PROCESSING OF RADIO-FREQUENCY ULTRASOUND SIGNALS FROM BLOOD AND CEREBRAL EMBOLI

Thesis submitted for the degree of
Doctor of Philosophy
at the University of Leicester

by

Joanne Cowe
Medical Physics Group
Department of Cardiovascular Sciences
University of Leicester

2006

ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106-1346

Processing of radio-frequency ultrasound signals from blood and cerebral emboli

Joanne Cowe

Abstract

Transcranial Doppler ultrasound (TCD) is commonly used to detect emboli in the cerebral circulation. However, current techniques to discriminate between signals from emboli and artefacts are subjective and ambiguous. The radio-frequency (RF) signal provides an extra dimension to the information available from conventional TCD systems, which may help to interpret complex events.

Artefacts generated by healthy volunteers and embolic signals recorded from a flow phantom were used to characterize the appearance of the two types of event. Characteristics of events, recorded during and immediately after carotid endarterectomy surgery, were compared to those from known sources. Additional information was provided by the RF signal, on events recorded during TCD monitoring, thus aiding classification. The RF signal may have a role as a gold standard for embolus detection.

Embolic signals appear as uniform and predictable shapes within the RF signal, enabling pattern recognition and image processing techniques to be used for their automated detection. Principal component analysis (PCA) has been used to characterize the typical variation in embolic signal shape, within the RF signal, using training sets of *in vitro* and *in vivo* data. PCA techniques were also utilised to discriminate between previously unseen embolic and artefact signals. The algorithms developed did not have the accuracy required for their use in a clinical setting but do have the potential to be developed further.

The RF signal allows us to track the motion of an embolus but the technique is limited by the narrow bandwidth of the transmitted pulse as this provides poor axial resolution. Mathematical modelling and *in vitro* experiments were carried out to assess the feasibility of using coded-excitation and pulse-compression within a TCD system to improve the bandwidth and hence the axial resolution. Quantitative and qualitative measures were used to compare the new processing method with that of a conventional, non-coded system. Results showed that it was feasible and beneficial to implement coded-excitation and pulse-compression in a TCD system.

Acknowledgements

The majority of this work was carried out as part of the "Ultrasonic Monitoring and Early Diagnosis of Stroke" (UMEDS) project. This was funded by the European Union (QLG1-CT-2002-01518).

I would like to thank WH Mess and APG Hoeks, University Hospital Maastricht, for their advice on the processing of TCD RF data and for providing the RF data used in the early part of this study.

I would also like to thank AR Naylor, NJM London, and R Sayers from the Vascular Surgery Group, Department of Cardiovascular Sciences, University of Leicester for facilitating the collection of the clinical data.

The staff in Medical Physics have, as a whole, been a great help. The work carried out in this project would have been impossible if it wasn't for John Gittins and Lingke Fan providing the equipment to enable the capture of the RF data. Steve Bentley and Harry Hall have done a good job in keeping me sane for the last few years. Also, many thanks go to Lorna Sweetman and John Gittins for putting up with many of my 'thick moments'!

I would also like to thank my husband, Glyn, for his patience, encouragement and suffering throughout the many highs and lows I have gone through during these past few years. His support has been fundamental to my completion of this project.

Last, but by no means least, many thanks to Prof. David Evans for being a terrific supervisor. His encouragement, teaching, and support have been second to none and his provision of countless opportunities for my personal and professional development has been very gratefully received.

Contents

List of Abbreviations

ACA	anterior cerebral artery
ACS	asymptomatic carotid stenosis
A/D	analogue-to-digital
AEDS	automatic embolus detection system
AF	atrial fibrillation
AM	amplitude modulation
ANN	artificial neural network
BDT	binary decision tree
BMF	blood mimicking fluid
CEA	carotid endarterectomy
CW	continuous wave
DAT	digital audio tape
EBR	embolus-to-blood ratio
FDA	(U.S.) Food and Drug Administration
FFT	fast Fourier transform
FM	frequency modulation
FP	false positive
FPGA	field programmable gate array
FWT	fast wavelet transform
ICA	internal carotid artery
IPEM	Institute of Physics and Engineering in Medicine
MBS	microbubble signals
MCA	middle cerebral artery
MEBR	measured embolus-to-blood ratio
MI	mechanical index
NHS	(UK) National Health Service
NLF	nonlinear forecasting
PC	principal component
PCA	principal component analysis
PM	phase modulation
PMD	power M-mode Doppler
PPV	positive predictive value
PRF	pulse repetition frequency
PW	pulsed wave
RF	radio-frequency
rms	root-mean-square
ROC	receiver operating characteristic
RSLL	range sidelobe level
SCS	symptomatic carotid stenosis
SNR	signal-to-noise ratio

SVM	support vector machines
TCD	transcranial Doppler
TMM	tissue mimicking material
TP	true positive
UMEDS	ultrasonic monitoring and early diagnosis of stroke
WDF	pseudo-Wigner distribution function

1 Introduction

1.1 Background

1.1.1 Stroke

Stroke is the third most common cause of death and the most common cause of adult disability in the UK. It is estimated to cost the NHS over £2.3 billion per year (Topic Working Group, 1999). There are two types of stroke – ischaemic and haemorrhagic. Ischaemic strokes are the most common, accounting for 86% of all incidents (Kelly-Hayes et al., 1995), and can be further divided into two sub-types – thrombotic and embolic. One in four strokes are embolic (Kelly-Hayes et al., 1995) and therefore research into ultrasonic detection of emboli has the potential to reduce death and disability, and to generate large financial savings.

An embolic stroke is caused when a clot, or other foreign body, travels through the bloodstream and subsequently becomes lodged in and obstructs a blood vessel in the brain. This restricts the flow of blood to the brain and results in almost immediate physical and neurological deficits. Emboli can originate from a number of sources such as the heart or plaques in arteries in the head or neck. For example, patients with atherosclerotic disease often develop plaques in the common carotid artery, primarily at the bifurcation into the internal and external carotid arteries. Embolic material can break off from the plaque and obstruct intracranial vessels resulting in stroke. Carotid endarterectomy (CEA) surgery is performed, in some cases, to remove the plaque. This can reduce the patient's risk of stroke but also carries a 2-4% risk of causing embolic stroke.

Transcranial Doppler ultrasound (TCD) is commonly used, in some centres, to detect embolic signals in the cerebral circulation. This can give an indication of the risk of stroke to the patient (Spencer, 1997). In Leicester, the routine TCD monitoring of patients undergoing CEA surgery has been associated with a 50% reduction in the incidence of intra-operative embolisation, a 60% sustained decline in the perioperative complication rate, and has abolished the incidence of post-operative thrombotic stroke (Naylor, 2000, Naylor et al., 2000). Although routine TCD monitoring has significantly reduced the risk to patients undergoing CEA, it is widely acknowledged that manual monitoring is tedious and identification of embolic signals challenging in a clinical environment. Such monitoring is, therefore, unlikely to be more widely used unless the techniques can be automated.

1.1.2 Transcranial Doppler Ultrasound

Transcranial Doppler ultrasound (TCD) most frequently uses pulsed wave (PW) Doppler. A simplified block diagram of such a system is illustrated in Figure 1-1.

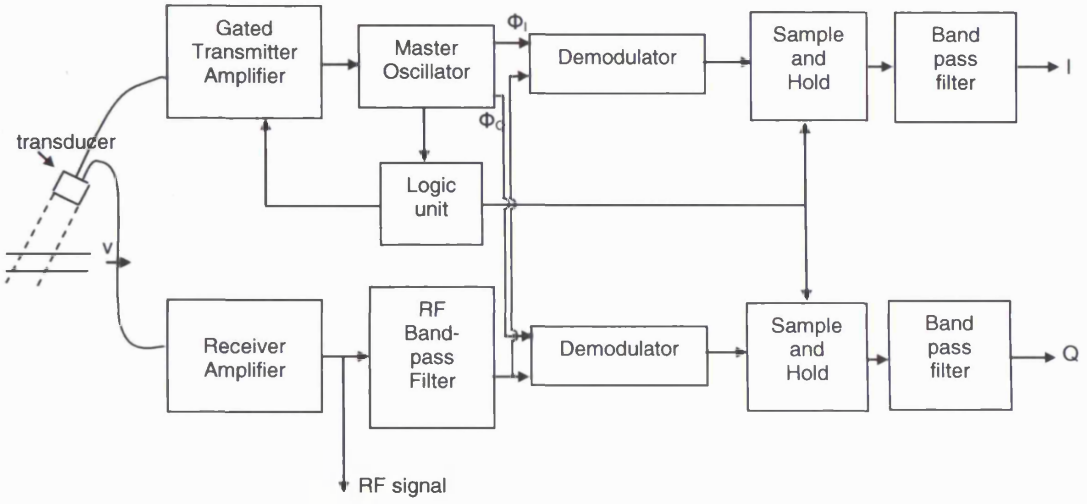

Figure 1-1 Simplified block diagram of PW Doppler system

When TCD is being used for embolus detection, the transducer is usually placed on the patient's temporal bone window. Its angle (and therefore the beam direction) is altered, together with the depth setting on the TCD machine, to obtain and monitor the blood flow signal from the middle cerebral artery (MCA). The signal received by the transducer is subject to a number of processing stages to obtain the in-phase and quadrature (I and Q) outputs shown in Figure 1-1. Performing a complex fast-Fourier transform (FFT) on the I and Q signals generates a directional Doppler signal which is usually illustrated as forward and reverse flow on a sonogram. A typical sonogram generated from an MCA blood flow signal is shown in Figure 1-2. Note that only forward flow is present, this is normal for the MCA.

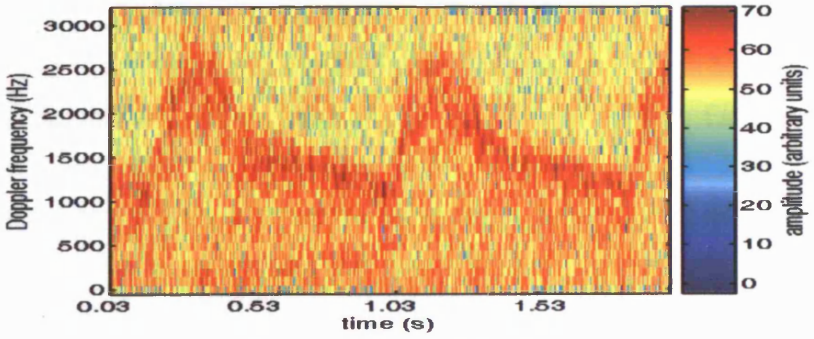

Figure 1-2 Typical sonogram from the MCA

An alternative to examining the Doppler signal is to extract the RF signal from the TCD system at the point shown in Figure 1-1. The RF signal will have undergone less processing than the Doppler signal thus reducing the likelihood of information loss and distortion. Additionally, any subsequent processing to be carried out on the RF signal can be made specific to the intended application e.g. embolus detection.

1.1.3 Detection of Emboli

1.1.3.1 History

Doppler ultrasound has been used to detect gaseous emboli in the blood since 1968 and, in the early years, was used as an early warning system in the study of decompression sickness (Spencer et al., 1990).

A few years later, Lubbers and Van Den Berg (1976) described the detection of emboli in blood flow using a continuous wave (CW) 2 MHz ultrasound system. They found that the dependence of scattering strength on embolus size matched with theory i.e. in general, the larger the embolus diameter, the larger the returned signal strength (apart from a dip at an approximate diameter of 600 μm in the case of red cell aggregate particles). They also found that solid particles returned a much smaller signal than gas bubbles. Finally, they concluded that the detection of emboli in blood was much easier using a transmit frequency of 2 MHz than using 9 MHz.

Padayachee et al. (1986) and Spencer et al. (1990) recognised the potential for Doppler ultrasound providing a technique to detect emboli in the cerebral circulation in a wide range of clinical situations. It was from this point on that the advancements in the techniques for the detection of emboli using TCD accelerated.

In 1994, Moehring and Klepper introduced the concept of the 'embolus-to-blood ratio' (EBR). This is the ratio of acoustic power backscattered from the embolus to that from the moving blood surrounding the embolus. They found that the EBR decreases with increasing transmit frequency and also with increasing vessel diameter (since the background blood signal then swamps the signal from the embolus). It was also noted that the EBR from a gaseous embolus is likely to be higher than that from a particulate embolus of the same size i.e. it is easier to detect gaseous emboli in the blood flow than particulate (probably more clinically significant) emboli.

Experts in the field of embolus detection by TCD met in 1998 and produced a consensus report (Ringelstein et al., 1998) describing limitations and problems of the technique, as well as providing guidelines for its use in clinical practice and scientific investigations. One conclusion from the consensus report was that, at that time, no automatic embolus detection system had the required sensitivity and specificity for clinical use.

More detailed reviews of the status of embolus detection were written by Evans (1999;2003), covering the scientific foundations of the subject, pitfalls of the technique and possible future directions of technical development. These also clarified that, in 1999, experienced human observers were still required for embolus detection/artefact rejection but, by 2003, a large amount of progress had been made in automated detection by Fan et al. (2001) and Cullinane et al. (2000).

As a result, to date, the majority of studies of embolic signals from TCD have used trained observers ('human experts') to monitor the number of emboli occurring over a given period. Observers are trained to classify occurrences of increased signal intensity, which are harmonic in tone, with a chirping or whistling quality (or with the characteristics of a 'snap', 'chirp' or 'moan' (Ackerstaff et al., 1995)), as emboli. This is in contrast to the noisy 'banging' quality of probe and patient movement artefacts (Spencer et al., 1990). These criteria are used to train new observers and are also included in automated systems, since these are designed to mimic the decision making of human experts. It should be noted, however, that these criteria cannot be objectively tested; therefore, on this basis, emboli and artefacts cannot be discriminated with 100% certainty.

The characteristics, described above, of signals from emboli and artefacts are heard in the demodulated Doppler signal and are accompanied by an increase in the intensity of the frequency domain signal displayed as a sonogram on TCD machines. An example of the appearance of an embolic event in the demodulated Doppler signal (in both the time and frequency domains) is shown in Figure 1-3 between approximately 0.97 and 1.02 s. The Doppler signal is essentially a one-dimensional signal indicating the presence of motion at a particular distance from the transducer.

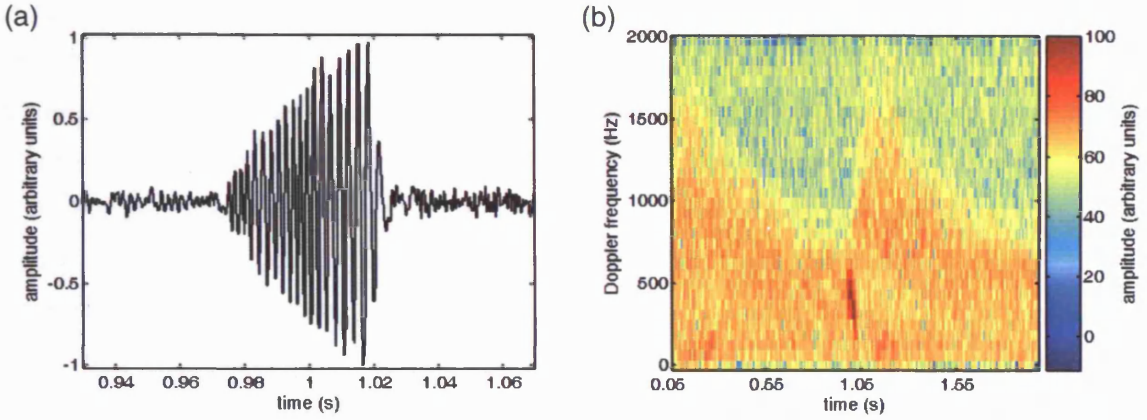

Figure 1-3 Example of an embolic event in the demodulated Doppler signal in (a) the time domain and (b) the frequency domain

There are a number of drawbacks associated with manual TCD monitoring techniques. Firstly, it is expensive as it requires an observer to monitor the outputs of the TCD system for the duration of the study, which may be several hours. However, this should be weighed against the cost saving achieved by reducing the risk of intra- and post-operative stroke and other complications. Additionally, the reliability of human expert decision making in a clinical setting is reduced in comparison to ideal laboratory conditions due to increased background noise and fatigue from long observation periods. In order for TCD monitoring for embolic signals to be a more widely used technique, these drawbacks should be reduced or overcome. An ideal solution would be to have a system that can automatically detect signals from emboli with at least as good reliability as the human expert. There have been many techniques described for the detection of embolic signals and some studies have been carried out on their automation. Some of these techniques will be discussed below.

1.1.3.2 Demodulated Doppler Signal Techniques

Single-Gate Techniques

Conventionally pulsed wave (PW) TCD systems extract the received ultrasound signal from one depth (or more accurately, one range/receive gate) and process this to display the sonogram and play the associated audio signal. This allows the observer to monitor the presence of motion (of blood and emboli) at that depth. There have been many techniques proposed to detect the presence of emboli in the circulation using a single gate and some of these will be discussed below.

In 1993, Markus et al. tested an off-line algorithm for the automatic detection of emboli on known events in an animal model and on events of a presumed nature in patients. The algorithm is based on recognising the characteristic narrow bell-shaped intensity increase in the frequency spectrum of an embolus in comparison to the bi-directional, wideband intensity increase of an artefact. In the patient studies, the algorithm achieved a sensitivity of 97.2% and specificity of 97%. The algorithm did not work well if the signals from an embolus were so large that receiver overload occurred, since the resulting signal closely resembled that of an artefact.

Three years later Van Zuilen et al. (1996) performed an *in vivo* study comparing the performance of three automatic embolus detection systems against the majority decision of a panel of four 'human experts'. Patients with symptomatic carotid atherosclerosis (i.e. producing particulate emboli with lower intensity signals) were used in this study. The same analysis is summarised by Mess et al. (1996b). The three automated systems used different algorithms for the differentiation of emboli and artefacts; the RB11 (EME) used the characteristic bell-shaped increase in the frequency spectrum caused by an embolus, the Pioneer Version 2.10 (EME) uses a simple algorithm based on the directionality of an intensity increase (with uni-directional more likely to be caused by an embolus and bi-directional more likely to be from an artefact), and the Embotec (STAC) used a trained neural network. The RB11 had a sensitivity of 70%, the Embotec 62%, and the Pioneer 44%. The agreement between human observers was high, emphasising the reproducibility of the technique (using TCD to detect emboli). The authors concluded that, based on the automated systems evaluated, assessment of signals by experienced observers was still necessary.

In 1998, Keunen et al. investigated the possibility of using nonlinear forecasting (NLF) to discriminate artefacts and emboli in the time-domain signal. This technique examines the periodicity of a signal and, based on this, tries to predict the next portion of the signal. If the prediction closely matches the signal that follows, a high prediction score is achieved. Keunen et al. found that the 20 emboli examined had high prediction scores whilst the 200 background signals and 100 artefacts did not. They concluded that emboli could be separated from artefacts and fluctuations in the background signal by NLF. However, the number of emboli used in their study is very small and no information is given regarding the magnitude of the embolic signals in comparison to the background signals and also regarding the duration of the embolic signals. It is intuitive that embolic signals with a low EBR have less chance of displaying periodicity than high intensity signals. Also, no information is provided on how the prediction score varies with the length of time window applied. This information would be necessary to fairly judge the validity of this technique.

Kemeny et al. (1999) evaluated a neural network based automatic embolus detection system (EMBotec V5.1 One). The same system was evaluated by Van Zuilen et al. (1996). They (Kemeny et al.) used two groups of patients in their study; one with arterial and one with cardiac embolic sources. They also produced artefact signals using normal volunteers. They recorded an artefact rejection rate of 85%, a sensitivity of 73.4% and a positive predictive value of 56.7%. As the authors acknowledged, neural networks suffer from several drawbacks; they can be over trained, they can repeat mistakes introduced by those doing the training, and they can make mistakes when encountering signals not encountered frequently during their training phase. The latter seems to have been the largest problem in the neural network system examined.

In 1999, Krongold et al. investigated the difference in performance of linear and quadratic time-frequency (Fourier transform) and time-scale (wavelet transform) embolus detectors when applied to the time-based signal. They found that time-scale detectors offered better performance than time-frequency detectors in terms of signal-to-noise ratio (SNR), probably due to the inverse relationship between Doppler frequency and time-duration of embolic signals. They also investigated the use of a time-scale-chirp detector and found this to have the overall best performance. However, the time-scale detectors could require ten times more computation than fast Fourier transform (FFT) based detectors and therefore the relatively small performance improvement has to be weighed against the additional implementation cost in consideration of its use in practice. Aydin et al. (1999) also investigated the use of the wavelet transform for describing embolic signals. They found that the temporal resolution and time localization of embolic signals was significantly better with a wavelet transform than a Fourier transform but that there were no significant differences between the EBRs found with the two techniques. Aydin et al. noted that the results obtained using the wavelet transform were heavily dependent on the choice of wavelet and that alternative wavelets may have resulted in improved EBRs. The wavelet chosen (Morlet wavelet) was different to that used by Krongold et al. (truncated bandpass Gaussian wavelet) and this may explain the difference in the findings of the two groups. Aydin et al. also acknowledged that the wavelet transform is computationally intensive but noted that a fast wavelet transform exists (analogous to the FFT) that is computationally efficient and can be implemented in real-time.

Markus et al. (1999) developed an automated embolus detection system using a 'frequency filtering' approach. This involved manual selection of a region of interest in the sonogram (i.e. a 2-D time-frequency region containing an embolus, artefact or Doppler speckle) and then applying a number of band-pass filters to the selected region of the

time-domain signal, using a sliding Hanning window. The frequency filtered signal shown in the paper is 2-D (time v's frequency) and highlights that the authors are effectively 'zooming-in' on the event selected in the sonogram, hence improving the time and frequency resolutions. The authors used the frequency filtering approach on a number of test events to extract characteristics of emboli, artefacts, and speckle and to set thresholds on parameters associated with these characteristics. Based on these characteristics, they developed an automated classification system. This used a 64-point FFT (32 bins of positive frequencies and 32 of negative, calculated every millisecond), instead of a number of band-pass filters, and calculated an amplitude threshold for each bin based on the surrounding points in the time and frequency domains. When the data in a given bin had amplitude above its threshold, it was subject to further analysis based on the event characteristics determined previously. When tested on an unseen data set (5-second data segments played in real-time), their system displayed a sensitivity of 91.1%, correctly rejected 200 episodes of artefact and misidentified 2 occurrences of speckle as emboli. It appears as though, by switching from 256 band-pass filters over a small frequency range in the data characterizing phase to a 64-point FFT over a larger frequency range in the automated detection phase, the frequency resolution will have significantly changed for the worse. Therefore, the accuracy of the parameters set in the characterization phase when applied to the new data set is questionable.

Cullinane et al. (2000) combined the parameters obtained from the first part of the study by Markus et al. (1999) with others from a nonlinear forecasting analysis to create a fuzzy logic based automated embolus detection system. This returns both an embolus probability and an artefact probability. Thresholds for each of these are set and used as a basis for rejecting artefacts/identifying emboli. Two patient groups were studied; in the carotid stenosis patient group, a sensitivity of 85.7% and specificity of 88.9% was achieved and, in the (post) carotid endarterectomy group, a sensitivity of 95.4% and specificity of 97.5% was achieved (an embolus probability threshold of 65% was used for both groups). In both groups, increasing the embolus probability threshold increased the specificity, but this resulted in a drop in sensitivity. The main reasons for emboli not being detected by the software were simultaneous occurrence with an artefact and low intensity signals.

The same automated system was further tested by Cullinane et al. (2002) on four patient groups: post carotid endarterectomy (postCEA), symptomatic carotid stenosis (SCS), asymptomatic carotid stenosis (ACS) and atrial fibrillation (AF). The system achieved the following sensitivities and specificities when compared to the decisions of a human expert: postCEA – sensitivity 95.8%, specificity 88.2%; SCS - sensitivity 98.4%, specificity 88.6%;

ACS - sensitivity 85.7%, specificity 13%; AF - sensitivity 54.8%, specificity 7%. It can be seen that the automated system did not perform that well in the ACS and AF groups when compared to the postCEA and SCS groups. It should be noted that, although the authors quote values for specificity, they have actually calculated the positive predictive value (PPV). This is quite low for the ACS and AF groups due to the low prevalence of embolic signals in these two groups compared with the postCEA and SCS groups. Again, the main reasons for emboli not being detected by the software were simultaneous occurrence with an artefact, low intensity signals and also possible overloading of the preamplifier resulting in bi-directional signals being misclassified as artefacts. Signals detected as emboli by the automated system but not by the human expert were re-examined by the expert. If these met the criteria as embolic signals they were then included in the analysis as true positives and were added to the human expert's count of embolic events. This is not consistent with protocols adopted at other research centres working in this field.

In 2003, Munts et al. assessed the offline performance accuracy of software (FS1, EME-Nicolet, Madison, WI, USA) incorporating the algorithms proposed by Markus et al. (1999) and Cullinane et al. (2000). It was tested on a 2-hour study tape made using data collected from 6 patients monitored post carotid endarterectomy surgery. Using a 7 dB intensity threshold the automated system achieved a sensitivity of 83%, with no false positives, when compared to the majority decision of a panel of five human experts ('gold standard' detected 93 emboli). When no intensity threshold was used, the sensitivity of the automated system was 73% and there were 3 false positives ('gold standard' detected 107 emboli). The authors noted that this software does not work well during CEA surgery when a proportion of emboli are likely to be gaseous, resulting in high intensity signals which can lead to receiver overload.

Markus and Reid (1999) showed that using the frequency filtering technique, discussed by Markus et al. (1999), can improve the EBR of an event by approximately 3 dB in the time-based signal. This has the potential for increasing the probability of detecting low intensity embolic signals originating from particulate emboli.

In 2000, Devuyst et al. proposed using a technique called matching pursuit to discriminate emboli from artefacts and to separate gaseous and particulate emboli using the time-based signal. This technique is similar to the wavelet transform and was shown to provide a very compact representation of the signals examined. Although the authors have managed to separate artefacts from emboli using the events' principal frequency and have separated gaseous from particulate emboli using parameters defining the regularity of the signal envelope and instantaneous frequency, they have only used 20 examples of each

of the three signal types. This is a really small sample size and their study should be expanded to provide a true validation of their technique. In addition, the authors acknowledge that the computational load of the matching pursuit technique is too large to make a real-time implementation of their technique practical.

A single gate, single channel (flow toward the transducer only), expert system based automated embolus detection system was described by Fan et al. (2001). Its performance was compared against three different 'gold standards' which adopt different combinations of decisions made by a panel of four 'human experts'. The expert system's performance was comparable to that of individual human experts when compared to the three gold standards, when signals with a signal-to-clutter ratio (SCR) greater than 5 dB were considered (achieving sensitivities between 86.7% and 95.8%, and specificities between 95.1% and 98.7%). When considering events with a SCR lower than 5 dB, the sensitivity of the automated system fell significantly. However, the inter-observer agreement on which events were embolic, for signals with an SCR lower than 5 dB, also dropped significantly, even under ideal laboratory conditions (as opposed to in a clinical setting). The implementation used was very computationally intensive therefore analysis was performed offline on short segments of data. Based on this system, Fan et al. (2004) designed and implemented a real-time automatic embolus detection system. Using the majority decision of a panel of three human experts as the gold standard, the authors tested the system on patients following CEA surgery. Using a measured embolus-to-blood ratio (MEBR) threshold of 7 dB, the automated system achieved a sensitivity of 93.6% and specificity of 99.2%; these were close to the results of the three human experts under ideal laboratory conditions.

In 2004, Marvasti et al. compared the performance of their new online automated embolus detection system, based on the fast wavelet transform (FWT), with the commercially available detection software FS-1 (Nicolet/EME Ltd.) that incorporates algorithms proposed by Markus et al. (1999) and Cullinane et al. (2000). In terms of sensitivity and specificity, the new system did not perform as well as FS-1 on a general data set. However, when presented with a data set composed of low intensity, short duration embolic signals the FWT system outperformed FS-1. The choice of wavelet is likely to affect the performance of the FWT-based system. Detection of emboli may be improved by development of a system that combines the better frequency resolution of the FFT and superior temporal resolution of the FWT.

A summary of the results, from the single-gate, automated embolus detection studies discussed above, is given in Table 1-1. It should be noted that not every performance

measure could be given for each study since this information was not always available in the corresponding reference. This highlights the difficulty in comparing the capability of multiple systems as not all authors use the same performance measures and, even when it appears that the same measure is being used, sometimes the method of calculation is different.

Study	Automated System	Patient Group	Amplitude threshold	Sensitivity (%)	Specificity (%)	PPV (%)	$\frac{FPs}{TPs+FPs}$ (%)
Markus et al., 1993	own	Symptomatic carotid artery stenosis & prosthetic mitral valves		97.2	97		
van Zuilen et al., 1996	RB11	Symptomatic carotid atherosclerosis		70			0 – 65
	Embotec			62			0 – 55
	Pioneer			44			0 – 87
Kemeny et al., 1999	Embotec	Arterial or cardiac embolic sources		73.4	85	56.7	
Markus et al., 1999	own	Symptomatic carotid artery stenosis	7 dB (intensity)	91.1	100 (artefacts) 98 (speckle)		
Cullinane et al., 2000	own	Carotid stenosis	7 dB (intensity)	85.7	88.9		
		Post CEA		95.4	97.5		
Cullinane et al., 2002	FS1	Post CEA	7 dB (intensity)	95.8	88.2*	88.2*	
		Symptomatic carotid stenosis		98.4	88.6*	88.6*	
		Asymptomatic carotid stenosis		85.7	13*	13*	
		Atrial fibrillation		54.8	7*	7*	
Munts et al., 2003	FS1	Post CEA	7 dB	83		100	
			none	73		96.3	
Fan et al., 2001	own	CEA and post-CEA	All SCRs	63.1 – 74.5	99.93 – 99.98		
			5 dB (SCR)	86.7 – 89.9	95.3 – 98.7		
			7 dB (SCR)	93 – 95.8	95.1 – 97.7		
Fan et al., 2004	own	Post CEA	7 dB (MEBR)	93.6	99.2		
Marvasti et al., 2004	own	Symptomatic carotid artery stenosis and post CEA		51 – 97.7	29.2 – 92.3		

** Authors of paper have quoted these figures as specificity but they have actually calculated the PPV. This is possibly also the case for the values given in Cullinane et al., 2000, but not enough information is present to conclude one way or the other*

Table 1-1 Summary of the performance of the single-gate, automated embolus detection systems discussed in this chapter

Multi-Gate Techniques

Many people believe that if more than one depth were to be monitored simultaneously, this would make the discrimination between emboli and artefacts more reliable. These are known as multi-gate techniques. Several methods employing multiple gates have been proposed including the dual-gate method, the arbitrary sample volume method, the four-gated method, and power M-mode Doppler (PMD). Each of these techniques will be discussed below.

Dual-Gate Method

One example of a multi-gate technique is the dual-gate method (also known as the coincidence method). This involves monitoring two gates simultaneously, both including the vessel of interest. Theoretically, if an embolus were present there would be an increase in intensity appearing sequentially in the gates within an expected time window defined by the blood velocity and sample volume separation. Alternatively, if the increase in intensity was present in both gates at the same time or only in one gate then this is characteristic of an artefact.

Smith et al. (1996) investigated the possibility of using a dual-gate system for emboli and artefact differentiation. They examined the time delay between two gates from embolic signals from 12 patients undergoing CEA and from artefacts from the same patients and from 2 healthy volunteers. The delays found for the two groups of event (emboli and artefacts) were separable, leading to the conclusion that this is a viable way for differentiating the two. The start time of an event was defined as the point where the signal amplitude exceeded the background Doppler signal amplitude by approximately 100%. This criterion may be difficult to achieve for low intensity events e.g. signals from small particulate emboli. It is noted that 5 out of the 138 emboli studied did not appear in the second gate. It is hypothesized that these may have followed a branch between the two gates. Smith et al. (1997) subsequently performed computer simulations to try and account for 'curious observations' during their dual-gate study. They found that the geometry of the middle cerebral artery (MCA), in relation to the insonating ultrasound beam, may account for discrepancies in embolus velocity when comparing a time-of-flight calculation with the Doppler-derived velocities, discrepancies in durations of embolic signals between the two gates, some emboli 'disappearing' before the second gate, and some emboli not appearing in the first gate but appearing in the second. They point out that, to maximise the chances of having emboli and artefact time delays that are separable, the maximum possible gate separation should be obtained since the apparent separation between gates is likely to be smaller than intuitively believed, giving a shorter

time delay. However, the maximum gate separation achievable is dependent on the MCA anatomy of the individual being examined.

Implementations of the dual-gate technique have been described by Mess et al. (1996a;1997;1999;2000). In 1996 they described preliminary results of measured time differences of the appearance of emboli in two sample volumes. This showed a large variation of between –3 and 31 ms. In 1999, they performed simulations to study the effects of MCA anatomy on time differences measured in a dual-gate system. They found that there were large variations in the time difference as a result of small changes to the insonation angle or sample volume depth. This variation was reduced by reducing the receive gate separation and using a shorter receive gate in the distal sample volume. They emphasise that TCD settings, which reduce variation in the time difference (i.e. ensuring the time differences are in a more predictable range), should be used if a dual-gate system is to be automated. Using the TCD settings derived from these simulations, Mess et al. evaluated a dual-gate automated embolus detection system in 2000. They used the pseudo-Wigner distribution function (WDF) for the measurements of the time difference as this is superior to the FFT in terms of temporal resolution. They concluded that the technique had potential but that a large number of emboli (28%) were missed by the automated system as they did not appear in both sample volumes. Of those appearing in both sample volumes, 87% of emboli were correctly identified by the automated system. However, when taking into account those emboli that only appeared in one of the sample volumes, the sensitivity was reduced to 66%. They also found that, with the exception of diathermy, the rejection of artefacts was very good.

Georgiadis et al. (1996) evaluated the dual gate technique *in vitro* and *in vivo* using a multigate probe. The time of onset of events was recorded manually from a DAT recorder display based on the appearance of the signal in the frequency domain. Overall, from their *in vivo* experiments, they recorded a sensitivity of 98.1% and specificity of 98.8%. The main reason for incorrect classification of emboli was their appearance in only one gate. A sample volume separation of 10 mm was used for all 19 subjects in their study. The study relating MCA anatomy to dual gate techniques by Mess et al. (1999), however, suggests that such a large separation would not be suitable for many people.

Georgiadis et al. (1998) tested automatic embolus detection software (TCD-8, version 8.00 T), which utilises the dual-gate approach, on patients undergoing coronary artery bypass surgery and cardiac valve replacement. When comparing results with a human observer gold standard they achieved a sensitivity of 64% and specificity of 78.5%. They believe that this was a satisfactory specificity. However, with 35 522 of the 44 933 high

intensity events being artefacts, a specificity of 78.5% resulted in 7643 false positives. This is very close to the 9411 actual embolic events recorded, making the reliability quite poor. They also encountered the problem of some emboli not appearing in both gates.

Droste et al. (1997) also tested automatic embolus detection software (TCD-8, version 8.00 K) employing the dual-gate technique. They considered normal subjects and patients with prosthetic heart valves. They reported a sensitivity of 74.3% and specificity of 59.9%.

Molloy and Markus (1996) examined the dual-gate technique, using a multidepth probe, *in vitro* and *in vivo*. Emboli from two groups of patients were examined – those with carotid artery stenosis and those with mechanical prosthetic heart valves. The time delay between an event appearing in each of the channels was measured from the time-domain signal. They achieved a specificity of 99.0%, a sensitivity of 75.2% for the heart valve patients and sensitivity of 92.6% for the carotid patients. The sensitivity for the valve patient group was reduced due to a large number of emboli not appearing in both gates in one patient. In the carotid patient group, the signals from emboli were, in general, much smaller than those from the heart valve group. This made it difficult, in some cases, to identify the embolus in the time-domain signal in both gates. This resulted in a decrease in the sensitivity in this group. This may prove to be a difficulty when automating this technique for use with carotid patients.

Devuyst et al. (2001) combined a dual-gate technique with the use of wavelets to determine the time delay of signals appearing in the two gates and to determine the main frequency of the signal in each of the gates. If possible, the frequency information was used to distinguish between artefacts (<250 Hz) and emboli (>500 Hz). For the region between 250 and 500 Hz, the time delay was used to make the decision (<4 ms is considered to be an artefact). A specificity of 98% and sensitivity of 97% were achieved. A combination of known artefacts, gaseous emboli and solid emboli were used in this study. This study showed that the combination of decision-making criteria might be more reliable than simply using the dual-gate time delay.

A dual-gate system was employed by Cullinane and Markus (2001) in their assessment of the use of a 1 MHz transducer for improved embolic signal detection. They found that the 1 MHz transducer improved embolic signal intensity compared with that from a 2 MHz transducer. However, the blood flow signal quality and the ease of finding the blood flow signal were both reduced with the 1 MHz transducer. The benefits of improved embolic signals are therefore outweighed when routine blood flow monitoring has to be performed

simultaneously. The authors suggest that a 1.5 MHz transducer may be a compromise but no study has been done to verify this.

Darbellay et al. (2004) combined a dual-gate TCD system with a binary decision tree (BDT) to discriminate emboli from artefacts. The criteria used to differentiate the two types of event were the frequency at the event's peak intensity, the time delay of the appearance of the events in the two gates, and the duration of the event. This achieved a sensitivity of 94% and specificity of 97% when compared against total agreement of a panel of three human experts. No rationale is offered giving possible reasons for the emboli not detected. One can only assume that they were due to typical reasons as offered by other authors such as emboli not appearing in both gates, receiver overload, low intensity signals etc. The authors also used three different classification techniques for separating gaseous from solid emboli: BDT, artificial neural networks (ANN), and support vector machines (SVM). BDT had a sensitivity and specificity of 81%, ANN had a sensitivity and specificity of 84%, and SVM had a sensitivity and specificity of 86%. Even although BDT is by far the least complex implementation, it demonstrated comparable results with the other two techniques. The authors state that, for the data collection for both stages of their study, each MCA (i.e. both sides of each of the 45 patients) was insonated at depths of 50 and 60 mm. It is very unusual that this many people in one study would have an anatomy allowing such a wide separation of gates and both gates to be placed so deep.

In conclusion, a number of people have considered the use of dual-gate techniques, with varying degrees of success. The success of using this technique is mainly dependent on the separation of the two gates, the MCA anatomy, and the nature of the emboli being examined. The majority of people have reported occurrences of emboli appearing in one gate but not the other, thus being incorrectly classified as artefacts. In order for such a technique to be automated and for users to have confidence in its results, it may be that it needs to be used in combination with additional discriminating criteria.

Arbitrary Sample Volume Method

As an alternative to placing the two gates over the vessel of interest, some authors have suggested placing one on the vessel and one outwith the vessel, close to the temporal bone. In this case, an artefact should appear in both gates almost simultaneously or only in the gate outwith the vessel, whereas an embolus will only appear in the gate placed over the vessel. The second gate in this technique has been referred to as a 'reference gate' or an 'arbitrary sample volume'.

Brucher et al. (1997) first proposed the idea of a 'reference gate'. They suggested that using their recommendations (including using the 'reference gate' technique), the sensitivity of automated embolus detection could be improved from 70-80% to more than 94%. However, there is no other data in this paper that backs up this statement.

The 'arbitrary sample volume' method is further explored by Georgiadis et al. (2000b). They used 250 artefacts and 250 emboli collected *in vivo* to determine a threshold of intensity difference between the two sample volumes at the time of an event, to best separate the two groups. This threshold was then applied to the intensities in the two gates of a further 1858 emboli and 1958 artefacts collected *in vivo* from a variety of patient groups. This resulted in a sensitivity of 96% and specificity of 98%. Georgiadis et al. (2000a) clinically tested the dual-gate and arbitrary sample volume methods on the same set of data. They found similar performance in the two techniques with a sensitivity and specificity of 90.7% and 91.3% respectively for the dual-gate approach and 88% and 91.9% for the arbitrary sample volume method. They suggest that a combination of the two techniques, i.e. having three sample volumes, two over the vessel of interest and one outwith the vessel, may provide an improvement on the results from the individual methods, but they have done no analysis to prove or disprove this.

Brucher and Russell (2002) used the 'reference gate' technique together with three other criteria to automatically differentiate emboli and artefacts. The other criteria were 'quarter Doppler shift', 'maximum duration limit', and 'bidirectional enhancement'. The first makes use of the Doppler shift characteristics from a scatterer when insonated simultaneously with two frequencies: 2 and 2.5 MHz. The authors state that the Doppler frequency shift caused by an embolus will be one quarter greater when insonated with 2.5 MHz than with 2 MHz, whereas an artefact will not show the same quarter shift. The second criterion is a threshold on the maximum time it would take an embolus to travel through the sample volume under observation. The third criterion is that artefacts are likely to appear in both the forward and reverse channels as opposed to emboli only appearing in one direction. They tested their system *in vitro* and *in vivo* and achieved a sensitivity of 98.6% and specificity of 98.9% *in vivo*. Russell and Brucher (2002) used the same system to automatically discriminate between solid and gaseous emboli. Since discrimination between different embolus compositions is not explored in this project, this paper will not be discussed further. However, Evans and Gittins (2005) concluded that there are significant differences in the beam shapes at 2 MHz and 2.5 MHz from the transducer used by Russell and Brucher. These differences will result in measurements of the ratios of the MEBRs at the two frequencies being unreliable making a technique to distinguish between gaseous and particulate emboli, based on this measurement, fallible. Differences

in MEBR using the two frequencies may also cause problems in the Brucher and Russell (2002) system for discriminating artefacts and emboli. For example, if a small particulate embolus is present, the MEBR at one frequency may be large enough to be above the 'high energy signal' threshold but the MEBR at the second frequency may not. In this case, it would not be possible to determine if a 'quarter Doppler shift' were present, hence leading to a possible misclassification as an artefact.

In summary, the 'arbitrary sample volume' method for embolus/artefact differentiation has shown to have similar performance to the dual-gate approach. It does not have the limitation of having to place two sample volumes over the vessel of interest and will overcome the difficulty found in the dual-gate method when an embolus only appeared in one of the gates.

Four-Gated Method

A four-gated automatic embolus detection system (TC4040, Nicolet-EME) is evaluated *in vivo* by Droste et al. (1999). The actual algorithms used for embolus detection/artefact rejection are not discussed and therefore it is impossible to comment on the validity of the technique. It is stated that some emboli were detected by the software in only one channel/depth gate and others in two, three or four channels. This implies that the detection of an event with an appropriate time delay between adjacent gates was not a criterion for classification as an embolus. However, in their discussion they indicate that an embolus signal with a low intensity increase in one gate may produce an increase lower than the threshold in another gate and would therefore escape automatic detection. The use of four gates requires careful selection of the monitoring depths to ensure all gates are on the same vessel and have adequate spacing so that time delays of intensity increases between gates allow emboli and artefacts to be discriminated. Overall they report a sensitivity of 86.5% and an artefact rejection rate of 62%. These results are not as good as in other techniques discussed.

Power M-Mode Doppler

Moehring and Spencer (2002) developed a new TCD modality, power M-mode Doppler (PMD), to address the problems of single-gate TCD systems only providing partial evidence of the presence of emboli, and dual-gate systems increasing the complexity of detecting emboli since the user has to simultaneously monitor two sonograms and/or time-based signals. PMD displays the signal power (colour-coded for flow direction with respect to the transducer) from 33 gates, each separated by 2 mm, in an M-mode format. It also shows the sonogram from a user-selected depth. The presence of an embolus is evident from a high power sloping track in the PMD display. This enables the path of an

embolus to be more easily traced. It also highlights the presence of emboli at depths other than that being monitored on the sonogram. The subjects used in this paper are all from groups likely to have gaseous, rather than particulate, emboli. It would be interesting to see how well the 'emboli tracks' can be distinguished from normal intensity variations in the background flow in the PMD display when particulate emboli are present, since these are known to result in lower intensity signals than gaseous emboli. The average Doppler shift frequency is also calculated at each gate depth. The clutter filtering process assigns Doppler shift frequencies lower than 200 Hz (corresponding to a velocity of 7 cm/s) zero power in the PMD display (i.e. colours them black). The authors found that all artefacts have mean velocities lower than this clutter cut-off and do so uniformly across depth. This results in a black stripe down the PMD display in the presence of an artefact. The authors intend using the characteristics described to develop an automated system for embolus detection using the PMD technology. However, care should be taken that slow moving emboli (with a Doppler shift frequency of 200 Hz) are not missed by being assigned zero power by the clutter filter. There is also the potential for emboli to be masked by zero power black stripes if an embolus and artefact occur simultaneously.

Saqqur et al. (2004) compared spectral TCD with PMD for the detection of microbubble signals (MBS). They found that more MBS were detected using PMD since bubbles in both the anterior cerebral artery (ACA) and the MCA could be detected.

Summary

A summary of the results from the multi-gate, automated embolus detection studies discussed above is given in Table 1-2. As with the single-gate system results, care should be taken when interpreting the table and comparing the performance of multiple systems, since methods of calculating the various parameters can vary. For example, some of the dual-gate system studies quote sensitivity based on only those emboli appearing in both gates whereas others regard the appearance of an embolus in only one gate as a failure to detect that as an embolus. The former approach will undoubtedly enhance the apparent performance of the system. If a parameter value is not shown in the table, it was not available in the corresponding reference.

Study	Automated System	Patient Group	Amplitude threshold	Sensitivity (%)	Specificity (%)	PPV (%)
Mess et al. (2000)	own	CEA & heart surgery		66	91	
Georgiadis et al. (1996)		Cardiac and carotid embolic sources		98.1	98.8	
Georgiadis et al. (1998)	TCD-8 (version 8.00 T)	Cardiac surgery	9 dB (intensity)	64	78.5	
Droste et al. (1997)	TCD-8 (version 8.00 K)	Prosthetic heart valves	6 dB	74.3	59.9	
Molloy and Markus (1996)		Carotid artery stenosis		92.6	99	
		Mechanical prosthetic heart valves		75.2	99	
Devuyst et al. (2001)	own	PFO & symptomatic ICA stenosis	9 dB (intensity)	97	98	99
Darbellay et al. (2004)	TCD-8 (version 8.00 K)	PFO & carotid stenosis	9 dB	94	97	
Brucher et al. (1997)	own	Artificial heart valves		>94		
Georgiadis et al. (2000b)	own	Cardiac or arterial embolic source		96	98	
Brucher and Russell (2002)	EmboDop	Heart valve & carotid stenosis	5 dB	98.6	98.9	
Droste et al. (1999)	TC4040	Mechanical heart valves, carotid artery disease	5 dB	86.5	62	

Table 1-2 Summary of the performance of the multi-gate automated embolus detection systems discussed in this chapter

1.1.3.3 RF Signal Techniques

As can be seen from above, a number of techniques for embolus detection using the demodulated Doppler signal from TCD systems have been proposed. There has been varied success when these systems have been tested *in vitro* and *in vivo*. An alternative approach is to take the raw backscattered radio-frequency (RF) signal i.e. before any processing has been performed. This allows frequencies other than those near the carrier frequency to be examined, a large number of depths compared to those contained within a conventional receive gate to be investigated, and the opportunity to apply signal processing techniques more appropriate for the detection of embolic signals. Several RF signal techniques have been proposed recently and these are discussed below.

Low frequency transducers (130 kHz and 250 kHz) were used by Palanchon et al. (2001) for the detection and characterization of circulating gaseous emboli. When the backscattered signal from an embolus contains second (and higher) harmonics, the embolus can be classified as gaseous since particulate emboli only scatter linearly. The transmission frequency should be close to the resonant frequency of the bubble to generate the harmonics in the backscattered signal. The resonance frequency of a bubble decreases with increasing bubble diameter (e.g. bubbles with diameters of 10 μm and

200 μm have resonant frequencies of 616 kHz and 32 kHz respectively). Therefore, using the frequencies 130 kHz and 250 kHz, the authors could size gaseous emboli between 19 and 62 μm. Larger bubbles behaved linearly and hence could not be discriminated from particulate emboli. The authors do not mention in this paper the discrimination of emboli from artefacts.

Two years later, Palanchon et al. (2003) expanded upon the ideas they presented in 2001 to utilize the subharmonic and ultraharmonic emissions in addition to the second (or higher) harmonic to detect and characterize gaseous emboli. They found that whilst bubbles close to the resonant size generate second or higher harmonic components, bubbles close to twice the resonant size generate a subharmonic or ultraharmonic component. Using a scanning frequency of 130 kHz they could classify bubbles smaller than 110 μm as gaseous emboli. The range of detectable bubble diameters is dependent on the magnitude of the transmitted acoustic pressure; the larger the acoustic pressure, the larger the bubble size range, but the lower the sensitivity of size estimation. By varying the transmit frequency and acoustic pressure, the size estimation could be improved and the range of bubble diameters that could be characterized, increased.

Palanchon et al. (2004) have also developed a dual-crystal transducer to expand on their findings of 2001 and 2003. It has a narrowband transmit capability and transmits at approximately 540 kHz. The receive part of the transducer is wideband and has a relatively constant detection level between 50 kHz and 5 MHz. Their findings using this transducer clarified their previous results on the detection and characterization of gaseous emboli using harmonic, subharmonic and ultraharmonic components of the received RF signal. They also suggest that, to increase the range of bubble sizes that can be characterized, a multifrequency transducer should be developed. However, using such a low transmit frequency is likely to cause problems *in vivo* when trying to locate and maintain a cerebral blood flow signal, as demonstrated by Cullinane and Markus (2001).

A neural network approach was discussed by El-Brawany and Nassiri in (2002). The first stage uses a predictive model of the blood backscatter signal. It compares its prediction of what should come next with the actual backscatter signal input and produces a prediction error. This error would be 0 if it was certain that the signal was from blood and 1 if it was certain that it was not. A cross-coherence spectrum is then calculated from the prediction error signal and the excitation pulse. The more similar the two signals, the higher the coherence value will be. The neural network then takes the backscattered signal, the prediction error and the coherence signals as inputs, and outputs a decision as to whether

an embolus was present or not. The neural network was trained with examples from blood-mimicking fluid, solid and gaseous emboli and wall motion artefacts. This system has only been tested on *in vitro* data and it is questionable as to how well it is likely to perform *in vivo*. For instance, an example of a wall motion artefact is shown to have lower amplitude in the coherence signal than a nearby embolic signal and, on this basis, the authors argue that this would enable the discrimination between artefacts and emboli. However the artefact shown has similar amplitude and appearance in the coherence signal to an embolus in another figure and therefore the discrimination between the two will not be as clear cut as the authors suggest. Also, the authors have not considered the appearance of other forms of artefact e.g. patient and probe movement. In addition, the authors have tried to eliminate the appearance of wall motion artefacts by adjusting the position of the start and end gates during acquisition. This will not be practicable in a clinical situation.

In 2002, Mess et al. showed that the clutter-filtered RF signal could be displayed as a 2-D amplitude plot of time versus depth. This allowed the depth at which the embolus appeared and disappeared to be determined, the path of the embolus to be visualized, and showed that a proportion of the gaseous emboli (but none of the particulate emboli) changed direction. The authors note that the direction and velocity of all gaseous and 71% of the particulate emboli could be determined from the amplitude plots but that the SNR of the remaining particulate emboli was too low to allow this. This paper offered an exciting new way to visualize emboli and leads the way for further work in this direction.

1.1.3.4 Summary

Many techniques have been proposed for the detection of emboli and their discrimination from the background blood flow and artefacts. The majority of these techniques have processed the demodulated Doppler signal to obtain characteristics of the various types of signal in the time and/or frequency domain. The automation of these classifiers has had varying degrees of success with the most promising described by Fan et al. (2001, 2004). However, there are some inherent difficulties with using the demodulated Doppler signal since the majority of TCD systems have been designed for the detection of blood flow and, therefore, the system parameters are not likely to be optimal for the detection of emboli. This may lead to signal overload when large emboli occur, resulting in possible misclassification as artefacts. Also, small emboli may be missed if an appropriate selection of window and window-overlap is not made (Markus, 1995; Aydin and Markus, 2000).

Recently, the processing of the raw RF signal has been introduced as a possible new technique for the detection of emboli. To our knowledge, only El-Brawany et al. (2002) have tried automating a system that uses the RF signal. However, the system they propose is likely to need considerable work before it can be tested on *in vivo* data.

Comparing studies by different authors is difficult due to different methods of calculating EBR, sensitivity and specificity, and the different thresholds (Markus and Molloy, 1997) and different 'gold standards' used. Ideally, authors should state the methods used and how parameters such as these have been calculated. Also, it would be ideal to have a gold standard that did not rely on the subjectivity of human experts.

1.2 Aims

The aim of this project was to identify new processing methods for TCD signals from blood and emboli. In particular, the RF signal was to be used instead of the conventional subject of processing, the demodulated Doppler signal. The goal was to investigate whether or not the RF signal would be a suitable means for discriminating emboli and artefacts, and to determine the advantages and disadvantages of using the RF signal instead of the demodulated Doppler signal.

2 Data Acquisition and Signal Processing

This chapter describes the data acquisition and signal processing techniques which were used throughout this project. Any additional techniques employed, specific to a particular area of the project, are described in the relevant chapter.

2.1 Data Acquisition

The RF, system clock, and receive (or transmit) gate signals were made externally available from a TCD system. These three signals were connected to a PC-housed 14-bit, 100 MS/s A/D converter (CompuScope 14100, Gage Applied Technologies, Inc., Lachine, QC, Canada). A trigger signal was also input to the A/D converter. This was either a manual input from the user or was provided by an automatic embolus detection system (AEDS) (Fan et al., 2004). A schematic of the acquisition system is shown in Figure 2-1. Variations from this schematic were used at different stages in this project. Any such variations are described in the relevant chapter.

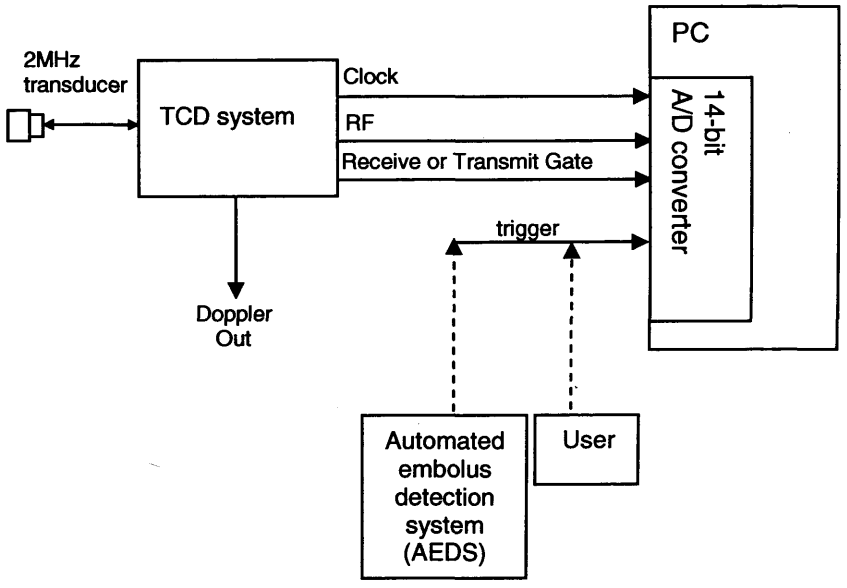

Figure 2-1 Schematic of RF acquisition system

Recordings were made from normal subjects, from patients undergoing CEA surgery and *in vitro* from a flow phantom. Where a human subject was used, a 2 MHz transducer was fixed over one of the subject's temporal bone windows so as to insonate the middle cerebral artery (MCA). The RF, receive gate, and system clock signals were continuously

input to the A/D converter and the data were stored in an on-board buffer. The TCD system's clock rate determined the sampling frequency, f_s, used by the A/D converter when digitising the RF and receive gate signals.

The AEDS provided a trigger signal on the occurrence of a suspected embolic event. A trigger could also be forced by the user so that the RF and receive gate data from artefacts, such as patient/probe movement, diathermy (electrocautery) etc., or 'embolic' signals from the flow phantom (when the AEDS was not being used) could be acquired. Following receipt of a trigger, all available RF and receive gate data, 1 s either side of the triggered event, were transferred from the A/D card to hard disk. An example of a small portion of such an acquisition is shown in Figure 2-2.

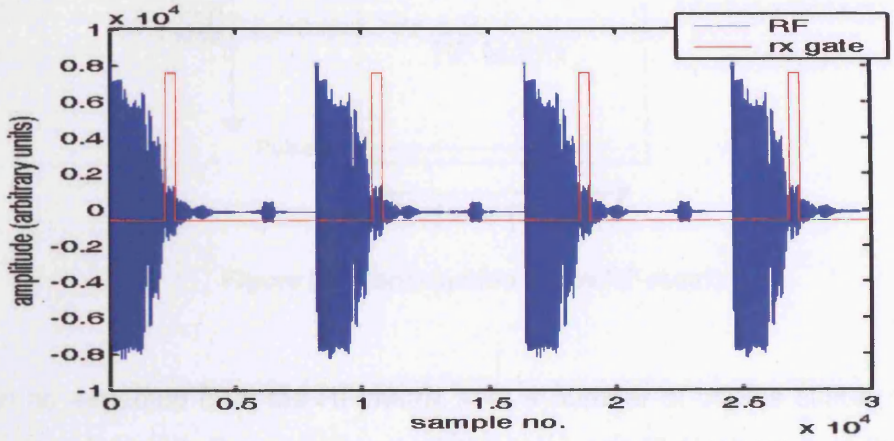

Figure 2-2 Portion of RF and receive gate signals acquisition

2.2 Signal Processing

The receive gate signal and depth setting on the TCD machine (or the transmit gate, where available) were used to determine the start of each transmit burst. The RF data were rearranged into a matrix with the received signal from each transmit burst forming one row. It is this RF-matrix that is the subject of processing for a large proportion of this project. The construction of the RF-matrix is illustrated in Figure 2-3. The abscissa of the RF-matrix is also referred to as fast time or depth and the ordinate as slow time.

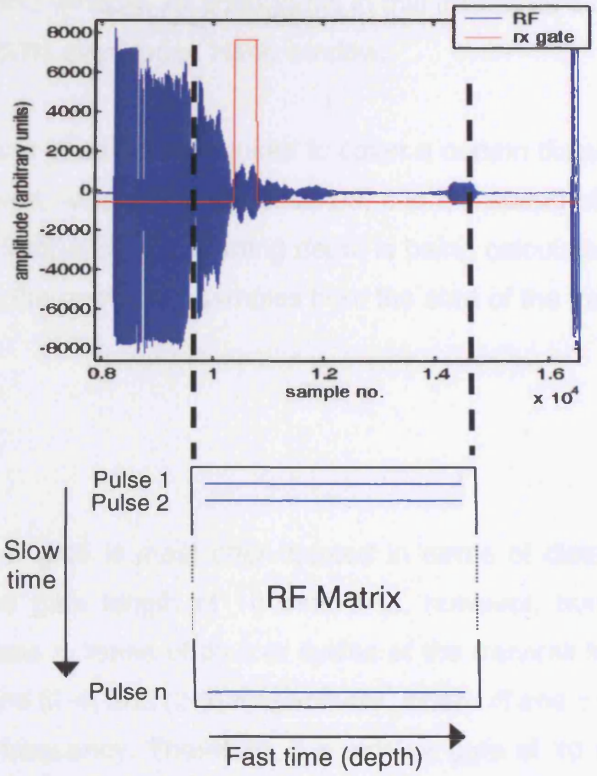

Figure 2-3 Construction of the RF-matrix

Data can be extracted from the RF-matrix from a number of depths simultaneously. The minimum usable depth, R_{min}, can be calculated using eqn (2-1) where τ_{tx} is the transmitted pulse duration (s) and c is the speed of sound in soft tissue (1540 m s^{-1}). The maximum detectable depth, R_{max}, can be determined from c and the pulse repetition frequency (PRF) using eqn (2-2).

$$R_{\min} = \frac{c\tau_{tx}}{2} \tag{2-1}$$

$$R_{\max} = \frac{c}{2 \cdot PRF} \tag{2-2}$$

The demodulated Doppler signal can be generated by extracting data from a receive gate, of a given starting depth and width, from every row in the RF-matrix and then demodulating each extracted gate with a 2 MHz sinusoidal carrier wave. Subsequently, each demodulated gate is integrated to give one value per pulse and the resulting signal is high-pass filtered to remove any low frequency clutter e.g. from wall motion. The sonogram can then be created from the demodulated Doppler signal using the MATLAB

command 'specgram'. Specgram computes the windowed discrete-time Fourier transform of a signal using a sliding window. All sonograms in this thesis have been produced using a 64-point FFT with a 50% overlapped Hann window.

To determine the number of samples required to cover a certain distance in the RF signal, eqn (2-3) should be used, where R is the distance, c is the speed of sound in soft tissue and f_s is the sampling frequency. If a starting depth is being calculated, it should be noted that eqn (2-3) provides the number of samples from the start of the transmit burst.

$$samples = \frac{2R}{c} \cdot f_s \tag{2-3}$$

The length of a receive gate is most often quoted in terms of distance e.g. many TCD systems use a receive gate length of 10 mm. It is, however, sometimes desirable to consider what this means in terms of time or cycles of the transmit frequency. These can be calculated using eqns (2-4) and (2-5) respectively, where R and c are as defined above and f_{tx} is the transmit frequency. Therefore, if a receive gate of 10 mm was used with a transmit frequency of 2 MHz, this may also be quoted as a receive gate of 13 μs or 26 cycles.

$$t = \frac{2R}{c} \tag{2-4}$$

$$cycles = \frac{2R}{c} \cdot f_{tx} \tag{2-5}$$

3 Characterizing the Appearance of Emboli and Artefacts in the RF Signal

3.1 Introduction

It was discussed in Chapter 1 how the criteria used to differentiate emboli and artefacts in the demodulated Doppler signal cannot be objectively tested and, as a result, when using this signal the two event types cannot be discriminated with 100% certainty. Additionally, since the Doppler signal is essentially one-dimensional, the information available about an event is limited. It has been shown by Mess et al. (2002) that, by acquiring the unprocessed radio-frequency (RF) ultrasound signals from a TCD machine, the trajectory of emboli can be visualized in a two-dimensional space of time and depth. By adding an extra dimension, the interpretation of an event potentially becomes much easier.

In a discussion on using a "human gold standard" for assessing the performance of automated embolic detection systems, Mess (2003) suggested that the detection of a moving intensity increase in the RF based amplitude plot 'could possibly perform as a better gold standard than the human observer expert panel'. It was also noted, however, that many of the particulate emboli could not be analysed in terms of velocity and flow direction in the RF plots. It will be shown in this chapter that the information contained in the two-dimensional RF signal may be adequate for consideration of its use as a more appropriate gold standard than that currently used. More specifically, it will be demonstrated that the RF signal may provide additional information on embolic events recorded during TCD monitoring with respect to their discrimination from artefacts, the number of events occurring during an emboli shower (a large number of emboli over a short period of time), and the depth at which events occur. A simple method for the enhancement of the signal-to-noise ratio (SNR) of an event, to aid its detection, will also be discussed.

3.2 Methods

This section is split into three areas; the data acquisition system, signal processing, and data collection. Additional details regarding the first two areas can be found in Chapter 2.

3.2.1 Data Acquisition System

A standard TCD machine (SciMed TC22, Bristol, UK) was modified to make the RF, system clock, and receive gate signals externally available. These three signals were connected, together with a trigger signal from an automatic embolus detection system (AEDS) (Fan et al., 2004), to the PC-housed 14-bit, A/D converter. The schematic of the acquisition system is shown in Figure 3-1.

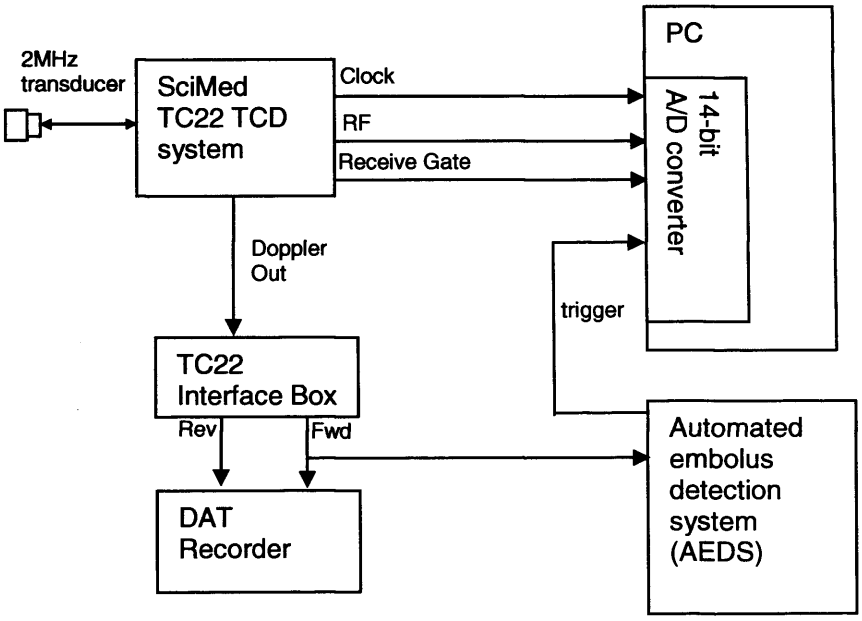

Figure 3-1 Schematic of RF acquisition system

The input range on the A/D converter was set so that data at potential depths of interest were not clipped. This is at the expense of the dynamic range being less well optimised for larger ranges, resulting in a poorer SNR at these depths.

Recordings were made from normal subjects and also from patients undergoing CEA surgery (see subsections on 'Artefact Characteristics' and 'Clinical Data' in the 'Data Collection' section, for further details). A 2 MHz (SciMed) probe was fixed over the patients' temporal bone windows, ipsilateral to the operated carotid artery, so as to insonate the middle cerebral artery (MCA). The clock rate on the TC22 system was 32.4 MHz. The pulse repetition frequency (PRF) of the system was dependent on the velocity scale used on the TCD machine at the time of acquisition. This scale is set by the vascular technician monitoring the patient. The PRF varied between 2,637 Hz and 7,910 Hz. The 'sample volume' setting on the TC22 system can be varied between 1 and 15 mm, but it is common practice by the monitoring technicians to set this at 10 mm. This

gives a transmit burst of approximately 13 μs (26 cycles) long and a receive gate of 10 mm. This setting has been used for all acquisitions. The start of the receive gate is positioned at the depth set on the TCD machine by the technician. The end of the receive gate is 10 mm deeper than the start. The depth setting on the TCD machine varied between 42 mm and 59 mm in the collection of the data used in this study.

3.2.2 Signal Processing

As an initial step, the depths present in the receive gate are extracted from each pulse in the RF-matrix and are processed to create the sonogram. The positions of any suspected emboli or artefacts can then be determined and the appropriate region in the RF-matrix can be selected for further processing. It should be noted that a sonogram could be created using the RF signal from any depth, not just the depth being monitored during and after surgery.

To examine potential emboli and artefacts in the RF, each column of the RF-matrix is clutter-filtered using a 2^{nd} order Butterworth high-pass filter with a cut-off frequency of 75 Hz (this is the same cut-off frequency on the wall-thump filter as set on the TCD machine during monitoring). A magnitude plot is then used to visualise the filtered RF-matrix over the desired range of depths and time interval. The magnitude plots are shown with respect to the root-mean-square (rms) noise level of that particular data acquisition. These were calculated from a region of the RF signal, from each acquisition, containing no suspected embolus or artefact.

In some cases, the SNR of the embolic signal may be low. This may occur if the embolus is small, if it is travelling in a low sensitivity region of the beam, or if the input range of the A/D converter is not optimally set for the depth range of the embolus. Despite the SNR being low, there will be a region of the received signal that, to a certain degree, looks like the transmitted pulse as a result of it being scattered from the embolus. Matched filtering is one commonly used technique to detect a known waveform in a noisy background. This involves cross-correlating a copy of the transmitted pulse with the received pulse. The result will contain a peak if any region of the received signal resembles the transmitted pulse. Since noise in the received signal will bear a resemblance to any target signal, the matched filter output will also contain noise (Smith, 1997), but this should be of noticeably lower amplitude than the peak located at the embolus position. Matched filtering has been performed on some data in this chapter, where an event is apparent in the sonogram but where no event is visible in the RF signal, to assess the benefits of the technique.

3.2.3 Data Collection

Data were collected *in vitro* to characterize the appearance of events from known sources. Data collected *in vivo* were compared to those from *in vitro* and observations on the nature of events contained within the signals were made. The various methods used for data collection are described below.

3.2.3.1 Test Object

An open-topped test object, made for quality assurance of imaging ultrasound systems, was used to characterize the return from a strong stationary reflector in the RF signal to aid the interpretation of subsequent signals from emboli and artefacts. The test object consisted of a Perspex tank filled with a mixture of 9.5% ethanol by volume in distilled water (Martin and Spinks, 2001). Within the tank there were a number of 0.5 mm diameter wires, on a plane parallel to the base of the tank, separated vertically and horizontally by 25 mm. The 2 MHz TCD transducer was clamped in place so that its face was just below the surface of the ethanol/water mixture, 54 mm from the first wire (hence 79 mm from the second and 104 mm from the third), such that the beam was perpendicular to the base of the tank. With the 'sample volume' setting at 10 mm on the TCD machine, a capture of the RF signal was made.

3.2.3.2 Artefact Characteristics

Four presumed healthy volunteers generated typically encountered artefacts through speech, groaning, cough, throat clearing, whistling and tapping the probe. A trigger was forced in the acquisition software to capture the RF at the time of these artefacts.

Another source of artefact is diathermy. This has quite a distinctive appearance on the sonogram but is an artefact that is known to generate errors in some automated embolic detection systems (Mess et al., 1997; 2000). Signals containing diathermy artefacts were acquired during the dissection phase of several CEAs by forcing a trigger in the acquisition software.

3.2.3.3 Emboli Characteristics

To compare the appearances of artefacts with those of emboli, several events were recorded from a flow phantom. The flow phantom was set up as shown in the schematic in Figure 3-2. The rubber acoustic absorber was required to attenuate the signal to a usable level. Since all items shown were clamped in place, any events must have occurred from

bubbles or clumps of blood mimicking fluid (BMF) (Ramnarine et al., 1998) passing the transducer on their way around the flow loop. These naturally occur in such a flow phantom for a period of time whilst the flow loop settles. Such events are embolic in nature. No bubbles or particles of known properties were injected into the flow loop to deliberately simulate emboli. No movement artefacts occurred. Events were recorded both from steady flow and from pulsatile flow; no difference in the appearance of 'embolic' events, between the two flow types, was observed.

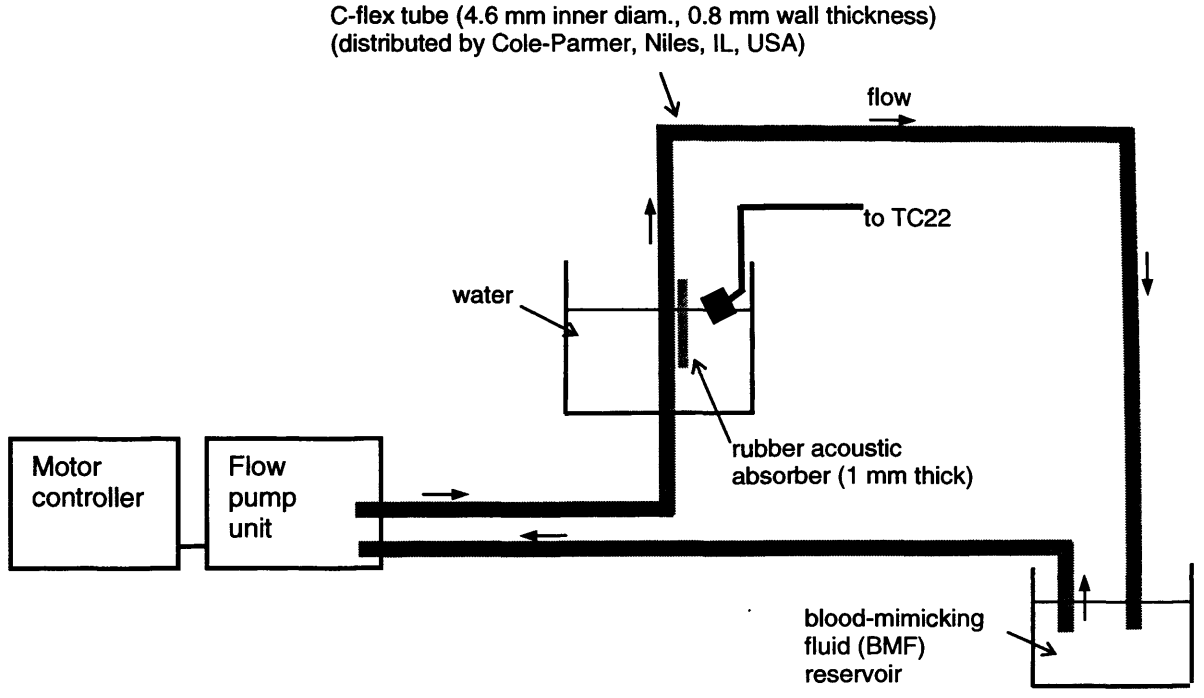

Figure 3-2 Schematic of flow phantom

3.2.3.4 Clinical Data

In vivo data were collected throughout the routine monitoring of 10 patients during and immediately after (for up to 3 h) CEA surgery. These data included presumed artefactual and embolic events, including emboli showers.

Matched filtering has been performed on some of the clinical data to aid the detection of events with a low SNR.

3.3 Results

3.3.1 Test Object

A magnitude plot of the unfiltered RF signal between 40 mm and 125 mm is shown in Figure 3-3. The depths of the wires relative to the face of the transducer correspond to the depths at the left side of the regions of increased intensity in Figure 3-3. The length of the increased intensity regions corresponds to the time taken for a complete transmitted pulse to reflect back from the respective wires. Bearing this in mind will assist in the interpretation of subsequent figures in this chapter.

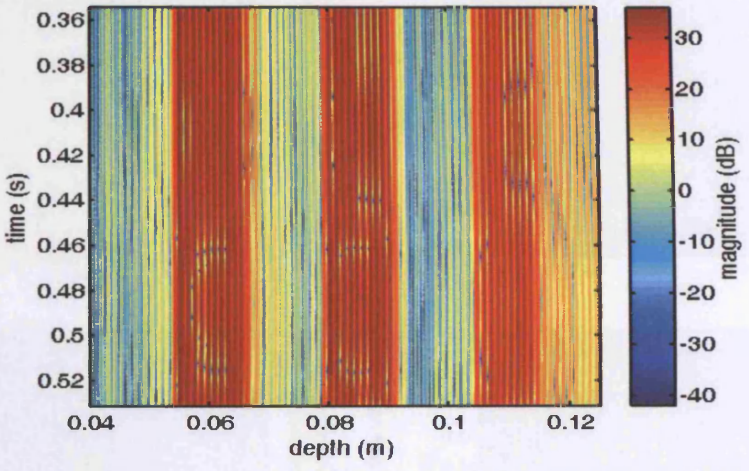

Figure 3-3 Magnitude plot of RF signal showing reflections from 3 stationary targets separated by 25 mm. The depths of the wires, relative to the face of the transducer, correspond to the depths at the left side of the regions of increased intensity (depth setting = 51 mm, PRF = 3955 Hz)

3.3.2 Artefact Characteristics

Figure 3-4 shows the sonograms and magnitude plots of the filtered RF signals for a small sample of the artefacts generated; Figure 3-4a illustrates an artefact generated from speech, Figure 3-4b from throat clearing, and Figure 3-4c from tapping the probe. These all appear in the sonogram as an increase in intensity at a low frequency and have a characteristic 'thump' when listening to the audio signal. When observing the magnitude plots of the filtered RF signal, it can be seen that whilst the appearance of the different artefacts are by no means identical, they all have the characteristic that there is an increased intensity over a large number of depths simultaneously. It should also be noted that different artefacts from the same source, e.g. probe tapping, may not necessarily be similar in appearance but will have the characteristic of having increased intensity over a large number of depths.

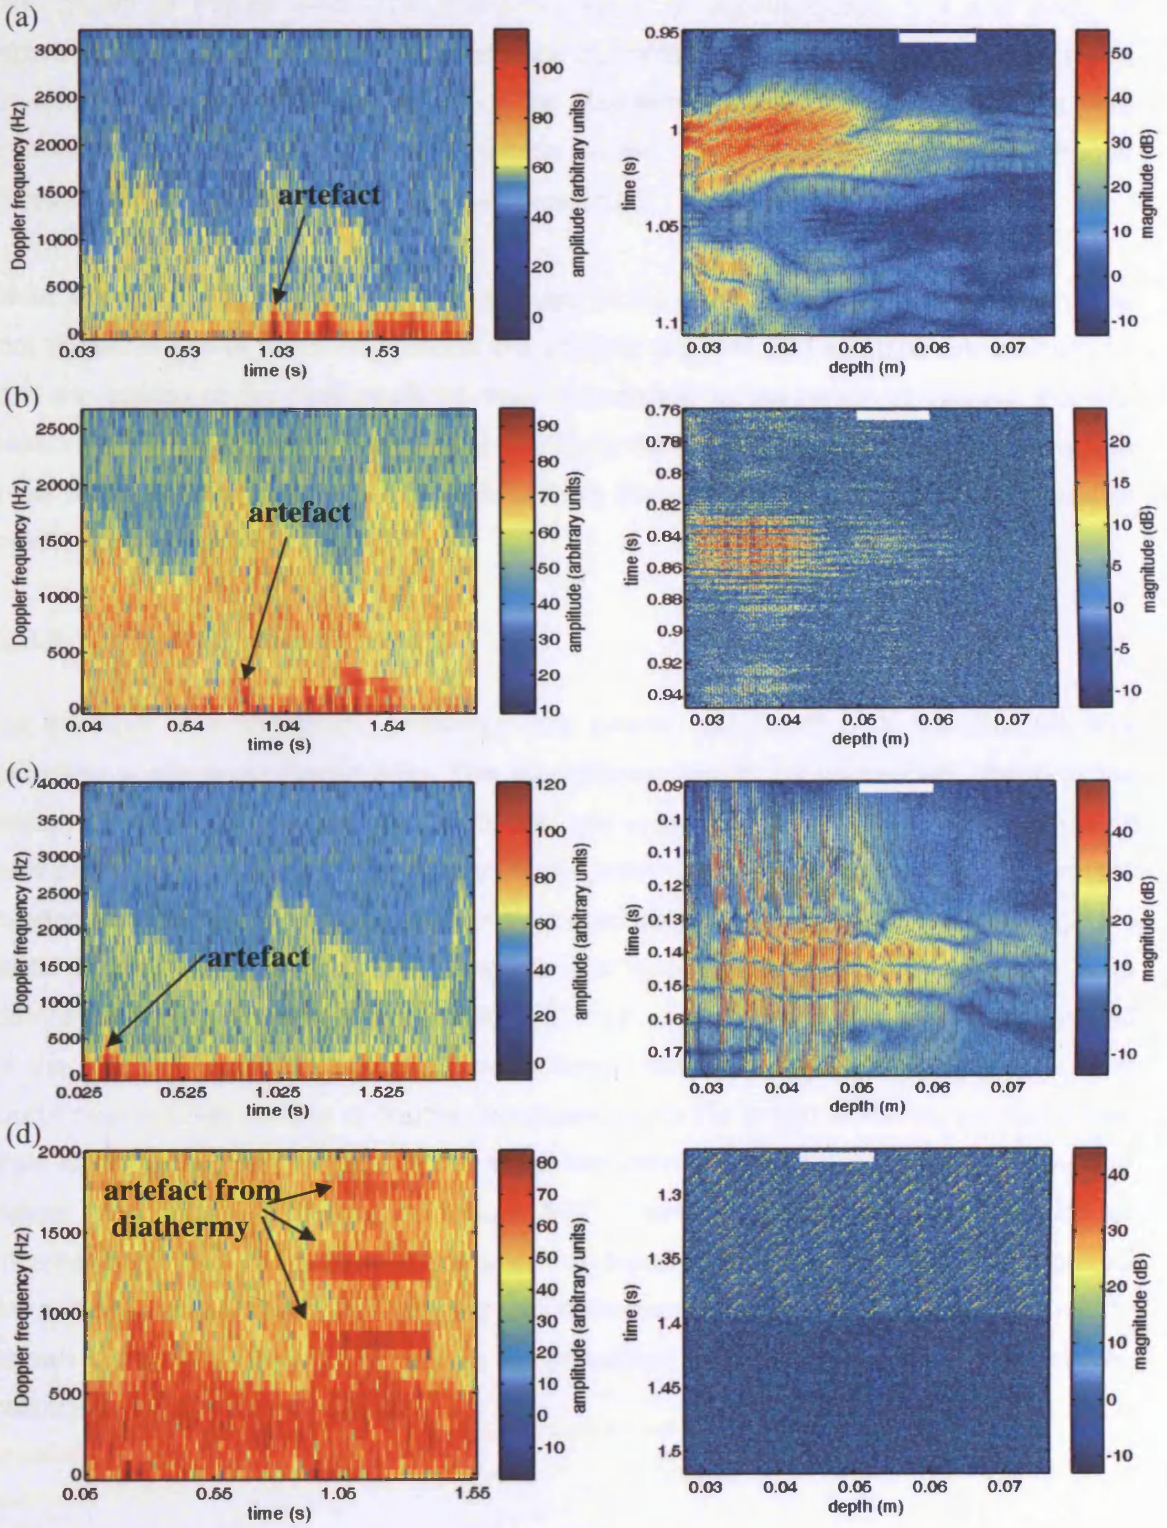

Figure 3-4 Sonograms and magnitude plots of filtered RF signal containing artefacts generated through (a) speech (depth setting = 55 mm, PRF = 6328 Hz), (b) throat clearing (depth setting = 50 mm, PRF = 5273 Hz), (c) tapping probe (depth setting = 50 mm, PRF = 7910 Hz), and (d) diathermy (depth setting = 42 mm, PRF = 3955 Hz). The white rectangles in the RF plots indicate the depths within the receive gate.

A sonogram and magnitude plot of the filtered RF signal, containing a diathermy artefact, are shown in Figure 3-4d. The diathermy starts at approximately 1 s and ends at approximately 1.4 s. Its appearance is quite distinctive and would be easy to identify. It should be noted that the amplitude of the diathermy artefact does not decrease with increasing depth unlike the other artefacts shown. The appearance of the diathermy artefact did not vary across the examples examined.

Note that the settling region of the filter, used in the generation of the Doppler signal, is not shown in any of the sonograms in this chapter (the first 200 samples are omitted) so that the scaling of the plots could be more appropriate for the region of interest. For this reason, none of the sonograms start at 0 s. Also, the white rectangle in the RF magnitude plots indicates the position of the gate used in that acquisition and hence also used to create the sonogram.

3.3.3 Emboli Characteristics

An example of a sonogram, containing two events, generated from the pulsatile flow phantom is shown in Figure 3-5a. The magnitude plots of the filtered RF signal, in the region of each event, are shown in Figure 3-5b and Figure 3-5c. The echoes apparent between 30 and 42 mm, in both magnitude plots, are due to the reflections from the rubber acoustic absorber. The region of increased intensity in the RF signal, resulting from each of the events shown in the sonogram (at approximately 1 s and 1.8 s), is quite different in appearance from that resulting from an artefact, as shown previously. Instead of the intensity increases occurring over a large number of depths simultaneously, they occur over a small number of depths corresponding to the length of the transmitted pulse. This is consistent with the size of the increased intensity regions in the plot showing the return from stationary reflectors (Figure 3-3). There is also a shift in the region of increased intensity to shallower depths with an increase in time. This can be interpreted as a movement of the object towards the transducer. The appearances of the events shown in Figure 3-5b and Figure 3-5c are consistent with the appearances of all events recorded from the flow phantom.

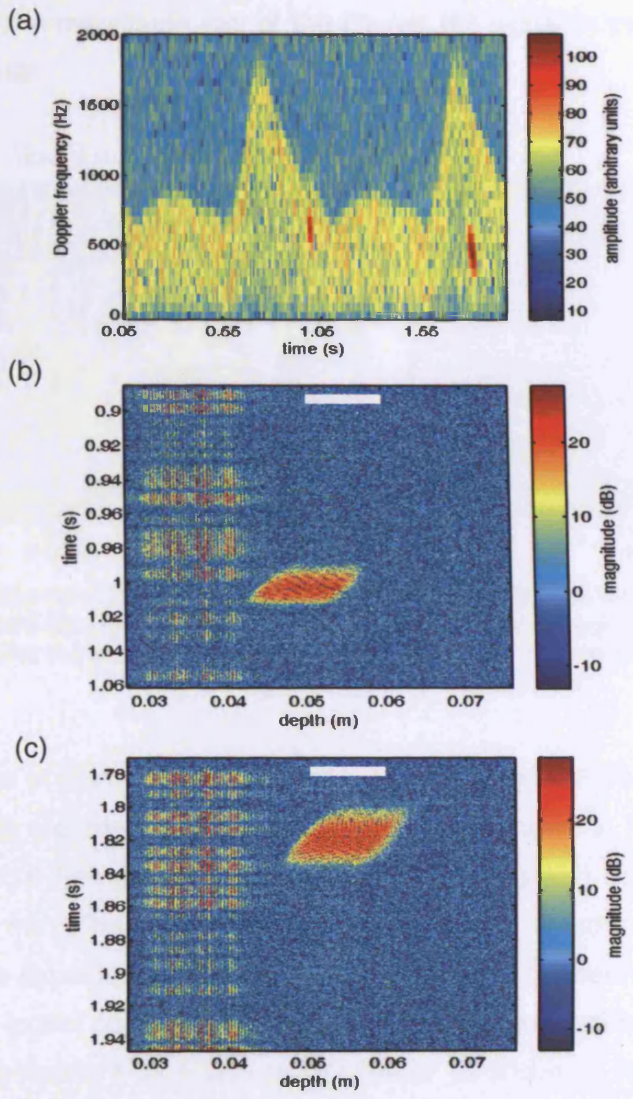

Figure 3-5 (a) Sonogram showing two events in signal captured from flow phantom, and (b) and (c) Magnitude plots of the filtered RF signal in the region of the two events, with the white rectangle indicating the depths within the receive gate (depth setting = 50 mm, PRF = 3955 Hz)

3.3.4 Clinical Data

3.3.4.1 Comparison of in vivo and in vitro Data

The signals recorded from the flow phantom can be compared with an event captured *in vivo* agreed to be an embolus by the AEDS and a human expert. Such an event is shown in the sonogram in Figure 3-6a as Event 2 (Events 1 and 3 will be discussed in the subsection on '*Improving the SNR of the RF signal*' and in the 'Discussion' section respectively). This signal was captured, during CEA surgery, just after the internal carotid artery (ICA) was reopened. The AEDS triggered on Event 2 and did not record Events 1

and 3 as emboli. The magnitude plot of the filtered RF signal in the region of Event 2 is shown in Figure 3-6b.

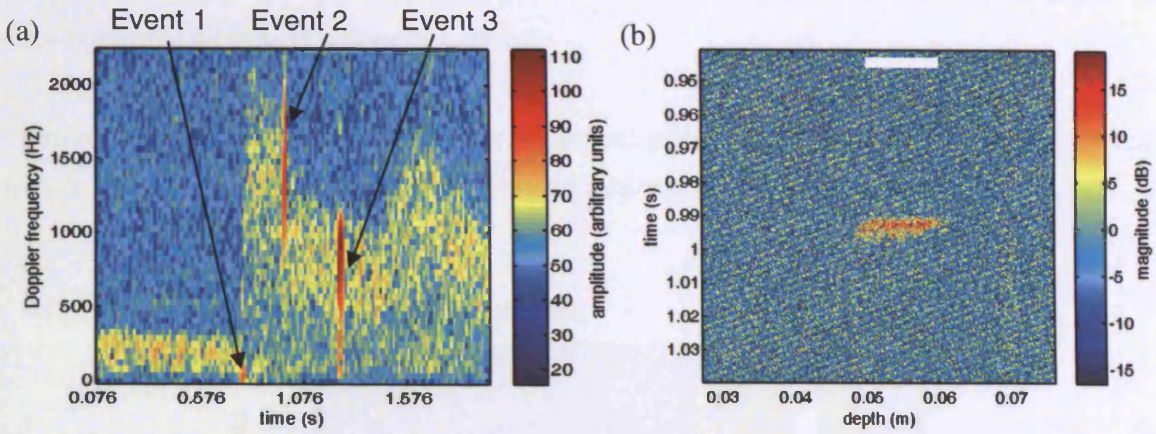

Figure 3-6 (a) Sonogram from signal captured *in vivo* during CEA surgery just after the ICA was reopened, and (b) Magnitude plot of filtered RF signal in region of Event 2, with the white rectangle indicating the depths within the receive gate (depth setting = 50 mm, PRF = 2637 Hz)

The characteristics of the increased intensity region in Figure 3-6b are similar to those of the embolic events recorded from the flow phantom as shown in Figure 3-5b and Figure 3-5c. In particular, it has duration of approximately 5 ms and covers a depth range of approximately 10 mm. The duration is much shorter than those from the phantom but is consistent with the expected duration of an embolus being between 2 and 100 ms (Evans, 1999). The depth extent corresponds to the length of the transmitted pulse and is, hence, consistent with the results from a stationary reflector as shown in Figure 3-3. There is also a shift in the region of increased intensity to shallower depths with an increase in time. This can be interpreted as a movement of the object towards the transducer.

3.3.4.2 *Using the RF Signal to Analyse an Emboli Shower*

A sonogram containing an emboli shower is shown in Figure 3-7a. These data were recorded intra-operatively at the reopening of the ICA. There are a few clearly separated emboli near the middle of the recording and an emboli shower at the end. It is clear that counting the number of emboli in a shower using the sonogram would be virtually impossible.

A small region of the emboli shower is illustrated in the demodulated Doppler signal in Figure 3-7b and in the magnitude plot of the filtered RF signal in Figure 3-7c. By having additional depth information in the RF signal, the identification of individual emboli is simplified in comparison to using the demodulated Doppler signal. When a cross-section at a particular depth is extracted from the RF plot, the result closely represents the shape

of the demodulated Doppler signal and the ease of identification of individual events is lost. Figure 3-7c also shows that, although the majority of events (circa 6 main events) occur at a depth of approximately 50 mm, there is also an event at a depth of approximately 42 mm at a time of 1.89 s.

Therefore, again, the RF signal has provided valuable additional information about these events, which was not available from the sonogram or the Doppler signal.

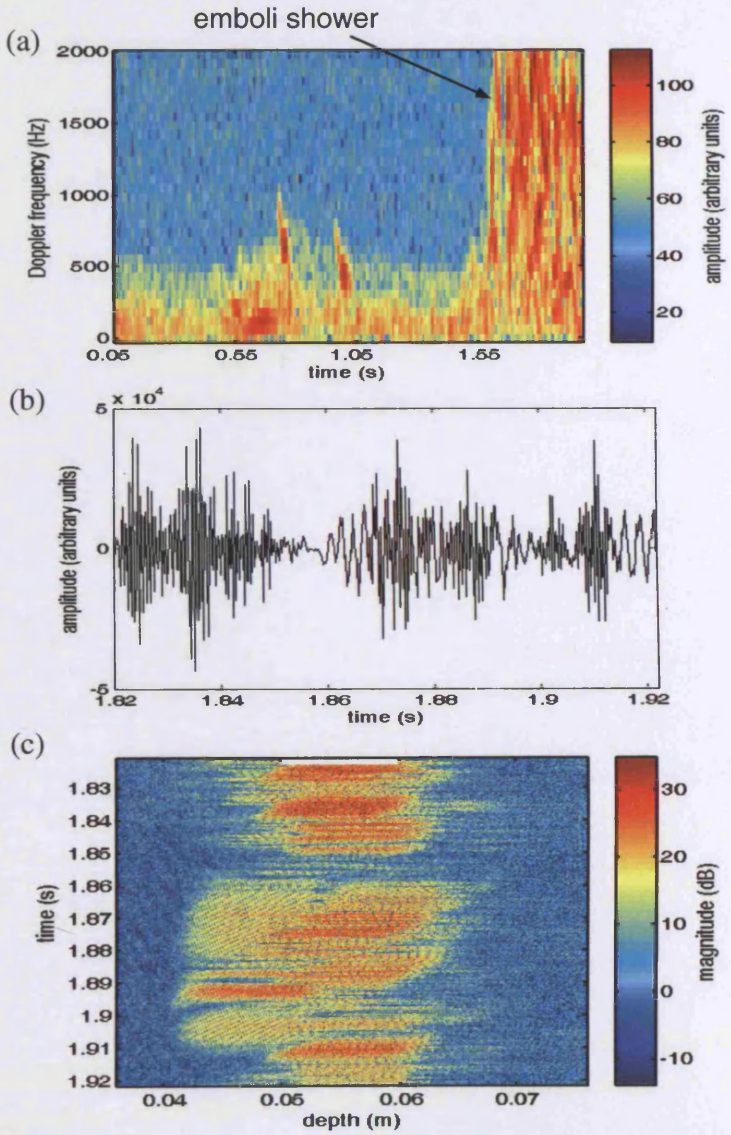

Figure 3-7 (a) Sonogram created from a 10mm range gate situated at 50 mm from data acquired intra-operatively at the re-opening of the ICA, (b) Demodulated Doppler signal over small region of emboli shower, and (c) Filtered RF in same region of emboli shower shown in (b), with the white rectangle indicating the depths within the receive gate (depth setting = 50 mm, PRF = 3955 Hz)

3.3.4.3 *Improving the SNR of the RF Signal*

An example of the benefits of matched filtering can be seen when analysing Event 1 from Figure 3-6. This event was considered to be an artefact by both the human expert and the AEDS given the characteristics evident in the sonogram (Figure 3-6a) and the Doppler signal. When examining the magnitude plot of the filtered RF signal in Figure 3-8a there appears to be no event present. Matched filtering was performed by cross-correlating a copy of the transmit pulse (Figure 3-9) with each line (row) of the portion of the RF-matrix shown in Figure 3-8a. The result is shown in Figure 3-8b. This indicates that there is an event with embolic characteristics present. At any one time, the peak in the matched filtered signal will be located close to the embolus depth (seen in previous cases as the left edge of the region of increased intensity in the filtered RF signal). This will be the point where the transmitted pulse will almost completely overlap with the pulse reflected by the embolus and so will give the highest correlation.

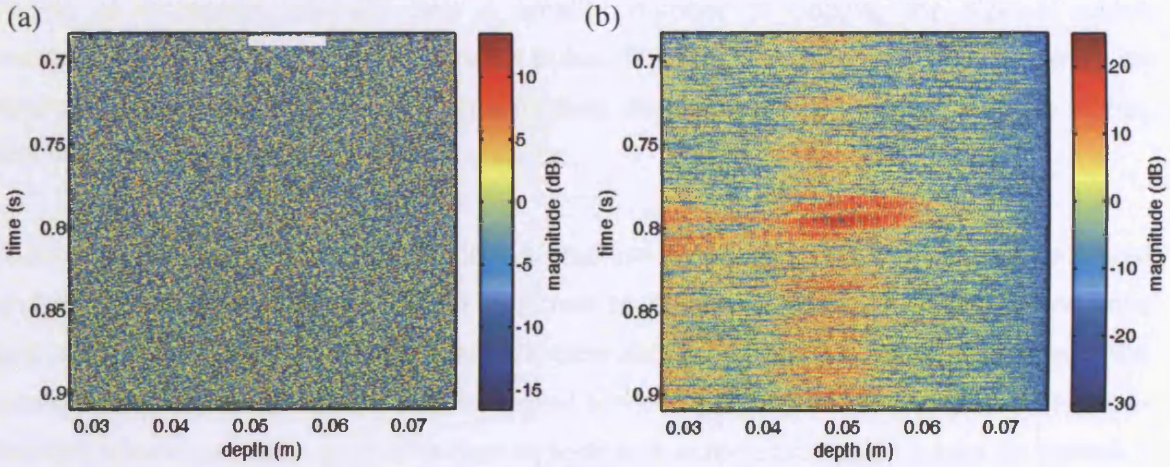

Figure 3-8 (a) Magnitude plot of filtered RF signal in region of Event 1 from Figure 3-6, with the white rectangle indicating the depths within the receive gate, and (b) Magnitude plot of matched filtered signal in region of Event 1 from Figure 3-6 (depth setting = 50 mm, PRF = 2637 Hz)

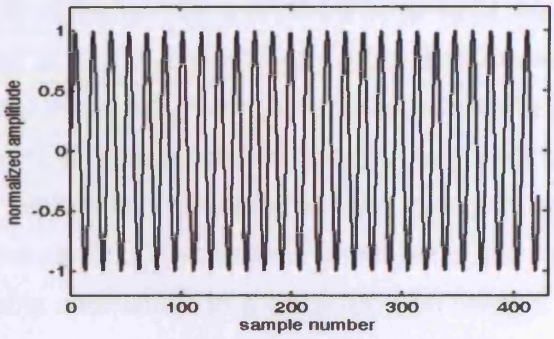

Figure 3-9 Transmitted pulse waveform

This effect has also been noticed with other acquisitions where there has been a very faint event in the sonogram and Doppler signal, but which has been completely buried in noise in the filtered RF signal. As a result of performing matched filtering on the filtered RF signal, the event not only becomes visible, but displays distinctive characteristics which allows the nature of the event (embolic or artefactual) to be hypothesized. Matched filtering has also been applied to events which already have an SNR adequate to enable the event to be detected and the nature determined. In these cases the application of matched filtering has proved to increase the SNR of the event significantly.

3.4 Discussion

It has been shown that emboli and artefacts have quite different characteristics in the magnitude plot of a filtered RF signal; artefacts have a region of increased intensity which covers a large number of depth samples simultaneously whereas an embolus will have a region of increased intensity over a smaller number of depths, the size of which corresponds to the length of the transmit pulse. The increased intensity resulting from an embolus may exhibit a shift in depth with time; depending on the duration of the signal, this shift in depth may not always be obvious.

Current automated embolus detection systems are designed to make decisions on signal characteristics, based on the criteria assumed by 'human experts' to distinguish embolic and artefactual events. Little can truly be known about the nature of these events from the one-dimensional demodulated Doppler signal or the sonogram and therefore the decision-making criteria currently used by human experts and automated systems may be flawed.

3.4.1 The RF Signal as an Alternative 'Gold Standard'

The performances of automatic embolic detection systems are often compared against the opinions from a panel of 'human experts'. If the majority of experts believe an event to be an embolus then that is how it is classified. There is more likely to be disagreement over low SNR events and these may then be incorrectly classified as artefacts. This may lead to an unfair assessment of an automatic system's performance. It is widely accepted that the 'human expert' cannot be considered to be a true gold standard, as no one truly knows what type of signal really corresponds to an embolus, but this continues to be used due to a lack of a suitable alternative. In a discussion on using a "human gold standard" for assessing the performance of automated embolic detection systems, Mess (2003) suggested that the detection of a moving intensity increase in the RF based amplitude plot

'could possibly perform as a better gold standard than the human observer expert panel', but also noted that many of the particulate emboli could not be analysed in terms of velocity and flow direction in the RF plots. It has been shown in this chapter, however, that the RF signal can provide additional information on an event. In particular, it has been shown that emboli and artefacts have distinctive characteristics when examining a magnitude plot of the filtered RF signal (or the result of matched filtering this signal with a copy of the transmit pulse), thus enabling a decision on the nature of an event to be made. It is still possible that, on this basis, the RF signal could be used as a gold standard for assessing the performance of automated embolus detection systems, as no events have been found that could not be analysed in this way.

In the analysis of our data, there have been several occasions where the AEDS decided an event was embolic but the human expert disagreed and vice versa. One example of each case is illustrated below.

The following data were collected from a patient post-operatively in the recovery room. The sonogram is shown in Figure 3-10a. There are three events present in this signal. The AEDS recorded Event 6 as an embolus and considered Events 4 and 5 to be artefacts. However, the human expert considered all three events to be artefacts, based on their sound and appearance. The magnitude plots of the filtered RF in the immediate vicinity of Events 5 and 6 are shown in Figure 3-10b and Figure 3-10c respectively. Event 4 is very similar in appearance to Event 5 and therefore has not been shown. It is evident from the RF plots that Event 5 has the hallmarks of an artefact with a region of increased intensity occurring over a large number of depths simultaneously. Event 6, on the other hand, appears to have the same characteristics of emboli shown previously i.e. increased intensity over a number of depths corresponding to the length of the transmitted pulse and duration of approximately 7 ms. The source of this event will remain unknown, but its RF characteristics suggest that it was probably an embolus.

Event 3 in Figure 3-6, captured during CEA surgery, just after the ICA was reopened, was considered to be an embolus by the human expert but was considered an artefact by the AEDS. The magnitude plot of the filtered RF signal in the region of this event is shown in Figure 3-11. From the RF plot, Event 3 looks more likely to be an embolus than an artefact based on its appearance in a low number of depths at any one time. The Doppler signal of this event is very large and overloaded the input of the AEDS resulting in the signal being clipped. This may have distorted the signal such that characteristics the AEDS recognises as being embolic were lost.

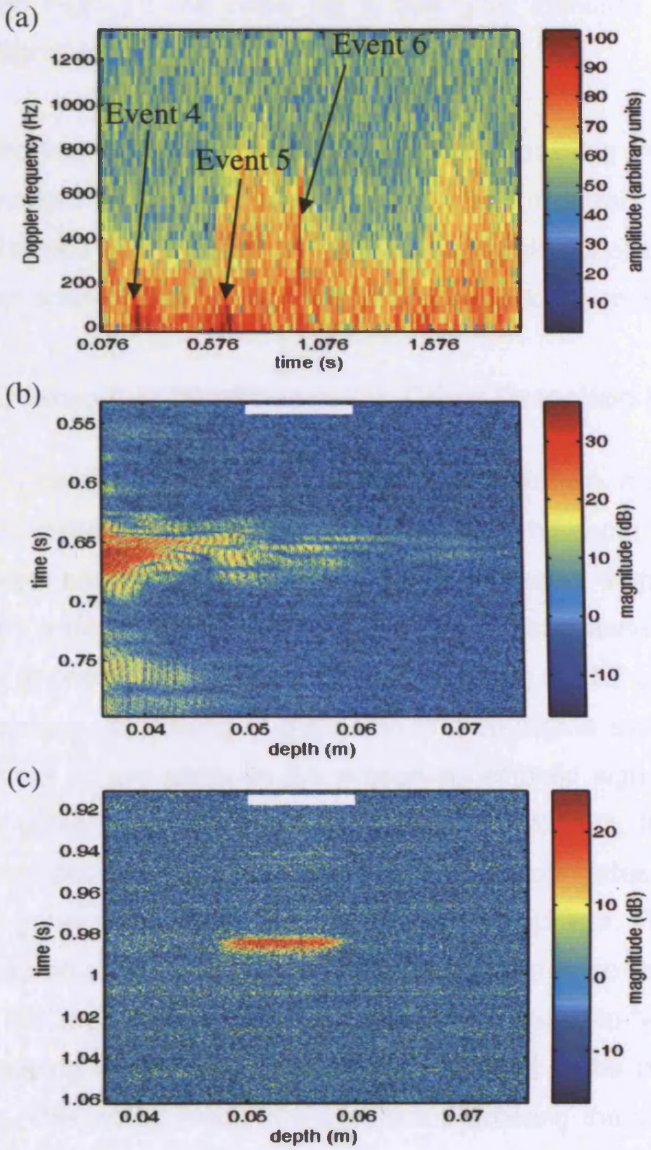

Figure 3-10 (a) Sonogram generated from data acquired post-operatively with 10 mm long receive gate positioned at 50 mm, (b) Magnitude plot of filtered RF in region of Event 5, and (c) Magnitude plot of filtered RF in region of Event 6. The white rectangles in the RF plots indicate the depths within the receive gate (depth setting = 50 mm, PRF = 2637 Hz)

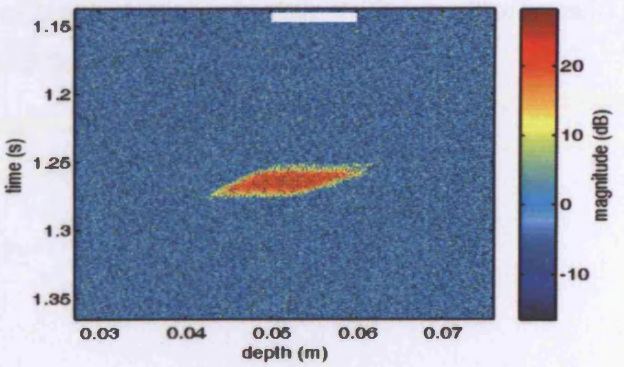

Figure 3-11 Magnitude plot of filtered RF signal in region of Event 3 from Figure 3-6, with the white rectangle indicating the depths within the receive gate (depth setting = 50 mm, PRF = 2637 Hz)

The above examples highlight the need for a true gold standard and illustrate the suitability of the RF signal for this role.

During an emboli shower, counting the number of emboli passing through the sample volume can be a problem for both the human expert and for current automated emboli detection systems (Russell and Brucher, 2005). It has been demonstrated in this chapter that the RF signal can enable closely spaced (in time) events to be analysed separately.

3.4.2 RF Signal Overcomes Weaknesses in Other Detection Systems

Whilst the source of individual *in vivo* events will remain unknown, it appears as though the RF signal can aid in the classification process. By using the depth extent of an event, the nature of that event can be inferred. This is the same theory employed by multigate systems i.e. that an artefact will occur at many depths simultaneously whereas an embolus will only be at one point at any one time. Examining the RF signal highlights the importance of appropriate positioning of the gates in a multigate system. For example, consider the placement of two gates in the presumed embolic signal shown in Figure 3-12. When 10 mm gates were positioned at 43 mm and 53 mm, the sonograms and Doppler signals show blood flow with very similar powers in both gates. Also, the embolus first appears in both gates at the same time and with similar power. Therefore, in a dual-gate system, this may be classified as an artefact since there is no evidence of travel of the event between the two gates. This may occur if the vessel-to-beam angle is quite large and the overlapping region of the two sample volumes is the portion of the beam insonating the MCA. This would result in the embolus entering the side of both sample volumes simultaneously. The likelihood of this occurring decreases as the length of the transmit burst or the size of the vessel-to-beam angle decrease. It has also been observed that there may be emboli at multiple depths at one time. An example of this is shown in Figure 3-13. Again, if the gates were not suitably placed, it may appear as though there was an artefact at approximately 0.79 s as there could be a simultaneous increase in power in both gates.

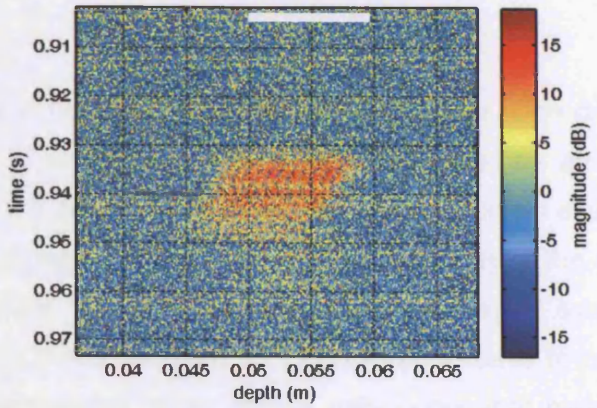

Figure 3-12 Magnitude plot of the filtered RF signal in the region of a presumed embolic event, with the white rectangle indicating the depths within the receive gate (depth setting = 50 mm, PRF = 2637 Hz)

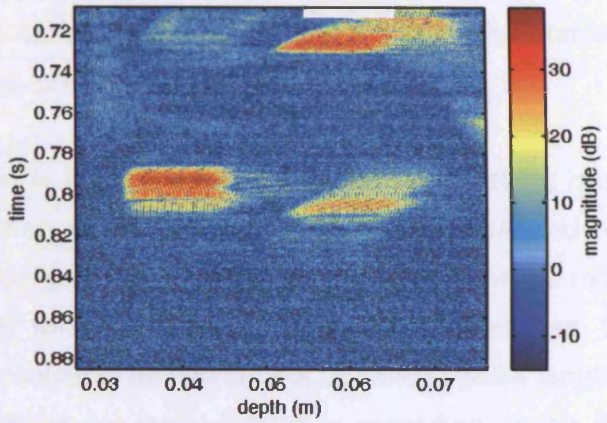

Figure 3-13 An example of two emboli occurring simultaneously at different depths, acquired intra-operatively at the re-opening of the ICA, with the white rectangle indicating the depths within the receive gate (depth setting = 54 mm, PRF = 3955 Hz)

The use of Power M-Mode Doppler (PMD) (Moehring and Spencer, 2002) assists in the visualization of the paths taken by emboli. It does, however, still rely on Doppler frequency characteristics of events to determine the nature of an event. We have shown (Event 1, Figure 3-6) that putative embolic events may exist at Doppler frequencies lower than 200 Hz. An example, showing artefactual events, in Moehring and Spencer (2002), employs a filter which displays mean Doppler frequencies lower than 200 Hz as a region of zero power. This is the means by which an artefact is identified – a region of zero power over all depths. However, if there is a slow moving embolus then this too will be coloured black and therefore, missed by observers. Also, if an artefact is present as a result of probe movement, there is likelihood that this may have a mean Doppler frequency greater than 200 Hz. As a result this will not be assigned a zero power and will not have the same appearance as the artefacts described. Using the RF signal places no reliance on the frequency characteristics of an event thus making it potentially more reliable.

3.4.3 Weaknesses

There are weaknesses in our current RF acquisition system. One drawback is that only RF data in the vicinity of events, suspected by the AEDS to be embolic, are captured. Therefore, no RF data will be available to provide clarification on events not detected by the AEDS. This could be resolved by increasing the sensitivity of the AEDS, resulting in a reduction in the specificity. Another shortcoming of the current system is that, due to volume of data, a single event can take up to 1 minute to save to hard disk. Therefore, whilst the event is being saved, other emboli may occur that may not be able to be acquired. Processor speeds are constantly increasing, however, and so this is not considered to be a long-term issue. The acquisition system may also be modified, in the future, such that a smaller duration of signal is transferred to hard disk for each event. This would reduce the saving time required and therefore the number of occasions that events were missed due to data being saved.

An additional weakness is that the sample volume length setting is set at 10 mm by the monitoring technicians so as to assist in locating the MCA and maintain a signal of adequate strength. This gives a long transmitted pulse, which results in poor spatial resolution. Therefore if two emboli were separated by less than 10 mm axially, their echoes would not be resolvable as two targets. A smaller pulse length was used by Mess et al. (2002) with a noticeable improvement in resolution in the RF magnitude plots. However, it appears as though their acquisition system was used solely to collect embolic signals in comparison to being an addition to a system, which is also required to perform conventional monitoring of the patient.

3.5 Conclusions

In conclusion, it has been shown that magnitude plots of the RF signal appear to provide a convenient way of identifying if an event is an embolus or artefact. It is also possible that these could become the new gold standard for assessing the performance of automated embolic detection systems.

It has been shown that low frequency events, assumed to be artefacts from their appearance in the sonogram and Doppler signal, may have embolic characteristics when analysing the RF signal. Therefore, those using TCD to detect emboli should be aware

that some emboli may produce much lower frequency signals than is traditionally believed.

It has been demonstrated that, due to the additional information present in the RF signal, the analysis of emboli showers may be easier. The RF signal clearly shows the depths at which the events occur as well as how many events are present.

Finally, by choosing the input range of the A/D card so that a larger number of depths can be examined for discrimination between artefacts and emboli, the input range may not be well matched to the signal level in the region of an embolus. This results in a low SNR making the identification of events difficult. However, by performing matched filtering on the RF signal, the SNR of an event can be greatly improved so that the nature of the event can be determined.

4 Automatic Detection of Emboli in the TCD RF Signal using Principal Component Analysis

4.1 Introduction

Transcranial Doppler ultrasound (TCD) is commonly used, in many centres, to detect embolic signals in the cerebral circulation during and after carotid endarterectomy (CEA) surgery. A trained observer monitors the signal, during this time, for occurrences of emboli. This is a very tedious task and, if performed for long durations, is likely to result in diminishing performance. The decision as to whether an embolus is present is very subjective and, therefore, two observers may not always agree on the nature of an event. To monitor every patient during and post-CEA surgery is also very costly. To solve these problems, an automated embolus detection system is required in which the discrimination between artefacts and emboli is based on clear differences between the two types of signal. This is not a simple task using the demodulated Doppler signal as there is a large overlap in the characteristics of emboli and artefacts. Although a number of systems using the demodulated Doppler signal are commercially available, doubts remain as to their reliability in all clinical situations. An alternative approach would be to use the raw radio-frequency (RF) signal from the TCD system. This would minimise the amount of signal loss and distortion that had occurred as a result of pre-processing and therefore has the potential to improve the ability to discriminate emboli and artefacts by providing additional information on these events.

In 2002, El-Brawany and Nassiri investigated the automatic detection of emboli using the RF signal and a neural network. The neural network was trained with examples from blood-mimicking fluid, solid and gaseous emboli and wall motion artefacts. This system has only been tested on *in vitro* data and considerable development would be required to use this system *in vivo*. In the same year, Mess et al. (2002) showed that the paths of emboli could be visualized in 2-D plots of the TCD RF signal. No attempt was made, using this data representation, to detect the occurrence of emboli automatically. It was shown in Chapter 3 that emboli can be differentiated from artefacts using 2-D plots of the RF signal; embolic signals appear over a small number of depth samples in comparison to artefact signals when examining magnitude plots of the RF-matrix. Furthermore, when viewed in the RF signal, emboli tend to appear as parallelograms of varying duration. The width of the parallelogram is determined by the transmitted pulse length, the slope by the velocity of the embolus and the PRF, and the duration is determined by the length of the sample volume. In contrast, artefacts appear in these magnitude plots with a variable, ill-defined

shape, across a large number of depth samples. An example of both types of signal is shown in Figure 4-1. The uniformity and predictability of the embolic signals' appearance have led to the investigation of the use of pattern recognition and image processing techniques for their automatic detection and discrimination from artefacts.

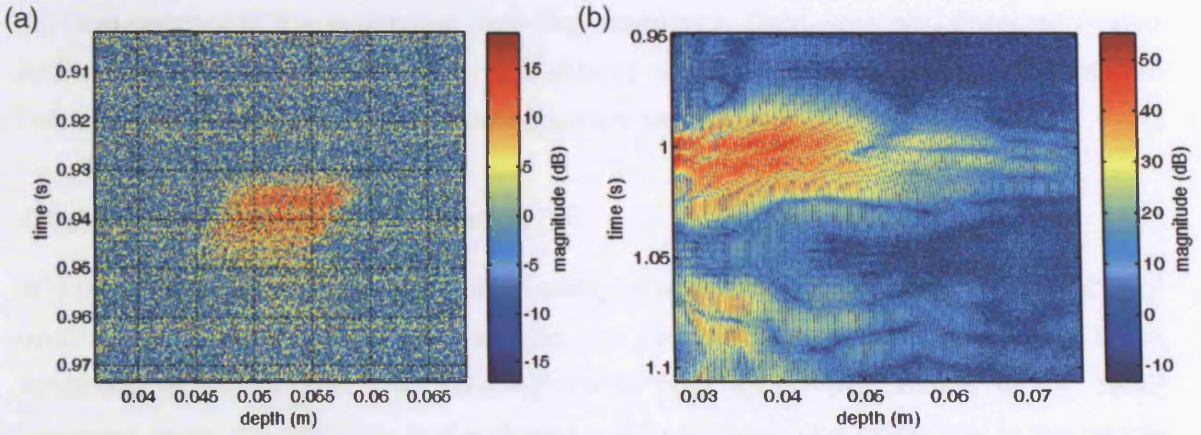

Figure 4-1 Typical appearance of (a) an embolus, and (b) an artefact in the RF signal

In particular, principal component analysis (PCA) has been applied to training sets of embolic events collected both *in vitro* and *in vivo*. This has enabled the typical variation of their shape and size to be determined. Algorithms have been developed, using the parameters determined from the training sets, to automatically differentiate emboli and artefacts in the RF signal. These algorithms have been tested using *in vitro* and *in vivo* data not included in the corresponding training set.

4.2 Methods

4.2.1 Data Acquisition System

The data acquisition system used for this part of the project is as described in Chapter 3. Additional details can also be found in Chapter 2. During the collection of the data used in this study, the PRF varied between 1978 Hz and 7910 Hz and the depth setting on the TCD machine varied between 42 mm and 55 mm.

Note that, due to the volume of data involved when using the RF signal, it was unfeasible to collect data continuously. Therefore, some mechanism was required to force a data capture when an embolus (or artefact) was suspected. The implementation used in this study meant that only emboli suspected by the AEDS were captured. This may introduce

a selection bias but it was an efficient way of building up a library of signals suitable for developing new detection algorithms.

4.2.2 Data Collection

Data were collected *in vitro* to obtain signals from events from known sources to use in the development of the automated detection algorithms. Data were also collected *in vivo* and were used together with the *in vitro* data to test and further develop the algorithms. The various methods used for data collection are described below.

4.2.2.1 Flow Phantom Emboli

In order to develop detection algorithms using events known to be embolic, a number of events were recorded from a flow phantom. The flow phantom was set up as shown in the schematic in Figure 4-2. The blood-mimicking fluid (BMF) (Ramnarine et al., 1998) circulates round the flow loop in the direction shown to simulate blood flow in the middle cerebral artery (MCA). A wall-less phantom (Ramnarine et al., 2001) was used to minimise the ultrasound distortion which can be present if tubing is used as a vessel channel. Since the transducer was clamped in place, any recorded events occurred from bubbles, clumps of BMF or small pieces of tissue mimicking material (TMM) passing the transducer on their way around the flow loop. The bubbles and clumps of BMF naturally occur in such a flow phantom for a period of time whilst the flow loop settles. Such events are embolic in nature. The small (≤ 1 mm) pieces of TMM were introduced to the BMF reservoir to mimic particulate emboli. No movement artefacts occurred.

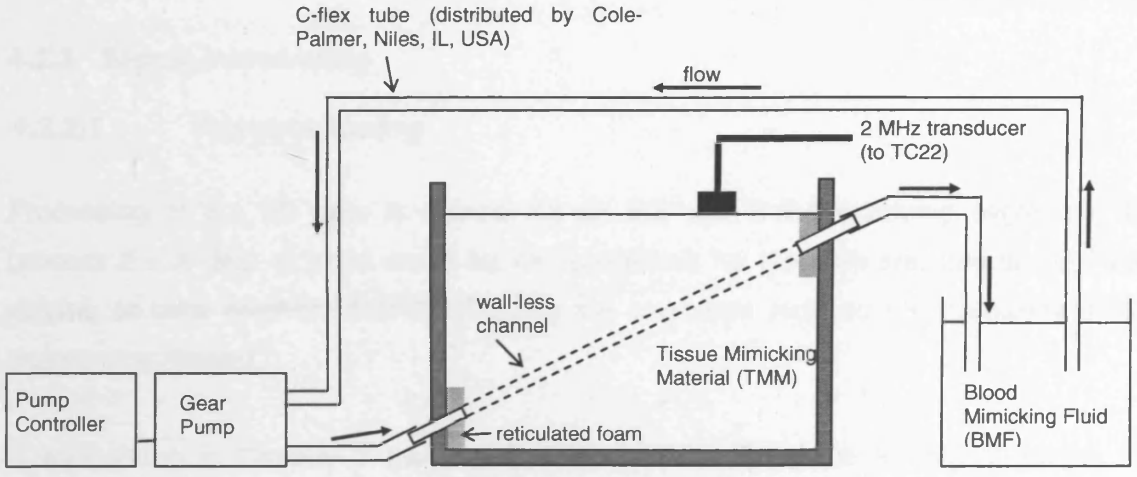

Figure 4-2 Schematic of flow phantom set-up

4.2.2.2 Clinical Data

In vivo data were collected during the routine monitoring of 15 patients undergoing and recovering from (immediately after surgery, for up to 3 h) CEA surgery. These data included presumed artefactual and embolic events (note that no definitive assessment can be made of *in vivo* events – they are labelled artefact or embolus based on human expert classification, which may be inaccurate, so are no more than presumed artefacts or emboli). During the data recordings, the transducer was fixed over the patients' temporal bone windows using a standard headband, ipsilateral to the operated carotid artery, so as to insonate the MCA.

4.2.2.3 Artefacts

Four presumed healthy volunteers generated typically encountered artefacts through speech, groaning, cough, throat clearing, whistling and tapping the probe. A trigger was forced in the acquisition software to capture the RF signal at the time of these artefacts. The transducer was held in place over the volunteers' temporal bone windows using a standard headband. The choice of which (left or right) MCA to insonate in the healthy volunteers was arbitrary.

Another source of artefact is diathermy (electrocautery). This has quite a distinctive appearance on the sonogram but is an artefact that is known to generate errors in some automated embolic detection systems (Mess et al., 1997; 2000). Signals containing diathermy artefacts were acquired during the dissection phase of several CEAs by forcing a trigger in the acquisition software.

4.2.3 Signal Processing

4.2.3.1 Pre-processing

Processing of the RF data is carried out on 300 ms, half-overlapping segments. To process 2 s of data at once would be unmanageable for the software, due to the large volume of data involved. Half-overlapping the segments reduces the possibility of an event being missed.

It was shown in Chapter 3 that the increase in intensity of the RF signal, during an artefact, occurs over a large range of depths simultaneously when compared to the appearance of an embolus. During an embolic event, the increase in intensity occurs in a number of depth samples corresponding to the transmitted pulse length, in this case

10 mm. The region of the RF matrix processed throughout this chapter was chosen to be between 926 (22 mm) and 3986 (95 mm) depth samples (or R_{max} [eqn (2-2)], whichever was smaller). This region was large enough to be able to distinguish artefacts and emboli, but small enough so as not to be too computationally intensive. Hereafter this extracted region will be referred to as the RF-matrix.

Each column of the RF-matrix is clutter-filtered using a 4^{th} order Butterworth high-pass filter with a cut-off frequency of 75 Hz (this is the same cut-off frequency on the wall-thump filter as set on the TCD machine during monitoring).

To improve the signal-to-noise ratio (SNR), matched filtering is performed on the clutter-filtered RF-matrix. This technique involves cross-correlating a copy of the transmitted pulse with each row of the filtered matrix. The result will contain a peak if any region of the received signal resembles the transmitted pulse. Since noise in the received signal will bear a resemblance to any target signal, the matched filter output will also contain noise (Smith, 1997), but this should be of noticeably lower amplitude than the peak located at an event's position.

Later stages of the processing perform more reliably if both RF-matrix dimensions are similar. Since the slow-time dimension is 300 ms, the RF-matrix is decimated by a factor of 10 (i.e. only every 10^{th} sample is used) in the fast-time dimension, giving approximately 300 samples. The envelope of the decimated array is found by taking the magnitude of the analytic signal (generated from the decimated array and its Hilbert transform). The resulting envelope is then smoothed using a 5 x 5 averaging filter.

Each value in the smoothed array is converted to an amplitude in dB. The amplitudes are calculated with respect to the root-mean-square (rms) noise level of that particular data acquisition. This makes more sense when setting a threshold for peak detection, since signals captured on different occasions will have slightly varying characteristics. To find the noise level, the rms value was calculated for a number of regions between 70 mm from the transducer and the end of the RF matrix. This range was chosen, as it is deep enough to be unlikely to contain signal from the MCA. The median rms value from these regions is taken to be the rms noise value for that RF-matrix. The median value is used instead of the mean, as this is less likely to be influenced by outliers from, for example, regions of artefact. Each value in the RF-matrix is divided by the rms noise value and converted to dB to obtain the new amplitude array.

Local maxima, which have a magnitude greater than a given threshold (denoted by '*ampl_thresh*'), are found in the amplitude array. Two values of *ampl_thresh* have been used in this study, 15 dB and 20 dB (note that this amplitude is not the measured embolus-to-blood ratio (MEBR) from the demodulated Doppler signal; an analysis of how the RF signal magnitude compares with the MEBR is described in the 'Quantifying Performance' subsection). Each local maxima with a magnitude above the threshold constitutes a candidate event.

For each candidate event, a circle of 22 equally spaced points, with radius 2, is centred on the local maxima. Each point in the circle is moved radially out from the centre until it reaches a point with magnitude of at least 10 dB lower than its associated local maxima. The contours described by these points are used to define emboli in the training set or are used to describe a previously unseen candidate event in a test set. In the latter case, further processing is carried out to classify the event as an embolus or an artefact. The processing sequence performed on these contours, either for training or for event classification, is illustrated in a flow diagram in Figure 4-3 with further details given in the respective subsection below.

4.2.3.2 Training

Training Data

The files containing flow phantom embolic events were processed to produce a sonogram and time-based demodulated Doppler signal. Based on these signals, a human expert identified 131 files containing only one embolic event. The start and end times of each event were recorded, together with the event duration. The events were organised in ascending order by their duration. Every third event from the ordered list was selected for the *in vitro* training set (giving 44 events). The remaining 87 events were used in the testing set. This method was used for the selection of training set events, as it was perceived that duration was likely to cause the largest variation in embolus shape within the RF signal.

A similar process was used to generate the *in vivo* training set. Eighty-nine *in vivo* files, each containing one embolic event (classification agreed by the AEDS and a human expert), were organised in ascending order by their duration. Every second event from the ordered list was selected for the training set (giving 44 events). The remaining 45 events were used in the testing set.

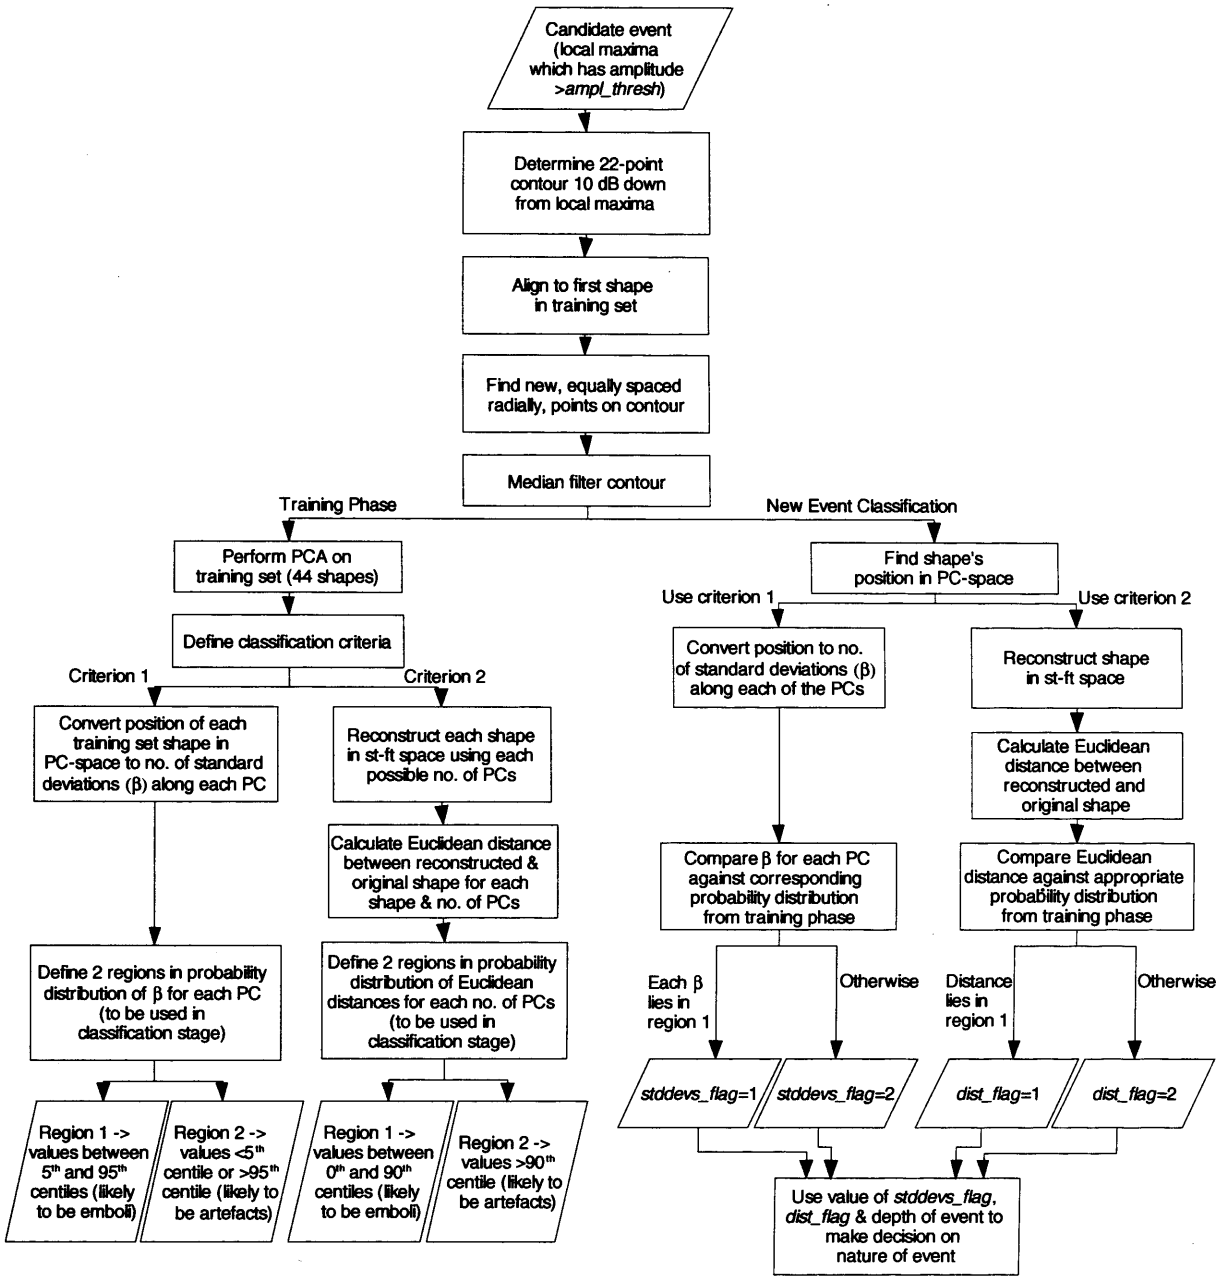

Figure 4-3 Flow diagram illustrating processing performed on event contours either for training or for event classification

Training Method

The co-ordinates of the 22-point embolus contour in the (pre-processed) RF-matrix were extracted for each of the 44 events in a given training set. The geometrical centres (average of the x and y co-ordinates of the 22 points on a contour) of all shapes in the training set were then aligned to the geometrical centre of the first shape. Note that this is an arbitrary choice – the centres could have been aligned to any shape in the training set. This step was necessary to ensure that translation was eliminated as a source of variation between the various shapes. In contrast, Cootes et al. (1992) aligned their training set by scaling, rotating and translating the shapes so that they corresponded to each other as closely as possible. It was not desirable to scale or rotate the emboli training set, since scale and rotation offer important modes of variation that help to define the limits on embolus shape. After alignment, a revised set of 22 points on each of the training set shapes' contours was found. These were equally spaced radially, using the geometrical centre of the aligned shape as a starting point. Since corresponding points across the various shapes were then at the same angle with respect to the centre point, it was possible to describe the points by their radial distance from the centre. This reduced the dimensionality and therefore required less computations than using the x and y (fast-time and slow-time) co-ordinates. Each contour was smoothed using a 3-point median filter; this reduced the chance of outlying points adversely influencing subsequent processing. The mean embolus shape was calculated from the smoothed contours and principal component analysis (PCA) was performed. The methodology for performing PCA is described in the appendix.

Each shape in our data set had 22 points and, therefore, performing PCA produced axes in 22-dimensional space. Each axis corresponds to a principal component (PC) and the new space produced will be referred to as PC-space (to distinguish it from the original slow time-fast time (st-ft) space in which the contours have been defined).

Varying the co-ordinates in PC-space, within suitable limits, can generate new examples of emboli shapes. Since the emboli training shapes' positions were not normally distributed along the PCs, it was not appropriate to fix the allowed variation at $\pm 3\sqrt{\lambda_i}$ (as suggested by Cootes et al. (1992)), where λ_i is the variance of the positions (b_i) of the training shapes in PC-space. Instead, the training sets were used to find more suitable variations in standard deviation for each of the PCs. The position of each training set shape was found in PC-space and was converted into a number of standard deviations along each of the PCs using eqn. (4-1), where β is a 2-D matrix of number of standard deviations, q varies between 1 and N_s, the number of shapes in the training set (44), i

varies between 1 and n, i.e. there is a value for each of the 22 PCs, and b and λ are as defined above (a derivation of each can also be found in the appendix).

$$\beta_{i,q} = \frac{b_{i,q}}{\sqrt{\lambda_i}} \tag{4-1}$$

Therefore, with 44 training set shapes each having 22 points on their contour, β will be a 22 x 44 array. The distribution of β for each PC, across all training set shapes, is examined and used to define two regions to be used in the classification phase; region 1 contains values between the 5th and 95th centile and region 2 is defined as any value less than the 5th centile or greater than the 95th centile. Events falling in region 1 are likely to be emboli and those in region 2 are likely to be artefacts.

A shape can be transformed back from PC-space to st-ft space using eqn (A-6) (see Appendix). If all 22 principal components (PCs) are used, the shape will be reconstructed exactly in st-ft space. The accuracy of the reconstruction decreases as the number of PCs used is decreased. However, it is usually found that the later PCs contribute little to the shape reconstruction and simply add noise and that a reasonable reconstruction can be made using the first few PCs. The reconstruction error can be estimated by calculating the Euclidean distance between the original (translated) and reconstructed shape in st-ft space using eqn (A-7) - the smaller the Euclidean distance, the more accurate the reconstruction. The training set was used to find acceptable values for the Euclidean distance. Each shape in the training set was reconstructed in st-ft space using each possible number of PCs (i.e. each shape was reconstructed 22 times, first using the first PC, then using the first 2 PCs and so on until all 22 PCs were used). The Euclidean distance between the reconstructed shape and the original translated shape was calculated using eqn (A-7), for all combinations of shape number and number of PCs used. Hence for each number of PCs there were 44 values of Euclidean distance. This is illustrated as a matrix in Figure 4-4. The probability distribution of Euclidean distances for each number of PCs (i.e. for each row in the matrix in Figure 4-4) was used to define two regions to be used in the classification stage; region 1 contains values between the 0th and 90th centile and region 2 contains values greater than the 90th centile (note that one set of *in vitro* results were obtained with the region cut-off point set at the 95th centile). Events falling in region 1 are likely to be emboli and those in region 2 are likely to be artefacts.

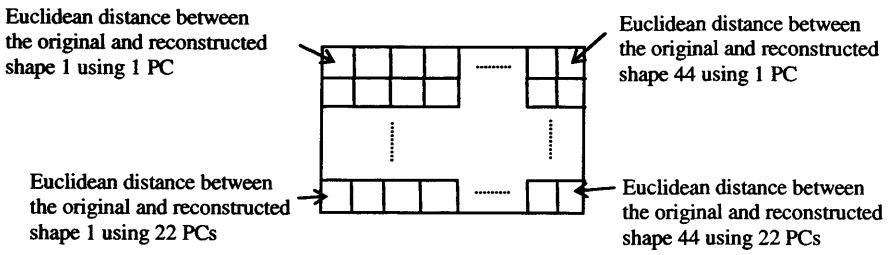

Figure 4-4 Illustration of matrix containing Euclidean distance between original and reconstructed shape for each training set shape and each possible number of PCs

The standard deviation and Euclidean distance parameters are processed in this way, i.e. using the 5th and 95th or the 0th and 90th centiles rather than using the minimum and maximum values, so that the limits are not so heavily influenced by outliers. This should reduce the potential for false positives, thus increasing specificity. The benefits of this greatly outweigh the disadvantages of the expected drop in sensitivity.

4.2.3.3 Event Classification

When a candidate event is identified, it is forwarded for classification as an embolus or artefact. The shape to be classified is translated such that its geometrical centre is aligned with that of the first shape in the training set. After alignment, a revised set of 22 (equally spaced radially) points on the shape's contour is found and the resulting shape is median filtered, as described in the 'Training Method' section. Two shape classification criteria have been developed and these have been tested both separately and combined.

The first criterion considers the shape's position in PC-space. Using eqn (A-5) with the first *t* PCs (where *t* can be any value up to 22 but is usually no greater than 5), the position of the translated shape in PC-space is found. This position is converted to the number of standard deviations along each of the PCs using eqn (4-1), where q=1 since only one shape is being used, and i varies between 1 and *t*. The value of β_i for each PC, is compared against the corresponding distribution of standard deviations from the training set (see 'Training' subsection) and the results are used to set a value for the variable *stddevs_flag*. If each β_i lies in region 1 in their corresponding distribution then *stddevs_flag*=1, otherwise *stddevs_flag*=2.

The second criterion utilises the shape's reconstruction error. Using the first *t* PCs, the shape is reconstructed in st-ft space using eqn (A-6). The Euclidean distance between the reconstructed shape and the original translated shape is calculated using eqn (A-7). This

is compared against the appropriate (to the number of PCs used) Euclidean distance distribution derived from the training set. A parameter *dist_flag* is set depending on the shape's position within this distribution i.e. if it lies within region 1, *dist_flag*=1, otherwise *dist_flag*=2.

The values of the two parameters, *stddevs_flag* and *dist_flag* are used to make a decision on the type of event present. If the appropriate flag to the criterion being used, or both flags if the rules are being combined, is set to 1 *and* the minimum depth of the event is between 3 and 6 cm, the event is classified as an embolus; otherwise the event is classified as an artefact.

Data (Euclidean distance, minimum depth, maximum depth, minimum time, maximum time, magnitude of peak in pre-processed RF-matrix) on emboli and artefacts are saved to separate files. Prior to saving, details of events found in the current 300 ms RF-matrix are checked against those in the file of the relevant event type and any events that have already been saved to that file are not saved again. Times of the saved events are manually checked against the embolus times extracted from the time-based demodulated Doppler signal, to verify the nature of the event.

4.2.4 Quantifying Performance

4.2.4.1 Performance Measures

Performance has been judged on sensitivity, specificity and positive predictive value (PPV). These are calculated using eqns (4-2) to (4-5), where TP is True Positive and FP is False Positive. When considering the system performance through testing on files containing both emboli and artefacts, it is considered that PPV is a better measure of performance than specificity given the high number of non-embolic events above *ampl_thresh* present.

$$\text{sensitivity} = \frac{\text{no. of TPs}}{\text{no. of events above } \textit{ampl\_thresh} \text{ classed as embolic}} \qquad (4\text{-}2)$$

$$\text{overall sensitivity} = \frac{\text{no. of TPs}}{\text{total no. of events present classed as embolic}} \qquad (4\text{-}3)$$

$$\text{specificity} = 1 - \frac{\text{no. of FPs}}{\text{no. of events above } ampl\_thresh \text{ classed as non-embolic}} \quad (4\text{-}4)$$

$$\text{PPV} = \frac{\text{no. of TPs}}{\text{no. of TPs} + \text{no. of FPs}} \quad (4\text{-}5)$$

4.2.4.2 *Comparing Signal Levels in the RF and Demodulated Doppler Signals*

The peak amplitude in the processed RF signal was compared with the measured embolus-to-blood ratio (MEBR) from the demodulated Doppler signal, for each of the 45 emboli in the *in vivo* test set. In each case, the demodulated Doppler signal was derived from the acquired RF signal in the position of the receive gate at the time of the acquisition. As noted in the 'Discussion' section, this is not necessarily the optimal receive gate position. The MEBR was calculated by manually selecting a window in the time-based signal between the start and end of the embolus. Two blood signal windows, one either side of the embolus, each of 100 ms duration, were also manually selected in the time-based signal. For each of the three windows (1 embolus and 2 blood), the mean signal level was removed from each point in that window. The two mean removed blood windows were concatenated to create one blood window. The rms value was found for the mean removed embolus window and the concatenated (mean removed) blood signals. The MEBR (in dB) was found using eqn (4-6).

$$MEBR = 20\log_{10}\left(\frac{\text{rms(embolus window)}}{\text{rms(blood window)}}\right) \quad (4\text{-}6)$$

4.3 Results

4.3.1 *In Vitro* Training Set

The cumulative variances accounted for by a number of the PCs, as a result of performing PCA on the *in vitro* training set, are shown in Table 4-1 as a percentage of the total variance.

The 5th and 95th centile positions on the standard deviation distributions and 90th centile position on the Euclidean distance distributions are shown in Table 4-2 for the first five PCs from the *in vitro* training set.

PCs	Cumulative variance accounted for after doing PCA on the *in vitro* training set (%)	Cumulative variance accounted for after doing PCA on the *in vivo* training set (%)
1	87.84	78.22
2	93.30	87.25
3	96.18	90.89
4	97.75	94.09
5	98.35	96.16
10	99.79	99.38
15	99.96	99.91
20	99.998	99.99
22	100	100

Table 4-1 Cumulative variances associated with various PCs as a result of performing PCA on both the *in vitro* and *in vivo* training sets

	In vitro training set			*In vivo* training set		
PC	5^{th} centile in standard deviation distribution	95^{th} centile in standard deviation distribution	90^{th} centile in Euclidean distance distribution	5^{th} centile in standard deviation distribution	95^{th} centile in standard deviation distribution	90^{th} centile in Euclidean distance distribution
1	-1.13	2.19	11.18	-1.22	1.69	32.03
2	-1.49	1.25	9.59	-2.28	1.21	19.94
3	-0.79	0.75	6.62	-1.38	1.11	18.99
4	-1.02	1.20	5.58	-1.13	1.83	16.16
5	-1.62	1.45	4.21	-1.49	2.26	11.15

Table 4-2 The 5^{th} and 95^{th} centile positions on the standard deviation distributions and 90^{th} centile position on the Euclidean distance distributions for the first five PCs from the *in vitro* and *in vivo* training sets

The positions of the *in vitro* training set shapes, in terms of number of standard deviations along the first four PCs, are shown in Figure 4-5. The position within the training set is identified for the most outlying shapes. The processed RF signals and 'event' contours for two of the outlying shapes (28 and 15) are shown in Figure 4-6. It was noted that the most outlying shapes in the plot of PC1 positions versus PC2 positions were located at the extremities of the distribution of durations of emboli in the training set.

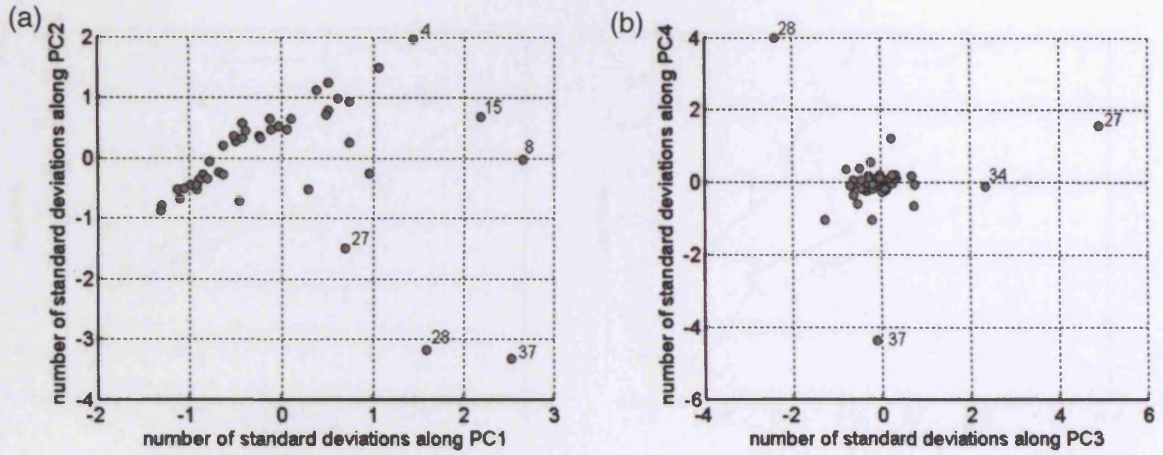

Figure 4-5 Positions of the *in vitro* training set shapes in terms of number of standard deviations along (a) PC1 and PC2, and (b) PC3 and PC4. The position within the training set is identified for the most outlying shapes.

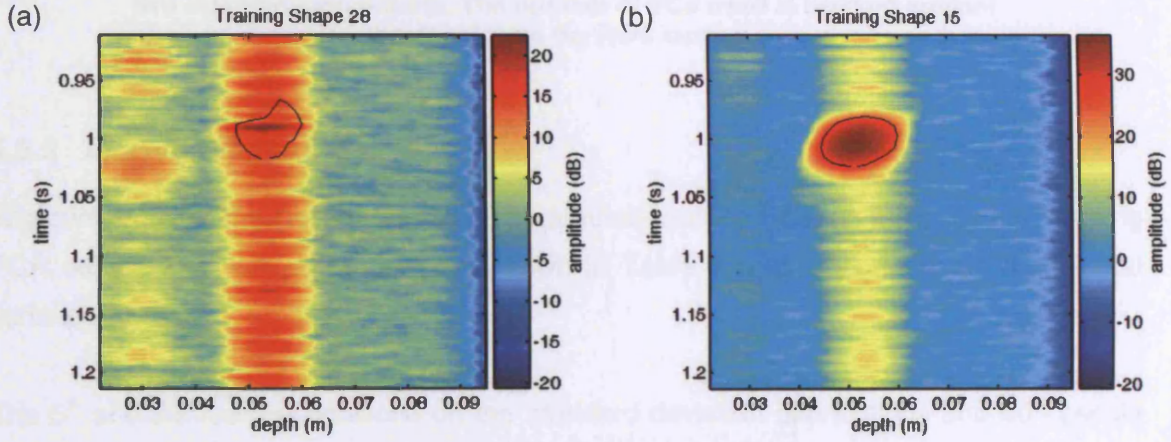

Figure 4-6 Processed RF signal and 'event' contour for *in vitro* training set shape (a) 28 and (b) 15

The results of testing the automated detection/classification algorithms on a set of 87 previously unseen (to the software) putative *in vitro* 'embolic' events are shown in Figure 4-7. Results are shown using both of the classification criteria separately and also using them combined. An ROC curve is also shown using the reconstruction error criterion with a 95th centile cut-off between the two regions. There were 291 candidate events present when using an *ampl_thresh* of 20 dB. Figure 4-7a shows conventional ROC curves using sensitivity and specificity. Figure 4-7b uses PPV instead of specificity. Note that, for both figures, the scale on each axis does *not* extend between 0 and 1. The ROC curves were generated by varying the number of PCs used. The number of PCs used is marked against each point on the ROC curves.

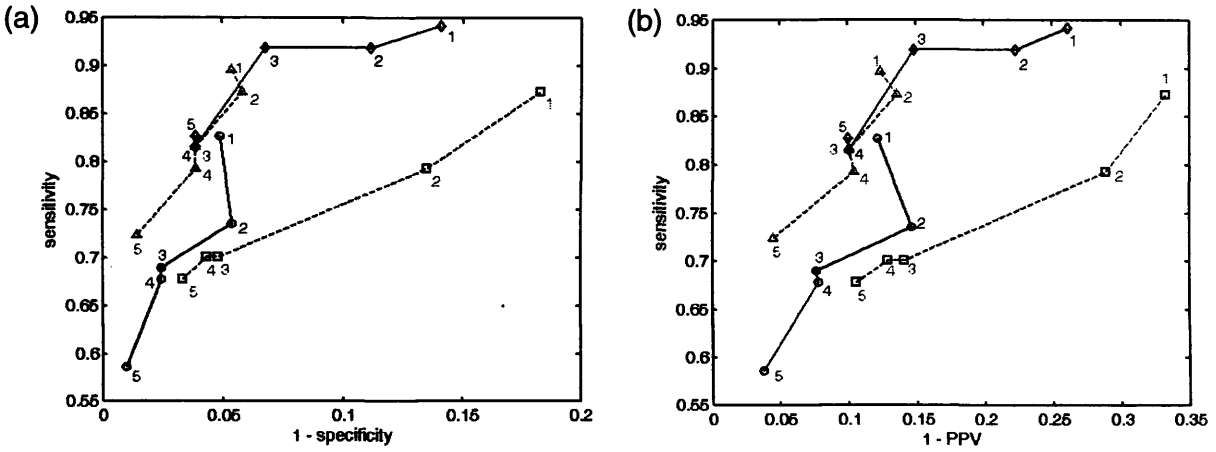

Figure 4-7 ROC curves showing results from testing detection/classification algorithms on 87 previously unseen *in vitro* 'emboli' using (a) specificity and (b) PPV, where □ are results using the position in PC-space criterion only, Δ are results using the reconstruction error criterion only (with a 90th centile cut-off) ◇are results using the reconstruction error criterion only (with a 95th centile cut-off) and, ○ are results using a combination of the two classification criteria. The number of PCs used is marked against each point on the ROC curves.

4.3.2 *In Vivo* Training Set

The cumulative variances accounted for by a number of the PCs, as a result of performing PCA on the *in vivo* training set, are shown in Table 4-1 as a percentage of the total variance.

The 5th and 95th centile positions on the standard deviation distributions and 90th centile position on the Euclidean distance distributions are shown in Table 4-2 for the first five PCs from the *in vivo* training set.

The positions of the *in vivo* training set shapes, in terms of number of standard deviations along the first four PCs, are shown in Figure 4-8. The position within the training set is identified for the most outlying shapes. The processed RF signals and 'event' contours for two of the outlying shapes (32 and 38) are shown in Figure 4-9. Similar to the *in vitro* case, it was noted that the most outlying shapes were all located at the extremities of the distribution of durations of emboli in the training set.

The mapping of peak RF amplitude to MEBR for the emboli in the *in vivo* test set is shown in Figure 4-10. The 1st order least squares best-fit line is shown as a solid line and the 95% confidence intervals are shown as dash-dot lines.

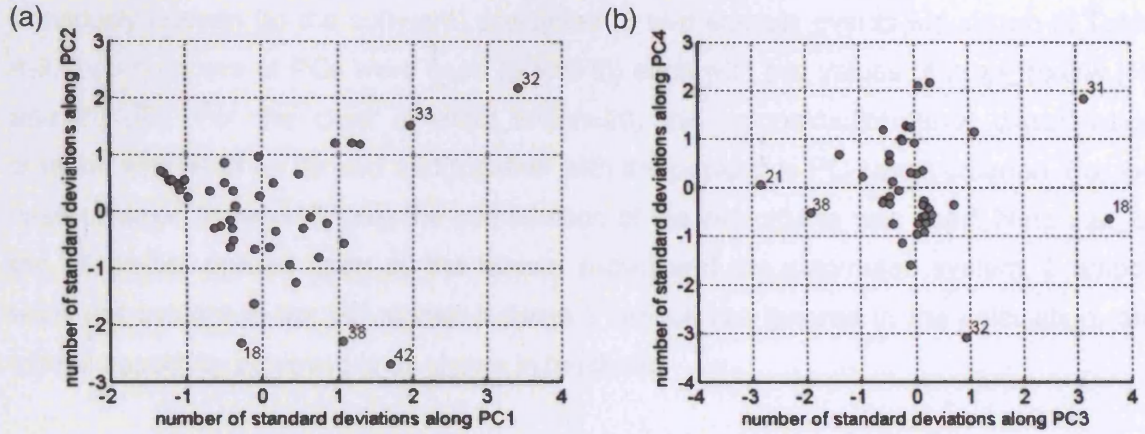

Figure 4-8 Positions of the *in vivo* training set shapes in terms of number of standard deviations along (a) PC1 and PC2, and (b) PC3 and PC4. The position within the training set is identified for the most outlying shapes.

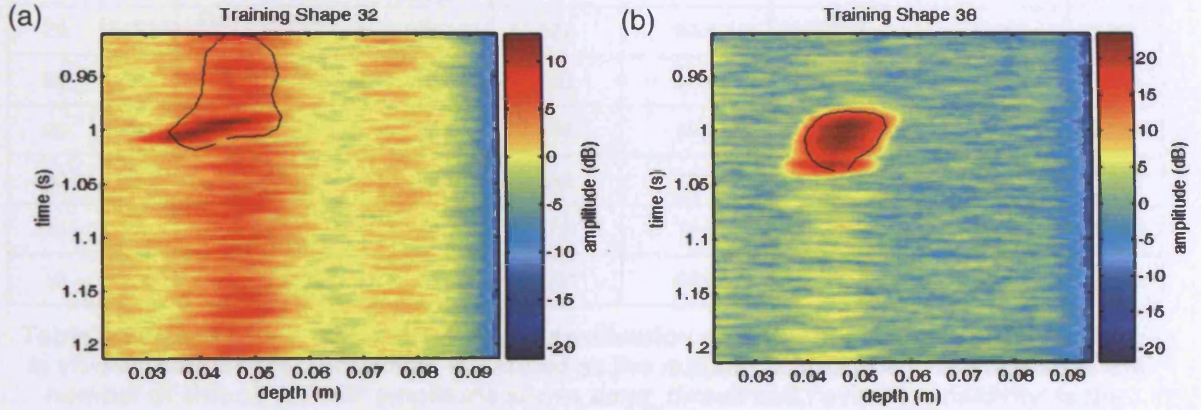

Figure 4-9 Processed RF signal and 'event' contour for *in vivo* training set shape (a) 32 and (b) 38

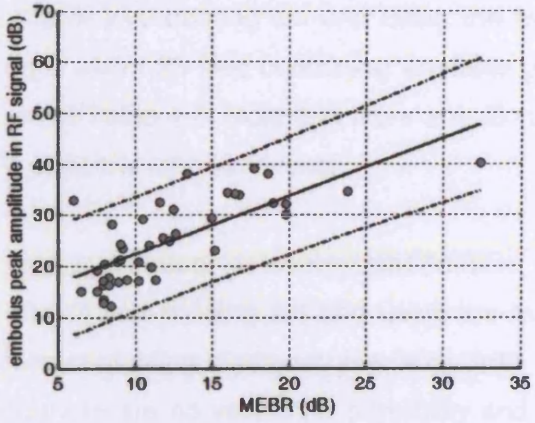

Figure 4-10 Mapping of peak embolus amplitude in the processed RF signal to the MEBR from the time-based demodulated Doppler signal, for emboli in the *in vivo* test set (— gives the 1st order least squares best-fit line, -·-·- shows the 95% confidence intervals)

The results of testing the automated detection/classification algorithms on a set of 45 previously unseen (to the software) presumed *in vivo* embolic events are shown in Table 4-3. Two numbers of PCs were used (3 and 5), each with two values of *ampl_thresh* (15 and 20 dB). For the case of *ampl_thresh*=20, the reconstruction error classification criterion was used on its own and together with the position in PC-space criterion. For the case of *ampl_thresh*=15, only the combination of the two criteria was used. Note that, of the 45 emboli agreed upon by the human expert and the automated system, 2 emboli were not evident in the RF signal. If these 2 emboli are ignored in the calculation, the overall sensitivity achieved is as shown in brackets.

ampl_ thresh (dB)	**No. of PCs**	**Classification criterion used**	**No. of emboli with RF amplitude above *ampl_ thresh***	**No. of candidate events**	**Sensitivity (%)**	**Overall Sensitivity (%)**	**Specificity (%)**	**PPV**
20	3	Reconstruction error	28	582	96.4	60.0 (62.8)	95.0	0.491
20	5	Reconstruction error	28	582	96.4	60 (62.8)	95.4	0.509
20	3	Two combined	28	582	89.3	55.6 (58.1)	96.8	0.581
20	5	Two combined	28	582	75.0	46.7 (48.8)	97.5	0.600
15	3	Two combined	40	1731	75.0	66.7 (69.8)	98.6	0.556
15	5	Two combined	40	1731	60.0	53.3 (55.8)	99.1	0.600

Table 4-3 Results from testing detection/classification algorithms on 45 previously unseen *in vivo* emboli, where 'sensitivity' is defined as the number of true positives relative to the number of emboli with RF amplitude above *ampl_thresh* and, 'overall sensitivity' is the number of true positives relative to the total number of emboli present

The performance of the automated detection/classification algorithms (using the parameters derived from the *in vivo* training set and using the two classification criteria combined), when tested on a set of 23 files containing artefacts generated by presumed healthy volunteers, is shown in Table 4-4. Note that there are no values for sensitivity and PPV as there were no TPs or actual embolic events.

The results of testing the automated detection/classification algorithms (using the parameters derived from the *in vivo* training set and using the two classification criteria combined) on a set of 9 files containing diathermy artefacts (with *no* emboli present) are shown in Table 4-5. Again, there are no values for sensitivity and PPV as there were no TPs or actual embolic events.

ampl_thresh (dB)	No. of PCs	Classification criterion used	No. of candidate events	Sensitivity (%)	Specificity (%)	PPV
20	3	Two combined	800	N/A	99.1	N/A
20	5	Two combined	800	N/A	99.4	N/A
15	3	Two combined	1495	N/A	99.3	N/A
15	5	Two combined	1495	N/A	99.6	N/A

Table 4-4 Results from testing the automated detection/classification algorithms on a set of 23 files containing artefacts generated by presumed healthy volunteers

ampl_thresh (dB)	No. of PCs	Classification criterion used	No. of candidate events	Sensitivity (%)	Specificity (%)	PPV
20	3	Two combined	0	N/A	N/A	N/A
20	5	Two combined	0	N/A	N/A	N/A
15	3	Two combined	73	N/A	89.0	N/A
15	5	Two combined	73	N/A	95.9	N/A

Table 4-5 Results from testing the automated detection/classification algorithms on a set of 9 files containing diathermy artefacts (with no emboli present)

4.4 Discussion

It has been shown that, using PCA, it is possible to characterize the variation in contour shapes defining emboli in the RF signal. *In vitro* emboli are clearer in the RF signal than *in vivo* emboli; this leads to more clearly defined contours and the ability to use fewer principal components to adequately define the possible variation in the contours' shapes. This is evident in the cumulative variances resultant from applying PCA to the two training sets, as shown in Table 4-1.

The 5^{th} and 95^{th} centile positions on the standard deviation distributions for the first 5 PCs (given in Table 4-2) illustrate that it is not suitable, in this case, to have the same allowable variation of $\pm 3\sqrt{\lambda}$ for each of the PCs. If a fixed value was used and this was greater than that obtained from one of the training set distributions, it is likely that there would be an increase in the number of true and false positives. Conversely, if the fixed value chosen was below that obtained from a training set distribution, the number of true and false positives are likely to decrease.

Table 4-2 also shows that as the number of PCs used is increased, the Euclidean distance between the original shape and reconstructed shape decreases. This is as

expected as the use of more PCs should lead to a more accurate reconstruction of the shapes. It is also clear that, for a given number of PCs, the reconstruction of the *in vitro* shapes is more accurate than the *in vivo* shapes. This is also what would be expected given that the cumulative variance accounted for by a given number of PCs is less for the *in vivo* data than the *in vitro* data.

The plots showing the training sets shapes' positions along the first 4 PCs (Figure 4-5 and Figure 4-8) illustrate that the *in vivo* shapes' position distribution is less concentrated into one region than that of the *in vitro* data set. This reflects the higher variability of *in vivo* embolus contour shapes observed in the RF plots. In both the *in vitro* and the *in vivo* cases, the most outlying shapes either had short or long durations with respect to the majority of the corresponding training set. This coincides with the hypothesis that one of the largest contributors to the variance between emboli contours would be the duration.

The mapping of peak embolus amplitude in the processed RF signal to MEBR for each of the emboli in the *in vivo* test set (Figure 4-10) shows that the expected MEBR for an *ampl_thresh* of 20dB is approximately 8dB. This is quite low and not dissimilar to MEBR thresholds used by other research groups, making the use of a 20 dB *ampl_thresh* a sensible choice. Using a 20 dB threshold was also preferable over using a 15 dB threshold given the number of candidate events present in each case; the lower the number of candidate events, the shorter the processing time. The variability in the mapping of RF amplitude to MEBR can be accounted for by two factors. The first is the variability in the measurement of the MEBR which is highly dependent on the location of the windows used to do the calculation. Secondly, the location of the receive gate may not necessarily incorporate the peak in the RF signal, resulting in a lower than expected MEBR for a given RF amplitude. An example of this is shown in Figure 4-11. The high-pass filtered RF signal in the region of the embolus is shown in Figure 4-11a, with the position of the receive gate indicated by dashed lines. The processed RF signal over the same region is shown in Figure 4-11b. It should be noted that, although the receive gate position is also shown in Figure 4-11b, the matched filtering causes a shift to the left of the peak of the signal which may lead a false belief that the embolus peak lies within the receive gate. It is clear that at the time of peak amplitude in the processed RF signal (between 0.95 and 0.96 s), the signal from the embolus (shown in Figure 4-11a) is only fractionally within the receive gate. This results in a low MEBR (6.03 dB) as shown by the demodulated Doppler signal in Figure 4-11c. If the receive gate is moved such that it lies between 5.5 cm and 6.5 cm, the MEBR is much higher (11.19 dB) as illustrated in Figure 4-11d. Therefore, the MEBR may not be calculated from optimum receive gate position, resulting in a lower than expected MEBR for a given peak amplitude in the RF signal.

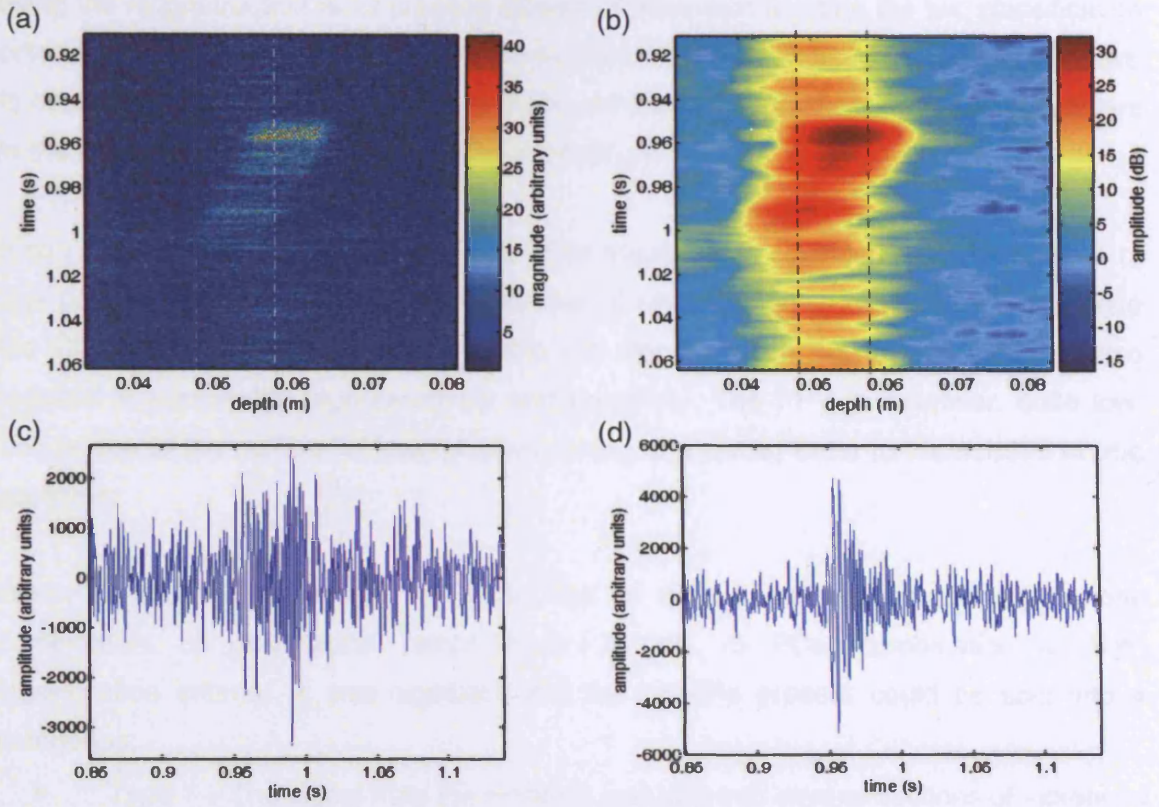

Figure 4-11 (a) High-pass filtered RF signal in region of signal from embolus with position of receive gate indicated by dashed lines, (b) fully processed RF signal in region of signal from embolus with position of receive gate indicated by dashed lines, (c) demodulated Doppler signal in region of embolus with receive gate positioned between 4.8 and 5.8 cm (MEBR = 6.03 dB) and, (d) demodulated Doppler signal in region of embolus with receive gate positioned between 5.5 and 6.5 cm (MEBR = 11.19 dB)

Testing the automated classification algorithm on an *in vitro* test set (see Figure 4-7), resulted in reasonably high sensitivity, specificity and PPV values. However, the events recorded *in vitro* are unlikely to show as high variability as those *in vivo* and it was therefore also necessary to test the detection/classification algorithms on *in vivo* data. It can be seen from the *in vitro* results that the choice of classification criterion involves a trade-off between sensitivity and specificity; using either of the two criteria on their own gives a higher sensitivity but lower specificity and PPV than using the two combined, for a given number of PCs. Relative to the combination of the two classification criteria, the reconstruction error criterion gives a larger increase in sensitivity with a smaller drop in sensitivity when compared with the position in PC-space criterion. When analysing the two curves generated using the reconstruction error criterion alone, it can be seen that widening region 1 in the training set distribution curve from the 90th to 95th centile results in an increase in sensitivity at the cost of a drop in specificity and PPV, for a given number of PCs. This was as expected.

Table 4-3 also shows an increase in sensitivity with a drop in specificity and PPV when using the reconstruction error criterion alone in comparison to using the two classification criteria combined, when tested on the *in vivo* data. The user of such a system would have to decide which was more important to them, sensitivity or specificity, and the parameters in the algorithms could then be adjusted accordingly.

It can also be seen from Table 4-3 that, when sensitivity is considered as the number of true positives as a percentage of the number of embolic signals present that are above the RF amplitude threshold used, testing the algorithms on the *in vivo* data set also resulted in reasonably high sensitivity and specificity. The PPV is, however, quite low. This is due to the number of false positives being of a similar order to the number of true positives.

When the results from testing the algorithms on *in vivo* data were examined, for one combination of parameters (*ampl_thresh* = 20 dB, 5 PCs, combination of both classification criteria), it was apparent that the 14 FPs present could be split into 4 categories:

- Type 1 – The signal from the embolus was split into several sections of increased intensity in the RF signal, resulting in it being classified as several (2 or more) emboli but visual inspection of the RF signal indicates that only one embolus was present. This was the source of 7 of the 14 FPs in the case examined. An example of this FP type is shown in Figure 4-12a, where the embolus was incorrectly classified as 2 emboli. Two possible reasons for the region of increased intensity being split into several sections are the presence of constructive and destructive 'interference' of signals returned from the embolus and, the embolus may be moving between areas of different beam sensitivity.
- Type 2 – False positives as a result of using half overlapping segments. An example of this FP type is shown in Figure 4-12b and Figure 4-12c. An FP occurs in the segment shown in Figure 4-12b. However, when the next, half overlapping, segment is examined (see Figure 4-12c) the same event is unlikely to be considered to be an embolus. This FP type has been seen to occur both at the start and end of segments and the algorithms could be modified in the future to better account for these. Four of the 14 FPs in the case examined were due to this FP type.
- Type 3 – An increase in intensity in the RF signal which may be from the blood or from a very short duration embolus that was not identified in the sonogram or time-based signal. This is analogous to a false positive arising from Doppler

speckle in a conventional, demodulated Doppler signal based system. An example of this FP type is shown in Figure 4-12d. It is impossible to say what the source of this event was but since it was not identified in the demodulated Doppler signal it has to be classified as an FP. There were 2 FPs of this type in the group of 14 FPs examined.

- Type 4 – False positives that are definitely part of an artefact signal. There was only 1 FP of this type in the 14 FPs examined. This is shown in Figure 4-12e. On visual inspection the FP is clearly part of an artefact signal but it is difficult to include logic in the algorithms to distinguish this from the situation where multiple emboli are occurring at the same time at different depths (as shown in Chapter 3).

Therefore, even although 14 FPs seems quite high, the nature of some of the events is such that the algorithms can be developed to eliminate them (e.g. type 1 and type 2) and some are impossible to determine if they are due to emboli or not (type 3). Incorporation of these observations, to reduce the number of FPs and hence improve the performance of the system, is planned.

Each training set consisted of 44 embolic events. It is difficult to comment on the effects of increasing the size of the training sets. It may be that if a larger number of events were used this would give a larger variation of emboli shapes. This in turn may lead to more emboli being classified but may also lead to more false positives. Alternatively, if the additional shapes had similar characteristics to those already in the training set, this would make the contours better defined by a smaller number of PCs which may lead to lower sensitivity but could increase the PPV.

The test sets were relatively small but the results from this pilot study demonstrate that, whilst the sensitivity and PPV are not yet adequate for use in clinical practice, the method of using PCA on event contours in the RF signal has potential and further development should be considered.

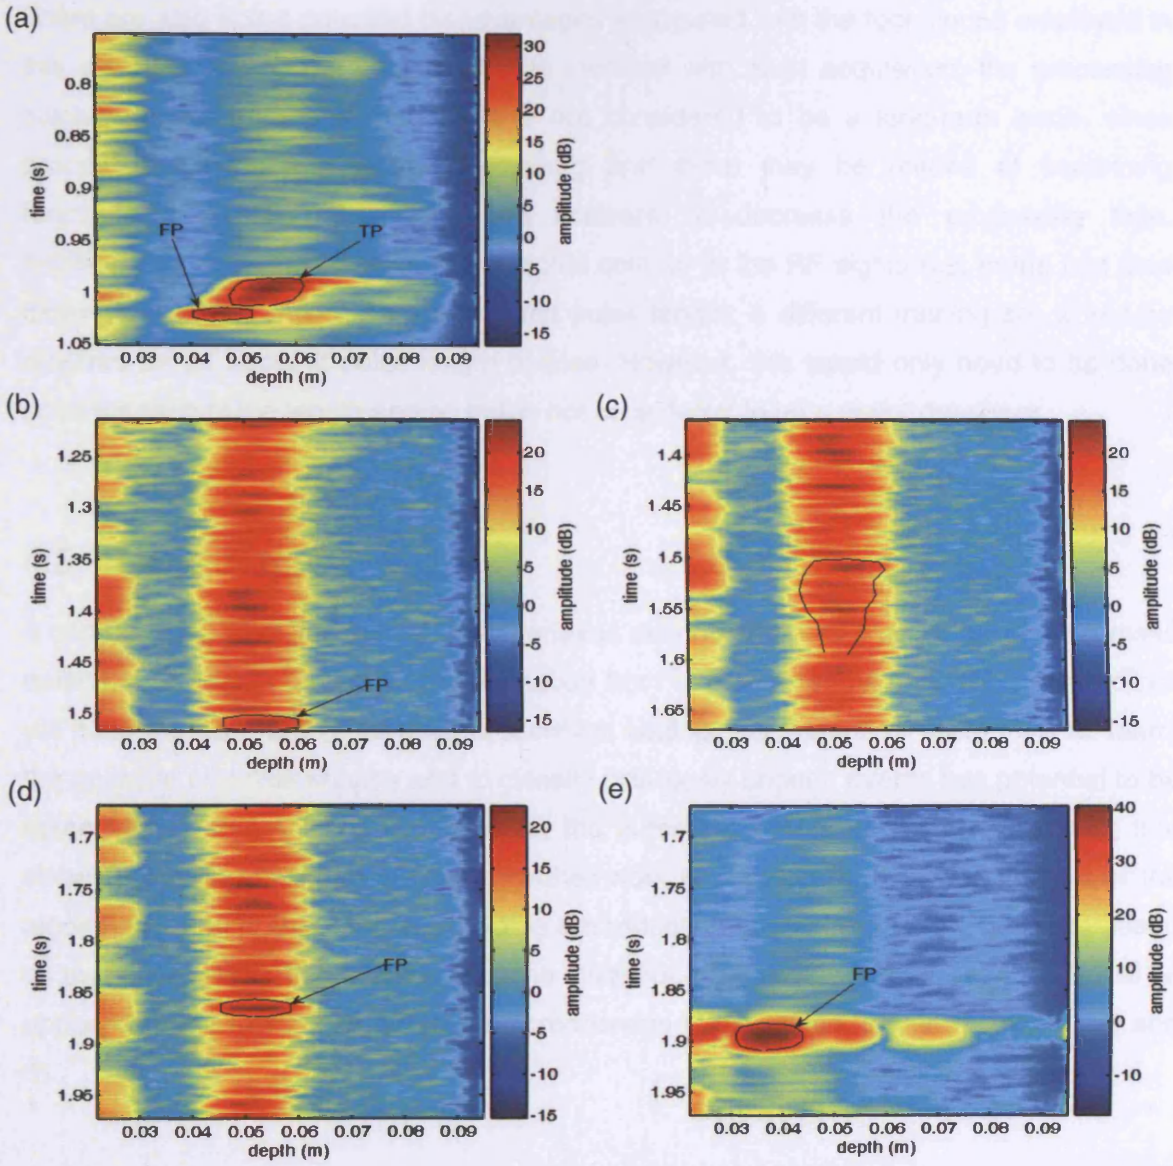

Figure 4-12 (a) Example of a 'Type 1' FP - The embolus signal is split into several regions of increased intensity resulting in its classification as 2 emboli. (b) Example of a 'Type 2' FP - An FP occurs at the end of the segment. When the next, half overlapped, segment is examined in (c) the same event is not classified as an embolus. (d) Example of a 'Type 3' FP - This increase in intensity in the RF signal may be from the blood or from a very short duration embolus that was not identified in the sonogram or time-based signal. (e) Example of a 'Type 4' FP - On visual inspection this is clearly part of an artefact signal.

The embolus detection method discussed in this chapter has the advantage that it is not dependent on signal characteristics such as frequency. Many automatic embolus detection systems use event frequency as a criterion for discriminating between emboli and artefacts but it has been shown (Chapter 3) that signals from emboli may have frequencies lower than that traditionally associated with these events, which may lead to their false classification as artefacts.

There are also some potential disadvantages associated with the techniques employed in this chapter. Due to the volume of data involved with each acquisition, the processing involved is far from real time. This is not considered to be a long-term issue, since processor speeds are always increasing and there may be means of combining functionality between hardware and software to decrease the processing time. Additionally, since the width of the embolus contour in the RF signal (i.e. in the fast time dimension) is dependent on the transmit pulse length; a different training set would be required for all transmit pulse length options. However, this would only need to be done once for each pulse length and so this is not considered to be a major drawback.

4.5 Conclusion

It can be concluded that, whilst the algorithms described in this chapter for the automatic detection of emboli (and their discrimination from artefacts) in the TCD RF signal do not yet have the accuracy to be used in a clinical setting, the method of using PCA to 'learn' the possible embolus shapes and to classify previously unseen events has potential to be taken further. It is a novel technique for the automatic detection of embolic signals that shows promising levels of success. Further work is planned for the improvement of the algorithms used in this chapter including the reduction in number of false positives based on the observations in this study, and the utilization of the pulse compressed RF signal as opposed to the raw RF signal (pulse compression is discussed further in Chapters 5 and 6).

5 Coded Excitation in TCD Ultrasound Systems to Improve Axial Resolution - Theoretical Considerations

5.1 Introduction

In previous chapters, it has been demonstrated that using the TCD RF signal may be a more reliable means of discriminating emboli and artefacts compared to the conventional subject of processing, the demodulated Doppler signal. Displaying magnitude plots of the RF signal allows us to track the motion of an embolus, but the technique is seriously limited by the narrow bandwidth of the transmitted pulse which leads to very poor axial resolution. When working with the RF signal, the axial resolution is defined by the transmitted pulse length (since no receive gate is used). Long transmit pulses are required to obtain an adequate signal-to-noise ratio (SNR) to locate and maintain a signal from blood flow in the middle cerebral artery (MCA). In our centre, pulses of approximately 13 μs are used, giving an axial resolution of 10 mm. This resolution is poor given the region of the MCA usually examined has length up to 20 mm and a diameter of approximately 3 mm. This has highlighted the need to improve the axial resolution, in order to track the progress of emboli as they propagate along the MCA and also to discriminate between closely spaced events.

To improve the resolution, the bandwidth of the transmitted pulse needs to be increased. This could be achieved by decreasing the transmitted pulse length. In doing so, the pressure amplitude over the duration of the pulse would need to be increased so as to achieve an adequate SNR. Increasing the pressure amplitude would have an adverse effect on the system safety, in particular on the mechanical index (MI). For example, the maximum MI from a standard TCD machine (SciMed TC22, Bristol, UK) was measured in our centre as 0.37, when using 13 μs duration pulses with a 2 MHz probe. If the pulse length were reduced by a factor of 5 to achieve a range resolution of 2 mm, it would be necessary to increase the transmitted pulse amplitude by a factor of 5 to maintain the same time-averaged power. This would also result in the maximum MI increasing by a factor of 5 to 1.85. This is very close the FDA's 'global maximum MI' limit of 1.9 (FDA, 1997) and it is, therefore, undesirable to operate with such high pulse amplitudes.

In addition to operating close to the FDA's maximum MI limit, increasing the pulse amplitude would have several other undesirable effects. Doubling the measured TCD MI

(0.37) would result in a value larger than the threshold MI (0.7) at which it is considered there may be a risk of cavitation (British Medical Ultrasound Society, 2000). This may be an issue if there is a risk of gaseous emboli passing through the sample volume. There is a similar danger of bubble destruction if contrast agents are being used and high pulse amplitudes are employed. In addition, non-linear effects will worsen as the pulse amplitude is increased. Therefore, any method that can improve the axial resolution without increasing the pulse amplitude would be desirable.

An alternative approach to improving the axial resolution is to use pulse compression. Pulse compression is achieved by applying frequency or phase modulation (FM or PM) to the transmitted signal to increase the bandwidth (amplitude modulation (AM) is also possible, but this is rarely used), and subsequently performing matched filtering on the received signal. The application of frequency or phase modulation to the transmitted signal is also known as coded excitation.

Pulse compression provides a resolution equivalent to using a non-coded, shorter pulse and also provides an improvement in the SNR. Radar and imaging ultrasound system designers usually wish to increase the total output power without lengthening the transmit pulse. This implies an increase in pressure amplitude which is also not desirable. Instead, they increase the pulse length and use pulse compression to recover the resolution previously obtainable. Conversely, for this application, it is desirable to decrease the effective pulse length without decreasing the overall output power or increasing the maximum pressure amplitude. If the length and amplitude of the transmit pulse are unaltered, pulse compression could be used to reduce the effective pulse length, thus improving the resolution.

Several other centres have considered the use of coded excitation in various ultrasound applications. Wilhjelm and Pedersen (1993a; 1993b; 1996) proposed a method for velocity estimation with a Doppler system transmitting linear FM signals. This had the disadvantage that it required separate transmit and receive transducers since the periods of transmit and receive overlapped, and also had a range resolution of the same order as conventional systems. Recently, Iwashita et al. (2003) and Moriya et al. (2005) employed a 'dual-chirp pulse' or 'complementary chirp signal' (i.e. a pulse which contains chirps with increasing frequency modulation and chirps with decreasing frequency modulation) to measure an ultrasound Doppler shift. Their system exploited the range-Doppler coupling effect (see the 'Theoretical Background' Section for definition) and used long pulses (100 μs or 1000 μs respectively). It is not clear from the text if the probe used in these systems had separate transmit and receive crystals. However, if a single crystal was used, the

minimum detectable depth would be large (77 mm or 770 mm respectively) due to the transmit pulse length and such a system would not be practical for use in TCD. Behar and Adam (2004), Pedersen et al. (2002; 2003), and Misaridis et al. (1999; 2000a; 2000b) evaluated the use of FM signals, together with matched and mismatched filtering, in imaging systems. They concluded that the higher SNR offered using coded excitation improves image quality and provides increased penetration in medical ultrasound imaging. Instead, the improved SNR can be used to allow imaging at higher frequencies, improving spatial resolution without loss of penetration. Bennett et al. (2004) considered the use of the fractional Fourier transform instead of matched filtering together with linear FM transmit signals. They found they could achieve similar, but not quite as good, results as when matched filtering was used but with the benefit that no *a priori* knowledge was required of the transmit signal. Nowicki et al. (2003) investigated the use of Golay codes (binary/phase modulated codes) for signal compression in diagnostic ultrasound imaging. They improved the SNR of their ultrasonographic images by a factor of three. They also imaged to 15 cm with no sacrifice in axial or contrast resolution and improved the contrast resolution of shallower structures. Chiao et al. (2005) compared the performances of bi-phase and chirp coding. In particular they quantified tradeoffs between penetration and resolution under frequency dependent attenuation, dynamic focusing, and nonlinear propagation and also compared the performance of the two forms of code with respect to image quality and system requirements. They concluded that both codes have a similar range sidelobe level (RSLL) due to nonlinear propagation and also that frequency-dependent attenuation helps to improve the RSLL. They found that Golay codes were more affected by dynamic focusing and motion than chirps but they have suggested means by which motion degradation can be reduced. This is also covered in more detail in Chiao (2001).

To date, to our knowledge, no one has considered the use of a single transducer TCD system using coded (either chirped or binary coded) transmit signals. Preliminary tests in our centre have shown that standard TCD transducers are likely to have too narrow a bandwidth to cope with phase modulation coded excitation (TCD transducers are generally designed to have narrow bandwidth for improved sensitivity). However, the bandwidth of some commercially available TCD transducers is sufficient to perform frequency modulation. As a result, frequency modulation coded excitation will be the focus of this study. It is proposed that such a system, when combined with a matched filtering/pulse compression scheme, could improve the axial resolution without decreasing the pulse length or increasing the pulse amplitude. It would also still be possible to process the returned signal to extract the conventional Doppler information thus not detracting from the requirements of a standard TCD system.

The aim of the first part of this study was to construct a mathematical model which would allow us to investigate the viability of using coded excitation in a TCD system. In particular, this model was used to assess a method of calculating the axial resolution from the envelope of the pulse-compressed signal versus the theoretical resolution. It was also used to establish if the potential side effects of pulse compression would be an issue in a TCD system when considering the likely pulse lengths and chirp bandwidths to be used. Experimental evaluation of a TCD system employing coded excitation is detailed in a subsequent chapter.

5.1.1 Theoretical Background

Matched filtering is carried out by performing a cross-correlation between each received signal and a copy of the transmitted signal. If the transmitted signal is coded e.g. contains frequency modulation, then the matched filtering process is also known as pulse compression. The waveform resulting from the cross-correlation of two chirps approximates a sinc pulse (the accuracy of the approximation is dependant on factors such as the pulse length, how rectangular the chirp spectrum is, how closely matched the two chirps are etc.). In addition to a high amplitude central peak (mainlobe), the sinc waveform contains peaks of diminishing amplitude extending out to $\pm\infty$; these are known as sidelobes. The mainlobe of a pulse-compressed signal is much narrower than the peak in the match-filtered signal when conventional non-coded pulses are used, or the peak in the received signal when no matched-filtering is used. Therefore, since the mainlobe is very narrow, 'targets' that are in close proximity to each other can be separately identified, when using pulse compression.

However, pulse compression does have some inherent potential side effects. The first of these is the sidelobes (also known as 'range sidelobes') present in the pulse-compressed signal. If the sidelobes are large they may be mistaken for an additional 'target' or they may mask the presence of a small nearby target. It is desirable to increase the ratio of mainlobe to sidelobe amplitude so that these problems are minimised. One method of decreasing the sidelobe level with respect to the mainlobe level is to apply a window to the copy of the transmit burst before performing the cross-correlation. A common choice of window for this application is the Dolph-Chebyshev window since this offers good compromise between reduction in sidelobe level and resulting SNR loss and increase in mainlobe width. This is the window used in this study. An illustration of a Dolph-Chebyshev window, in both the time and frequency domains, is given in Figure 5-1. A Dolph-Chebyshev window is defined by its duration and Fourier transform sidelobe

magnitude (expressed as dB below the mainlobe magnitude). The window shown in Figure 5-1 has the same duration (τ_{tx}) as the transmitted pulse, and a Fourier transform sidelobe magnitude of 60 dB; the sampling frequency, f_s, was 32.5 MHz. Note that the frequency axis in Figure 5-1b has a log scale.

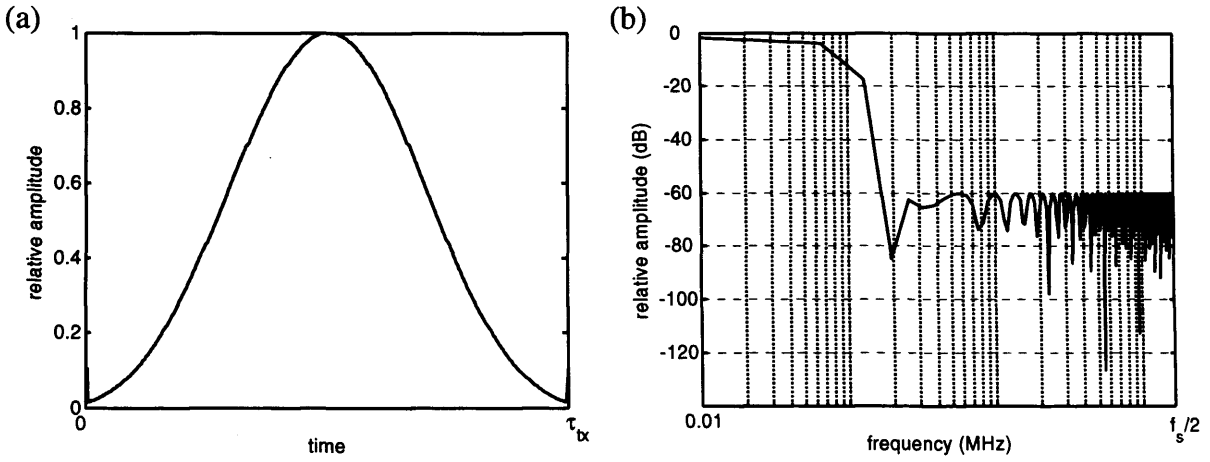

Figure 5-1 Illustration of a Dolph-Chebyshev window with 60 dB sidelobe magnitude in (a) the time domain and (b) the frequency domain

Another potential pulse compression side effect is range-Doppler coupling. When a target is moving, a 'true' Doppler shift is introduced between the transmitted and received pulses. This results in the two spectra being misaligned. When matched filtering is performed on these misaligned spectra, a peak occurs in the pulse-compressed signal that is not at the true target range. This is known as range-Doppler coupling.

Frequency dependent attenuation and non-linear propagation may also be issues since the transmitted pulse will contain a range of frequencies. If these frequency components are affected in different ways, this will result in the matched filtering process being less effective.

5.2 Methods

5.2.1 Range Resolution

As stated above, the waveform resulting from the cross-correlation of two chirps approximates a sinc pulse. It is usually considered that the usable portion of the sinc pulse is half the null-to-null points of the mainlobe. The null-to-null distance is $2/B_{tx}$, where B_{tx} is the bandwidth of the transmitted chirp. Therefore, the usable portion of the pulse-compressed signal would be the inverse of B_{tx} (Skolnik, 1981, Pedersen et al., 2003).

Converting to units of distance, this results in a range resolution, R_{res}, given by eqn (5-1) (Mahafza, 2000), where c is the speed of sound (taken to be 1540 m s^{-1} for soft tissue).

$$R_{res\_theoretical} = \frac{c}{2B_{tx}} \tag{5-1}$$

Note that eqn (5-1) has the same form as that used to calculate the resolution when a conventional, non-coded pulse is used, where the bandwidth would be equal to the inverse of the duration of the transmitted pulse. When a chirp is used, the duration of the pulse has no bearing on the range resolution; instead the bandwidth is the difference between the minimum and maximum frequencies present within the chirp.

It is not always practicable to use half the null-to-null distance of the mainlobe as these points may not be easy to determine. Instead, in this study, the range resolution was calculated as the width of the pulse-compressed mainlobe in the signal envelope at the point where the amplitude of the mainlobe is half (-6dB) of the maximum (see Figure 5-2). Range resolution is measured in units of distance. This can be calculated from the number of samples (as illustrated in Figure 5-2) using eqn (5-2), where f_s is the sampling frequency.

$$R_{res\_measured} = \frac{samples}{f_s} \cdot \frac{c}{2} \tag{5-2}$$

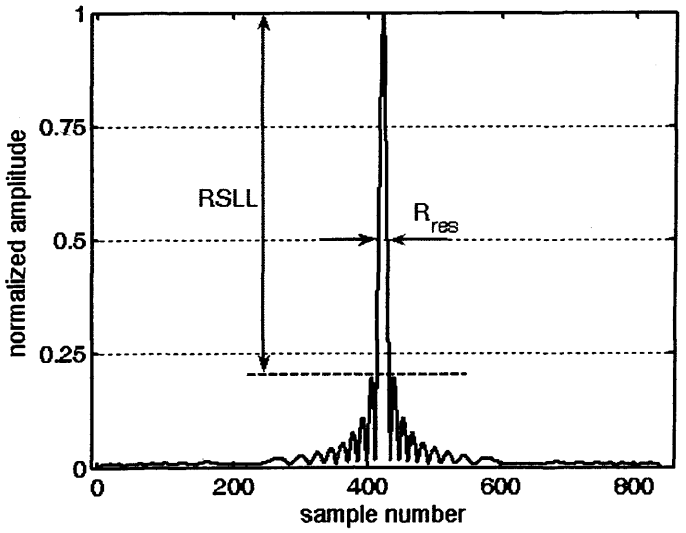

Figure 5-2 Illustration of points in pulse-compressed signal used to calculate parameters R_{res} and RSLL

5.2.2 RSLL

Range sidelobe level (RSLL) was also measured from the pulse-compressed signal envelope and is defined as the ratio of the maximum mainlobe amplitude to the maximum sidelobe amplitude. The points used in the measurement of RSLL are illustrated in Figure 5-2. RSLL is measured in dB.

5.2.3 The Model

A mathematical model was generated to compare the suggested method of calculating the range resolution against the theoretical method. The same model was used to assess the severity of the potential side effects of pulse compression when used in a TCD system, considering typical values for parameters such as pulse length, chirp bandwidth and target velocity.

In the model, the RF blood signal was modelled as an array of normally distributed (in the frequency domain) random numbers of length such that a signal was present from 0 to 10 cm (or from 0 to 100 cm when a long transmit pulse was used). This method was based on the signal generation scheme used by Hoeks et al. (1993). A target was inserted into the signal at a specified depth (by adding the transmitted chirp to the blood signal with the start sample of the transmitted chirp placed at the required depth and the remaining samples in the transmitted pulse being added to deeper samples in the blood signal); in this case a depth of 4 cm was arbitrarily chosen. A sampling frequency of 32.5 MHz was used.

The transmitted chirp had a centre frequency, f_c, of 2 MHz. The duration and bandwidth of the chirp was varied depending on the simulation being performed. These are detailed further, below.

The mean and standard deviation of the range resolution and RSLL were calculated over 100 model runs, for each of two pulse lengths (9.9 μs and 250 μs) and for several bandwidths between 0.2 and 2 MHz. This was repeated with a Dolph-Chebyshev window applied to the transmit pulse before cross-correlation. The Dolph-Chebyshev window sidelobe level was varied between 40 and 140 dB. When a window was applied, the range resolution is given as a ratio of the resolution with a window to the resolution without, and the RSLL is given as the difference in RSLL with and without a window applied.

Non-linear propagation and frequency dependent attenuation were found to have minimal effect in simulations using chirps by Chiao and Hao (2005). The simulations performed by Chiao and Hao used chirps with a centre frequency of 6.1 MHz; this is a higher frequency than that used for TCD and so is likely to worsen any effects present (Muir and Carstensen, 1980). Simulations performed by Rao (1994) also showed that, for a typical TCD centre frequency and tissue attenuation coefficient, frequency dependent attenuation had little effect on the pulse-compressed signal. Non-linear propagation is known to be a greater problem for large amplitude pulses. However, since we are not proposing to alter the transmitted pulse amplitude from conventional, non-coded, TCD operation, the effects of non-linear propagation have not been included in this model. Frequency dependent attenuation was included in the model using an attenuation coefficient of 0.7 dB cm^{-1} MHz^{-1}. The mean range resolution and RSLL were calculated over 100 model runs using a pulse duration of 9.9 μs and various chirp bandwidths between 0.2 and 2 MHz. These results were compared to results generated when no frequency dependent attenuation was included within the model.

The theoretical depth error due to range-Doppler coupling can be calculated using eqn (5-3) (Hunt et al., 2004), where ε_{depth} is the depth error in metres, τ_p is the pulse (chirp) duration in seconds, f_c is the chirp centre frequency, v_{tgt} is the target velocity in m s^{-1} and B_{tx} is the bandwidth of the transmitted chirp.

$$\varepsilon_{depth} = \frac{\tau_p f_c v_{tgt}}{B_{tx}} \tag{5-3}$$

The range-Doppler coupling effect was included in the model by adding a Doppler shift to the target signal that was inserted into the simulated blood signal. Doppler shifts were calculated for the start and end frequencies of the transmit chirp and added to the original frequencies. This gave the start and end frequencies of the chirp inserted into the simulated blood signal. The mean and standard deviation of the depth error were calculated over 100 model runs for various target velocities between 0 and 5 m s^{-1} (chirp bandwidth was 0.5 MHz) and for various chirp bandwidths between 0.2 and 2 MHz (target velocity was 0.6 m s^{-1}). The depth error was calculated by subtracting the 100 values obtained for target depth for a given target velocity, from the mean depth over 100 runs at a target velocity of 0 m s^{-1}. These simulations were repeated using two different pulse lengths, 9.9 μs and 250 μs.

The differences in predicted depth error from the model and using eqn. (5-3) were subsequently examined for varying chirp duration (between 5 and 250 μs) for various target velocities (0.6, 2 and 5 m s^{-1}, chirp bandwidth of 0.5 MHz) and for various chirp bandwidths (0.5, 1, 1.5 and 2 MHz, target velocity of 0.6 m s^{-1}). The difference between the model and theoretical results were measured as a percentage of the theoretical depth error using eqn (5-4).

$$\text{error between model and theory} = \frac{\varepsilon_{depth} - \text{model depth error}}{\varepsilon_{depth}} \times 100 \qquad (5\text{-}4)$$

5.3 Results

The theoretical range resolutions are shown, together with the mean resolutions obtained from the model, in Figure 5-3a for various chirp bandwidths. The model results are shown for two different chirp durations; the theoretical calculation does not vary with chirp duration and so only one theoretical curve is shown. The standard deviations of the model results are shown in Figure 5-3b.

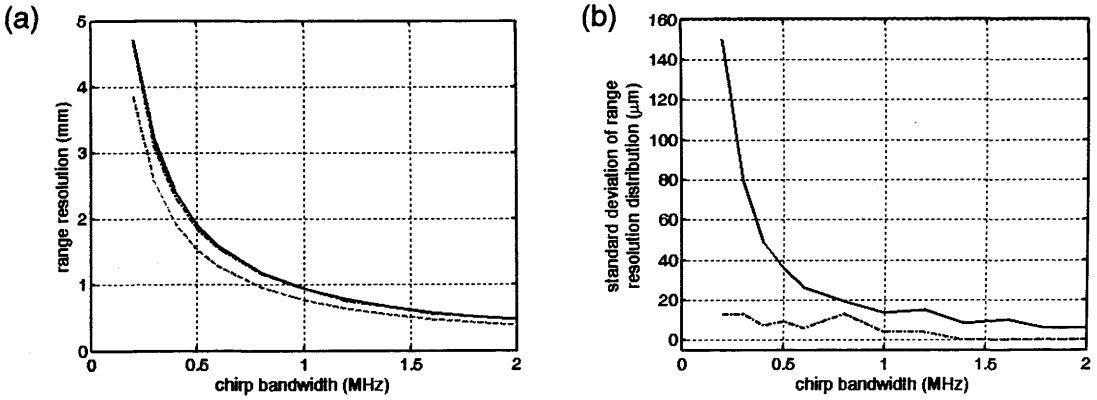

Figure 5-3 (a) (---) Theoretical range resolution [mm] calculated using eqn. (5-1) and mean range resolution [mm] over 100 model runs using a (—) 9.9 μs chirp and a (-----) 250 μs chirp, and (b) Standard deviation [μm] of the range resolution over 100 model runs, using a (—) 9.9μs chirp and a (-----) 250 μs chirp

The effects on range resolution, of windowing the transmit pulse before performing the cross-correlation, are shown in Figure 5-4 for varying Dolph-Chebyshev window sidelobe levels. Figure 5-4a shows the mean and standard deviation of the range resolution ratio over 100 runs at each window level, for a chirp bandwidth of 0.5 MHz, chirp duration 9.9 μs. Figure 5-4b shows the mean over 100 runs at each window sidelobe level, for various chirp bandwidths between 0.2 and 2 MHz using a chirp duration of 9.9 μs. Figure

5-4c and Figure 5-4d show the same as Figure 5-4a and Figure 5-4b but using a 250 μs chirp.

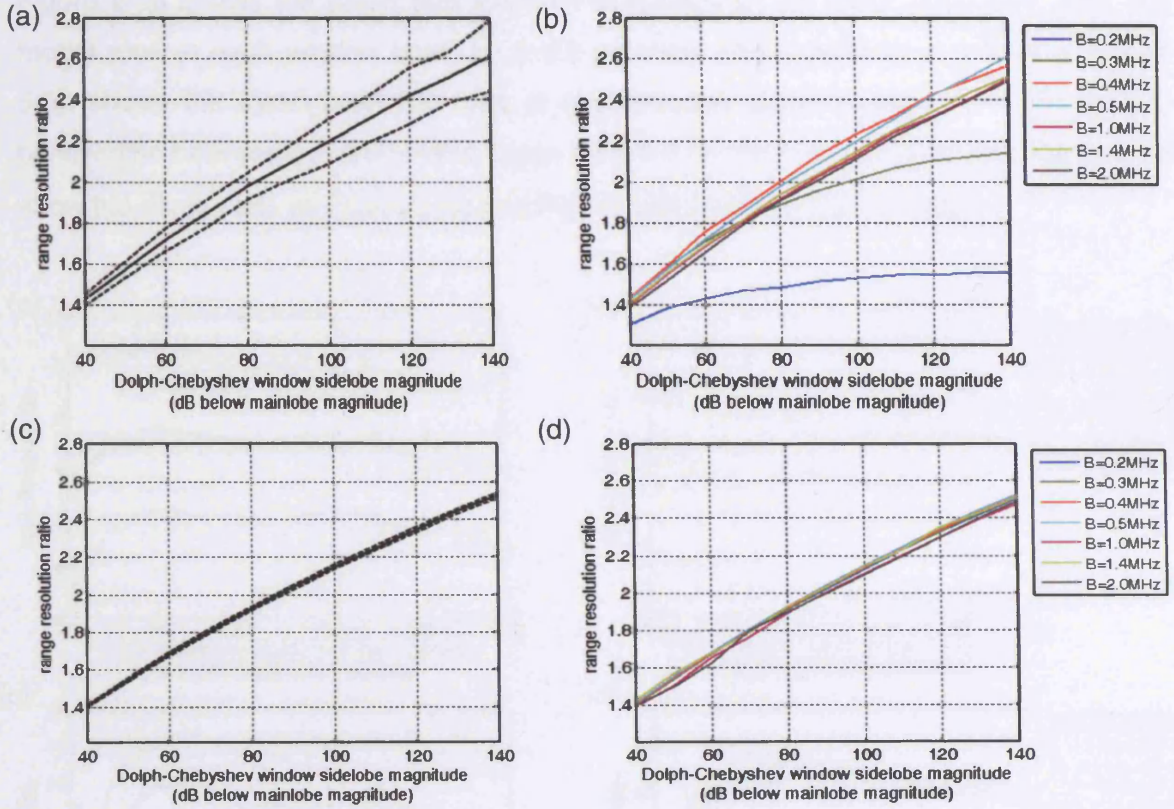

Figure 5-4 (a) Mean (—) ± one standard deviation (---) of the range resolution ratio over 100 runs, using a 9.9 μs chirp, bandwidth = 0.5MHz, (b) Mean of the range resolution ratio over 100 runs, using a 9.9 μs chirp, for various chirp bandwidths. (c) and (d) are as per (a) and (b) but using a 250 μs chirp.

The mean RSLL obtained over 100 model runs, for different chirp bandwidths and for two different chirp durations, is shown in Figure 5-5a. The standard deviations of the RSLL results are shown in Figure 5-5b.

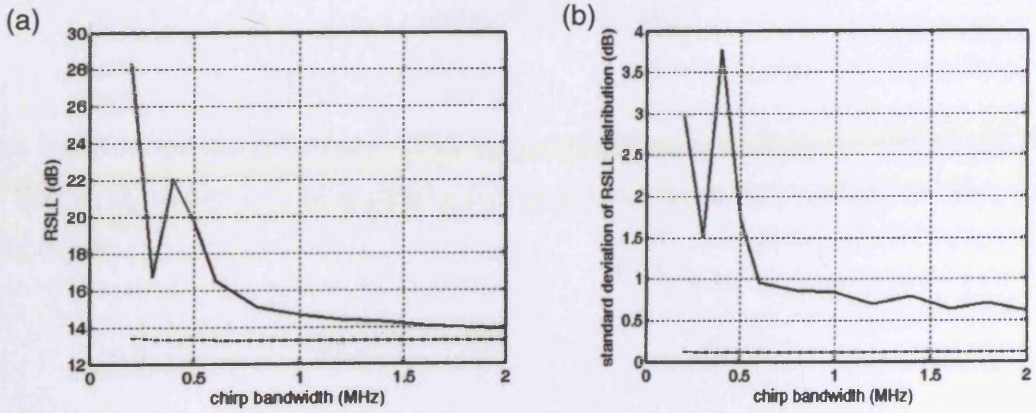

Figure 5-5 (a) Mean and (b) standard deviation of RSLL over 100 runs, using a (—) 9.9 μs chirp and a (----) 250 μs chirp

The effects on RSLL, of windowing the transmit pulse before performing the cross-correlation, are shown in Figure 5-6 for varying Dolph-Chebyshev window sidelobe levels. Figure 5-6a shows the mean and standard deviations of the RSLL difference over 100 model runs at each window level, for a 9.9 μs chirp with bandwidth of 0.5 MHz. Figure 5-6b shows the mean over 100 runs at each window sidelobe level, for various chirp bandwidths between 0.2 and 2 MHz, again for a 9.9 μs chirp. Figure 5-6c and Figure 5-6d show the same plots as Figure 5-6a and Figure 5-6b but for a 250 μs chirp.

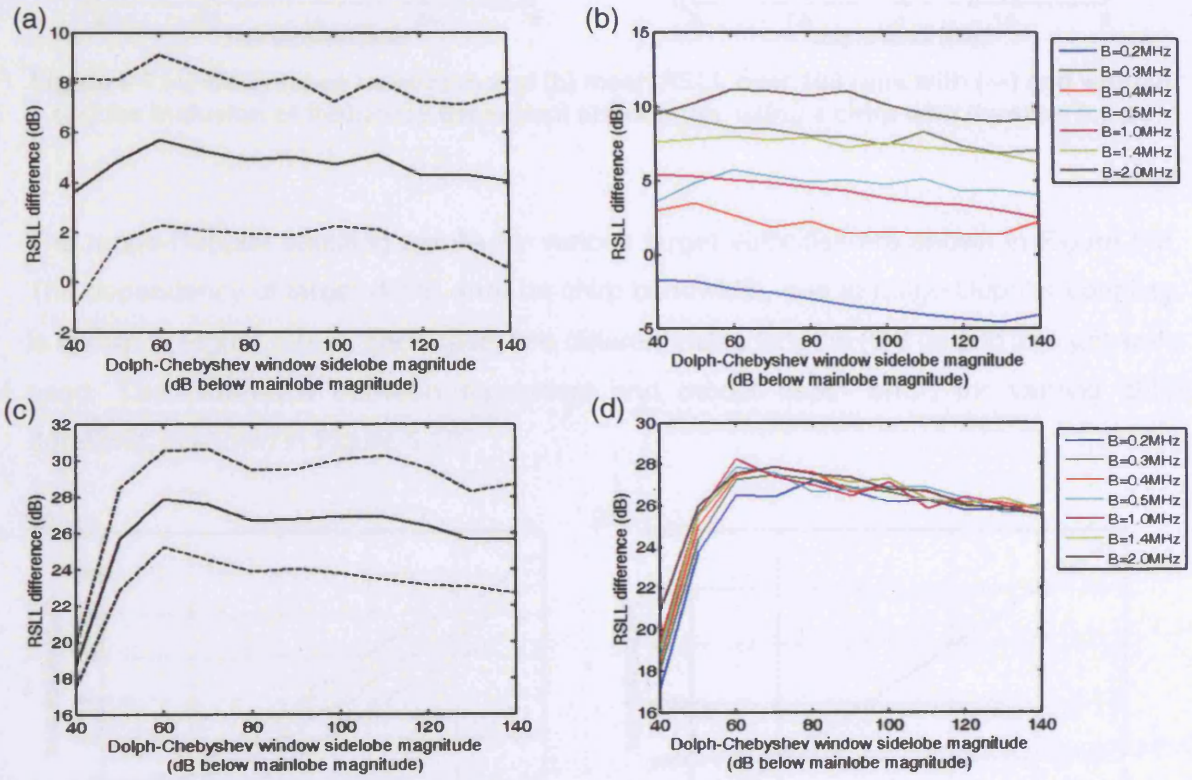

Figure 5-6 (a) Mean (—) ± one standard deviation (---) of the 'RSLL difference' over 100 runs, using a 9.9 μs chirp with bandwidth = 0.5 MHz, (b) Mean of the 'RSLL difference' over 100 runs, using a 9.9 μs chirp, for various chirp bandwidths, (c) and (d) as per (a) and (b) but using a 250 μs chirp.

The mean range resolution and RSLL over 100 model runs, with and without the inclusion of frequency dependent attenuation, are shown in Figure 5-7. A chirp duration of 9.9 μs was used.

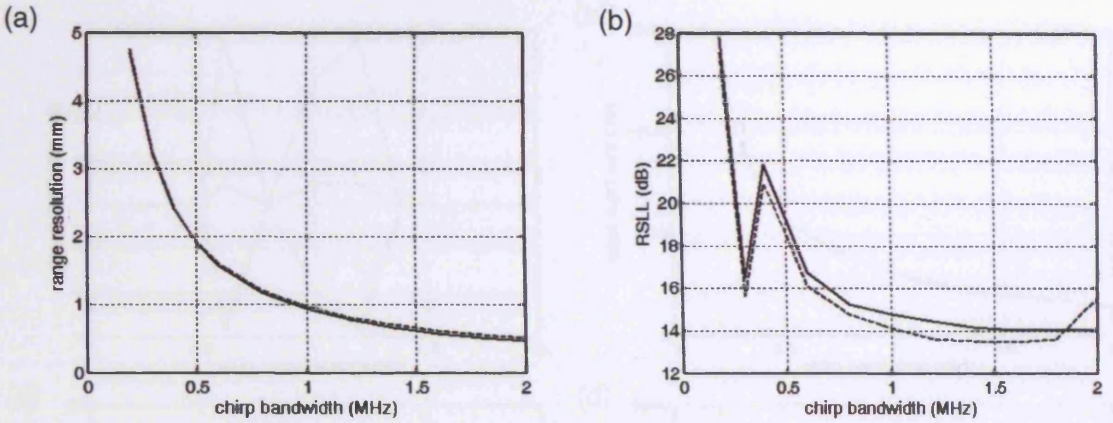

Figure 5-7 (a) Mean range resolution and (b) mean RSLL over 100 runs with (---) and without (—) the inclusion of frequency dependent attenuation, using a chirp with duration 9.9 μs

The range-Doppler coupling results for various target velocities are shown in Figure 5-8. The dependency of target depth error on chirp bandwidth, due to range-Doppler coupling, is shown in Figure 5-9. In each case, two different pulse lengths (9.9 μs and 250 μs) were used. The difference between theoretical and model depth error, for varying chirp durations, is shown in Figure 5-10.

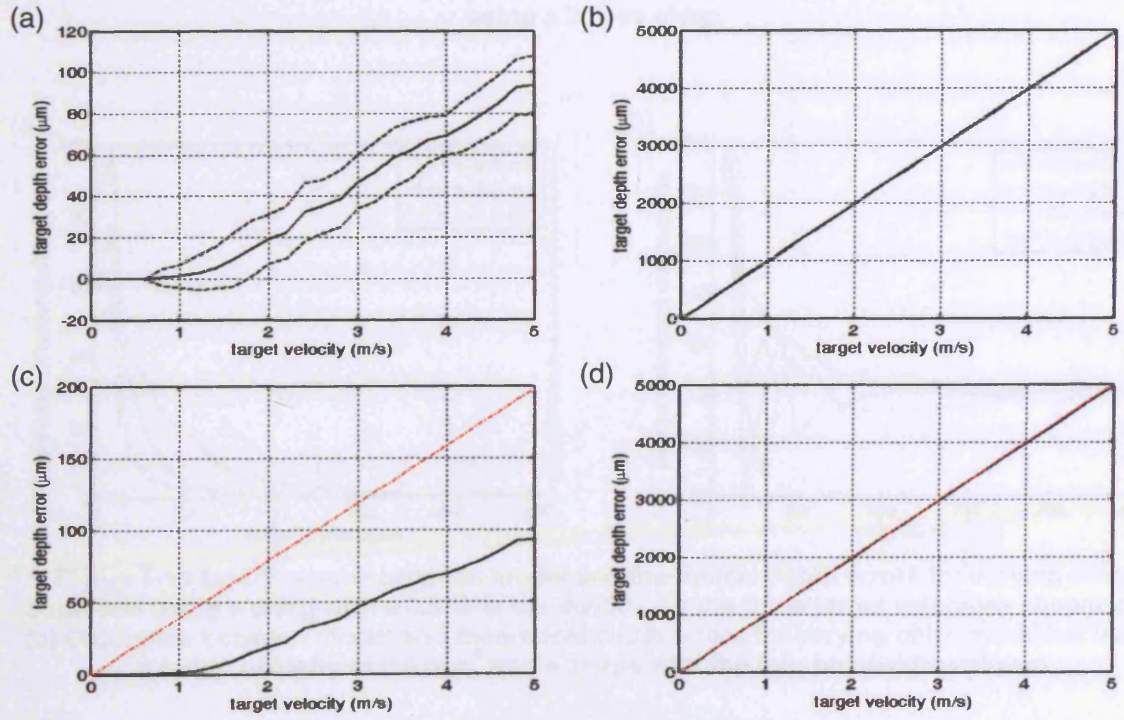

Figure 5-8 Mean (—) ± one standard deviation (---) of the target depth error over 100 model runs, using (a) a 9.9 μs chirp and (b) a 250 μs chirp with 0.5MHz bandwidth. (c) (-·-·-) Theoretical depth error calculated using eqn. (5-3) and (—) Mean target depth error over 100 model runs using a 9.9 μs chirp with 0.5 MHz bandwidth, (d) as per (c) but using a 250 μs chirp.

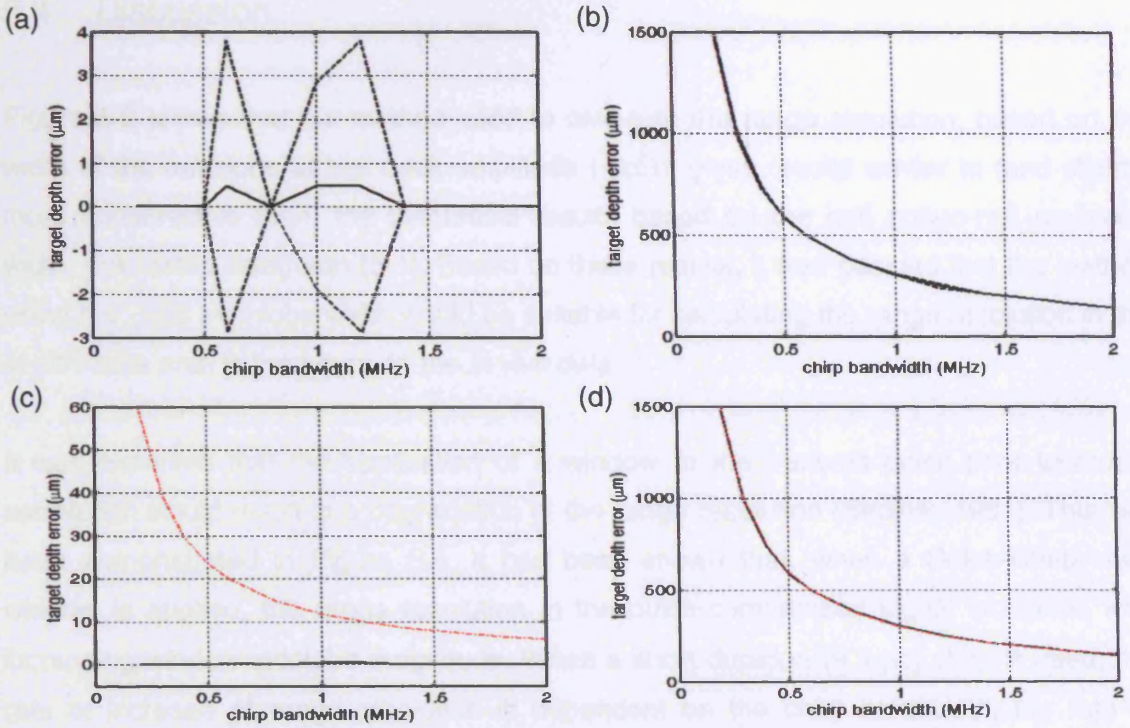

Figure 5-9 Mean (—) ± one standard deviation (---) of the target depth error over 100 model runs, using (a) a 9.9 μs chirp and (b) a 250 μs chirp, target velocity = 0.6 m s^{-1}. (c) (-·-·-) Theoretical depth error calculated using eqn. (5-3) and (—) Mean of the target depth error over 100 model runs using a 9.9 μs chirp, target velocity = 0.6 m s^{-1}, (d) as per (c) but using a 250 μs chirp.

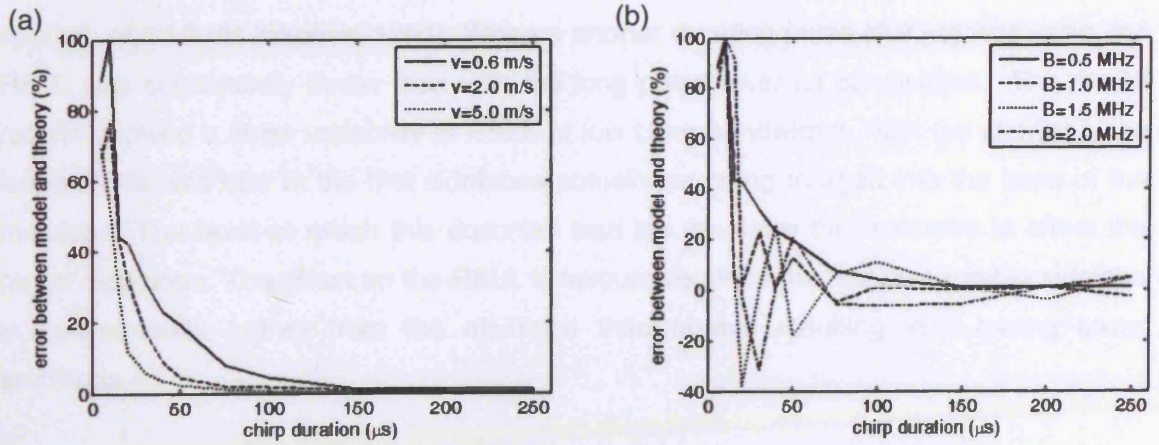

Figure 5-10 (a) Difference between model and theoretical depth errors for varying chirp durations using a chirp with a 0.5 MHz bandwidth, for the three target velocities shown and, (b) Difference between model and theoretical depth errors for varying chirp durations using a target velocity of 0.6 m s^{-1} and a chirps with the four bandwidths shown

5.4 Discussion

Figure 5-3 shows that the method used to calculate the range resolution, based on the width of the mainlobe at half peak amplitude (-6dB), gives results similar to (and slightly more conservative than) the theoretical results based on the half null-to-null mainlobe width, calculated using eqn (5-1). Based on these results, it was decided that the method using the - 6dB mainlobe width would be suitable for calculating the range resolution in the *in vitro* data and, in the future, in the *in vivo* data.

It was expected that the application of a window to the transmit pulse prior to cross-correlation would result in a degradation of the range resolution (Skolnik, 1981). This has been demonstrated in Figure 5-4. It has been shown that, when a Dolph-Chebyshev window is applied, the range resolution in the pulse-compressed signal increases with increasing window sidelobe magnitude. When a short duration (9.9 μs) chirp is used, the rate of increase of range resolution is dependent on the chirp bandwidth; the rate of increase rises for bandwidths up to 0.4 MHz and then is similar for bandwidths between 0.4 and 2 MHz.

When a long duration chirp (250 μs) was used, the mean RSLL over all bandwidths was 13.28 dB (see Figure 5-5a). This agrees well with the theoretical RSLL, with no window applied, of 13.2 dB (Skolnik, 1981). When a shorter duration pulse (9.9 μs) was used, the RSLL was consistently better than with the long pulse, over all bandwidths. The model results showed a large variability in RSLL at low chirp bandwidths, with the shorter pulse length. This was due to the first sidelobes sometimes being merged into the base of the mainlobe. The level at which this occurred was too far down the mainlobe to affect the range resolution. The effect on the RSLL is favourable since the first measurable sidelobe is consequently further from the mainlobe than normal resulting in it having lower amplitude.

It can be seen from Figure 5-6 that, when a Dolph-Chebyshev window with sidelobe magnitude of 60dB or greater is used prior to pulse compression, the sidelobes in the pulse compressed signal can be reduced to a level approximately 26 dB lower than the non-windowed case, when a long duration (250 μs) chirp is used. This is comparable to the theoretical difference of 26.8 dB (Skolnik, 1981). However, when a shorter pulse (9.9 μs) is used the change in RSLL is not as large. This may be for two reasons. Firstly, the short pulse will approximate a square wave in the frequency domain to a lesser

degree than the longer pulse. As a result, the pulse-compressed signal will less accurately approximate a sinc function in the short pulse case. This gives a better RSLL without windowing and therefore the improvement in RSLL will not be as large when the window is applied. Secondly, in the short pulse case, the sidelobes in the pulse-compressed signal (with no window applied) are already quite close to the level of the background blood signal. Therefore, applying a window may have little effect if the sidelobe level cannot be decreased much further.

The peak amplitude in the pulse-compressed signal was smaller when using a shorter pulse. This was as expected and the result is due to a shorter pulse giving rise to a smaller pulse compression ratio (Skolnik, 1981)). The peak amplitude obtained using the shorter pulse is, however, more than adequate to be able to detect events in the pulse-compressed signal.

It can be seen from Figure 5-7 that the presence of frequency dependent attenuation degrades the range resolution and the RSLL minimally. The degradation increases slightly with increasing bandwidth. This is to be expected since, as the bandwidth increases, the difference in attenuation between the start and end of the chirp is greater resulting in a larger variation from the transmit pulse; this, in turn, would result in a poorer correlation. However, despite this degradation, the benefits in terms of using pulse compression to improve axial resolution are still significant.

When the short pulse (9.9 μs) was used, the model results consistently underestimated the theoretical depth error given by eqn (5-3). This is evident from Figure 5-8c and Figure 5-9c. However, it can be seen in Figure 5-8d and Figure 5-9d that when the long pulse (250 μs) was used, the model and theoretical results matched closely. Figure 5-10 shows that the model and theoretical results for depth error are most closely matched for chirp durations greater than 100 μs. It can be seen in Figure 5-10a that the difference between the two predictions decreases with increasing target velocity. In Figure 5-10b it is clear that the magnitude of the difference between the model and theoretical results is not as dependent on chirp bandwidth as it is on target velocity.

It can be seen from Figure 5-8a that any error in the depth of the peak in the pulse-compressed signal is small when using a 9.9 μs chirp and when considering the velocities of blood likely to be encountered from a TCD signal (<2 m s^{-1}). When the velocity is 2 m s^{-1}, the mean depth error is approximately 20 μm. Note that, although velocities of 2 m s^{-1} could be encountered, these are rare and so this is used as a worst-case scenario.

Given that the theoretical range resolution of a pulse-compressed signal, using a chirp bandwidth of 0.5 MHz, is 1.54 mm (see eqn (5-1)), it is clear that a depth error of 20 μm would be inconsequential. For target velocities lower than 2 m s^{-1}, the expected depth error is less than 20 μm. Equation (5-3) shows that the depth error is proportional to the chirp length and this is reflected in the higher depth errors shown in Figure 5-8b where a pulse length of 250 μs was used.

5.5 Conclusion

The potential side-effects of pulse compression such as range sidelobes, range-Doppler coupling and frequency dependent attenuation have been considered. Their effects have been examined in a mathematical model designed to assess the viability of using pulse compression in a TCD system. The results have demonstrated that the magnitude of the pulse compression side-effects are small for typical TCD pulse lengths, target velocities and suitable chirp bandwidths and that the improvement in axial resolution outweighs any of the negative effects. Therefore, no reason has been found to indicate that pulse-compression should not be used within a TCD system. Experimental evaluation of a TCD system employing coded excitation is detailed in Chapter 6.

6 Coded Excitation in TCD Ultrasound Systems to Improve Axial Resolution - Experimental Evaluation

6.1 Introduction

Current techniques to track the motion of emboli using the transcranial Doppler (TCD) radio-frequency (RF) signal (see Chapter 3) are limited by the narrow bandwidth of the transmitted pulse since this leads to poor axial resolution. The bandwidth could be increased by shortening the duration of the transmitted pulse but this would also require an increase in pulse amplitude which is undesirable from a safety standpoint. It is proposed that coded excitation and pulse compression could be used to increase the bandwidth without the requirement to increase the pulse amplitude.

Mathematical modelling has shown, in Chapter 5, that the magnitudes of the possible disadvantages to using pulse compression in a TCD system are minimal when compared to the improvement in axial resolution. These promising modelling results led to *in vitro* experiments being carried out to assess if implementing such a system were practical and potentially beneficial.

The aim of this study was to assess TCD system performance using coded-excitation against using conventional, non-coded pulses. The two processing methods were compared qualitatively and quantitatively for different chirp bandwidths, with and without the inclusion of a temporal bone sample. Measures of performance included range resolution and range sidelobe level, the ability to track the motion of emboli, the ability to locate vessel position and whether or not the conventional Doppler signal could still be extracted.

6.2 Methods

The *in vitro* work described in this chapter was performed using a 2 MHz ultrasound system, built in-house. A pulse repetition frequency (PRF) of 10.1 kHz was used for all acquisitions. The processing within this system is FPGA based and so can be programmed to transmit any pulse shape. Chirped pulses and conventional non-coded pulses were transmitted alternately. This allowed the same 'embolic' event to be compared using the two different types of processing. The resulting PRF of each of the

'split' signals was approximately 5 kHz. A transmit burst of approximately 13 μs (26 cycles) long and a receive gate of 10 mm were used by the ultrasound system. The end of the receive gate is positioned at the depth set on the TCD machine by the technician. The start of the receive gate is 10 mm shallower than the end. The depth setting on the TCD machine was fixed at 55 mm in the collection of the in vitro data used in this study.

The frequency response of the 2 MHz transducer was measured using a radiation force balance (Perkins, 1989). This enabled the transducer's centre frequency and -3 dB bandwidth to be determined and therefore appropriate chirps designed.

Chirp bandwidths equal to or greater than the transducer's –3 dB bandwidth were used in order to reduce the range sidelobe level in the pulse-compressed signal (Misaridis and Jensen, 1999). The transmitted chirp was amplitude tapered using a Tukey window with a ratio of taper to constant sections of 0.15. The amplitude tapering reduces the ripples present in the spectrum of the chirps and hence will reduce the amplitude of sidelobes in the pulse-compressed signal (Misaridis and Jensen, 1999).

A number of 'embolic' events were recorded from a flow phantom. The flow phantom was set up as shown in the schematic in Figure 6-1. The blood-mimicking fluid (BMF) (Ramnarine et al., 1998) was circulated round the flow loop in the direction shown to simulate blood flow in the MCA. A wall-less phantom (Ramnarine et al., 2001) was used to minimise the ultrasound distortion which can be present if tubing is used as a vessel channel. Since the transducer was clamped in place, any recorded events occurred from bubbles or clumps of BMF passing the transducer on their way around the flow loop. No movement artefacts occurred. Data were recorded both with and without a temporal bone sample (male cadaver, 83 years) between the transducer and the tissue mimicking material (TMM).

Data recordings were also made from normal subjects to investigate the appearance of artefacts in the pulse-compressed signal. The 2 MHz probe was fixed over the subjects' temporal bone windows, on an arbitrary choice of side, so as to insonate the MCA. The subjects generated typically encountered artefacts such as tapping the probe, coughing and speech.

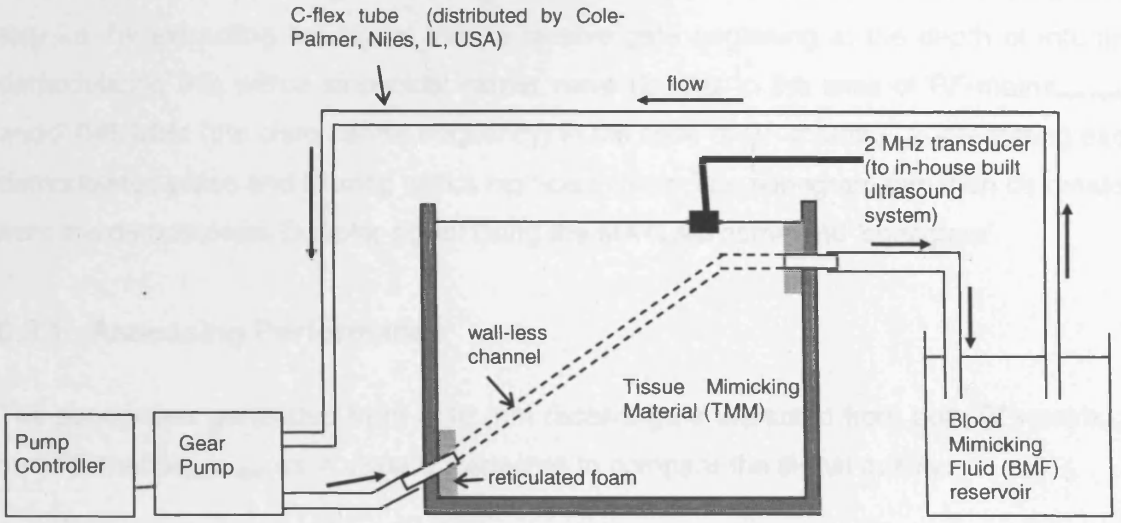

Figure 6-1 Schematic of flow phantom

The RF, system clock, and transmit gate signals were made externally available from the ultrasound system. These three signals were connected to the PC-housed, 14-bit A/D converter. A commercially available 2 MHz TCD transducer (Skidmore Medical, Bristol, UK) was used.

The RF, transmit gate, and system clock signals were continuously input to the A/D converter and the data were stored in an on-board buffer. The clock rate (32.5 MHz) determined the sampling frequency, f_s, used by the A/D converter when digitising the RF and transmit gate signals.

The user initiated a trigger signal to the A/D converter on the occurrence of a suspected embolic event (in the case of the flow phantom) or artefact (in the case of the normal subjects). Following receipt of this trigger, all available RF and transmit gate data, 1 s either side of the time of triggering, were transferred from the A/D card to hard disk.

The transmit gate signal was used to determine the start of each transmit burst. The RF data were rearranged into a matrix with the received signal from each transmit burst forming one row. This matrix was divided into two, with the odd rows forming one matrix and the even rows forming another. One of these matrices corresponds to the signals received when transmitting a chirp, RF-matrix$_{chirp}$, and the other corresponds to the signals received when transmitting a conventional non-coded pulse, RF-matrix$_{non\text{-}coded}$. It is these RF-matrices that are the subject of processing for the rest of this chapter.

A demodulated Doppler signal can be obtained from either of the RF-matrices in the usual way i.e. by extracting the signal from a receive gate beginning at the depth of interest, demodulating this with a sinusoidal carrier wave (2 MHz in the case of RF-matrix$_{non\text{-}coded}$ and 2.045 MHz (the chirp centre frequency) in the case of RF-matrix$_{chirp}$), integrating each demodulated pulse and filtering with a high-pass filter. The sonogram can then be created from the demodulated Doppler signal using the MATLAB command 'specgram'.

6.2.1 Assessing Performance

The sonograms generated from a 10 mm receive gate extracted from both RF-matrix$_{chirp}$ and RF-matrix$_{non\text{-}coded}$ were visually inspected to compare the signal quality.

A pulse-compressed matrix was produced by cross-correlating each row of RF-matrix$_{chirp}$ with a copy of the transmitted chirp. Sonograms were generated from 2 mm gates extracted from the pulse-compressed matrix (demodulated using the centre frequency of the transmitted chirp) and compared with the 10 mm receive gate sonograms from RF-matrix$_{chirp}$ and RF-matrix$_{non\text{-}coded}$ in terms of signal quality. Sonograms from each of the three sources were also used to assess the ability to locate vessel position and the associated demodulated Doppler signals were used to assess the ability to distinguish between artefacts and emboli using the time-based signal.

The ability to follow the embolus 'track' was assessed qualitatively in three different matrices; the high-pass filtered RF-matrix$_{non\text{-}coded}$ (each column of RF-matrix$_{non\text{-}coded}$ is clutter-filtered using a 4th order Butterworth high-pass filter with a cut-off frequency of 75 Hz), the match-filtered matrix (each row of the high-pass filtered RF-matrix$_{non\text{-}coded}$ is cross-correlated with a copy of the conventional, non-coded transmit pulse) and finally the pulse-compressed matrix (the high-pass filtered RF-matrix$_{chirp}$ is cross-correlated with a copy of the transmitted chirp). Magnitude plots of these matrices are shown with respect to the root-mean-square (rms) noise level of that particular data acquisition. These were calculated from a region of the appropriate RF signal containing no suspected embolus or artefact.

The appearance of artefacts and emboli are compared visually in the pulse-compressed signal to assess if they can still be discriminated between reliably according to the criteria defined in Chapter 3.

The appearance of *in vitro* 'emboli' in these three matrices (the high-pass filtered RF-matrix$_{non\text{-}coded}$, the matched filtered matrix and the pulse-compressed matrix) was used to

quantify the range (axial) resolution of the different techniques. The range resolution was calculated in the same way as in the model described in Chapter 5 i.e. the width (in units of distance) of the event peak in the signal envelope at the point where the amplitude of the peak is half (-6dB) of the maximum. This method is used for the signals in each of the three matrices. The mean and standard deviation of the axial resolution were calculated over 13 embolic events using two different chirp bandwidths (0.59 MHz and 0.8 MHz). Each bandwidth was tested with and without the temporal bone sample placed between the transducer and the TMM.

The mean and standard deviation of the range sidelobe level (RSLL) were also calculated for each case (i.e. two bandwidths, each with and without the use of temporal bone) from the pulse-compressed matrix. The RSLL is defined as the ratio of maximum mainlobe amplitude to maximum sidelobe amplitude and is measured in dB.

The range resolution and RSLL in the pulse-compressed matrix are calculated with and without shaping of the transmit chirp with a Dolph-Chebyshev window (60 dB sidelobe level) prior to cross-correlation.

6.3 Results

The frequency response of the 2 MHz transducer is shown in Figure 6-2. The curve has a relative power of –3 dB in 4 places; 1.75 MHz, 2.15 MHz, 2.26 MHz and 2.34 MHz. The two outer points (1.75 and 2.34 MHz) were used to determine the centre frequency, 2.045 MHz, of the transmitted chirps. Using the outer points gives the transducer a –3 dB bandwidth of 0.59 MHz.

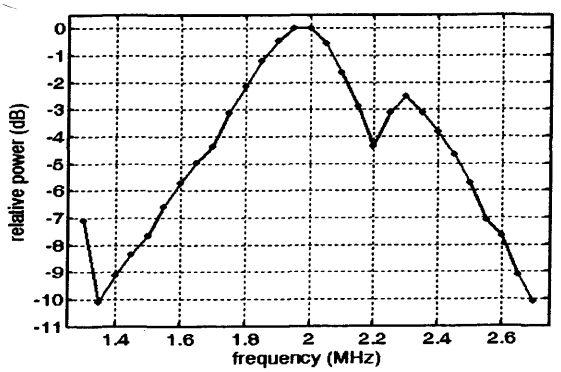

Figure 6-2 Frequency response of the 2 MHz transducer, measured using a radiation force balance

An example of the sonograms generated from a 10 mm receive gate extracted from RF_matrix$_{non\text{-}coded}$ and RF_matrix$_{chirp}$, both from the same acquisition, are shown in Figure 6-3. Sonograms generated from a number of 2 mm receive gates, extracted from the pulse-compressed matrix (of the same data acquisition as shown in Figure 6-3) are shown in Figure 6-4.

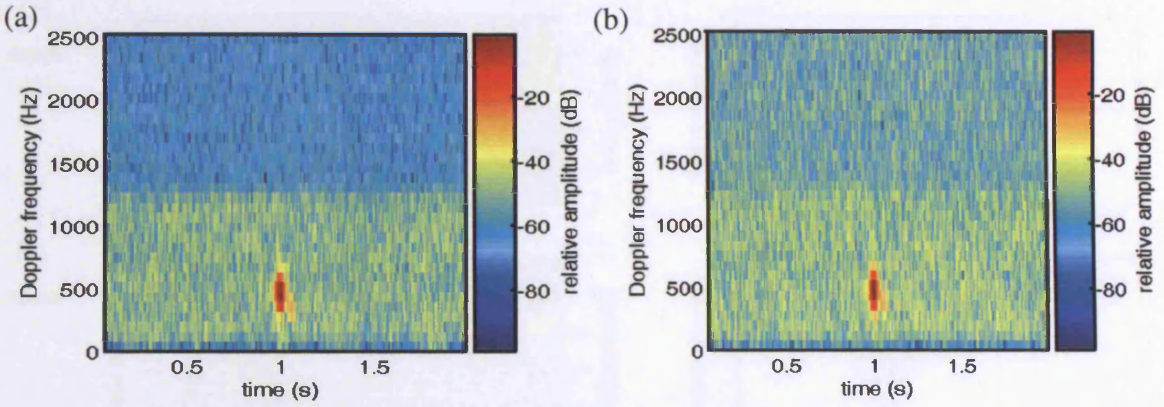

Figure 6-3 Sonograms generated from a 10 mm gate extracted from (a) RF-matrix$_{non\text{-}coded}$ and (b) RF-matrix$_{chirp}$. The transmitted chirp had a bandwidth of 0.8 MHz. In both cases the receive gate was placed between 45 and 55 mm.

Sonograms generated from a number of 2 mm receive gates, extracted from the pulse-compressed matrix of a signal acquired from a healthy volunteer at the time they were generating a speech artefact, are shown in Figure 6-5.

The clutter-filtered RF-matrix$_{non\text{-}coded}$, the match-filtered matrix and the pulse-compressed matrix, in the region of the embolus shown in Figure 6-3 and Figure 6-4, are shown in Figure 6-6a, Figure 6-6b and Figure 6-6c respectively. The sonogram, demodulated Doppler signal and magnitude RF plots for an embolus with much lower amplitude, are shown in Figure 6-7.

Figure 6-8 illustrates the magnitude plot of the pulse-compressed matrix in the region of a probe-tapping artefact, acquired from a healthy volunteer.

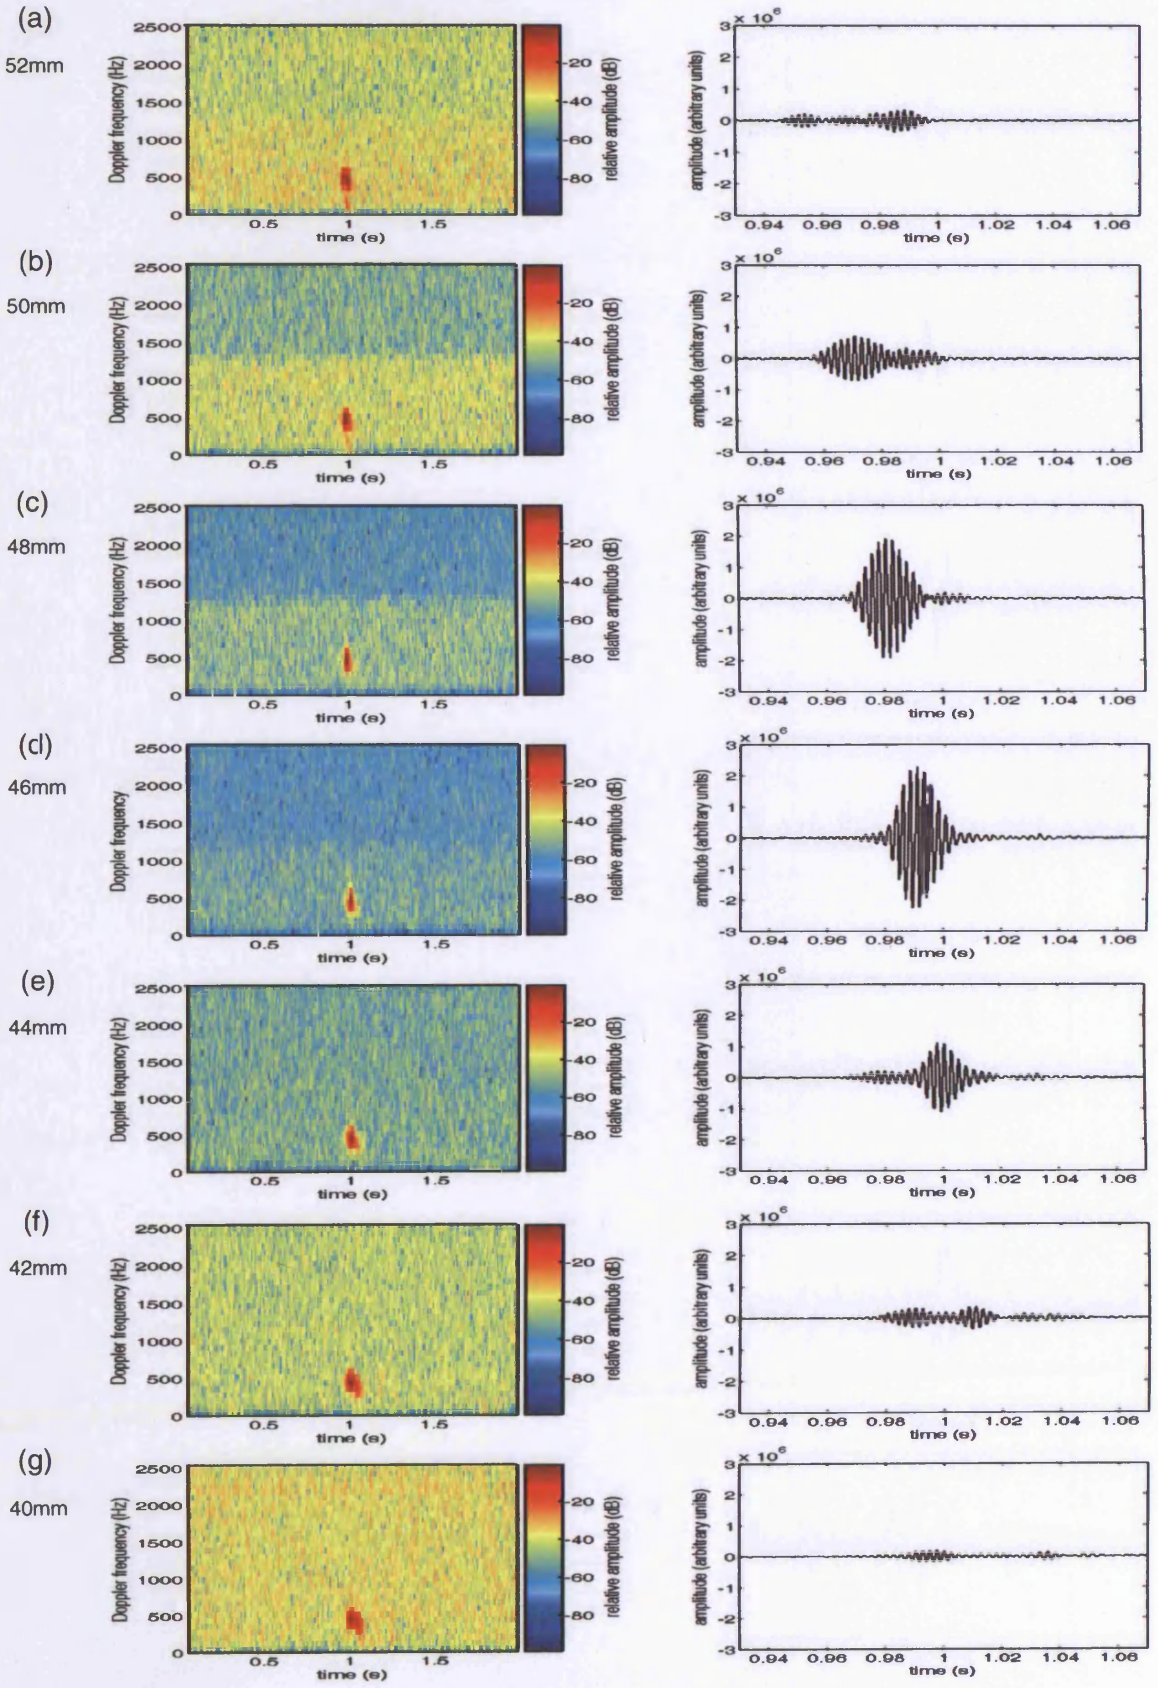

Figure 6-4 Sonograms and demodulated Doppler signals from a number of 2 mm gates extracted from the pulse-compressed matrix. These were generated from the same data acquisition illustrated in Figure 6-3. A chirp with bandwidth 0.8 MHz was used.

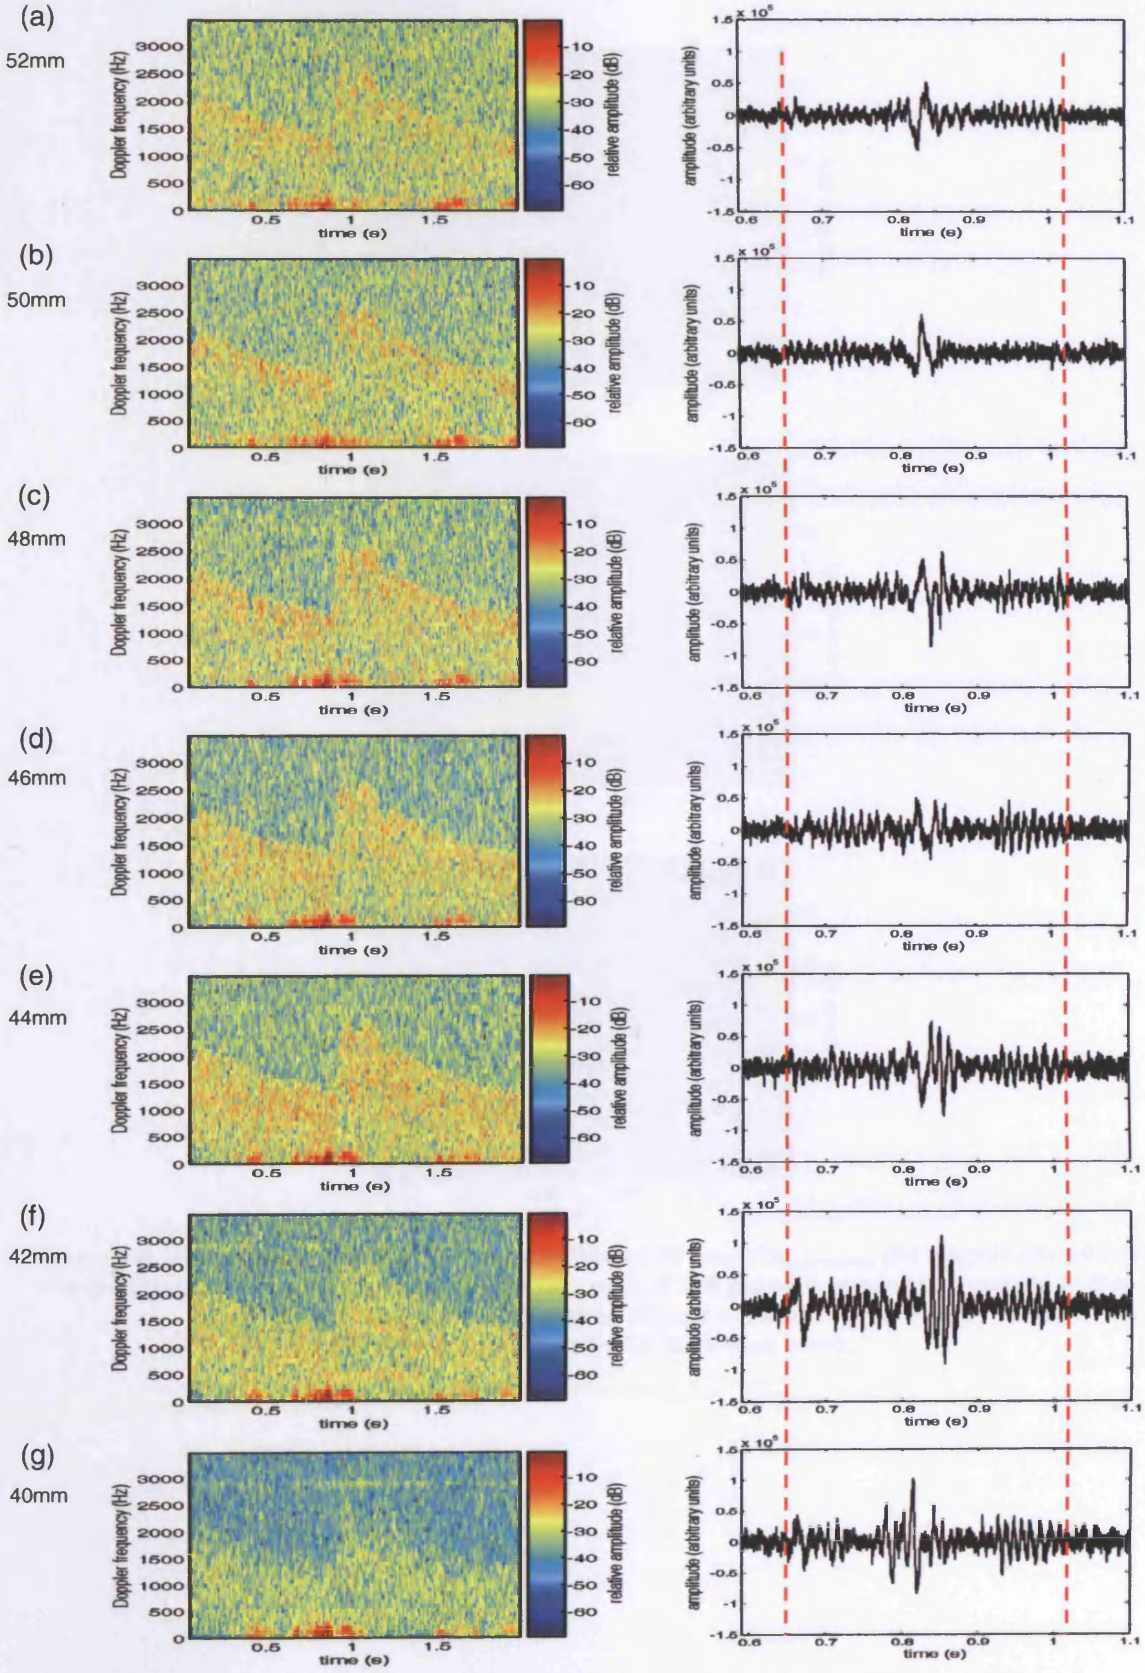

Figure 6-5 Sonograms and demodulated Doppler signals from a number of 2 mm gates extracted from the pulse-compressed signal, illustrating speech artefacts generated by a healthy volunteer. A chirp with bandwidth 0.8 MHz was used.

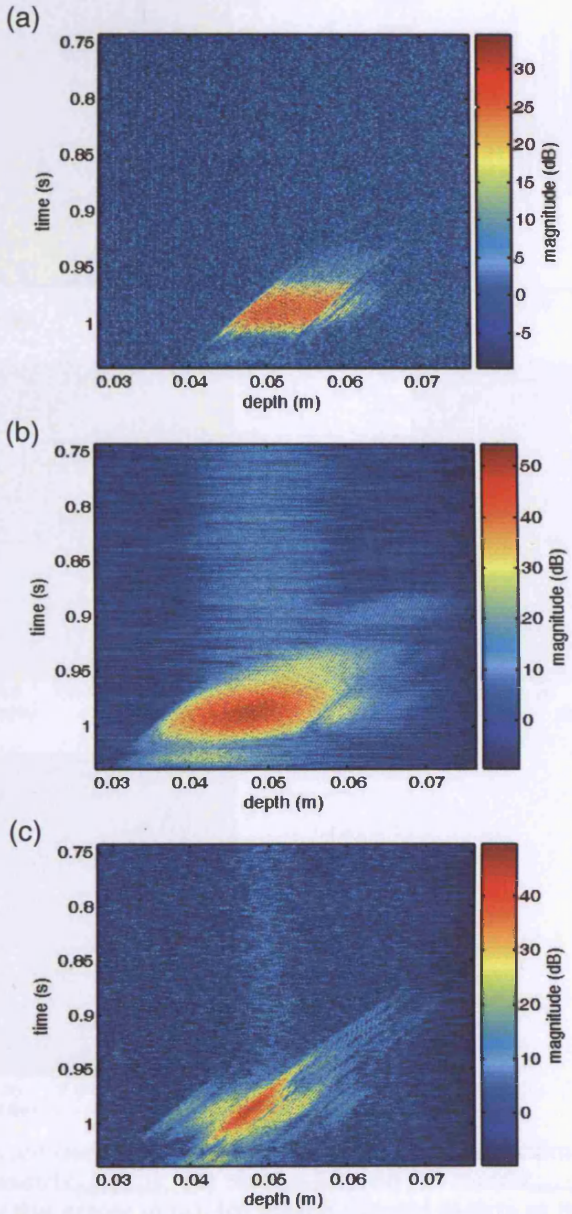

Figure 6-6 (a) Magnitude plot of the clutter-filtered RF-matrix$_{non\text{-}coded}$, (b) magnitude plot of the match-filtered matrix and (c) magnitude plot of the pulse-compressed matrix, in the region of the embolus shown in Figure 6-3 and Figure 6-4. A chirp with bandwidth 0.8 MHz was used.

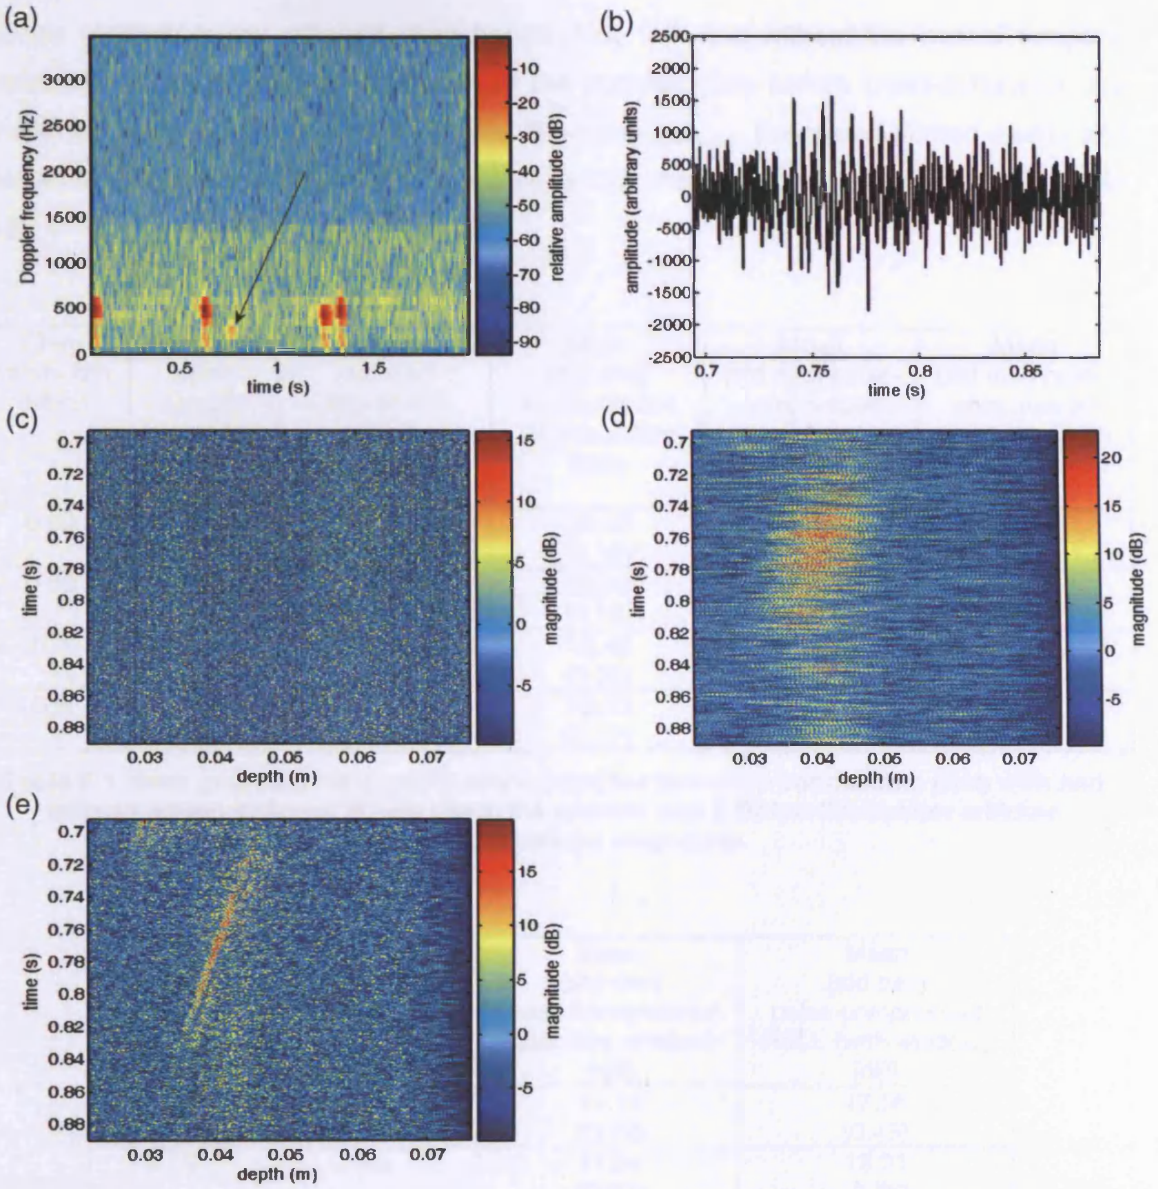

Figure 6-7 (a) Sonogram generated from RF-matrix$_{non\text{-}coded}$, (b) demodulated Doppler signal generated from RF-matrix$_{non\text{-}coded}$,, (c) clutter-filtered RF-matrix$_{non\text{-}coded}$ in the region of the embolus indicated by the arrow in (a), (d) match-filtered matrix in the region of the embolus indicated by the arrow in (a) and, (e) pulse-compressed matrix in the region of the embolus indicated by the arrow in (a). A chirp with bandwidth 0.8 MHz was used.

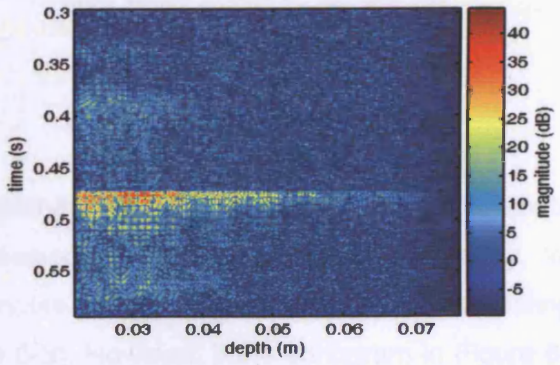

Figure 6-8 Magnitude plot of the pulse-compressed matrix in the region of a probe-tapping artefact. A chirp with bandwidth 0.8 MHz was used.

Range resolutions for different chirp bandwidths, with and without the use of temporal bone, and with and without windowing of the transmit chirp before cross-correlation are shown in Table 6-1 for the clutter-filtered RF-matrix$_{non\text{-}coded}$, the match-filtered matrix and the pulse-compressed matrix. RSLLs are shown for the pulse-compressed matrix in Table 6-2.

Chirp bandwidth (MHz)	Temporal bone sample used	Mean (std dev) filtered RF resolution [mm]	Mean (std dev) match-filtered RF resolution [mm]	Mean (std dev) pulse-compressed resolution (no window) [mm]	Mean (std dev) pulse-compressed resolution (with window) [mm]
0.59	No	10.01 (0.23)	10.42 (0.12)	2.14 (0.10)	3.46 (0.06)
0.59	Yes	9.91 (0.09)	10.03 (0.19)	2.51 (0.07)	3.32 (0.05)
0.8	No	9.86 (0.05)	10.42 (0.05)	1.49 (0.009)	2.54 (0.02)
0.8	Yes	9.85 (0.12)	10.00 (0.13)	1.65 (0.11)	2.87 (0.07)

Table 6-1 Mean (std dev) range resolutions [mm] for two chirp bandwidths each with and without temporal bone. Where used, the window was a Dolph-Chebyshev window with 60 dB sidelobe magnitude.

Chirp bandwidth (MHz)	Temporal bone sample used	Mean (std dev) pulse-compressed RSLL (no window) [dB]	Mean (std dev) pulse-compressed RSLL (with window) [dB]
0.59	No	13.19 (0.79)	17.56 (0.43)
0.59	Yes	11.94 (0.63)	18.31 (8.76)
0.8	No	12.90 (0.20)	19.58 (0.54)
0.8	Yes	11.73 (3.18)	18.22 (0.56)

Table 6-2 Mean (std dev) RSLLs [dB] for two chirp bandwidths each with and without temporal bone. Where used, the window was a Dolph-Chebyshev window with 60 dB sidelobe magnitude.

6.4 Discussion

Demodulating the signal within a receive gate extracted from RF-matrix$_{chirp}$ (without having carried out pulse-compression), with the chirp centre frequency, results in a sonogram with a poorer signal-to-noise ratio (SNR) than when demodulating a signal from RF-matrix$_{non\text{-}coded}$ (see Figure 6-3). However, if the sonogram in Figure 6-3a is compared with that in Figure 6-4c, similar SNRs are evident. It is therefore feasible to demodulate a gate

extracted from the pulse-compressed matrix with the chirp centre frequency whilst still maintaining an adequate SNR. Using the pulse-compressed signal also permits the use of a much narrower receive gate when compared to using the conventional RF-matrix$_{non\text{-}coded}$.

It is evident from the sonograms generated from the pulse-compressed matrix (see Figure 6-4) that the SNR of the blood signal is different in each of the 2 mm receive gates. From this diagram, we can determine that the best quality signal was obtained from a receive gate located between the depths 48mm and 50 mm (Figure 6-4c). This is a much more accurate ascertainment of the vessel position than can be achieved when using conventional techniques with a 10 mm range gate.

The demodulated Doppler signals shown in Figure 6-4 allow the movement of an embolus with time to be identified. This can be compared with the stationarity of artefacts, as illustrated in Figure 6-5. This is the principle behind current multi-gate systems used to automatically discriminate between emboli and artefacts. It may be that, using the system described in this chapter i.e. extracting a number of small, non-overlapping gates from the pulse-compressed signal, multi-gate systems could be made more reliable; it would reduce the likelihood of an embolus appearing in two gates simultaneously (see Chapter 3), it would allow a shorter distance to be monitored with more gates i.e. increasing the possibility of multiple gates insonating the MCA with a separation adequate to measure a time delay between gates (Mess et al., 1999), and it may also provide a tool for investigating 'appearing' and 'disappearing' emboli (Smith et al., 1997).

Similarly, the magnitude plots in Figure 6-6 show that the path of an embolus can be more easily identified from the pulse-compressed matrix than from either the clutter-filtered RF-matrix$_{non\text{-}coded}$ or from the match-filtered matrix. The significance of this is highlighted in Figure 6-7 where a much lower amplitude embolic signal is examined. The embolus is identifiable in the sonogram and in the demodulated Doppler signal but not in the clutter-filtered RF-matrix$_{non\text{-}coded}$. The match-filtered matrix highlights the presence of an event with amplitude larger than that of the background but it is impossible to define the embolus path from this signal. However, using the pulse-compressed matrix in Figure 6-7e, the path of the embolus is clearly depicted.

The improved resolution not only allows the better location of vessels and emboli paths but may also allow us to see two or more emboli travelling side-by-side in close proximity. This may be the case in the magnitude plot of the pulse-compressed matrix shown in Figure 6-9. It is not clear from the sonogram, demodulated Doppler signal, clutter-filtered RF-matrix$_{non\text{-}coded}$ or the match-filtered matrix in Figure 6-9a - Figure 6-9d how many

emboli are present. However, when the same signal is examined in the pulse-compressed matrix in Figure 6-9e (the pulse-compressed matrix is also shown in Figure 6-9f, with a linear scale), it appears as though two emboli are travelling side-by-side. The distance between the peaks is 1.9 mm and the channel diameter in the flow phantom is approximately 4 mm; it is therefore feasible for two emboli to travel through the channel simultaneously. Both peaks are of similar magnitude so it is unlikely that one peak is a sidelobe of the other.

It can be seen from Figure 6-8 that the appearance of artefacts in the pulse-compressed matrix are as described in Chapter 3 i.e. they occur over a large number of depth samples simultaneously. All artefacts recorded in this study conformed to these findings. It has also been shown that embolic signals only appear in a few depth samples at any one time in the pulse-compressed matrix. It is therefore still possible to discriminate between emboli and artefacts in the pulse-compressed signal, using the criteria defined for the RF signal in Chapter 3.

Table 6-1 shows that the there is a considerable improvement in axial resolution in the pulse-compressed signal, when compared to the clutter-filtered RF-matrix$_{\text{non-coded}}$ and the match-filtered matrix. The presence of bone reduces the resolution improvement but not to a large degree.

Table 6-2 shows that the RSLL level is quite low in the pulse-compressed signals. This may lead to the masking of signals from smaller emboli. The RSLL level is improved when a window is applied to the transmit chirp before cross-correlating. However, the application of the window also has a detrimental effect on the axial resolution (see Table 6-1). This agrees with the findings from the model simulations (see Chapter 5).

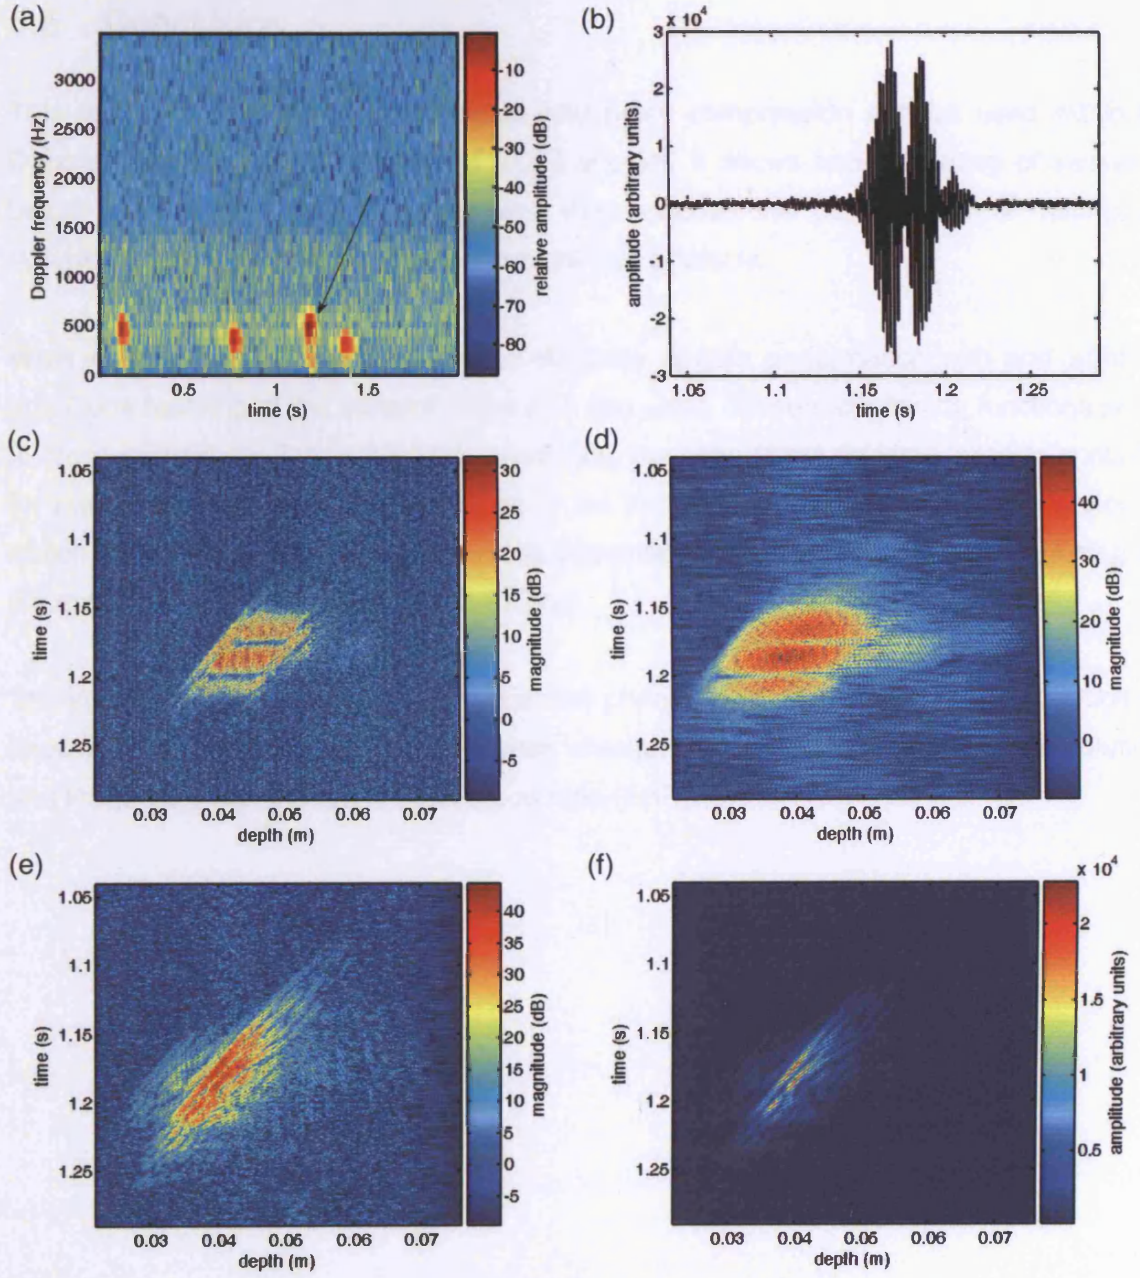

Figure 6-9 (a) Sonogram generated from RF-matrix$_{non\text{-}coded}$, (b) demodulated Doppler signal generated from RF-matrix$_{non\text{-}coded}$,, (c) clutter-filtered RF-matrix$_{non\text{-}coded}$ in the region of the embolus indicated by the arrow in (a), (d) match-filtered matrix in the region of the embolus indicated by the arrow in (a), (e) pulse-compressed matrix in the region of the embolus indicated by the arrow in (a), and (f) as per (e) but using a linear amplitude scale. A chirp with bandwidth 0.8 MHz was used.

6.5 Conclusion

This study shows that coded-excitation and pulse compression can be used within a Doppler (and, in particular, within a TCD) system. It allows better locating of vessels, better tracking of embolus position and may improve the performance of multi-gate systems used to discriminate between emboli and artefacts.

Work in this area will be continued to compare system performance with and without amplitude tapering of the transmit pulse and also using different windowing functions prior to cross-correlation. This will include windowing the copy of the transmit pulse to account for the non-uniform frequency response of the transducer. This should result in a closer match of the two signals used in the cross-correlation and hence reduce the broadening of the mainlobe in the pulse-compressed signal.

The study will also be extended using a flow phantom with the facility to inject emboli of known sizes. This will allow us to assess whether parameters such as range resolution and RSLL vary with the embolus-to-blood ratio (EBR).

7 Summary and Conclusions

7.1 Summary

A summary of past and present techniques for the detection of emboli is given in Chapter 1. This showed that the reliability of current automated techniques is highly varied and the means by which their performance is assessed (i.e. which performance measures are used) lacks consistency across the field. It was also noted that no suitable gold standard exists by which to assess the performance of automatic systems. In recent years some centres have considered the use of the TCD RF signal, as an alternative to the demodulated Doppler signal, for embolus detection. However, the use of the RF signal is still in its infancy; this paved the way for further work to be done in this area.

It was shown in Chapter 3, using events from known sources, that emboli and artefacts can be easily distinguished in the TCD RF signal. As a result, the RF signal may provide a more suitable gold standard, for differentiating between emboli and artefacts, than that currently used – the human expert. It was also found that the RF signal was a useful means for analysing emboli showers since closely spaced (in time) events could be separately identified; this is one aspect of embolus detection that is very challenging using the demodulated Doppler signal.

A method of automating the detection of emboli and their discrimination from artefacts, using the TCD RF signal, was proposed in Chapter 4. Principal components analysis (PCA) was used to characterize the typical variation in embolic signal shape, using training sets of *in vitro* and *in vivo* data. PCA techniques were then employed to discriminate between previously unseen embolic and artefact signals. Whilst the results demonstrated that this method was not yet accurate enough to be used in clinical practice, the method did show potential and should be further developed in the future.

Theoretical aspects of using coded excitation and pulse compression to improve the axial resolution in a TCD system is discussed in Chapter 5. It was found that the magnitude of the pulse compression side-effects (e.g. range sidelobes, range-Doppler coupling and frequency dependent attenuation) were small for typical TCD pulse lengths, target velocities and suitable chirp bandwidths; the improvement in axial resolution outweighed any of the negative effects.

In vitro experiments evaluating a TCD system employing coded excitation were detailed in Chapter 6. These showed that using coded excitation and pulse compression within a TCD system was feasible and that this scheme allowed better locating of vessels and better tracking of embolus position than conventional processing. Using such a system may also improve the performance of multi-gate systems used to discriminate between emboli and artefacts.

7.2 Conclusions

Embolus detection is a vast area of study. Research groups have been working for many years on several aims; the discrimination of emboli and artefacts, the sizing of emboli, and distinguishing between gaseous and particulate embolic events. The majority of previous work has been undertaken using the demodulated Doppler signal in both the time and frequency domains. As a result, a large international knowledge-base has been built up regarding the information contained in this signal on the occurrence of the various event types. Despite the large amount of work carried out on embolus detection, there are still no techniques universally accepted to answer each of the aims above. A novel approach was required that may provide additional, previously unavailable, information. Consequently, processing of the TCD RF signal was investigated in this project.

The TCD RF signal has proven to be extremely useful for the discrimination of emboli and artefacts. The analysis of emboli showers has also been made significantly easier in the 2-D RF plots. The RF signal may have a role as a new gold standard for embolus detection.

Principal component analysis (PCA) was a useful tool in characterizing the variability in embolus shapes from *in vitro* and *in vivo* training sets. The identification of new events as either emboli or artefacts, using PCA, was also successful but not to a degree of accuracy which would allow the technique to be used in clinical practice. However, the technique did show promise and the algorithms could be developed further to improve their sensitivity, specificity and PPV.

Coded excitation and pulse compression have been used within radar systems and imaging ultrasound systems for several years. The work carried out in this project has demonstrated that this processing could also be employed within a TCD system to improve the axial resolution. This would allow better locating of vessel position, better tracking of emboli and may enhance the performance of multi-gate systems designed to

automatically discriminate between emboli and artefacts. It also has the potential to allow monitoring of more closely spaced (spatially) events than was previously possible using conventional processing schemes e.g. two emboli travelling side-by-side, or two closely located vessels.

Overall, using the TCD RF signal has proven to be beneficial both in embolus detection and as a tool for the development of new processing techniques.

7.3 Future Work

Initially, more work should be carried out on the optimisation of parameters for using coded excitation and pulse compression within a TCD system. When the effects of changing parameters within such a system are fully understood, there are several studies which could be undertaken.

Firstly, a pulse compression TCD system could be used together with a flow phantom which has the facility to inject emboli of known sizes. The change in achievable axial resolution with change in embolus size could be investigated and this may be an alternative approach to the sizing of emboli in the cerebral circulation.

In addition, since narrower receive gates can be used within a pulse compression system in comparison to a conventional TCD system, a study into whether the reliability of automated embolus detection using the dual-gate method could be improved, should be performed.

It has been demonstrated that the appearance of emboli in the pulse-compressed RF-matrix is clearer and even more distinctive than in the original RF-matrix. Therefore, this opportunity should be taken to develop new automated techniques for the differentiation of emboli and artefacts, based on the additional information obtained regarding the characteristics of these events.

8 Publications & Presentations Arising from the Work

The following publications have arisen from this work:

- Cowe and Evans (2004) (abstract)
- Cowe et al. (2005a) (abstract)
- Cowe et al. (2005b)
- Cowe et al. (2006a) (abstract)
- Cowe et al. (2006b) (abstract)
- Cowe and Evans (2006)

The following presentations have arisen from this work:

- Poster presentation at 9th Meeting of the European Society of Neurosonology and Cerebral Hemodynamics. Wezlar, Germany, May 2004
- Oral presentation at the IPEM 'The Physics and Technology of Medical Ultrasound' Conference. York, UK, March 2005
- Oral presentation at 10th Meeting of the European Society of Neurosonology and Cerebral Hemodynamics. Padova, Italy, May 2005
- Oral presentation at 13th New England Doppler conference. Seattle, USA, June 2005
- Poster presentation at 11th Meeting of the European Society of Neurosonology and Cerebral Hemodynamics. Düsseldorf, Germany, May 2006
- Oral presentation at Euroson 2006, 18th European Congress of Ultrasound. Bologna, Italy, September 2006

Appendix – Principal Component Analysis

Principal component analysis (PCA) is a mathematical procedure that transforms a set of variables into a new set of variables called principal component coefficients. The main aim of PCA is to reduce the dimensionality (i.e. the number of variables) of the data whilst at the same time preserving most of the original variation. PCA produces a set of orthogonal axes in n-dimensional space such that the maximum variance in the data occurs along axis 1, the most remaining variance occurs along axis 2 etc. Each shape in our data set has 22 points and, therefore, performing PCA on this data set produces axes in 22-dimensional space. Each axis corresponds to a principal component (PC).

Therefore, given our set of aligned shapes, where $\mathbf{r}_i$ is a vector which describes the n points of the i^{th} shape in the set and each point is a radial distance from the geometrical centre of the shape, PCA can be performed as described in this appendix.

$$\mathbf{r}_i = \left(r_{i_0}, r_{i_1}, r_{i_2}, \ldots, r_{i_{n-1}}\right)^T$$

The mean shape is found using eqn (A-1), where N_s is the number of shapes in the set.

$$\bar{\mathbf{r}} = \frac{1}{N_s}\sum_{i=1}^{N_s}\mathbf{r}_i \qquad \text{(A-1)}$$

The deviations from the mean, $d\mathbf{r}_i$, are calculated for each shape in the training set using eqn (A-2).

$$d\mathbf{r}_i = \mathbf{r}_i - \bar{\mathbf{r}} \qquad \text{(A-2)}$$

The covariance matrix, **S**, is given by eqn (A-3).

$$\mathbf{S} = \frac{1}{N_s - 1}\sum_{i=1}^{N_s} d\mathbf{r}_i d\mathbf{r}_i^T \qquad \text{(A-3)}$$

The modes of variation of the points of the shape, i.e. the way in which the points in the shape move together, are described by the unit eigenvectors ($\mathbf{p}_i$) of **S**, such that eqn (A-4) holds, where λ_i is the i^{th} eigenvalue of **S** and $\lambda_i \geq \lambda_{i+1}$.

$$\mathbf{S}\mathbf{p}_i = \lambda_i \mathbf{p}_i \qquad \text{(A-4)}$$

The eigenvectors of the covariance matrix corresponding to the largest eigenvalues describe the most significant modes of variation. The proportion of the total variance given by each eigenvector is equal to its corresponding eigenvalue.

The position of a shape in the n-D PC-space, **b**, can be found using eqn (A-5), where **P** is the matrix of the eigenvectors:

$$\mathbf{P} = (\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_t, \ldots, \mathbf{p}_n)$$

$$\mathbf{b} = \mathbf{P}^T(\mathbf{r} - \bar{\mathbf{r}}) \qquad \text{(A-5)}$$

The position in PC-space can be approximated using the first *t* eigenvectors.

By reversing this process, i.e. using eqn (A-6), a shape can be reconstructed in st-ft space using a chosen number of PCs. The more PCs used, the more accurate the reconstruction. If the shape being reconstructed closely matches the shapes within the training set, it will be possible to reconstruct it using relatively few PCs.

$$\mathbf{r}_{new} = \bar{\mathbf{r}} + \mathbf{P}\mathbf{b} \qquad \text{(A-6)}$$

The reconstruction error can be estimated using a Euclidean distance between points of the original shape, **r**, and the reconstructed shape, $\mathbf{r}_{new}$. This is given by eqn (A-7).

$$Euclid\_dist = \sqrt{\sum_{j=0}^{n-1} (\mathbf{r}_j - \mathbf{r}_{new_j})^2} \qquad \text{(A-7)}$$

The variance of $\mathbf{b}_i$ over the training set is λ_i. Cootes et al. (1992) state that new examples of the shapes can be generated by varying $\mathbf{b}_i$ within suitable limits and that these limits are likely to be in the order shown in eqn (A-8), since most of the population will lie within three standard deviations of the mean. However, the fixed value of 3 standard deviations has not been used for this application. Full details are given in the 'Methods' and 'Results' sections of Chapter 4.

$$-3\sqrt{\lambda_i} \le b_i \le 3\sqrt{\lambda_i} \qquad \text{(A-8)}$$

Bibliography

Ackerstaff RGA, Babikian VL, Georgiadis D, Russell D, Siebler M, Spencer MP, Stump D. Basic Identification Criteria of Doppler Microembolic Signals: Concensus Committee of the 9th International Cerebral Hemodynamics Symposium. Stroke 1995; 26: 1123.

Aydin N and Markus HS. Optimization of Processing Parameters for the Analysis and Detection of Embolic Signals. European Journal of Ultrasound 2000; 12: 69-79.

Aydin N, Padayachee S, Markus HS. The Use of the Wavelet Transform to Describe Embolic Signals. Ultrasound Med Biol 1999; 25: 953-958.

Behar V and Adam D. Parameter Optimization of Pulse Compression in Ultrasound Imaging Systems With Coded Excitation. Ultrasonics 2004; 42: 1101-1109.

Bennett M, McLaughlin S, Anderson T, McDicken N. Filtering of Chirped Ultrasound Echo Signals With the Fractional Fourier Transform. Proc 2004 Ultrasonics Symposium 2004; 2036-2040.

British Medical Ultrasound Society. Guidelines for the Safe Use of Diagnostic Ultrasound Equipment. BMUS Bulletin 2000.

Brucher R and Russell D. Automatic Online Embolus Detection and Artifact Rejection With the First Multifrequency Transcranial Doppler. Stroke 2002; 33: 1969-1974.

Brucher R, Russell D, Frenkenberger H. Automatic Cerebral Embolus Detection Using Dual-Gate Doppler. New trends in cerebral Hemodynamics and Neurosonology 1997; 364-373.

Chiao,RY. Method and Apparatus for Flow Imaging Using Golay Codes, US Patent 6,312,384 (Issued 2001)

Chiao RY and Hao X. Coded Excitation for Diagnostic Ultrasound: a System Developer's Perspective. IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control 2005; 52: 160-170.

Cootes, T. F., Taylor, C. J., Cooper, D. H., and Graham, J. Training Models of Shape from Sets of Examples. Proceedings of the British Machine Vision Conference 1992; 9-18.

Cowe J and Evans DH. A New Approach to Automatic Detection of Emboli Based on the Cross Correlation of RF Data (Abstract). Cerebrovascular Diseases 2004. 17(Suppl 4): 16-17.

Cowe J and Evans DH. Automatic Detection of Emboli in the TCD RF Signal Using Principal Component Analysis. Ultrasound Med Biol 2006; In Press.

Cowe J, Gittins J, and Evans DH. RF Signals Aid Classification of Embolic Events Recorded During TCD Monitoring (Abstract). Cerebrovascular Diseases 2005a. 19(Suppl 1): 22.

Cowe J, Gittins J, and Evans DH. Improving the Axial Resolution in a TCD System Using Coded-Excitation. Cerebrovascular Diseases 2006a. 21(Suppl 3): 21-22.

Cowe J, Gittins J, and Evans DH. Improving the Axial Resolution in a TCD System Using Pulse Compression. Giornale Italiano di Ecografia 2006b. 9 (3 (Suppl 1)): 65.

Cowe J, Gittins J, Naylor AR, Evans DH. RF Signals Provide Additional Information on Embolic Events Recorded During TCD Monitoring. Ultrasound Med Biol 2005b; 31: 613-623.

Cullinane M, Kaposzta Z, Reihill S, Markus HS. Online Automated Detection of Cerebral Embolic Signals From a Variety of Embolic Sources. Ultrasound Med Biol 2002; 28: 1271-1277.

Cullinane M and Markus HS. Evaluation of a 1 MHz Transducer for Transcranial Doppler Ultrasound Including Embolic Signal Detection. Ultrasound Med Biol 2001; 27: 795-800.

Cullinane M, Reid G, Dittrich R, Kaposzta Z, Ackerstaff RGA, Babikian VL, Droste DW, Grossett D, Siebler M, Valton L, Markus HS. Evaluation of New Online Automated Embolic Signal Detection Algorithm, Including Comparison With Panel of International Experts. Stroke 2000; 31: 1335-1341.

Darbellay GA, Duff R, Vesin J-M, Despland P-A, Droste DW, Molina C, Serena J, Sztajzel R, Ruchat P, Karapanayiotides T, Kalangos A, Bogousslavsky J, Ringelstein EB, Devuyst G. Solid or Gaseous Circulating Brain Emboli: Are They Seperable by Transcranial Ultrasound? Journal of Cerebral Blood Flow and Metabolism 2004; 24: 860-868.

Devuyst G, Darbellay GA, Vesin J-M, Kemeny V, Ritter M, Droste DW, Molina C, Serena J, Sztajzel R, Ruchat P, Lucchesi C, Dietler G, Ringelstein EB, Despland P-A, Bogousslavsky J. Automatic Classification of HITS into Artifacts or Solid or Gaseous Emboli by a Wavelet Representation Combined With Dual-Gate TCD. Stroke 2001; 32: 2803-2809.

Devuyst G, Vesin J-M, Despland P-A, Bogousslavsky J. The Matching Pursuit: A New Method of Characterizing Microembolic Signals? Ultrasound Med Biol 2000; 26: 1051-1056.

Droste DW, Dittrich R, Hermes S, Kemeny V, Schulte-Altedorneburg G, Hansberg T, Ringelstein EB. Four-Gated Transcranial Doppler Ultrasound in the Detection of Circulating Microemboli. European Journal of Ultrasound 1999; 9: 117-125.

Droste DW, Hagedorn G, Nötzold A, Siemens H-J, Sievers HH, Kaps M. Bigated Transcranial Doppler for the Detection of Clinically Silent Circulating Emboli in Normal Persons and Patients With Prosthetic Cardiac Valves. Stroke 1997; 28: 588-592.

El-Brawany MA and Nassiri DK. Microemboli Detection Using Ultrasound Backscatter. Ultrasound Med Biol 2002; 28: 1439-1446.

Evans DH. Detection of Microemboli. In: Babikian VL, Wechsler LR, Toole JF, eds. Transcranial Doppler Ultrasonography, 2nd edn. Boston: Butterworth Heinemann, 1999: 141-155.

Evans DH. Ultrasonic Detection of Cerebral Emboli. Proc 2003 Ultrasonics Symposium 2003; 316-326.

Evans DH and Gittins J. Limits of Uncertainty in Measured Values of Embolus-to-Blood Ratios in Dual Frequency TCD Recordings Due To Nonidentical Sample Volume Shapes. Ultrasound Med Biol 2005; 31: 233-242.

Fan L, Evans DH, Naylor AR. Automated Embolus Identification Using a Rule-Based Expert System. Ultrasound Med Biol 2001; 27: 1065-1077.

Fan L, Evans DH, Naylor AR, Tortoli P. Real-Time Identification and Archiving of Micro-Embolic Doppler Signals Using a Knowledge-Based DSP System. Med Biol Eng Comput 2004; 42: 193-200.

FDA. Information for manufacturers seeking marketing clearance of diagnostic ultrasound systems and transducers. 1997

Georgiadis D, Goeke J, Hill M, König M, Nabavi DG, Stögbauer F, Zunker P, Ringelstein EB. A Novel Technique for Identification of Doppler Microembolic Signals Based on the Coincidence Method In Vitro and In Vivo Evaluation. Stroke 1996; 27: 683-686.

Georgiadis D, Uhlmann F, Astler M, Cencetti S, Zierz S. Automated Identification of Doppler Microembolic Signals: Comparison of Two Techniques. Neurological Research 2000a; 22: 738-740.

Georgiadis D, Uhlmann F, Lindner A, Zierz S. Differentiation Between True Microembolic Signals and Artefacts Using an Arbitrary Sample Volume. Ultrasound Med Biol 2000b; 26: 493-496.

Georgiadis D, Wenzel A, Zerkowski HR, Zierz S, Lindner A. Automated Intraoperative Detection of Doppler Microembolic Signals Using the Bigate Approach. Stroke 1998; 28: 137-139.

Hoeks APG, Arts TGJ, Brands PJ, Reneman RS. Comparison of the Performance of the RF Cross Correlation and Doppler Autocorrelation Technique to Estimate the Mean Velocity of Simulated Ultrasound Signals. Ultrasound Med Biol 1993; 19: 727-740.

Hunt SM, Oppenheim M, Close S, Brown PG, McKeen F, Minardi M. Determination of the Meteoroid Velocity Distribution at the Earth Using High-Gain Radar. Icarus 2004; 168: 34-42.

Iwashita, Kazuyuki, Moriya, Tadashi, Tagawa, Norio, and Yoshizawa, Masasumi. Doppler Measurement Using a Pair of FM-Chirp Signals. IEEE Ultrasonics Symposium 2003; 1219-1222.

Kelly-Hayes M, Wolf P, Kase CS, Brand FN, McGuirk JM, D'Agostino RB. Temporal Patterns of Stroke Onset: The Framingham Study. Stroke 1995; 26: 1343-1347.

Kemeny V, Droste D, Hermes S, Nabavi DG, Schulte-Altedorneburg G, Siebler M, Ringelstein EB. Automatic Embolus Detection by a Neural Network. Stroke 1999; 30: 807-810.

Keunen RWM, Stam CJ, Tavy DLJ, Mess WH, Titulaer BM, Ackerstaff RGA. Preliminary Report of Detecting Microembolic Signals in Transcranial Doppler Time Series With Nonlinear Forecasting. Stroke 1998; 29: 1638-1643.

Krongold BS, Sayeed AM, Moehring MA, Ritcey JA, Spencer MP, Jones DL. Time-Scale Detection of Microemboli in Flowing Blood With Doppler Ultrasound. IEEE Transactions on Biomedical Engineering 1999; 46: 1081-1089.

Lubbers J and van den Berg JW. An Ultrasonic Detector for Microgasemboli in a Bloodflow Line. Ultrasound Med Biol 1976; 2: 301-310.

Mahafza BR. Radar Systems Analysis and Design Using Matlab.: Chapman and Hall/CRC, 2000

Markus H. Importance of Time-Window Overlap in the Detection and Analysis of Embolic Signals. Stroke 1995; 26: 2044-2047.

Markus H, Cullinane M, Reid G. Improved Automated Detection of Embolic Signals Using a Novel Frequency Filtering Approach. Stroke 1999; 30: 1610-1615.

Markus H, Loh A, Brown MM. Computerized Detection of Cerebral Emboli and Discrimination From Artifact Using Doppler Ultrasound. Stroke 1993; 24: 1667-1672.

Markus HS and Molloy J. Use of a Decibel Threshold in Detecting Doppler Embolic Signals. Stroke 1997; 28: 692-695.

Markus HS and Reid G. Frequency Filtering Improves Ultrasonic Embolic Signal Detection. Ultrasound Med Biol 1999; 25: 857-860.

Martin K and Spinks D. Measurement of the Speed of Sound in Ethanol/Water Mixtures. Ultrasound Med Biol 2001; 27: 289-291.

Marvasti S, Gillies D, Marvasti F, Markus H. Online Automated Detection of Cerebral Embolic Signals Using a Wavelet-Based System. Ultrasound Med Biol 2004; 30: 647-653.

Mess WH, Titulaer BM, Ackerstaff RGA. An In Vivo Model to Detect Microemboli With Multidepth Technique. Preliminary Results. Cerebrovascular Disease 1996a; 6.

Mess WH, Titulaer BM, Ackerstaff RGA. Discrimination and Characterization of Emboli: Old and New Aspects. New trends in cerebral Hemodynamics and Neurosonology 1997; 355-363.

Mess WH, van Zuilen EV, Jansen C, Tweel IVD, Ackerstaff RGA. Reliability of Automated Microemboli Detection. Neurology 1996b; 46: A301-A302.

Mess WH, Willigers JM, Ledoux LAF, Ackerstaff RGA, Hoeks APG. Microembolic Signal Description: a Reappraisal Based on a Customized Digital Postprocessing System. Ultrasound Med Biol 2002; 28: 1447-1455.

Mess WH (2003) On the Characteristics and Detectability of Intracranial Microembolic Signals. Ph.D. Dissertation, University of Maastricht.

Mess WH, Titulaer BM, Ackerstaff RGA. Middle Cerebral Artery Anatomy and Characteristics of Embolic Signals: a Dual Gate Computer Simulation Study. Ultrasound Med Biol 1999; 25: 531-539.

Mess WH, Titulaer BM, Ackerstaff RGA. A New Algorithm for Off-Line Automated Emboli Detection Based on the Pseudo-Wigner Power Distribution and the Dual Gate TCD Technique. Ultrasound Med Biol 2000; 26: 413-418.

Misaridis T, Gammelmark K, Jorgensen CH, Lindberg N, Thomsen AH, Pedersen MH, Jensen JA. Potential of Coded Excitation in Medical Ultrasound Imaging. Ultrasonics 2000a; 38: 183-189.

Misaridis T and Jensen JA. An Effective Coded Excitation Scheme Based on a Predistorted FM Signal and an Optimized Digital Filter. Proc 1999 Ultrasonics Symposium 1999; 1589-1593.

Misaridis T, Pedersen MH, Jensen JA. Clinical Use and Evaluation of Coded Excitation in B-Mode Images. Proc 2000 Ultrasonics Symposium 2000b; 1689-1693.

Moehring MA and Klepper JR. Pulse Doppler Ultrasound Detection, Characterization and Size Estimation of Emboli in Flowing Blood. IEEE Transactions on Biomedical Engineering 1994; 41: 35-44.

Moehring MA and Spencer MP. Power M-Mode Doppler (PMD) for Observing Cerebral Blood Flow and Tracking Emboli. Ultrasound Med Biol 2002; 28: 49-57.

Molloy J and Markus HS. Multigated Doppler Ultrasound in the Detection of Emboli in a Flow Model and Embolic Signals in Patients. Stroke 1996; 27: 1548-1552.

Moriya,T, Tanahashi,Y, and Yoshizawa,M. Ultrasound Measurement Apparatus, US Patent 6918875 (Issued 2005)

Muir TG and Carstensen EL. Prediction of Nonlinear Acoustic Effects at Biomedical Frequencies and Intensities. Ultrasound Med Biol 1980; 6: 345-357.

Munts AG, Mess WH, Bruggemans EF, Walda L, Ackerstaff RGA. Feasibility and Reliability of On-Line Automated Microemboli Detection After Carotid Endarterectomy. A Transcranial Doppler Study. European Journal of Vascular and Endovascular Surgery 2003; 25: 262-266.

Naylor AR. Making Carotid Surgery Safer. British Medical Bulletin 2000; 56: 539-548.

Naylor AR, Hayes PD, Allroggen H, Lennard N, Gaunt ME, Thompson MM, London NJM, Bell PRF. Reducing the Risk of Carotid Surgery: A 7-Year Audit of the Role of Monitoring and Quality Control Assessment. Journal of Vascular Surgery 2000; 32: 750-759.

Nowicki A, Secomski W, Litniewski J, Trots I. On the Application of Signal Compression Using Golay's Codes Sequences in Ultrasound Diagnostic. Archives of Acoustics 2003; 28: 313-324.

Padayachee TS, Gosling RG, Bishop CC, Burnand K, Browse NL. Monitoring Middle Cerebral Artery Blood Velocity During Carotid Endarterectomy. British Journal of Surgery 1986; 73: 98-100.

Palanchon P, Bouakaz A, Klein J, DE Jong N. Subharmonic and Ultraharmonic Emissions for Emboli Detection and Characterization. Ultrasound Med Biol 2003; 29: 417-425.

Palanchon P, Bouakaz A, Klein J, DE Jong N. Emboli Detection Using a New Transducer Design. Ultrasound Med Biol 2004; 30: 123-126.

Palanchon P, Bouakaz A, van Blankenstein JH, Klein J, Bom N, DE Jong N. New Technique for Emboli Detection and Discrimination Based on Nonlinear Characteristics of Gas Bubbles. Ultrasound Med Biol 2001; 27: 801-808.

Pedersen MH, Misaridis T, Jensen JA. Clinical Comparison of Pulse and Chirp Excitation. Proc 2002 Ultrasonics Symposium 2002; 1632-1635.

Pedersen MH, Misaridis T, Jensen JA. Clinical Evaluation of Chirp-Coded Excitation in Medical Ultrasound. Ultrasound Med Biol 2003; 29: 895-905.

Perkins MA. A Versatile Force Balance for Ultrasound Power Measurement. Physics in Medicine and Biology 1989; 34: 1645-1651.

Ramnarine KV, Anderson T, Hoskins PR. Construction and Geometric Stability of Physiological Flow Rate Wall-Less Stenosis Phantoms. Ultrasound Med Biol 2001; 27: 245-250.

Ramnarine KV, Nassiri DK, Hoskins PR, Lubbers J. Validation of a New Blood-Mimicking Fluid for Use in Doppler Flow Test Objects. Ultrasound Med Biol 1998; 24: 451-459.

Rao NAHK. Investigation of a Pulse Compression Technique for Medical Ultrasound: a Simulation Study. Med Biol Eng Comput 1994; 32: 181-188.

Ringelstein EB, Droste DW, Babikian VL, Evans DH, Grosset DG, Kaps M, Markus HS, Russell D, Siebler M. Consensus on Microembolus Detection by TCD. Stroke 1998; 29: 725-729.

Russell D and Brucher R. Online Automatic Discrimination Between Solid and Gaseous Cerebral Microemboli With the First Multifrequency Transcranial Doppler. Stroke 2002; 33: 1975-1980.

Russell D and Brucher R. Embolus Detection and Differentiation Using Multifrequency Transcranial Doppler. Stroke 2005; 36: 706.

Saqqur M, Dean N, Schebel M, Hill MD, Salam A, Shuaib A, Demchuk AM. Improved Detection of Microbubble Signals Using Power M-Mode Doppler. Stroke 2004; 35: e14-e17.

Skolnik MI. Introduction to Radar Systems. In: Cerra FJ, ed.: McGraw Hill, 1981

Smith JL, Evans DH, Fan L, Bell PRF, Naylor AR. Differentiation Between Emboli and Artefacts Using Dual-Gated Transcranial Doppler Ultrasound. Ultrasound Med Biol 1996; 22: 1031-1036.

Smith JL, Evans DH, Naylor AR. Signals From Dual Gated TCD Systems: Curious Observations and Possible Explanations. Ultrasound Med Biol 1997; 23: 15-24.

Smith SW. The Scientist and Engineer's Guide to Digital Signal Processing. San Diego: California Technical Publishing, 1997

Spencer MP. Transcranial Doppler Monitoring and Causes of Stroke From Carotid Endarterectomy. Stroke 1997; 28: 685-691.

Spencer MP, Thomas GI, Nicholls SC, Sauvage LR. Detection of Middle Cerebral Artery Emboli During Carotid Endarterectomy Using Transcranial Doppler Ultrasonography. Stroke 1990; 21: 415-423.

Topic Working Group. NHS R&D Strategic Review: Coronary Heart Disease and Stroke. 1999

van Zuilen EV, Mess WH, Jansen C, van der Tweel I, van Gijn J, Ackerstaff RGA. Automatic Embolus Detection Compared With Human Experts: A Doppler Ultrasound Study. Stroke 1996; 27: 1840-1843.

Wilhjelm JE and Pedersen PC. Target Velocity Estimation With FM and PW Echo Ranging Doppler Systems - Part I: Signal Analysis. IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control 1993a; 40: 366-372.

Wilhjelm JE and Pedersen PC. Target Velocity Estimation With FM and PW Echo Ranging Doppler Systems - Part II: Systems Analysis. IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control 1993b; 40: 373-380.

Wilhjelm JE and Pedersen PC. Analytical and Experimental Comparisons Between the Frequency-Modulated Frequency-Shift Measurement and the Pulsed-Wave Time-Shift Measurement Doppler Systems. J Acoust Soc Am 1996; 100: 3957-3970.