# Ultra-fast Imaging of Two-Phase Flow in Structured Monolith Reactors; Techniques and Data Analysis

University of Cambridge
Department of Chemical Engineering

**Jonathan Jaime Heras**
**Churchill College**

A dissertation submitted for the Degree of
Doctor of Philosophy

**September 2006**

## Preface

The work presented in this dissertation was carried out in the Department of Chemical Engineering, University of Cambridge between October 2002 and December 2005. All of the work described in this dissertation is believed to be original, except where explicit reference is made to other authors. This dissertation is all my own work and includes nothing that is done in collaboration, except where specifically stated in the text. This dissertation has not been previously submitted in part, or in whole, to any other universities for degrees, diplomas or other qualifications. This dissertation does not exceed 150 figures, nor 65,000 words.

## Acknowledgement

I wish to express my gratitude to Professor Lynn Gladden for offering me the privilege to work in the Magnetic Resonance Research Centre (MRRC) and for her guidance and advice during the course of my research.

I would like to acknowledge the support of the members of the Magnetic Resonance and Catalysis group, especially Dr. Mick Mantle and Dr. Andrew Sederman for sharing their nuclear magnetic resonance (NMR) experience with me, Dr. Matthew Lim and Dr. Mark Sains for helpful discussions and helping me hold together my sanity, and the rest of the group for making the MRRC such a pleasant place to work. Also, I would like to thank Dr. Paul Alexander of the Cavendish Laboratory for his data analysis program IDA, without which the data analysis methods implemented in this thesis would not have been possible.

I would also like to acknowledge the Engineering and Physical Sciences Research Council (EPSRC) which financed my research and study in Cambridge.

Jon, Cambridge, 2006

## Summary

This thesis will address the use of nuclear magnetic resonance (NMR) and magnetic resonance imaging (MRI) techniques to probe the "monolith reactor", which consists of a structured catalyst over which reactions may occur. This reactor has emerged as a potential alternative to more traditional chemical engineering systems such as trickle bed and slurry reactors. However, being a relatively new design, its associated flow phenomena and design procedures are not rigorously understood, which is retarding its acceptance in industry. Traditional observations are unable to provide the necessary information for design since the systems are opaque and dynamic. Therefore, NMR is proposed as an ideal tool to probe these systems in detail.

The theory of NMR is summarised and the development of novel NMR techniques is presented. Novel techniques are validated in simple systems, and tested in more complex systems to ascertain their quantitative nature, and to find their limitations. These techniques are improvements over existing techniques in that they either decrease the acquisition time (allowing the observation of dynamically-changing systems) or allow us to probe systems in different ways to extract useful information. The goal of this research is to better understand the flow phenomena present in such systems, and to use this information to design better, more efficient, more controllable industrial reactors.

The analysis of the NMR data acquired is discussed in detail, and several novel image-processing techniques have been developed to aid in the quantification of features within the images, and also to measure quantities such as holdup and velocity. These novel techniques are validated, and then applied to the systems of interest.

Various configurations of monolith reactor, ranging from low flow rate systems to more challenging (and more industrially relevant) turbulent systems, are probed using these methods, and the contrasting flow phenomena and performance of these systems are discussed, with a view to optimisation of the choice of design parameters.

# Chapter 1 – Introduction

---

**CHAPTER 1 – INTRODUCTION**



## 1.1. An overview of the project

The main interest of the Magnetic Resonance and Catalysis Group (MRRC) in the Department of Chemical Engineering, University of Cambridge is to apply nuclear magnetic resonance (NMR) techniques to advance the knowledge of various aspects of chemical engineering. NMR is a powerful analytical technique because it is inherently non-invasive and non-intrusive. It can be both spatially and chemically selective. Its power lies in the large array of contrasting mechanisms available, for example relaxation and diffusion weighting. Chemical reactor engineering is one of the key areas of interest in the research group. Quantitative information relevant to chemical reactor study that can be obtained from NMR includes:

- Bulk and local chemical composition measurement
- 3D structural information, even in opaque systems
- Transport parameters such as velocity, acceleration, diffusion and dispersion

From the list above, it is clear why NMR is an ideal tool in the study of chemical reactor engineering. A single technique provides all the information needed to study such reactors. However, NMR does have its disadvantages such as high initial capital and maintenance cost, and limitations in terms of sample size and suitable materials for study (i.e. samples of non-metallic construction). In general, NMR is not envisioned as a piece of process equipment, but rather as a research tool to help understand flow phenomena and to aid the design of industrial equipment.

With our expertise in NMR and data analysis techniques, we would like to answer three fundamental questions relevant to chemical reactor engineering:

- Given that conventional magnetic resonance imaging (MRI) techniques are "slow", can we increase the speed of data acquisition?
- Can we use and develop traditional techniques to obtain information on dynamic systems?
- Can we develop data analysis tools to extract the maximum possible amount of information from data acquired with a particular NMR technique?

Significant progress has been made on these three questions. The main system of interest has been the "monolith reactor", which utilises a structured catalyst for reaction. The techniques and discussions developed in this thesis are also relevant to other areas of research. The organisation of the work presented in this thesis will now be outlined.

## 1.2. The organisation of the thesis

This thesis is organised in fifteen chapters. A brief summary of each chapter is presented below:

**Chapter 1** **Introduction**

An introduction to this project and why NMR is used.

**Chapter 2** **Principles of Nuclear Magnetic Resonance**

NMR is a vast subject. In this chapter, principles of traditional NMR techniques which are relevant to this research are presented. This chapter is useful for readers who are not familiar with the subject.

**Chapter 3** **The monolith reactor**

So-called "monolith reactors" are emerging as an industrial alternative to trickle beds and slurry reactors. This chapter consists of a literature review on the current state of research regarding structured monolith reactors.

**Chapter 4 Tomographic methods**

An outline of tomographic measurement techniques is presented, with a focus on MRI and a discussion of its applicability to our systems of interest.

**Chapter 5 Novel MRI techniques**

Some novel, rapid MRI techniques which have been developed as part of this work are presented. These are specifically intended for probing dynamic systems such as the monolith reactor. The quantitative nature and limitations of these new techniques are highlighted.

**Chapter 6 Single-phase flow in a monolith**

The novel MRI techniques of Chapter 5 are applied to single-phase flow in a monolith reactor, to investigate their quantitative nature in systems of porous media.

**Chapter 7 Image analysis techniques**

This chapter covers some of the fundamentals of image analysis including image manipulation, Fourier space processing, texture analysis, distribution quantification and feature tracking. Novel analysis tools for processing MRI data are presented. These techniques are evaluated and later applied to monoliths in Chapters 9-11.

**Chapter 8 Monolith configurations and holdup-measurement techniques**

This chapter presents a reference of the various monolith reactor configurations, and also presents NMR techniques for measuring the holdup in these process configurations.

**Chapter 9 Bubble and slug size analysis**

This chapter presents the application of image analysis techniques to MRI data with the aim of measuring bubble and slug size distributions in the monolith reactor. The process is outlined, and observations of the flow phenomena are discussed.

**Chapter 10 Quantification of flow maldistribution**

This chapter presents the application of novel processing techniques developed in Chapter 7, with the aim of quantifying the quality of fluid distribution over the monolith cross-section.

**Chapter 11 Two-phase velocity measurement in the monolith reactor**

This chapter presents the application of novel MRI velocity measurement techniques to two-phase flow configurations. The applicability of these techniques to such dynamic systems is discussed, and results are presented where appropriate.

**Chapter 12 Summary of experimental findings**

This chapter is a summary of the findings of Chapters 8-11 for easy reference. The performance of the novel MRI techniques is discussed, and then the observations of flow phenomena in the various process configurations are summarised.

**Chapter 13 Deconvolution techniques**

A review of deconvolution literature is presented, highlighting some of the procedures for post-processing convolved data (such as may be obtained in MRI experiments). These techniques are applied to data from monolith systems in Chapter 14.

**Chapter 14 Propagator deconvolution**

The results of the novel line-excitation technique of Chapter 5 are processed using a novel deconvolution algorithm, with the intention of extracting propagator-type information. The procedure is outlined, and results are presented and discussed.

**Chapter 15 Conclusions and further work**

Overall conclusions from this research are drawn and further work is proposed.

## 1.3. Nomenclature

| Abbreviations | Definition |
|---|---|
| MRI | Magnetic resonance imaging |
| MRRC | Magnetic Resonance Research Centre |
| NMR | Nuclear magnetic resonance |
| RARE | Rapid acquisition with relaxation enhancement |

# Chapter 2 – Principles of Nuclear Magnetic Resonance

In this chapter, the basic theory of nuclear magnetic resonance (NMR) and magnetic resonance imaging (MRI) is presented. Since NMR is a very wide field, this chapter focuses specifically on the fundamentals of NMR and the techniques relevant to this research.

**CHAPTER 2 – PRINCIPLES OF NUCLEAR MAGNETIC RESONANCE**



## 2.1. NMR background

The first experimental observations of the nuclear magnetic resonance (NMR) phenomenon were made independently in 1946 by Bloch at Stanford (Bloch, 1946) and Purcell at Harvard (Purcell *et al.*, 1946), and they received the Nobel prize in

physics in 1952. They had discovered that certain nuclei could absorb energy in the radiofrequency range of the electromagnetic spectrum when placed in a magnetic field, and upon their return to equilibrium they would re-emit this energy.

Lauterbur (1973) presented a paper describing a new imaging technique that utilised a weak magnetic field with a spatially-varying gradient in addition to the main magnetic field. This allowed him to spatially locate two tubes of water by a back-projection method. Sir Peter Mansfield was credited with further developing these magnetic field gradients for the use of MRI in medicine, and recently in 2003 both Lauterbur and Mansfield were awarded a Nobel Prize for their contribution.

In a time span of less than sixty years, NMR has transformed from an obscure physicists' curiosity into one of the most powerful analytical tools available today and NMR is now routinely used by research scientists. Since the first magnetic resonance image (MRI) was taken in 1973, MRI has been developed rapidly and has become an indispensable tool in medicine. Applications of NMR and MRI to engineering disciplines, such as chemical engineering (Gladden, 1994; Gibbs and Hall, 1996) are relatively recent.

There are many excellent books on basic NMR and MRI theory (Fukushima and Roeder, 1983; Harris, 1986; Callaghan, 1991; Haacke *et al.*, 1999). In this chapter the basic NMR principles will first be illustrated, followed by discussions of the NMR techniques relevant to this research, namely various forms of MRI.

### 2.1.1. NMR fundamentals

The NMR process stems from the interaction of nuclear spin with an external magnetic field. This section will cover the basic quantum mechanical principles, and describe the effects of an external magnetic field on the nuclear spin.

Nuclei are composed of neutrons and protons. Nuclei with even number of neutrons and protons will have no spin (i.e. intrinsic angular momentum) and are hence NMR inactive, e.g. $^{12}C$, $^{16}O$ will give no NMR signal. Nuclei with odd numbers of neutrons

and protons will have an integer spin quantum number, $I$, for example: $^{2}$H, $^{14}$N (for which $I = 1$). Finally, nuclei with a mix of odd and even numbers of neutrons and protons will have a half-integer spin quantum number, for example: $^{1}$H, $^{13}$C (for which $I = ½$). In this research, all the work is done on $^{1}$H nuclei only.

The $^{1}$H nucleus is so useful because it has a very high natural abundance – the $^{1}$H isotope has an abundance of 98.4%, whereas $^{13}$C accounts for only 1.1% of all C nuclei. The sensitivity of the $^{1}$H nucleus is also very high, around 50 times that of the $^{13}$C nucleus, and moreover $^{1}$H is present in water in high concentrations, which makes flow imaging relatively straightforward, and eliminates the need for expensive isotope-enriched fluids.

In the absence of a magnetic field, the spin angular momentum vector has no preferred direction. However, when atoms with unpaired nuclear spins are placed in a magnetic field $\mathbf{B_0}$ (say in the $z$-direction), they will either align with or against the static magnetic field in low and high energy states. Under thermal equilibrium, the nuclei will spread out amongst the permitted energy levels according to the Boltzmann distribution given in Eq. 2.1.

$$\frac{N_\beta}{N_\alpha} = \exp\left(\frac{-\hbar\gamma\mathbf{B}_0}{k_B T}\right) \qquad \textbf{(Eq. 2.1)}$$

where $N_\alpha$, $N_\beta$ are the population of nuclei in the upper or lower energy states respectively, $k_B$ is the Boltzmann constant, $T$ is the absolute temperature, $\hbar$ is the Planck constant divided by $2\pi$, and $\gamma$ is the gyromagnetic ratio. The gyromagnetic ratio is nucleus specific, and has a value of 42.58 MHz/Tesla for $^{1}$H.

The spin-angular momentum vector of spin-½ nuclei has two permitted directions, which can be represented pictorially, as shown in Figure 2.1.

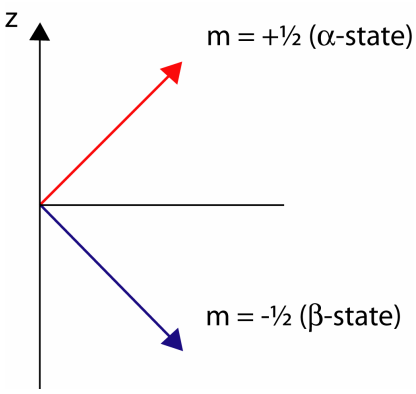


Figure 2.1 – Possible orientations of the angular momentum vectors for a spin-½ nucleus. This splitting of the states is also known as Zeeman splitting.

For $^{1}$H nuclei in a 4.7 T static magnetic field at a temperature of 298 K, the ratio between nuclei in the upper and lower states is 0.99997, i.e. there are slightly more nuclei in the lower energy state. This small excess in the lower energy state produces the net magnetisation in the sample. It might be assumed that no significant signal would be detected under these conditions since the spin excess is millions of times smaller than the total number of spins. However, the density of these spins is so great (1 ml of water contains approximately $1\times10^{23}$ $^{1}$H nuclei), that we obtain a measurable net magnetisation vector, **M**. This may be viewed as a bulk property, and hence we may visualise the system in classical terms.

A transition between the two energy states of Figure 2.1 is possible by the absorption of a photon of the correct energy corresponding to the energy difference between the two states. A particle in the lower state absorbs the photon, and ends up in the higher energy state. If left to thermal equilibrium, photons will be re-emitted as the spins return from the high to the low energy state, and this constitutes the NMR signal we measure. The energy ($\Delta E$) and frequency ($\nu$) of the photos are defined in Eq. 2.2.

$$\Delta E = h\nu = \hbar\gamma\mathbf{B}_0 \text{, and so } \nu = \frac{\gamma\mathbf{B}_0}{2\pi} \qquad \textbf{(Eq. 2.2)}$$

For example, for $^{1}$H nuclei in a static magnetic field of 4.7 T, electromagnetic radiation tuned to a frequency of 200 MHz is needed to cause a transition from the lower to the higher energy level. This corresponds to a region within the radio frequency range of the electromagnetic spectrum.

## 2.1.2. Spin excitation and the rotating frame

The motion of the net magnetisation vector, **M,** may be treated using classical mechanics. The torque is equated to the rate of change of angular momentum within the static magnetic field $\mathbf{B_0}$, given by Eq. 2.3.

$$\frac{dM}{dt} = \gamma \mathbf{M} \times \mathbf{B_0} \qquad \textbf{(Eq. 2.3)}$$

The solution to Eq. 2.3 implies that there is a precession of the magnetisation vector about the magnetic field at a rate of $\omega_0 = \gamma \mathbf{B_0}$, which is known as the Lamor frequency (Callaghan, 1991).

The application of an oscillating transverse magnetic field, $\mathbf{B_1}$, causes the magnetisation vector **M** to simultaneously precess about the longitudinal polarising field $\mathbf{B_0}$ at $\omega_0$ and about the radio frequency (RF) field $\mathbf{B_1}$ at $\omega_1$. As this precession is rather complex, it is common to express spin vectors in a rotating frame which rotates about $\mathbf{B_0}$ at the Lamor frequency, $\omega_0$, and has an effective longitudinal field of $(\mathbf{B}_0 - \omega/\gamma)$. Thus, if an RF pulse is applied on resonance, i.e. $\omega = \omega_0$, **M** will precess about $\mathbf{B_1}$ as shown in Figure 2.2.

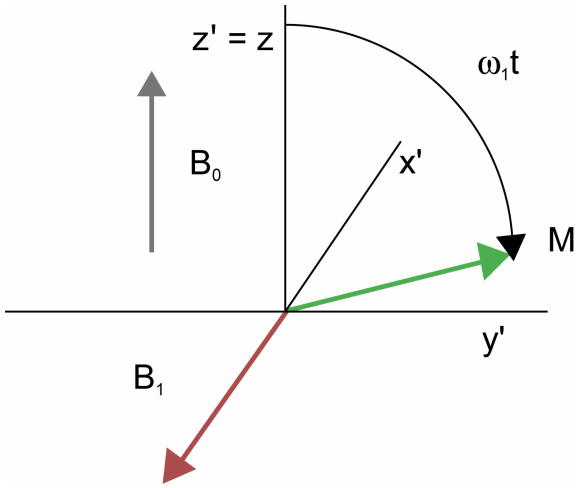


Figure 2.2 – The evolution of magnetisation under the influence of $\mathbf{B_0}$ and $\mathbf{B_1}$, observed in the rotating frame with transformed axes x', y', z'.

An RF pulse with appropriate amplitude (power) and duration which rotates **M** completely into the transverse plane is termed a 90° or π/2 pulse. If the duration of

the RF pulse is doubled, **M** is rotated through π radians and this pulse is called a 180° or π pulse. RF pulses which have sufficient bandwidth to perturb all the spins within the sample are called "hard" or non-selective pulses, whilst those which only perturb a select group of spins are called "soft" or selective pulses.

### 2.1.3. Signal detection and post-processing

After the spins have been excited with a RF pulse, they undergo free precession in the transverse plane as they return to the lower energy equilibrium state. An oscillatory voltage is induced in the RF coil, which acts as a receiver as well as the transmitter of RF excitation pulses. The induced voltage measured is on the order of μV, and its amplitude is proportional to the transverse magnetisation that exists within the RF coil. This induced signal is termed as free induction decay (FID) and it is dampened by $T_2$-relaxation (Section 2.1.6). The mathematical description of an FID following a 90° pulse is given by Eq. 2.4.

$$M(t) = M_0 \exp(i\varpi_0 t)\exp\left(-\frac{t}{T_2}\right) \qquad \textbf{(Eq. 2.4)}$$

An on-resonance spin system yields a simple decaying exponential. Figure 2.3 shows an example of an off-resonance FID. The return to equilibrium is termed "relaxation", and there are two important physical phenomena which govern this relaxation rate; (i) spin-lattice ($T_1$), and (ii) spin-spin ($T_2$) interactions. These are covered in Sections 2.1.5 and 2.1.6.

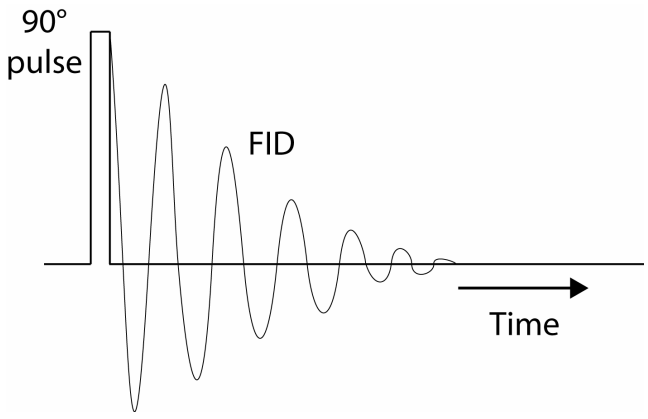


Figure 2.3 – An example of the free-induction decay (FID) resulting from the "pulse-acquire" sequence.

The FID contains signal information in the time domain. By performing a Fourier transformation, the same information can be represented in the frequency domain, which is generally more useful. The basic Fourier transformation equations are shown in Eq. 2.5 and Eq. 2.6. Since we are dealing with digitised data, in reality we use a discrete fast Fourier transform (Cooley and Turkey, 1965).

$$F(v) = \int_{-\infty}^{\infty} f(t)\exp(-i2\pi vt)dt \qquad \textbf{(Eq. 2.5)}$$

$$f(t) = \int_{-\infty}^{\infty} F(v)\exp(+i2\pi vt)dv \qquad \textbf{(Eq. 2.6)}$$

Figure 2.4 shows a typical NMR spectrum, whose peak has a Lorentzian line shape. The width of the peak is typically characterised by the Full-Width-Half-Maximum (FWHM) of the spectrum, given by $1 / \pi T_2^*$. ($T_2^*$ is defined in Section 2.1.6).

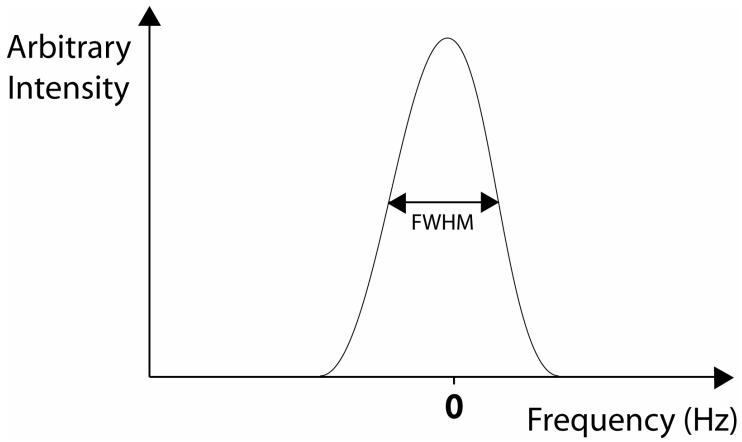


Figure 2.4 – A typical NMR absorption spectrum. The FWHM of the spectrum is given by $1 / \pi T_2^*$.

## 2.1.4. Chemical Shift

The phenomenon of chemical shift arises because of the shielding of nuclei from the external magnetic field by circulating electrons. Each unique position within a molecular structure will experience a slightly different local electronic environment, and will therefore experience different degrees of shielding. This effectively shifts the local Lamor frequency to a slightly different value, and this small effect is sufficient to be measured. The chemical shift of a particular nucleus in a molecule is

normally measured in terms of the difference in resonance frequency between the nucleus of interest, $v$, and a reference nucleus, $v_{ref}$, in dimensionless form (Eq. 2.7).

$$\delta = \frac{10^6 (v - v_{ref})}{v_{ref}} \qquad \textbf{(Eq. 2.7)}$$

The reference nucleus for $^1H$, $^{13}C$ and $^{29}Si$ NMR spectroscopy is normally tetramethylsilane (TMS), $Si(CH_3)_4$ i.e. the chemical shift of the $^1H$ nuclei in TMS is defined as zero. So $\delta$ is a molecular property independent of the static magnetic field strength and is normally quoted in parts per million (ppm). Figure 2.5 shows a spectrum of methanol to illustrate the chemical shift phenomenon. This technique is the basis of NMR spectroscopy, but is not the main focus of this work.

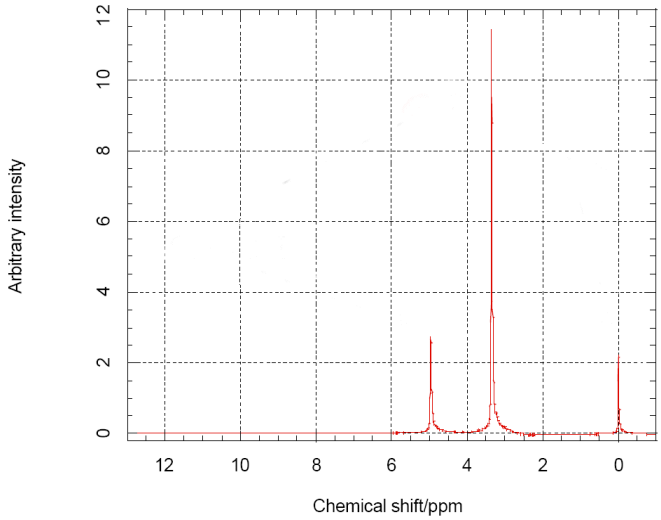


Figure 2.5 – $^1H$ spectrum of methanol taken at 200 MHz. Methanol has two functional groups: methyl ($CH_3$) and hydroxyl (OH). The chemical shifts of $CH_3$ and OH are 3.35 and 4.95 ppm respectively, with respect to the proton resonance of TMS (0 ppm). Image adapted from Yuen (2003).

## 2.1.5. Spin-lattice relaxation ($T_1$)

In NMR there are two distinct types of nuclear spin relaxation ($T_1$ and $T_2$ relaxation) which return a perturbed system to its thermal equilibrium state. This rate of regrowth of the longitudinal signal is characterised by a time constant $T_1$, known as the "spin-lattice" relaxation time. This type of relaxation stems from an exchange of energy between the spin system and its surroundings (referred to as the lattice by

convention). There are various mechanisms causing $T_1$ relaxation (Fukushima and Roeder, 1983), but they will not be discussed here. For liquid samples, the recovery of the $M_z$ magnetisation following a 90° excitation pulse is described by the differential equation in Eq. 2.8.

$$\frac{dM_z}{dt} = -\frac{(M_z - M_0)}{T_1} \qquad \textbf{(Eq. 2.8)}$$

where $M_0$ is the equilibrium magnetisation. This equation may be easily solved with appropriate initial conditions to yield Eq. 2.9.

$$M_z(t) = M_z(0)\exp\left(-\frac{t}{T_1}\right) + M_0\left[1 - \exp\left(-\frac{t}{T_1}\right)\right] \qquad \textbf{(Eq. 2.9)}$$

where $M_z(0)$ is the initial magnetisation along the $z$-axis, $M_0$ is the equilibrium magnetisation along the $z$-axis and $T_1$ is the characteristic spin-lattice relaxation time constant. At room temperature, typical values of $T_1$ relaxation constants for protons lie between 0.1 and 10 seconds for pure liquids and for liquids confined in porous solids. One method for measuring $T_1$ is the "inversion recovery" method (Vold *et al.*, 1968), shown in Figure 2.6.

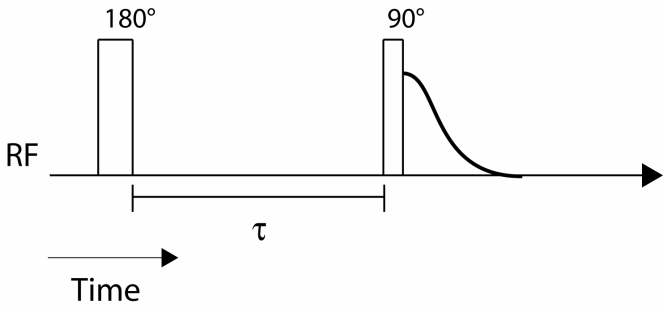


Figure 2.6 – The inversion recovery pulse sequence. The 90° pulse is also known as the "read" pulse.

The spin system is subjected to a complete inversion by means of a 180° pulse, in such a way that the magnetisation remains along the $z$-axis during the evolution period, $\tau$. At the end of $\tau$, the magnitude of the magnetisation is read out with a 90° pulse. For perfect inversion, $M_z(0) = -M_0$ and Eq. 2.9 reduces to Eq. 2.10.

$$M_z(\tau) = M_0\left[1 - 2\exp\left(-\frac{\tau}{T_1}\right)\right] \qquad \textbf{(Eq. 2.10)}$$

In a typical $T_1$ experiment, a number of spectra are acquired with different values of $\tau$. The data are then fitted with Eq. 2.10 to determine the value(s) of $T_1$ which characterise the system. It is interesting to note that in inversion recovery, the magnetisation initially has a negative value and becomes positive when $\tau$ is large. This implies that at a certain value of $\tau$, the observed signal, $M_z(\tau)$, will be zero. This particular $\tau$ is known as the "nulling point" and it is determined by setting Eq. 2.10 equal to zero. The nulling point is of use, particularly if we have more than one chemical species present and we wish to quell the signal from one of them; by selecting an evolution time which corresponds to the nulling point of the species we wish to suppress, we remove all of the signal corresponding to that species, while retaining a significant proportion of the other species, illustrated in Figure 2.7.

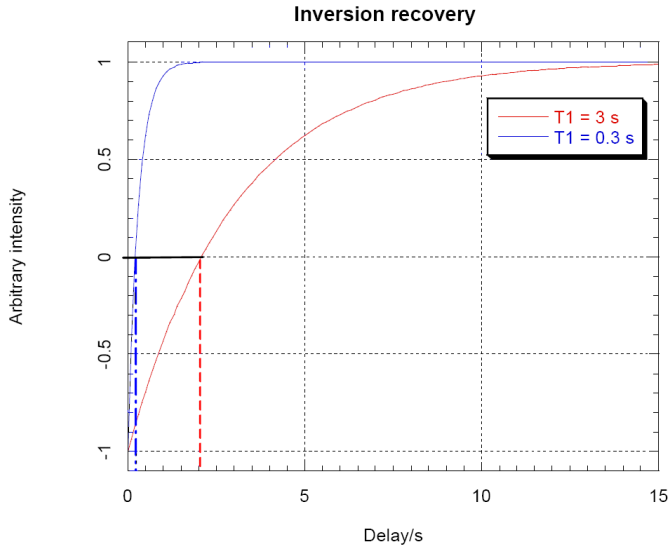


Figure 2.7 – Recovery of longitudinal magnetisation in inversion recovery as a function of delay time, $\tau$. Recovery curves are shown for two chemical species with different spin-lattice relaxation times. Nulling points are identified with dotted lines. Reproduced from Yuen (2003).

It is important to note that the amount of signal available in a NMR experiment is directly proportional to the amount of initial longitudinal magnetisation. Therefore sufficient time should be allowed for $T_1$ relaxation between successive repetitions of

RF excitations. Typically, a rest-period of $5\times T_1$ is recommended between experiments to allow full relaxation (~99%) to be achieved. Insufficient $T_1$ relaxation will lead to $T_1$-weighting of the data.

## 2.1.6. Spin-spin relaxation ($T_2$) and the spin echo

Within the magnetic field in which the sample is placed, the spin system is exposed to local field inhomogeneities which cause the spins to rotate at the Lamor frequencies corresponding to the particular local magnetic field of the spins. This slight difference causes dephasing of the component spins of the original magnetisation vector $\mathbf{M_0}$. The spin-spin relaxation time, $T_2$, characterises the loss of phase coherence within the spin system. For a simple liquid, $T_2$ relaxation can be modelled as exponential decay, as in Eq. 2.11.

$$M_{x,y}(t) = M_{x,y}(0)\exp\left(-\frac{t}{T_2}\right) \qquad \textbf{(Eq. 2.11)}$$

In addition to this effect, molecular interactions also cause a loss of phase coherence. This type of relaxation stems from random events which tend to realign the nuclei within a sample to re-establish thermal equilibrium.

The total apparent $T_2$ relaxation, $T_2^*$, is the sum of the spin-spin relaxation process ($T_2$) and the coherent loss associated with magnetic field inhomogeneity ($T_2$'). These are combined in the manner show in Eq. 2.12.

$$\frac{1}{T_2^*} = \frac{1}{T_2'} + \frac{1}{T_2} \qquad \textbf{(Eq. 2.12)}$$

The Hahn echo pulse sequence (Hahn, 1950) shown in Figure 2.8 reverses the effect of magnetic field inhomogeneity ($T_2$'). After initial excitation by a 90° pulse, the magnetisation vectors fan out (dephase) because of the magnetic field inhomogeneity and thus the total transverse signal decreases. In order to recover signal lost to $T_2$' relaxation, a 180° pulse is applied after a delay $\tau$, which inverts the phase of the magnetisation, and the dephasing of the spins is reversed. Since the rate at which phase is accumulated by the spins is unchanged, after a further delay $\tau$, a spin echo (SE) is formed as all the spins return to the original, coherent direction. The

dephasing caused by the magnetic field homogeneity is then said to be refocused. This process is illustrated in Figure 2.9. The amplitude of the echo is modulated by the molecular $T_2$ relaxation, which cannot be reversed. The spin echo is a valuable method of retaining signal which would otherwise decay, and is the basis of many MRI techniques.

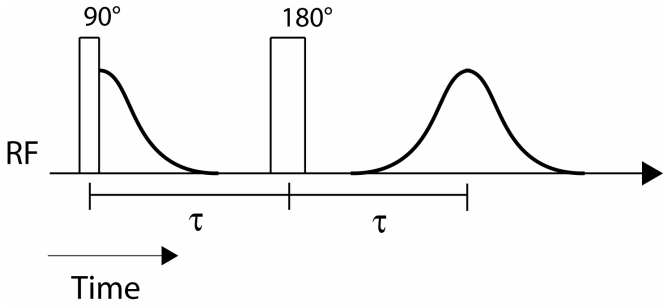


Figure 2.8 – The Hahn echo pulse sequence. A spin echo is formed at time $2\tau$ after the initial excitation.

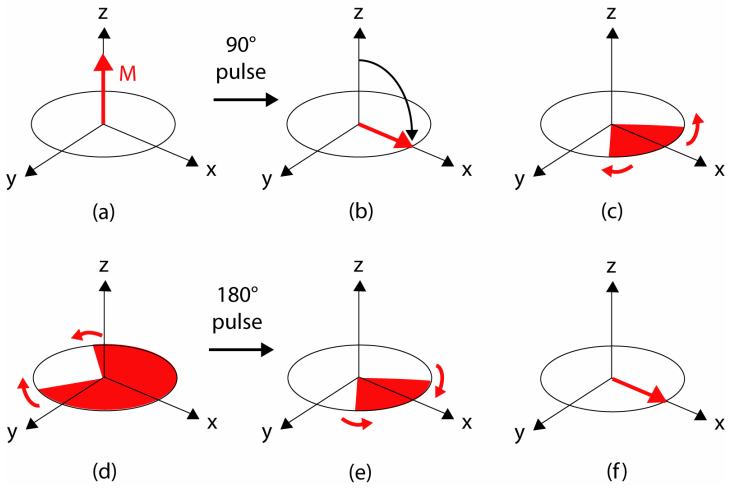


Figure 2.9 – Pictorial explanation of the Hahn echo: (a) the system is at equilibrium; (b) after a 90° excitation pulse, causing the magnetisation vector to rotate into the transverse plane; (c) the spins dephase due to local differences in Lamor frequency; (d) after a delay, $\tau$, a 180° pulse is applied which reverses the phase of all the magnetisations but does not alter the precession frequencies; (e) the spins are now moving back towards coherence and (f) after a further delay $\tau$, the spins form a spin echo.

Values of $T_2$ can be measured by performing a number of Hahn echo experiments with different $\tau$ delays and then fitting the data to Eq. 2.11.

It is important to note that if there is substantial molecular diffusion during the $2\tau$ delay, then the Hahn echo experiment will not completely refocus all of the phase coherence loss, since the spins may have moved to an environment of different local field strength during this time. Constant refocusing with a much shorter $\tau$ alleviates this effect and this sequence is known as the Carr-Purcell-Meiboom-Gill (CPMG) sequence (Carr and Purcell, 1954; Meiboom and Gill, 1958), and is shown in Figure 2.10.

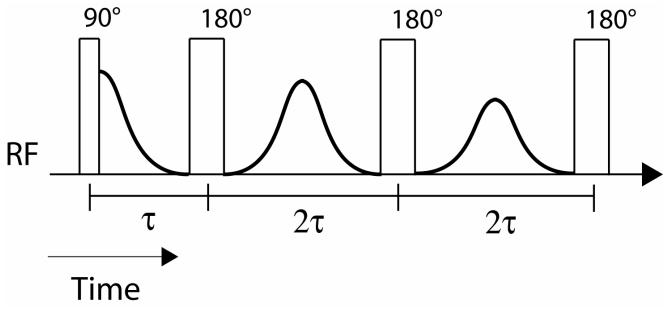


Figure 2.10 – The CPMG pulse sequence; a spin echo is formed after each 180° pulse. The amplitude of each echo is $T_2$ modulated.

## 2.1.7. Stimulated echo

If we have a sample with a very short $T_2$, the echo formed by a spin echo will be of low signal strength due to $T_2$ relaxation. However, if we have a considerably longer $T_1$ (as is the case for many liquids), we may store the magnetisation is such a way that the signal is $T_1$ attenuated rather than $T_2$ attenuated. This RF pulse sequence is called a stimulated echo (StE). Here three 90° pulses are applied to the sample, shown in Figure 2.11.

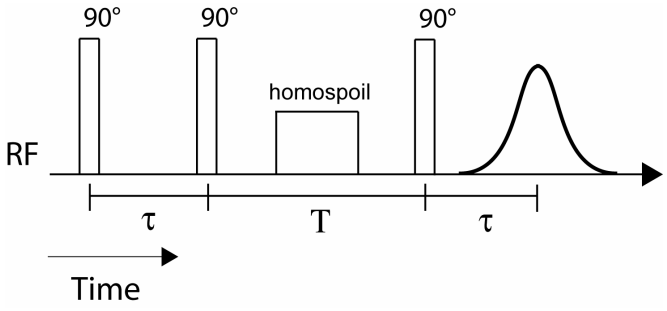


Figure 2.11 – The stimulated echo pulse sequence, a stimulated echo is formed at time $2\tau$ + T after the initial excitation.

In a stimulated echo experiment, the first 90° RF pulse rotates all the magnetisation into the transverse plane. The second 90° RF pulse rotates one component of the magnetisation back to the longitudinal plane for storage. The final 90° RF pulse rotates the "stored" component back into the transverse plane again. Counting from the first 90° RF pulse, a stimulated echo is formed after a total time $2\tau$ + T has elapsed. The advantage of using a stimulated echo (rather than a spin echo technique) is that the spins undergo $T_1$ relaxation during the storage time T, and therefore the stimulated echo yields a stronger signal for systems where $T_1 > T_2$. (The spins will still undergo $T_2^*$ relaxation during time $\tau$).

As mentioned earlier, signals from FIDs occur after every 90° pulse (Section 2.1.3). In practice it is impossible for all the spins to be in exactly the same environment, and so the application of a perfect 90° pulse is impossible. Hence, for every pair of pulses of any angle there is a possibility of a spin echo. For example, in the pulse sequence shown in Figure 2.11, there will be three FIDs and three additional SEs formed (Figure 2.12). This stray magnetisation is undesirable but can be destroyed by applying a strong homospoil gradient pulse during the storage time T. This causes rapid dephasing of any magnetisation in the transverse plane, suppressing unwanted echoes.

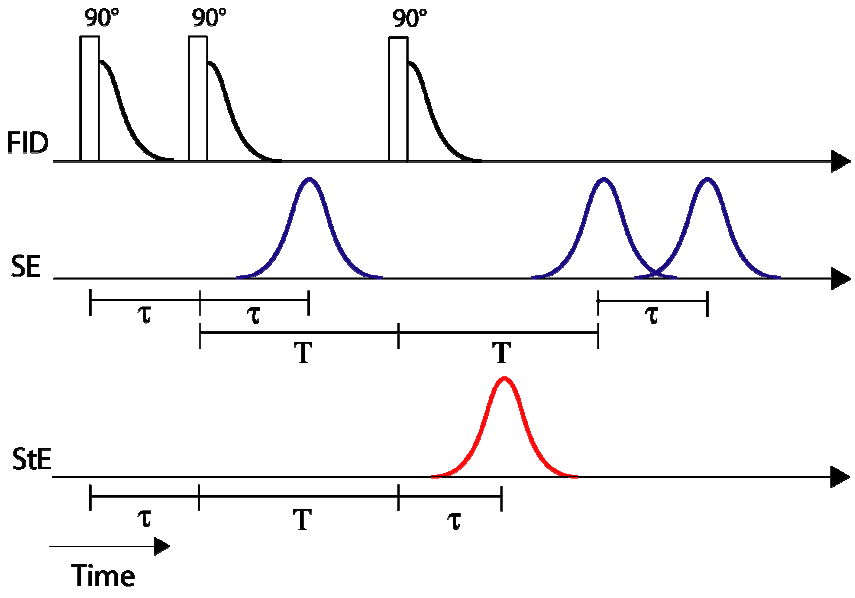


Figure 2.12 – The three FIDs, three SEs and a StE formed with three 90° pulses. The time between the first and second 90° pulses is $\tau$, and the time between the second and third 90° pulses is T. Free induction decays (FIDs) are labelled in black, spin echoes (SE) are labelled in blue, and the stimulated echo (StE) is labelled in red.

## 2.1.8. Magnetic susceptibility

The static magnetic field $\mathbf{B_0}$ of a modern superconducting NMR spectrometer is approximately homogeneous for a simple single-phase liquid sample. In the case of multiple phases in porous media however, the local magnetic environment may vary within the sample because of the presence of paramagnetic, ferromagnetic or diamagnetic species, or simply because of the presence of different materials with different inherent magnetic permittivities. Magnetic susceptibility is discontinuous at the interface between different phases, for example gas-liquid or liquid-solid. These susceptibility differences give rise to local gradients in the magnetic field, and hence nuclei in different parts of the sample will resonate at different frequencies – this will broaden the line width of the spectrum from a narrow peak of perhaps a few Hertz to hundreds of Hertz. In MRI, large susceptibility differences can lead to distortion of the image, or a loss in resolution. A method to overcome the problem is to apply a large enough imaging gradient which increases the spectral bandwidth available per voxel (volume element) and reduces the distortion arising from magnetic susceptibility. With reference to the monolith reactor, which is the focus of this research, we have a highly non-homogeneous system with multiple interfaces. This problem is further exacerbated in the presence of two-phase flow due to the presence of moving gas-liquid interfaces. Some pulse sequences are more robust to the effects of magnetic susceptibility than others, and this will be discussed later (Section 4.2.2). A discussion of magnetic susceptibility is given by Callaghan (1991).

## 2.1.9. NMR spectrometer (hardware)

Modern NMR spectrometers operate in the so-called "pulse" or Fourier transformation (FT) mode, in which all frequencies are excited simultaneously. The separation of frequency components is performed with Fourier transformation during post-processing. A modern FT imaging spectrometer has four major components:

- **The magnet**, which provides the necessary static polarising field, $\mathbf{B_0}$. The most important criteria for a magnet are: the homogeneity and stability of the field strength with respect to both time and position. The homogeneity of the field can be improved by a process called "shimming". During shimming, the

operator adjusts the amount of current passing through small gradient coils installed inside the magnet. These gradient coils create small and well-defined magnetic field gradients superimposed upon the main $\mathbf{B_0}$ field, and carefully chosen combinations of these small field gradients can compensate for inhomogeneities. Low-field machines may be constructed from permanent or electro-magnets, but above one or two Tesla they would consume so much energy that superconducting magnets become the only realistic option. Currently the highest field strength commercially available is 21.2 T ($^1$H resonance frequency of 900 MHz). Higher field strengths give better signal-to-noise and improved spectra, however they also amplify the problems of magnetic susceptibility – hence when studying porous systems there is a trade off between better signal-to-noise, and reducing image artefacts. The materials of the sample are also important, and need to be carefully chosen. For example, metallic materials will distort the magnetic field lines, and result in a non-homogeneous $\mathbf{B_0}$ field through the sample. This is then manifested as a distortion artefact in the acquired data.

- **The gradient system** which surrounds the sample and the RF coil. This reduces the volume available for the sample, but the presence of gradients is necessary for imaging. The purpose of the gradient system is to superimpose a small magnetic field gradient onto the static magnetic field, resulting in the spatial variance of the magnetic field which allows slice selection (Section 2.2.2), frequency encoding (Section 2.2.4) and phase encoding (Section 2.2.5). New, rapid MRI techniques tend to make ever greater demands on the gradient coils and this frequent gradient-switching causes substantial heating in the coils and the generation of eddy currents, which can degrade the quality of spectra and images. The eddy-current problem is reduced by two methods: firstly through "pre-emphasis" which involves over-driving the gradient coils at the leading edge of the gradient pulse and then reversing the polarity of the coils at the trailing edge. This has the effect of generating eddies first in one direction, then in reverse and the result is that the effects of eddy currents are largely compensated for. The second method involves the use of "ramped" gradient pulses, in which the change in magnetic field is more gentle than for the case of a conventional "square" gradient pulse, thereby ensuring weaker

eddy currents. The gradient pulses are therefore trapezoidal, but this is not a problem since the strength of the pulse is determined by the area under the shape, and there is no requirement for it to be rectangular.

- **The radio frequency system** has the primary functions of transmitting and receiving NMR signals. The radio frequency system is an integrated, computer-controlled transceiver unit which includes a waveform generator, a pulse generator and an amplifier. The amplified radio frequency pulses are sent to the NMR probe, where the sample is located within the RF coil.

- **The computer system** is the brain of a modern NMR spectrometer. Its primary function is to command and control the spectrometer (shaping and firing the gradient and RF pulses at precise times and instructing the receiver to acquire data). It also has the functions of storing, retrieving and post-processing of the data (for example performing Fourier transformation). Figure 2.13 is a photograph of the NMR spectrometer on which the majority of the work reported in this thesis was performed; a 4.7 T, 200 MHz, super-wide bore Bruker spectrometer.

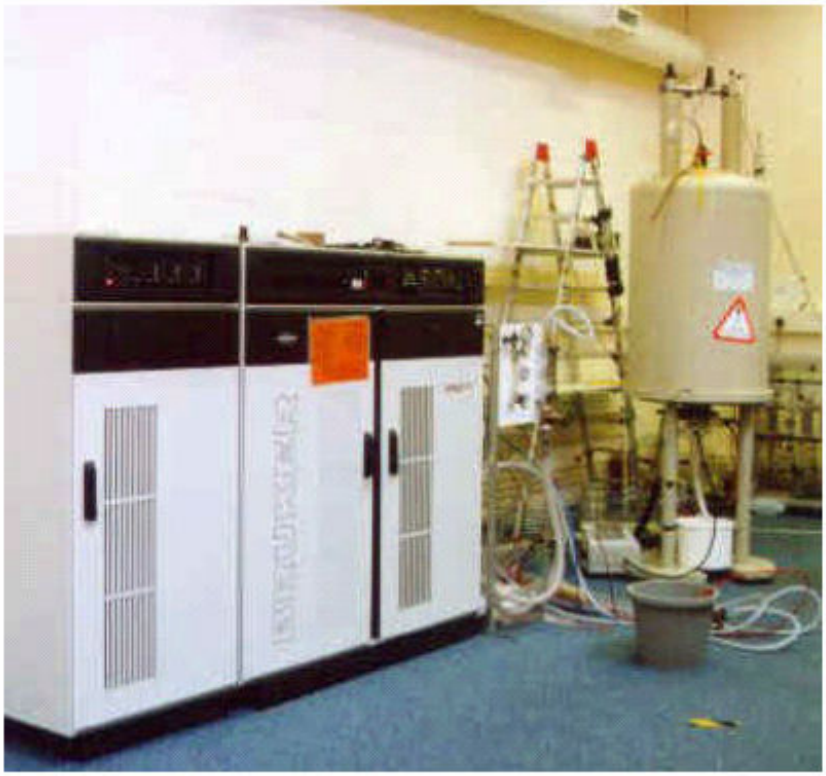

Figure 2.13 – The Bruker spectrospin DMX-200 NMR spectrometer in the Magnetic Resonance Research Centre, University of Cambridge.

## 2.2. Magnetic Resonance Imaging

This section will introduce the theory of imaging methods using NMR, and describe the form and function of the various magnetic field gradients and pulse sequences used for this purpose. Various concepts in MRI will be discussed: slice selection, **k**-space, frequency and phase encoding, classical and fast imaging techniques.

### 2.2.1. Magnetic field gradients

The idea of using additional magnetic field gradients to achieve spatial resolution and slice selectivity is a simple but important one. Lauterbur and Mansfield realised that if a spatially varying field were applied across the object, the Lamor frequencies would also be spatially varying. They showed that the different frequency components could then be separated (via a Fourier transform) to obtain spatially resolved information about the object. This key point of spatially-encoding the data opened the door to the development of MRI techniques.

In magnetic resonance imaging (MRI), a linearly varying magnetic field gradient, **G**, is applied independently of the much larger polarising field, $\mathbf{B_0}$. This then causes the Lamor frequency of the sample to become spatially dependent, shown in Eq. 2.13 and Figure 2.14.

$$\omega(\mathbf{r}) = \gamma(\mathbf{B_0} + \mathbf{G}\bullet\mathbf{r}) \qquad \textbf{(Eq. 2.13)}$$

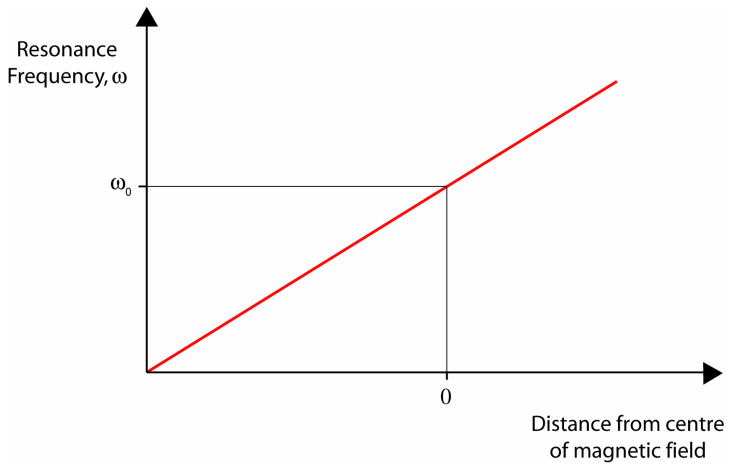


Figure 2.14 – A magnetic field gradient introduces spatially varying resonance frequencies. The isocentre of the gradient is the point where the resonance frequency is unaffected by the applied gradient.

## 2.2.2. Slice selection

Slice selection is achieved through the combined use of a magnetic field gradient and a soft pulse with an adjustable transmission (or carrier) frequency. A soft pulse is an RF pulse which has a limited bandwidth. Since off-resonant nuclear spins are not affected, this pulse will only excite nuclear spins which correspond to the RF pulse bandwidth. Hence, we may selectively excite a physical slice within our sample. The ideal excitation profile is the "hat" function which excites spins within a well-defined region, but none elsewhere. The Fourier transformation of a hat function is a sinc function, $\left(\sin(at)/at\right)$, where $a$ is an arbitrary constant and $t$ is time. However, in practice the sinc function must be truncated, and this results in lobes and undulations in the Fourier transform. A Gaussian pulse shape is often preferred since in the Fourier domain it remains a Gaussian. These pulses are illustrated in Figure 2.15.

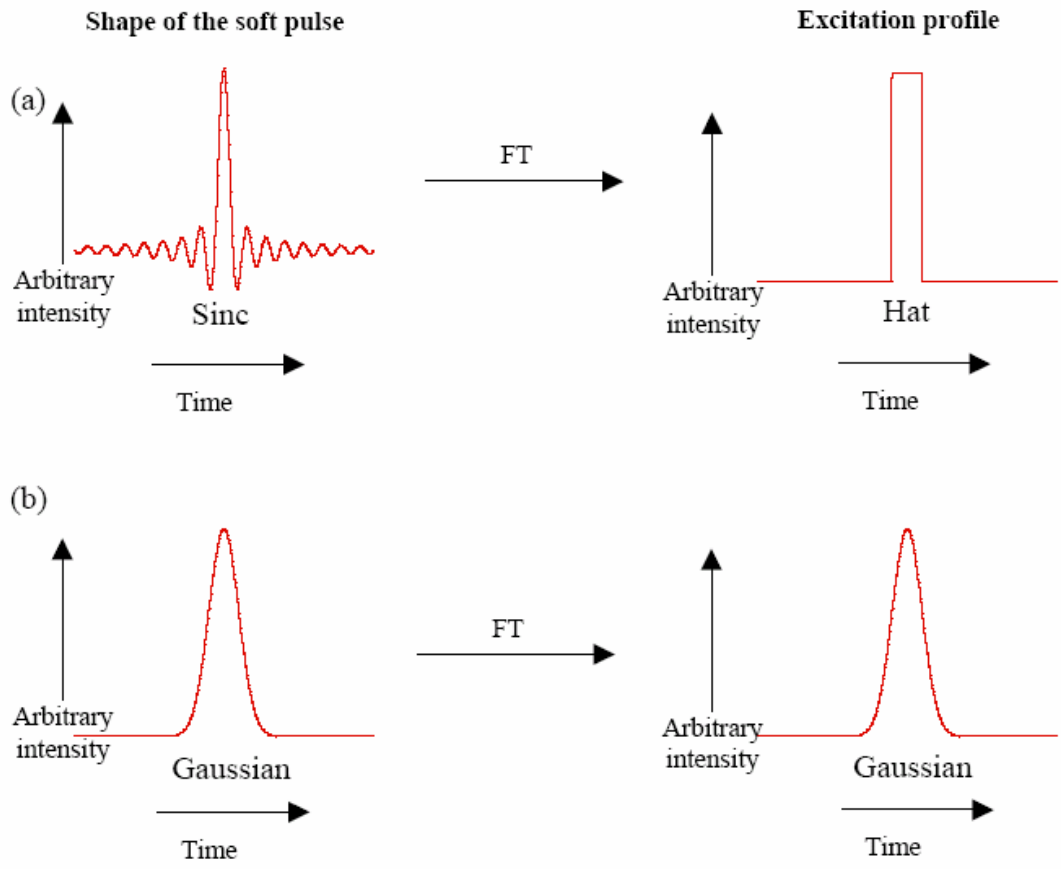


Figure 2.15 – (a) A sinc pulse gives a rectangular excitation profile, but in reality the pulse is truncated at some point which introduces oscillations in the profile; (b) a Gaussian pulse gives a Gaussian excitation profile. Reproduced from Yuen (2003).

For a given slice thickness, the slice gradient strength is adjusted accordingly to match the bandwidth of the soft pulse. To select a slice with thickness $\delta_s$ using a soft pulse of bandwidth $f_s$, the strength of the slice gradient $\mathbf{G_s}$ is calculated with Eq. 2.14.

$$\mathbf{G_s} = \frac{2\pi f_s}{\delta_s \gamma} \qquad \textbf{(Eq. 2.14)}$$

For example, if a 500 μs pulse with a bandwidth of 2500 Hz is to excite a slice of 1 mm, a slice gradient of strength 5.88 G cm$^{-1}$ is required. From Eq. 2.14, doubling the slice thickness simply means halving $\mathbf{G_s}$.

The product $\mathbf{G_s}\gamma$ gives the frequency range per unit length. Using the previous example where $\mathbf{G_s}$ is 5.88 G cm$^{-1}$, $\mathbf{G_s}\gamma$ will be 25,000 Hz cm$^{-1}$. This means to select a slice 1 cm from the centre of the slice gradient, a transmission (resonance) frequency offset of 25,000 Hz is needed. Thus, different positions in the sample can be selected without physically moving the sample. This is shown in Figure 2.16.

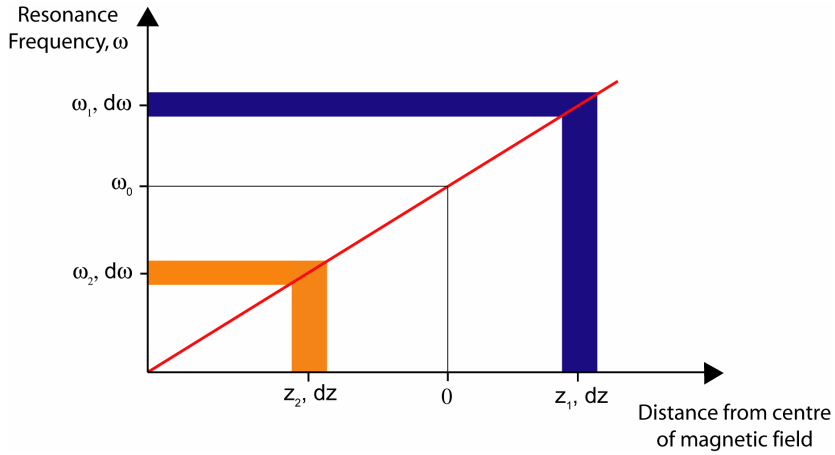


Figure 2.16 – Diagram explaining slice selection, the red line is the magnetic field gradient; the soft pulse has a bandwidth of d$\omega$ which manipulates a slice of thickness d$z$. To select a slice at position $z_1$, the soft pulse is transmitted at $\omega_1$. To select a slice at position $z_2$, the soft pulse is transmitted at $\omega_2$.

## 2.2.3. k-space vector formalism

Mansfield and Grannel (1973) introduced the concept of **k**-space to help simplify the description of MRI experiments. The **k**-vector has units of reciprocal space, and is defined as in Eq. 2.15.

$$\mathbf{k} = \frac{\gamma \mathbf{G} t}{2\pi} \qquad \textbf{(Eq. 2.15)}$$

Eq. 2.15 shows that **k**-space can be traversed either in gradient strength or time, but the direction of the traverse is determined by the direction of the gradient **G** alone. The signal resulting from the magnetic field gradient can thus be written in terms of **k**-space as Eq. 2.16.

$$S(\mathbf{k}) = \iiint \rho(\mathbf{r}) \exp\left[i 2\pi \mathbf{k} \bullet \mathbf{r}\right] d\mathbf{r} \qquad \textbf{(Eq. 2.16)}$$

The Fourier transformation of Eq. 2.16 gives the spin density distribution in real space of Eq. 2.17.

$$\rho(\mathbf{r}) = \iiint S(\mathbf{k}) \exp\left[-i 2\pi \mathbf{k} \bullet \mathbf{r}\right] d\mathbf{k} \qquad \textbf{(Eq. 2.17)}$$

Eq. 2.16 and Eq. 2.17 show that $S(\mathbf{k})$ and $\rho(\mathbf{r})$ are mutually conjugate. In the next two sections, methods of sampling **k**-space will be discussed, namely frequency encoding and phase encoding. For further reading on **k**-space sampling strategies, see Hennig (1999).

## 2.2.4. Frequency encoding

One method of traversing **k**-space is via "frequency encoding". Here a single line in **k**-space is obtained in the presence of a read gradient $\mathbf{G_R}$, shown in Figure 2.17. Different parts of the sample will have different Lamor frequencies and all the signals are acquired together in the form of a spin echo. In terms of **k**-space the first magnetic field gradient (also known as the dephasing read gradient) in Figure 2.17a brings the spins to the "positive" leading edge of **k**-space (Figure 2.17b). The 180° pulse serves two purposes: (1) refocusing spins which have been dephased due to magnetic field inhomogeneity (Section 2.1.6) and (2) inverting the spins through the centre of **k**-space to the "negative" leading edge of **k**-space. The second magnetic

field gradient is then applied to traverse **k**-space from the negative to the positive leading edge again, while data are acquired.

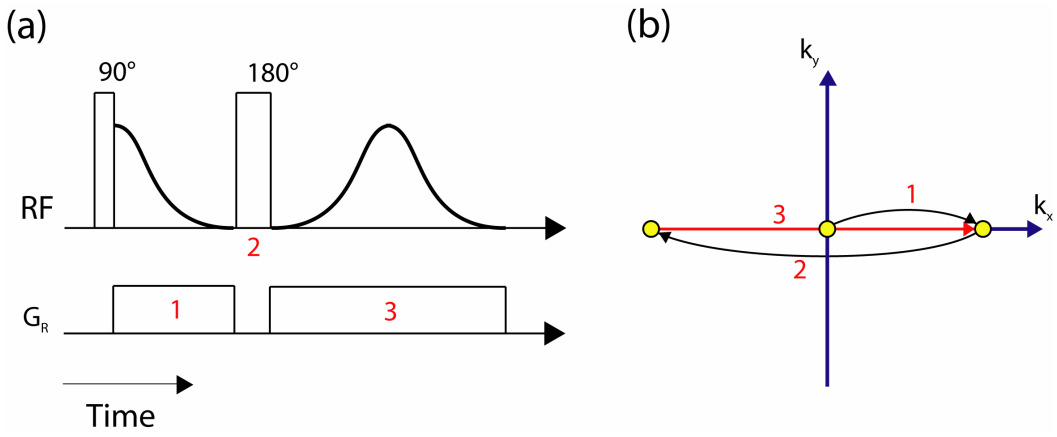


Figure 2.17 – (a) Pulse sequence demonstrating frequency encoding ($G_R$ is applied in the x-direction); (b) **k**-space diagram explaining the evolution in **k**-space. Data are acquired during the red arrow.

The data acquired by applying the pulse sequence shown in Figure 2.17 are a 1D projection of spins along the axis of the applied frequency encoding gradient. The inverse of the spacing between each point in **k**-space governs the field-of-view (FOV), whereas $1/\mathbf{k}_{max}$ defines the resolution, shown in Eq. 2.18.

$$\mathrm{FOV} = \frac{1}{\Delta \mathbf{k}} = \frac{2\pi}{\gamma \mathbf{G}_R t_d} \qquad \textbf{(Eq. 2.18)}$$

where $t_d$ is the dwell time which is related to the sweep width. In MRI, the sweep width is chosen based on two criteria: (1) sufficient bandwidth for each pixel – for example if our spectra has a line-width of 100 Hz and we desire a resolution of 256 pixels, we will require a sweep width of at least 25,600 Hz to prevent signal from neighbouring pixels overlapping and therefore reducing the apparent resolution. (2) The duration of the data acquisition, which ultimately affects the echo time. Typically a sweep width of around 150,000 Hz is used in the experiments performed in this thesis ($t_d$ = 5 µs). Once the dwell time is fixed, $\mathbf{G}_R$ is then adjusted accordingly to give the desired FOV (Eq. 2.18). For example, for a field of view of 1 cm acquired with a sweep width of 100,000 Hz, a frequency encoding gradient $\mathbf{G}_R$ of 23.5 G cm$^{-1}$ is required.

## 2.2.5. Phase encoding

The second method of traversing **k**-space is via "phase encoding". During phase encoding, a gradient pulse is switched on for a brief period of time. A phase angle is imparted onto the sample which varies according to the position of the nuclei within the sample. The phase shift induced in an individual spin is given by Eq. 2.19.

$$\phi = \gamma G t \qquad \textbf{(Eq. 2.19)}$$

To acquire successive points in **k**-space, we fix the duration of the phase encoding gradient pulse and ramp the strength between successive repetitions, illustrated in Figure 2.18.

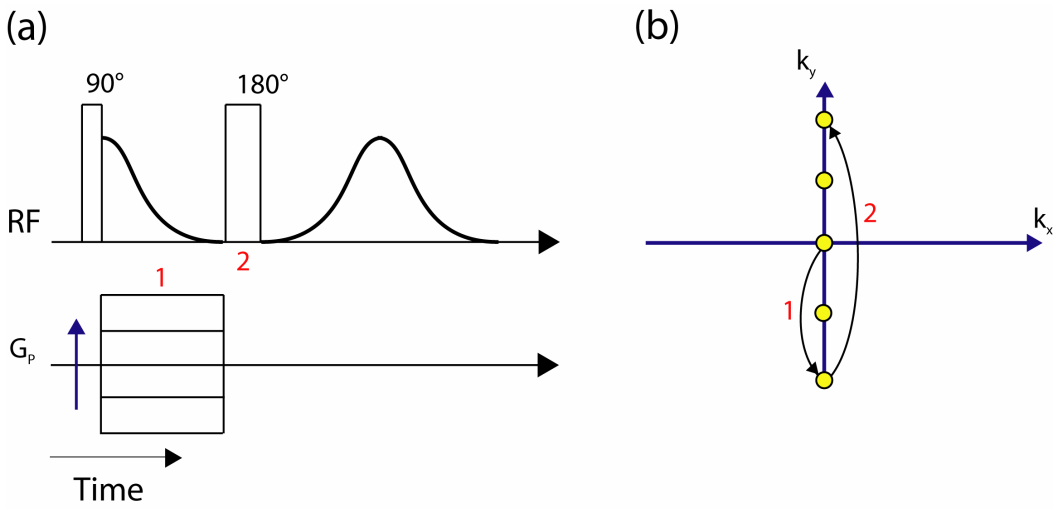


Figure 2.18 – (a) A pulse sequence demonstrating the concept of phase encoding. In this example, $\mathbf{G_P}$ is applied in the y-direction. The blue arrow indicates that the phase encoding gradient is ramped after each repetition; (b) **k**-space diagram explaining the evolution in **k**-space. A single point in **k**-space is collected for each repetition.

The phase encoding gradient is normally ramped from the negative to the positive leading edge of **k**-space, acquiring data one point at a time. As mentioned before, the inverse of the spacing between each point in **k**-space governs the FOV, as in Eq. 2.20.

$$\mathrm{FOV} = \frac{1}{\Delta \mathbf{k}} = \frac{2\pi}{\gamma \Delta \mathbf{G_P} t_p} \qquad \textbf{(Eq. 2.20)}$$

where $t_p$ is the duration of the phase encoding gradient. The important parameter is the step size of the phase encoding gradient, $\Delta\mathbf{G_P}$. For example, a field of view of

1 cm acquired with a phase encoding gradient of length 1 ms, $\Delta\mathbf{G_P}$ must have a value of 0.235 G cm$^{-1}$. If 128 pixels are required, the phase encoding gradient must be ramped from -14.9 to +14.9 Gcm$^{-1}$ in 128 steps, given by Eq. 2.21.

$$\mathbf{G}_{\mathbf{P},\max} = \frac{\Delta\mathbf{G_P}(\mathrm{NOP}-1)}{2} \qquad \textbf{(Eq. 2.21)}$$

where NOP is the number of pixels in the phase encoding direction. The factor of 2 takes into account the fact that the phase encoding gradient is ramped from negative to positive.

## 2.2.6. Gradient echo

The gradient echo is similar to the spin echo (Section 2.1.6), but instead of using RF pulses to refocus spins, it relies upon the gradients themselves to generate an echo.

The gradient echo does not refocus dephasing due to $T_2$', and so the relaxation contrast will be different from that obtained by the spin echo technique. Thus, attenuation or phase shift effects due to susceptibility or variations in the polarising field will persist. This results in less robust imaging sequences, but a major advantage of the gradient echo is for medical applications where we desire to avoid excessive power dissipation in the patient, an effect associated with using hard RF pulses.

We have already seen that if we apply a linearly increasing magnetic field gradient along the length of our sample, each spin position will experience a slightly different magnetic field strength and therefore rotate at a different Lamor frequency. When observing the spins as a whole, the spins will therefore dephase due to magnetic field inhomogeneities, just as for $T_2$' dephasing. If we then apply the same gradient strength in the opposite direction, the spins will rephase and we thus obtain a gradient echo (Figure 2.19). This is the basis of several imaging sequences, most notably EPI (Section 2.3.2). However, while the gradient echo reverses the phase induced by the application of gradients, it does not refocus the dephasing due to $T_2$'.

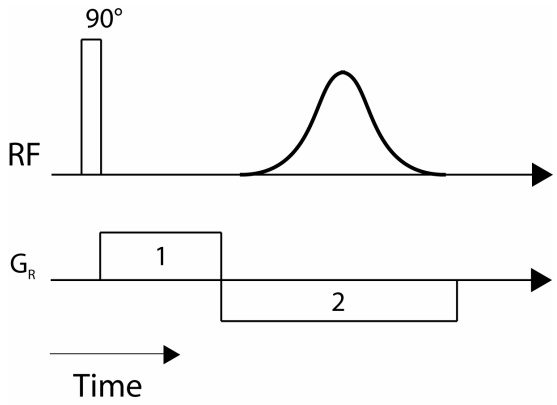


Figure 2.19 – Illustration of a gradient echo sequence: (1) dephasing gradient; (2) refocusing gradient.

## 2.2.7. A classical 2D imaging method

So far the discussion has been focussed on 1D imaging (or profiling). 2D and 3D magnetic resonance imaging are possible by combining the concepts discussed, and a classic imaging method will be outlined in this section.

A 2D image is typically obtained by combining frequency and phase encoding within the pulse sequence (Figure 2.20). Successive lines in **k**-space are acquired (as with the frequency encoding method), and the phase encoding method is used to move the acquisition line upwards in our 2D **k**-space raster. This method of scanning 2D **k**-space is known as "spin-warp" imaging (Edelstein *et al.*, 1981). In this particular version, slice selection is achieved with a soft (selective) refocusing RF pulse. A typical 2D single-slice image takes a few minutes to acquire. It would also be possible to traverse the 2D **k**-space raster by using phase encoding in both dimensions, however this would take far longer due to increased number of repetitions required, arising from the recovery time we must wait between each acquisition to recover spin magnetisation. The only advantage of phase-phase encoding is that each point in **k**-space is acquired at the same echo time, whereas in frequency-phase encoding the echo time will vary in the read-direction according to the length of the read gradient. Phase-phase acquisition is therefore of use for systems with very short $T_2$ times since the echo time is kept very short, maximising the strength of the signal acquired.

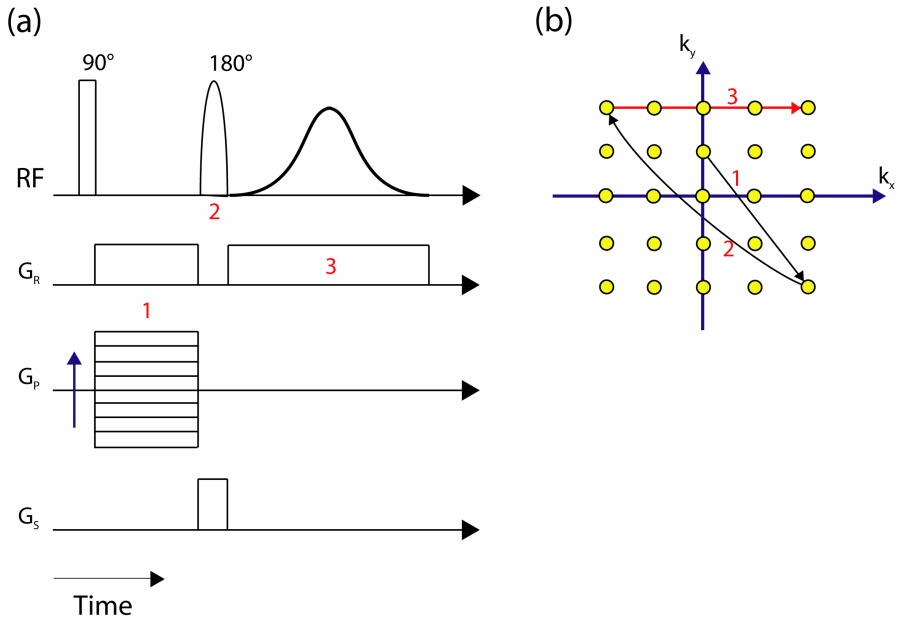


Figure 2.20 – (a) an example of a "spin-warp" 2D single-slice imaging pulse sequence, the blue arrow indicates that the phase encoding gradient is ramped after each repetition; (b) the corresponding **k**-space trajectory diagram during one repetition. Data are acquired during the red arrow.

A fully resolved 3D image can be obtained by adding a second phase encoding gradient orthogonal to the frequency and phase encoding directions of the 2D sequence (Figure 2.21).

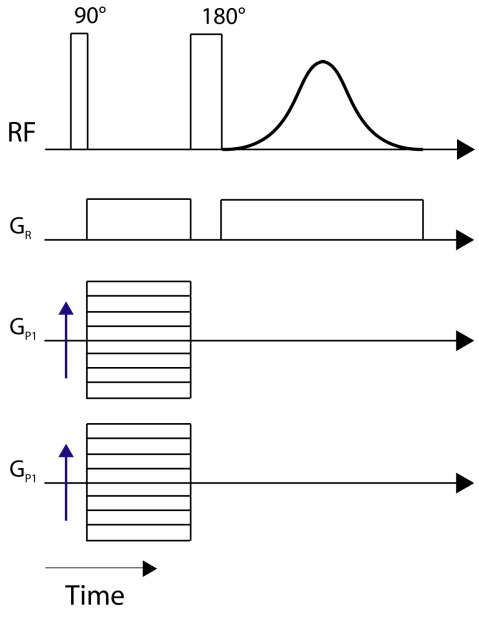


Figure 2.21 – An example of a "spin-warp" 3D imaging sequence.

In this type of imaging, no soft pulses are required. All the spins contribute to each point in **k**-space. Because of the additional phase encoding gradient, this type of experiment usually takes several hours to complete.

## 2.3. Fast imaging methods

Fast imaging is perhaps one of the most interesting areas of MRI development. The thrust for such single-shot acquisitions has been driven in medical applications by cardiac imaging and functional imaging, and in chemical engineering by the desire to study dynamic systems where steady-state cannot be assumed. By careful manipulation of the pulse sequence we can reduce imaging times and begin to probe such dynamic systems. This section will outline some of the approaches to accelerating the imaging process, and outline a few traditional pulse sequences. Novel techniques developed as a part of this work, as well as other fast techniques developed in the Department of Chemical Engineering at the University of Cambridge are covered later in Chapter 5.

The data acquisition time, $t_{scan}$, of an MRI image is calculated from Eq. 2.22.

$$t_{scan} = \frac{TR \times N(\mathrm{2D}) \times N(\mathrm{3D}) \times NS}{N(S)} \qquad \textbf{(Eq. 2.22)}$$

where $TR$ is the repetition time, $N$(2D) and $N$(3D) are the number of two- and three-dimensional phase encoding steps, $NS$ is the number of signal averages, and $N(S)$ is the number of lines of **k**-space collected per repetition. In the classical imaging method described in Section 2.2.7, only one line of **k**-space is acquired per repetition. There are various methods to reduce $t_{scan}$:

- **Low flip-angle excitation pulse and gradient echo** – $TR$ can be reduced, however to avoid introducing severe $T_1$-weighting of the image low pulse angles of around 5-15° are used, such that the system will not become saturated by rapid excitations i.e. an equilibrium excitation is reached, hence it is desirable to acquire several dummy scans at the beginning of a sequence since the system will not yet have reached equilibrium. A gradient echo is typically used to reduce the total echo time and the amount of RF heating

caused by rapid excitations (which is more important in medical applications). The disadvantage of using a gradient echo is that the image is now $T_2^*$ rather than $T_2$ weighted, as discussed earlier (Section 2.2.6). Examples of this type of method are FLASH (Fast Low Angle Single Shot) and SNAPSHOT (Haase, 1990).

- **Segmented scanning** – multiple lines in **k**-space are collected per repetition i.e. $N(S) > 1$. This is achieved by collecting an echo train in which each echo has a different phase encoding. There are two methods of generating an echo train: (1) CPMG-type pulse sequence or (2) gradient-echo type sequence. Images obtained using the former method are $T_2$-weighted, whereas the latter is $T_2^*$-weighted. An example of these respective methods are RARE (Rapid Acquisition with Relaxation Enhancement) (Hennig *et al.*, 1986), and EPI (Echo Planar Imaging) (Callaghan, 1991).

- **Partial Fourier imaging** – instead of collecting the whole **k**-space raster, only half of the **k**-space is collected. This works because **k**-space is inherently symmetrical, and so the second half of **k**-space can be inferred and the whole raster is Fourier transformed. This reduces the number of phase encoding steps for a given image, but in doing so the signal-to-noise ratio is sacrificed (Stark and Bradley, 1999).

### 2.3.1. Rapid acquisition with relaxation enhancement (RARE)

RARE is a fast, single-excitation imaging technique used in the work reported in this thesis, and is capable of acquiring a 2D image in under one second. In RARE, a CPMG echo train is collected and each echo is phase-encoded differently. Since a CPMG train is used, the image is $T_2$ rather than $T_2^*$ weighted. RARE may be used either as a single-shot or multi-shot method depending on the $T_2$ of the sample. In the single-shot method, typically 64 or 128 echoes are collected.

One notable feature of RARE is the "rewinding" phase encoding gradient at the end of each echo acquisition, and this is what gives the "relaxation enhancement" suggested in its title. This brings the magnetisation back to the centre of **k**-space and thus any errors in phase encoding are not accumulated. The single slice 2D RARE pulse sequence shown in Figure 2.22 uses soft pulses only. With RARE, a single 2D image can typically be acquired in hundreds of milliseconds rather than minutes.

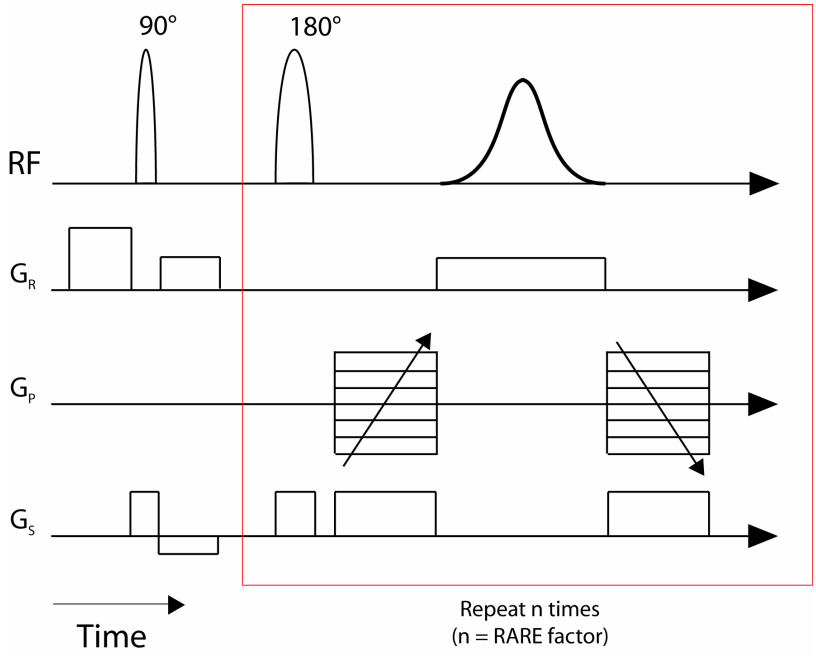


Figure 2.22 – A single slice 2D RARE sequence.

Since many points in **k**-space are acquired with a single excitation, RARE is highly sensitive to $T_2$ attenuation and may not be applicable to certain short-$T_2$ systems. Also, the large number of successive refocusing RF pulses causes considerable energy dissipation within the sample, and therefore RARE is avoided in medical applications.

## 2.3.2. Echo Planar Imaging (EPI)

It is rare to collect all the **k**-space points following a single RF excitation, due to the problems associated with acquiring many lines of **k**-space before the signal is lost from $T_2$ or $T_2^*$ decay. EPI is a gradient echo based technique, and uses a series of phase gradient "blips" to navigate **k**-space. Each line of **k**-space is read by the application of a read gradient to induce a gradient echo, then a phase gradient blip is

applied to move to the next line of **k**-space, and a negative read gradient is then applied to generate another gradient echo, and so on, until all of **k**-space is traversed. This series of short phase blips coupled with the gradient echo allows the acquisition of **k**-space very rapidly (Figure 2.23). Since the entire image is acquired from a single acquisition, there is no $T_1$-weighting in the image – contrast stems from spin density and $T_2^*$.

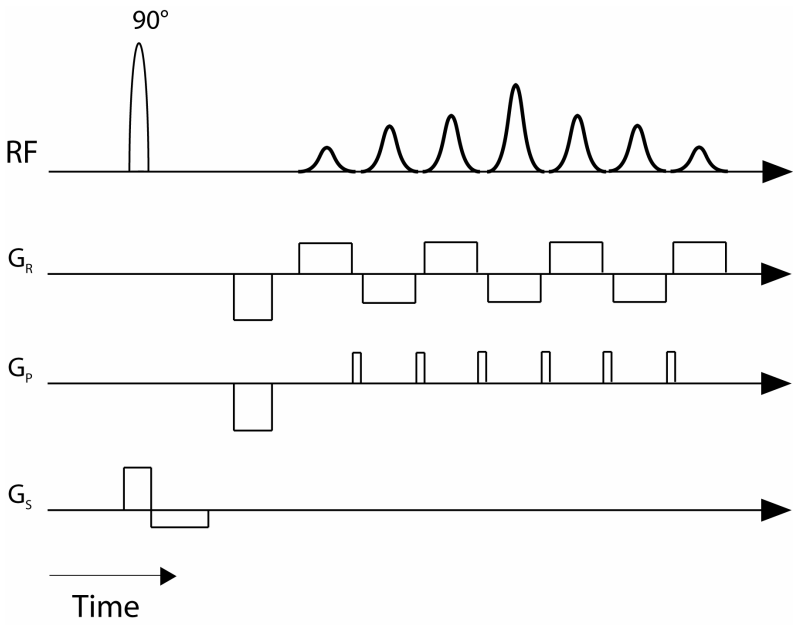


Figure 2.23 – Illustration of a 2D blipped EPI sequence.

The EPI imaging method is not as robust as RARE to imaging artefacts, and in particular to magnetic susceptibilities within the sample. RARE contains a phase-rewind gradient between each echo in the train, whereas EPI uses a series of blips to scan across **k**-space one step at a time. Hence, any errors in timing, pulse configuration, or inhomogeneities in the sample will build up and manifest as imaging artefacts. This means that in certain systems EPI will not necessarily be a quantitative analytical technique. Despite its limitations, EPI is much faster than RARE (20 ms compared to 150 ms for a 64×64 raster) and can be an invaluable tool, particularly for obtaining qualitative information on dynamic systems. It is widely used in medical imaging, for example, where the human body can be considered to be largely homogeneous in terms of magnetic susceptibilities.

## 2.4. MRI of transport processes

Two kinds of motion may be studied with NMR and they are: coherent motion, for example velocity; and incoherent motion, for example self-diffusion. The details of how velocity can be measured with NMR are discussed in this section. Diffusion is not covered, but the principles behind the method are the same.

As shown previously, when a pulsed magnetic field gradient is applied, a phase shift is imparted to the nuclear spins as given by Eq. 2.23.

$$\phi(\mathbf{r}) = \gamma \mathbf{g} \delta \mathbf{r} \qquad \textbf{(Eq. 2.23)}$$

where **r** is the position of the spins, **g** and $\delta$ are the strength and duration of the applied gradient respectively. After an observation delay, Δ, the same magnetic field gradient is applied, but reversed in direction. The net phase shift induced is given by Eq. 2.24.

$$\phi_n(\mathbf{r}) = \gamma \mathbf{g} \delta (\mathbf{r} - \mathbf{r}') \qquad \textbf{(Eq. 2.24)}$$

If the spins are stationary, i.e. **r** = **r'**, there will not be a net phase shift. For coherent motion, the displacement term (**r** – **r'**) in Eq. 2.24 will be proportional to the linear displacement, and can be expressed in terms of spin velocity, **v**, and the observation delay, Δ, giving Eq. 2.25.

$$\phi_n(\mathbf{r}) = \gamma \mathbf{g} \delta \Delta \mathbf{v}(\mathbf{r}) \qquad \textbf{(Eq. 2.25)}$$

It is clear that the phase shift induced, $\phi_n(\mathbf{r})$ is directly proportional to the velocity of the spin, and this gives NMR the ability to measure velocity quantitatively. The basic pulsed gradient spin echo (PSGE) sequence (Stejskal and Tanner, 1965) for achieving velocity contrast is shown in Figure 2.24.

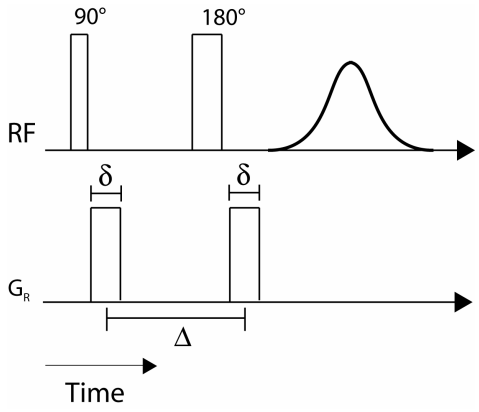


Figure 2.24 – The basic pulsed gradient spin echo (PGSE) sequence. Δ is the observation time and $\delta$ is the duration of the pulsed field gradient.

## 2.4.1. Velocity imaging

Spatially-resolved velocity maps of nuclear spins can be obtained by combining the PGSE sequence (Figure 2.24) with a suitable imaging sequence. An example of a velocity imaging pulse sequence is shown in Figure 2.25.

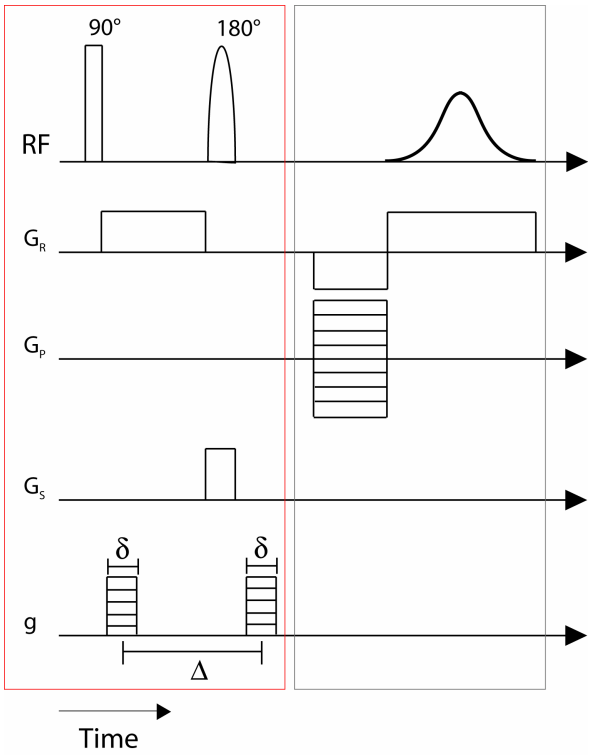


Figure 2.25 – An example of a velocity imaging pulse sequence. The red box marks the pre-conditioning, velocity-encoding section; the grey box marks the imaging section of the pulse sequence.

Since the phase of an MRI signal is arbitrary, a minimum of two velocity-encoded images are needed to determine the difference in phase shift, i.e. we acquire a zero-flow acquisition, and measure the phase shift relative to this for a flowing acquisition. To improve the accuracy of the measurement, normally four velocity-encoded images are taken (Sederman *et al.*, 1998). The difference in phase shift $\phi_d$ in each pixel is converted to velocity $\mathbf{v}(\mathbf{r})$ with Eq. 2.26.

$$\mathbf{v}(\mathbf{r}) = \frac{\phi_d(\mathbf{r})}{\delta \Delta \mathbf{g}_{\mathbf{inc}} \gamma} \qquad \textbf{(Eq. 2.26)}$$

where $g_{inc}$ is the increment size of the velocity gradients. The maximum distinguishable phase shift is $2\pi$, beyond which the measured phase would wrap around to zero. Therefore when performing a velocity imaging experiment, it is important to adjust the parameters to ensure that the phase shift induced by the maximum velocity is less that $2\pi$, otherwise the velocity image may contain a "folding over" artefact[1].

It should be noted that when gradient coils are operated during the velocity-encoding of a flowing system, eddy currents are generated. These eddy currents cause the spins to incur systematic phase-shift errors (as mentioned previously in Section 2.1.9). However, the zero-flow velocity-encoded image will measure the effect of gradient inhomogeneities or eddy currents in the presence of zero flow, and subtracting this data from all flowing data sets will obtain the correct phase shift for the experiments.

Higher orders of motion than velocity, such as acceleration, jerk and snap may also be measured (Callaghan, 1991). The relative phase of the NMR signal (which is what is actually measured) is given by Eq. 2.27.

$$\phi = \gamma \left[ x_0 \int G(t)dt + v \int tG(t)dt + \frac{a}{2} \int t^2 G(t)dt + \ldots \right] \qquad \textbf{(Eq. 2.27)}$$

The integrals are the successive moments of the applied magnetic field gradients, and hence can be manipulated in an experiment. Thus, the zeroth moment ($\int G(t)dt$) can be set to zero by choosing a pair of gradients that have opposite signs, and the resultant phase will be dominated by the velocity term (as is seen in the PGSE

[1] In fact it is sometimes possible to correct for this fold-over artefact, and some methods for this are discussed in Section 7.7.

diagram of Figure 2.24). In principle then, it is possible to generate acceleration (or higher order) images, although each successive term that is nulled makes the experiment more delicate and difficult, and the signal-to-noise ratio will decrease.

## 2.4.2. Average displacement propagator

Following a similar argument in the definition of **k**-space imaging (Section 2.2.3), a dynamic reciprocal space vector, **q**, is defined in Eq. 2.28 (Callaghan, 1991).

$$\mathbf{q} = \frac{\gamma \mathbf{g} \delta}{2\pi} \qquad \textbf{(Eq. 2.28)}$$

A dynamic displacement is defined as **R** = **r**' – **r**, and the average displacement propagator, $\overline{P}_s(\mathbf{R}, \Delta)$, is defined as the average probability that any molecule in the sample will move by a dynamic displacement **R** over the observation time Δ. The acquired normalised signal is given by Eq. 2.29.

$$E_{\Delta}(\mathbf{q}) = \int \overline{P}_s(\mathbf{R}, \Delta) \exp(i 2\pi \mathbf{q} \bullet \mathbf{R}) d\mathbf{R} \qquad \textbf{(Eq. 2.29)}$$

Thus, $\overline{P}_s(\mathbf{R}, \Delta)$ can be obtained from the measurement of $E_{\Delta}(\mathbf{q})$. This is achieved by running a series of experiments using the PGSE pulse sequence (Figure 2.24) with different values of **g** (also known as **q**-gradient steps). The data are then subsequently Fourier transformed with respect to **q**. Effectively the experiment tags the molecules by manipulating their nuclear spins, which is equivalent to a traditional tracer experiment.

The concept of **q**-space can be extended to velocity imaging (as an alternative to the phase-shift method described in (Section 2.4.1). It may be shown that after Fourier transformation in **q**-space using multiple **q**-contrasted images, the velocity in each pixel is given by Eq. 2.30 (Callaghan and Xia, 1991).

$$\mathbf{v}(\mathbf{r}) = \frac{2\pi k_v(\mathbf{r})}{\delta \Delta \mathbf{g}_{\mathbf{inc}} \gamma N} \qquad \textbf{(Eq. 2.30)}$$

where $k_v(\mathbf{r})$ and $N$ are the position of the peak and the digital array size in Fourier transformation respectively. Normally this method requires at least 12 **q**-gradient steps (typically 16) with zero-filling to 256 steps for effective Fourier transformation.

The method has the drawback of extreme time consumption, but is more accurate than a simple PGSE velocity experiment.

## 2.5. Conclusions

In this chapter, basic NMR theory has been introduced. This was then followed by discussions of various traditional MRI techniques used in this thesis. In Chapter 3, the focus will turn to the main application of this thesis, the monolith reactor, and Chapter 4 will discuss the application of NMR to the study of such porous systems.

## 2.6. Nomenclature

| Symbols | Definition |
|---|---|
| a | Arbitrary constant |
| $\mathbf{B_0}$, $\mathbf{B_1}$ | Static polarising magnetic field, and RF magnetic field |
| $E$ | Energy |
| $f_{sw}$, $f_s$ | Sweep width, bandwidth of a soft pulse |
| $\mathbf{G}$, $\mathbf{G_S}$, $\mathbf{G_R}$, $\mathbf{G_P}$, $\mathbf{G_{P,max}}$, $\Delta\mathbf{G_P}$ | Magnetic field gradient, slice gradient, frequency encoding gradient, phase encoding gradient, maximum phase encoding gradient, step size of the phase encoding gradient |
| $\mathbf{g}$, $\mathbf{g_{inc}}$ | Dynamic magnetic field gradient, increment size of the dynamic magnetic field gradient |
| $i$ | Imaginary number |
| $I$ | Spin quantum number |
| $h$ | Planck constant |
| $\hbar$ | Planck constant divided by $2\pi$ |
| $\mathbf{k}$, $\Delta\mathbf{k}$ | Reciprocal space vector, spacing between **k**-space points |
| $k_v$ | Position of the peak in **k**-space |
| $k_B$ | Boltzmann constant |
| $\mathbf{M}$ | Net magnetisation vector |
| $M$, $M_0$, $M_x$, $M_y$, $M_z$ | Magnetisation, equilibrium magnetisation, magnetisation in $x$, $y$ and $z$-axes |
| $N$ | Digital array size |
| $N_\alpha$, $N_\beta$ | Population in $\alpha$ and $\beta$ states respectively |
| $\mathbf{P}$ | Spin angular momentum |
| $P$ | Magnitude of spin angular momentum |
| $\mathbf{q}$ | Dynamic reciprocal space vector |
| $\mathbf{r}$ | Position |
| $\mathbf{R}$ | Dynamic displacement |
| Re | Reynolds number |
| $S$, $S_0$ | Signal intensity, signal intensity immediately after excitation |
| $t$, $t_{RF}$, $t_d$, $t_p$, $t_{scan}$ | Time, duration of the RF pulse, dwell time, duration of the phase encoding gradient and data acquisition time |
| T | Temperature, or storage time in the stimulated echo pulse sequence (depending on the context) |
| $T_1$ | Spin-lattice relaxation constant |
| $T_2$, $T_2^*$, $T_2'$ | Spin-spin relaxation constant, apparent spin-spin relaxation constant, effect of magnetic field inhomogeneity |
| $TR$ | Repetition time |
| $\mathbf{u}$, $\mathbf{v}$ | Velocity |
| x,y,z | Spatial axis in laboratory frame |

| | |
|---|---|
| x',y' | Spatial axis in rotating frame |
| | |
| $\Delta$ | Observation time |
| $\phi$, $\phi_n$, $\phi_d$ | Phase angle, net phase shift, phase shift difference |
| $\sigma$ | Shielding factor |
| $\boldsymbol{\mu}$ | Magnetic dipole moment |
| $\rho$ | Material density |
| $\boldsymbol{\rho}$ | Spin density |
| $\tau$, $\tau_n$, $\tau_1$, $\tau_2$ | Delay, nulling delay, delays between RF pulses |
| $\delta$ | Chemical shift or duration of **q**-gradient (depending on the context) |
| $\delta_s$ | Slice thickness |
| $\theta$ | Angle |
| $\gamma$ | Gyromagnetic ratio |
| $\omega_0$, $\omega_1$, $\Delta\omega$ | Lamor frequency, frequency of the RF pulse and frequency offset |
| $\upsilon, \upsilon_{ref}$ | Frequency and reference frequency |

| **Abbreviations** | **Definition** |
|---|---|
| CPMG | Carr-Purcell-Meiboom-Gill |
| EPI | Echo-planar imaging |
| FID | Free induction decay |
| FLASH | Fast low angle single shot |
| FOV | Field of view |
| FT | Fourier transform |
| FWHM | Full-width-half-maximum |
| GE | Gradient echo |
| MRI | Magnetic resonance imaging |
| NMR | Nuclear magnetic resonance |
| NOP | Number of pixels |
| PGSE | Pulsed gradient spin echo |
| ppm | Parts per million |
| RARE | Rapid acquisition with relaxation enhancement |
| RF | Radio frequency |
| SE | Spin echo |
| StE | Stimulated echo |
| TMS | Tetramethylsilane |

## 2.7. References


Bloch, F. (1946). "Nuclear Induction." Physical Review **70**(7-8): 460-474.

Callaghan, P. T. (1991). Principles of Nuclear Magnetic Resonance Microscopy. Oxford, Clarendon Press.

Callaghan, P. T. and Y. Xia (1991). "Velocity and Diffusion Imaging in Dynamic NMR Microscopy." Journal of Magnetic Resonance **91**: 326-352.

Carr, H. Y. and E. M. Purcell (1954). "Effects of Diffusion on Free Precession in Nuclear Magnetic Resonance Experiments." Physical Review **94**: 630-639.

Cooley, J. W. and J. W. Turkey (1965). "An Algorithm for the Machine Calculation of Complex Fourier Series." Mathematics of Computation **19**: 297-301.

Edelstein, W. A., J. M. S. Hutchinson*, et al.* (1981). "Human Whole-body NMR Tomographic Imaging: Normal Sections." British Journal of Radiology **54**: 149-151.

Fukushima, E. and S. B. W. Roeder (1983). Experimental Pulse NMR - A Nuts and Bolts Approach, Perseus Books Group.

Gibbs, S. J. and L. D. Hall (1996). "What roles are there for magnetic resonance imaging in process tomography?" Measurement Science & Technology **7**(5): 827-837.

Gladden, L. F. (1994). "Nuclear-Magnetic-Resonance in Chemical-Engineering - Principles and Applications." Chemical Engineering Science **49**(20): 3339-3408.

Haacke, M. E., R. W. Brown*, et al.* (1999). Magnetic Resonance Imaging: Physical Principles and Sequence Design, John Wiley and Sons Inc.

Haase, A. (1990). "Snapshot Flash MRI - Applications to T1, T2, and Chemical-shift Imaging." Magnetic Resonance in Medicine **13**: 77-89.

Hahn, E. L. (1950). "Spin Echoes." Physical Review **80**: 580.

Harris, R. K. (1986). Nuclear Magnetic Resonance Spectroscopy. Harlow, Longman Group UK.

Hennig, J., A. Nauerth*, et al.* (1986). "Rare Imaging - a Fast Imaging Method for Clinical Mr." Magnetic Resonance in Medicine **3**(6): 823-833.

Hennig, J. (1999). "K-space sampling strategies." European Radiology **9**(6): 1020-1031.

Lauterbur, P. C. (1973). "Image Formation by Induced Local Interactions - Examples Employing Nuclear Magnetic Resonance." Nature **242**: 190-191.

Mansfield, P. and P. K. Grannel (1973). "NMR "Diffraction" in Solid." Journal of Physics C **6**: L422-426.

Meiboom, S. and D. Gill (1958). "Modified Spin-echo Method for Measuring Nuclear Relaxation Times." Review of Scientific Instruments **29**: 688-691.

Purcell, E. M., H. C. Torrey*, et al.* (1946). "Resonance Absorption by Nuclear Magnetic Moments in a Solid." Physical Review **69**(1-2): 37-38.

Sederman, A. J., M. L. Johns*, et al.* (1998). "Structure-flow correlations in packed beds." Chemical Engineering Science **53**(12): 2117-2128.

Stark, D. D. and W. G. Bradley (1999). Magnetic Resonance Imaging. St Louis, Mosby.

Stejskal, E. O. and J. E. Tanner (1965). "Spin Diffusion Measurements: Spin Echoes in the Presence of a Time-dependent Field Gradient." Journal of Chemical Physics **42**: 288-292.

Vold, R. L., J. S. Waugh*, et al.* (1968). "Measurement of Spin Relaxation in Complex Systems." Journal of Chemical Physics **48**: 3831-3832.

Yuen, E. H. L. (2003). Aspects of Nuclear Magnetic Resonance Studies of Transport and Reaction in Packed Bed Reactors. Dept of Chemical Engineering. Cambridge, University of Cambridge**:** 257.

# Chapter 3 – The Monolith Reactor

Monolith structures have emerged as candidates for the new generation of reactors utilising structured packings. This chapter contains a detailed literature review of the current knowledge on the design and operation of monolith reactors in chemical engineering applications. Their performance will be compared to more traditional reactors (e.g. slurry, packed bed), and the flow phenomena observed in monoliths will be discussed.

**CHAPTER 3 – THE MONOLITH REACTOR**

## 3.1. Introduction

In chemical engineering, there is always a strong drive towards smaller, more compact, more efficient processes, as these incur much reduced costs in terms of plant footprint and energy consumption. It is far better to deploy reactor designs with "clean" technology than to rely on end-of-pipe solutions, expensive purifications or recycles to overcome low conversions and poor selectivity. Smaller reactors mean reduced residence times, allowing fast adaptation to altered process conditions and feeds. By making plants smaller they are intrinsically safer, storing reduced quantities of hazardous chemicals. Also start-up and shutdown times will be reduced, making batch processes more effective. Recent improvements to packed-bed reactors and other mature technologies have been incremental and have not resulted in a large increase in performance. Multi-functional and structured reactors will play a greater role in process intensification, blurring the borders between catalyst and reactor. Please refer to Stankiewicz and Moulijn (2000) and Dautzenberg and Mukherjee (2001) for more details on process approaches.

An important development in the area of process intensification is that of the monolith reactor. This design has several advantages over similar reactor designs such as trickle beds and slurry reactors, though one of the main drawbacks is the lack of detailed knowledge in optimisation of their operation. The role of NMR and MRI as research tools is to aid the understanding and quantification of the flow characteristics and complex local behaviour of these systems, and then apply this to aid the optimisation of large-scale processes used in industry.

This chapter will begin by describing the construction of the monolith reactor, and illustrate its performance compared to traditional reactor types with examples from the literature. Particular characteristics of monolith reactors, such as flow regimes and mass transfer properties will be discussed for different configurations of monolith reactors, followed by a brief summary of the current knowledge on modelling monolith processes.

## 3.2. Monolith reactors

Monolith reactors are well established in the automotive industry as exhaust gas converters, and are probably the most popular catalytic reactor worldwide. The first vehicle catalytic converters were available on cars in 1975, and today several hundred million have been fitted, mostly in the USA and Europe. In industry they are used for exhaust gas treatment and gas combustion (Pfefferle and Pfefferle, 1987) due to their ability to efficiently reduce exhaust emissions, as illustrated in Figure 3.1. As such, the gas-solid monolithic reactor is a well-proven technology.

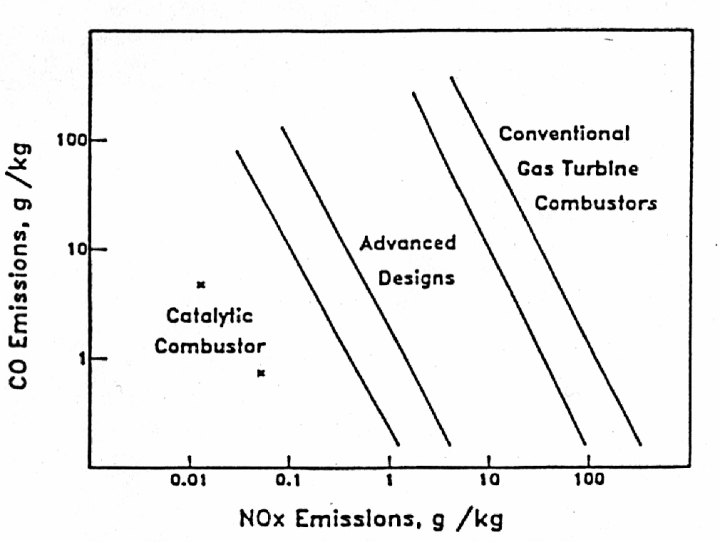


Figure 3.1 – Emission performance. Reproduced from Pfefferle and Pfefferle (1987).

The monolith reactor is widely used commercially and in industry, and has the potential for use in many different areas (Kapteijn *et al.*, 2001; Nijhuis *et al.*, 2001). This is reflected in the substantial literature available on the subject. Some examples of structured monolith reactors are shown in Figure 3.2.

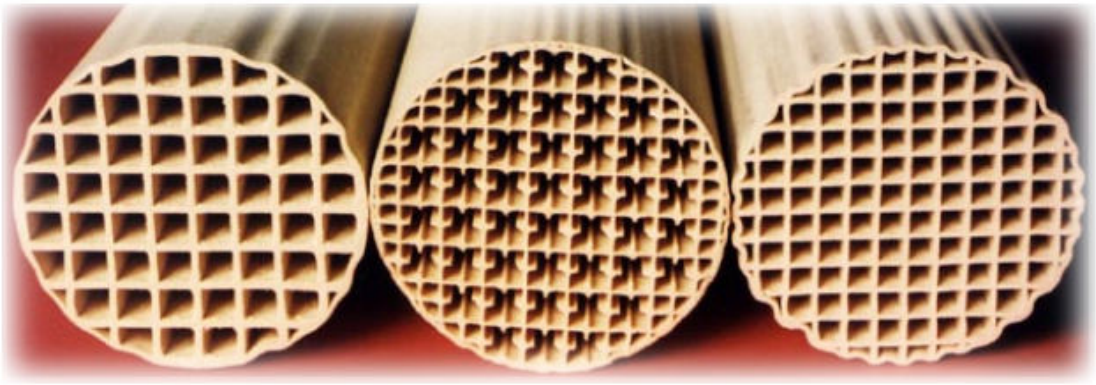

Figure 3.2 – Three geometries of monolith structure, showing (left-to-right) square channels, finned channels, and higher density square channels. The cross-sections are approximately 5 cm in diameter. Reproduced from Heibel *et al.* (2001).

However, there is a lack of experience and knowledge with gas-liquid flows through these structures, and hence a lack of verified mathematical models for their design, the application of which could be potentially very great in industry. The pace of research on monoliths as multiphase reactors has accelerated in the last decade, and the aim of this work is to better understand the local flow phenomena, and to relate this to reactor performance. A summary of some of the main advantages and uses of the monolith reactor are presented in the following section. Further reading is available in Cybulski and Moulijn (1994), Cybulski and Moulijn (1998) and Roy *et al.* (2004a).

## 3.3. Design of monoliths and monolith reactors

Although monoliths are sophisticated structures with complex geometries, high cell densities and specific structural requirements, they are commercially available at reasonable costs. They are at present still more expensive than traditional fixed bed reactors using traditional catalyst particles. The technology behind monolith reactors is still relatively new, with the first conference on monolithic catalysts organised in 1995 in St Petersburg. As such, many patents are still in place, though the first patents are now close to expiration which should bring the costs down significantly in future years. Also, the growth in the market for structured reactors should bring the bulk cost down. It should be noted that this high cost of monoliths is partly due to the fact that they are designed to meet the demands of high-temperature combustion operations – using cheaper, low-temperature materials will reduce costs where the

reactor conditions may not be as harsh as for automotive applications. There are many inorganic materials that can be extruded in the monolith production process that are possible contenders.

The explosion in combustion emission-reduction processes has been driven primarily by legislation, however many chemical processes will be driven by the technical and economic advantages of monolithic systems, particularly where they offer a process advantage such as better selectivity or higher yield, or the simplification of the design or operation of a unit.

In choosing design parameters for a monolith reactor, we have a number of options:

- substrate materials and geometry
- catalyst loading
- fluid distribution and flow regimes

These aspects will be addressed in the following sections.

## 3.3.1. Substrate materials

This section will first cover some of the methods used to create the structures specific to monolith reactors, and then discuss some of the possible substrate materials which are typically used in the construction of monolith reactors.

### 3.3.1.1. Monolith fabrication

Monoliths are generally made from extruded cordierite, a highly anisotropic ceramic crystalline phase that has high thermal expansion anisotropy, leading to orientation during extrusion and low overall expansion. Cordierite is a ceramic material of magnesium oxide, silicon dioxide and aluminium oxide in the ratio 2:5:2. Cordierite monoliths are used in automotive emission control, diesel particulate filters, indoor air purification and water filtration, amongst other uses. Cordierite has several key characteristics which make it suitable for use in industrial conditions (Williams, 2001), including good thermal shock resistance due to a low thermal expansion, a porosity and pore size distribution suitable for ease of washcoat application and

washcoat adherence, sufficient refractoriness because of the high melting point of around 1450-1730°C, sufficient strength for survival in a harsh environment, and compatibility with washcoats and catalysts. They typically have an axial crushing strength around 10-35 MPa, cell densities between 9 and 1600 cpsi (cells per square inch), an open area of 50-90%, and open porosity of 30-40%, and pore sizes of 5-15 μm (Cybulski and Moulijn, 1994; Woehl and Cerro, 2001). This ceramic support can subsequently be washcoated with a high surface area catalyst support material and impregnated, or alternatively extruded directly into catalytic bodies using appropriate materials such as zeolites, or $Al_2O_3$. Hence, a simple structured catalytic reactor, with a very well defined geometry can be achieved.

The combination of thin walls, straight channels and high cell density characterising a monolith reactor gives a large overall surface area for catalysis with a low backpressure. Typical geometries of monolith catalysts include round, square, hexagonal and triangular channels (Thulasidas *et al.*, 1995a; Zhao and Bi, 2001). More complex structures have been experimented with, such as channels with internal fins to stabilise counter-current film flow (Lebens *et al.*, 1999b).

High cell density structures are more difficult to manufacture and washcoat with catalyst, they are also associated with higher pressure drops (though still low compared with packed bed reactors), and they are more prone to flow heterogeneities than low cell density monoliths. The advantage of going to higher density cells is that there is a greater surface area available for reaction. In practice, densities above 600 cpsi are rarely used. Since they are the most widely applicable design, only square channels of densities 200-400 cpsi have been investigated in this research.

Cross-sectional diameters of up to 35 cm have been fabricated for use in heavy duty vehicles. Much larger diameter stacks are formed from square blocks stacked together, such as are found in industrial exhaust treatment towers. Lengths of monolith above two metres are impractical for normal production, so again stacks of monoliths must be used for heavy duty gas filters, or where a long residence time is required. It has been proven that when stacking monoliths, a graded cell configuration can improve performance i.e. having sections with large cells at the inlet, intermediate cells in the centre, and fine cells downstream (Hawthorn, 1974).

Stacks of monoliths which are rotated relative to each other can also help to redistribute any flow maldistribution (Satterfield and Ozel, 1977). When stacking monoliths, care must be taken – when liquid leaves the monolith it is accelerated under gravity, forming a *vena contracta*, and so the second monolith section must be placed very close such that menisci are formed, thereby directing the flow from the first to the second section.

### 3.3.1.2. Materials of fabrication

Monoliths may be formed from many different materials, the most common being ceramic, carbon, or metallic. The fabrication and advantages of these fabrication materials will now be outlined.

Ceramic monoliths are formed by a simple process, but care must be taken to ensure product quality and integrity. The raw materials for the support are mixed and then plasticised with water or other solvents. Extrusion is done with a series of unique dies developed by the manufacturers, and the extrudate is dried and fired to attain the desired properties. Enough moisture must be driven off without cracking the monolith body. Advances in ceramic processing have led to the development of higher cell density materials with improved strength and thermal properties.

Extruded carbon monoliths may be formed in a similar manner (Garcia-Bordeje *et al.*, 2002), but with the use of phenolic resins in the extrusion step, followed by curing, then carbonising in nitrogen and activation in $CO_2$ or steam. Carbon monoliths are used for water purification, VOC emission control, air purification, and other uses. Carbon is relatively inert, so the undesired side reactions catalysed by the support surface hardly occur. Recovery of noble metals is very easy by burn-out of the carbon support. Properties of the carbon are adjustable over a wide range of values, such as surface area, pore-size distribution and isoelectric point. They can either be formed from full-body carbon extrusion, or as a coating on a cordierite substrate; these coated monoliths are significantly stronger than all-carbon monoliths, however care must be taken to ensure that the coating adheres well to the support to avoid flaking caused by the shrinking of the coating during carbonisation. This can cause loss of active

material and damage to downstream equipment, so it is important to obtain a suitable formulation of washcoat to avoid such problems.

Metal monoliths may be fabricated from crimping a sheet of metal foil on a pair of rollers having teeth of a predefined profile (usually sinusoidal or triangular) and this is then combined with a flat foil sheet, and the two are rolled around a spindle until the right diameter is reached. Fabrication is completed by the welding of the outside edge (Cybulski and Moulijn, 1998). Variation of the cell density is achieved by varying the pitch and width of the profile on the crimping rolls. Metal substrates have the benefits of high heat conduction, which may be useful if the reaction being performed is exothermic and requires cooling. Ceramic monoliths typically have a heat conductivity of around 0.1 $Wm^{-2}K^{-1}$, whereas for metals this is nearer to 50 $Wm^{-2}K^{-1}$. Flytzani-Stephanopoulos and Voecks (1980) studied steam reforming of hydrocarbons in coated metallic monoliths, and heat transfer was improved to such an extent that unreacted hydrocarbons exiting the bed were five times lower than those in packed beds. Another advantage of metal supports is that they can be manufactured with thinner walls, thereby giving a large surface area with reduced backpressure. Typical ceramic automobile catalysts have walls around 150 μm thick; with metal supports, 40 μm is possible and 50 μm is the standard. Metal monoliths will also conduct electricity – this can be used to electrically heat the monolith, which may improve the reaction kinetics – for example, in automobile applications the catalyst takes time to warm up, leading to sub-optimal performance on start-up. An electrical current could also be used to drive a chemical reaction, or to drive an electro-osmotic flow in the case of very narrow channels (~10s of μm) where pressure driven flow is impractical.

Less common monolith materials which are used for gas-phase applications, such as paper or laminated fibreglass, will disintegrate in liquid systems and so are not used.

### 3.3.2. Catalyst loading

As we have seen, monolith catalyst substrates may be manufactured from metal, ceramic or carbon. Typically, the surface area of these supports is low (around 0.7 $m^2/g$) and so they are then washcoated with a highly porous catalyst support

material to increase the available surface area (to around 240 $m^2/g$). The monolith is then impregnated with catalyst material, as illustrated in Figure 3.3. $Al_2O_3$ is used as a high surface area washcoat, though in the case of slow chemical reaction compared to diffusion, a thin washcoat may be inadequate and in this case catalyst substrates may be made entirely of $Al_2O_3$.

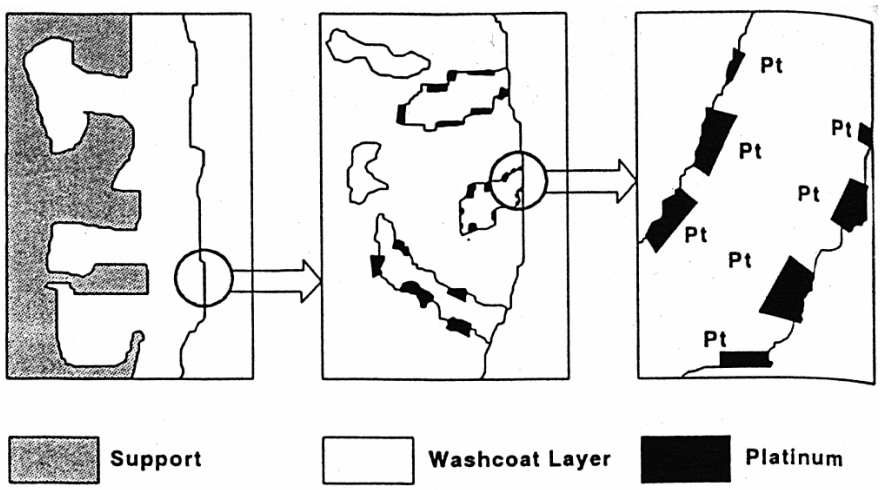


Figure 3.3 – Coated catalyst for automotive emissions control. Reproduced from Cybulski and Moulijn (1994).

The heat treatment and composition of the monolith can directly affect the wall morphology and hence dictate the catalyst loading and pore structure characteristics of the reactor. The work of de Lathouder *et al.* (2004) measured residence time distributions (RTDs) in monoliths with differing pore structures, and analysed their activities. They report that the monoliths with larger pore structures allowed high and accessible catalyst loadings. Dow Chemicals have produced a monolith with significantly permeable walls, allowing diffusion of material between neighbouring channels, thereby reducing the negative effects of stagnant channels. However they found that there was a wider velocity distribution and the monolith behaved less like plug-flow, which may discount its potential advantages (Bakker *et al.*, 2005).

The activated monolithic catalyst units are simple to handle, and can be easily stored and transferred to a reaction vessel when needed. In this way, catalysts may be stored on the shelf until required.

A comparison of the performance of rounded and square channels was presented by Perez-Cadenas *et al.* (2005). They found that during washcoating, the round channels acquired an even coating of catalyst, whereas the square channels tended to be heterogeneous, with more coating accumulating in the corners. They report that the rounded channels gave better activity and selectivity under reaction conditions (hydrogenation of esters). Liu and Roy (2004) report that circular channels gave 2-5 times better performance than square channels, which they attributed to a more uniform liquid film in circular geometries. Hayes and Kolaczkowski (1994) report that for square channels, the coating is 10-15 times thicker in the corners than in the mid-side. This uneven coating was then prone to cracking during later heat treatment preparations, as also reported by Kapteijn *et al.* (2005). Kolb and Cerro (1991) present some techniques for improving the coating of the inside of square capillaries.

## 3.4. Potential advantages

This section will outline some of the advantages of monolith reactors, over traditional choices such as slurry reactors and packed bed reactors. Reactor choice is driven by technical and economic factors, especially where they offer advantages such as better selectivity or higher yield, or the simplification of the design or operation of a unit. Monolith reactors are an attractive alternative to slurry reactors or trickle bed reactors (also called packed bed reactors) which are commonly used for three-phase chemical reactions. Some typical reactor parameters are summarised for comparison in Table 3.1 and Table 3.2.

Monolith reactors have the advantage that the mass and heat transfer zones can be separated and scaled independently, giving an extra degree of flexibility in reactor design and control over traditional reactors. The channel opening and catalyst wall thickness can be varied independently, whereas in a packed bed the only geometry parameter is the pellet size. There may also be advantages in terms of mass transfer in two phase systems due to high catalyst wetting, high effectiveness factors, and thin films that are conducive to high mass transfer rates. The following sections will compare the specific characteristics and advantages of slurry and packed bed reactors with those of the monolith reactor.

Table 3.1 – Comparison between typical ceramic monoliths and packed bed reactors for gas phase reactions. Reproduced from Cybulski and Moulijn (1998).

| | **Monolith** | **Packed Bed** |
|---|---|---|
| Particle / channel diameter (mm) | 1 - 4 | 1.0 - 5.0 |
| Voidage | 0.65 | 0.42 |
| Specific surface area ($m^2/m^3$) | 1200 - 3800 | 1000 - 3000 |
| Volume fraction of catalyst | 0.05 - 0.25 | 0.55 - 0.60 |
| Diffusion Length (μm) | 25 - 150 (washcoat thickness) | 100 - 2500 (depending on catalyst structure) |
| Flow characteristics | Re = 50 - 600 | $Re_p$ = 50 - 500 |
| Pressure drop (air at STP, kPa/m) | 0.1 at Re = 300 | 27 at $Re_p$ = 150 |

Table 3.2 – Comparison between typical monolith, trickle bed and slurry reactors for gas-liquid reactions. Adapted from Edvinsson and Cybulski (1995).

| | **Monolith** | **Trickle Bed** | **Slurry** |
|---|---|---|---|
| Particle / channel diameter (mm) | 1.0 - 4.0 | 1.0 - 6.0 | 0.001 - 0.200 |
| Volume fraction of catalyst | 0.05 - 0.25 | 0.55 - 0.6 | 0.005 - 0.01 |
| Superficial gas velocity (m/s) | 0.01 - 1.0 | 0.05 - 1.5 | Stirrer rate 500 - 1000 rpm |
| Superficial liquid velocity (m/s) | 0.01 - 0.20 | 0.005 - 0.05 | |
| Pressure drop (kPa/m) | 3.0 | 50.0 | 6.0 |
| Mass transfer, liquid/solid ($s^{-1}$) | 0.03 - 0.09 | 0.06 | 1 - 4 |
| Mass transfer, gas/liquid ($s^{-1}$) | 0.05 - 0.30 | 0.06 | 0.1 - 1 |

### 3.4.1. Slurry reactors

Comparing monolith and slurry reactors, it is found that slurry reactors can obtain high conversion and selectivity but have the problem of catalyst separation and attrition, which is avoided in the monolith reactor. Precious metal catalysts are not typically used in slurry reactors due to the difficulties in recovering the catalyst. The monolith has equally high selectivities as a slurry reactor, but without the complication of catalyst particle handling. Machado *et al.* (2005), present an

economic comparison of slurry and monolith reactors, and report that for fast reaction chemistries, monolith reactors can achieve volumetric activities many times higher than a slurry reactor. On the other hand, if the process is kinetically limited, mass-transfer effects are removed and the rate of reaction becomes dependent directly on the mass of catalyst present. In this special case, the slurry reactor (and equally a trickle bed) will outperform the monolith reactor due to the higher load and lower catalyst cost. The authors also take into account the reduction in catalyst handling for monolith systems, increased control, increased safety, increased hygiene and lower catalyst cost, and weigh this against the higher manufacturing cost of monoliths. They emphasise the importance of catalyst life, stating that monolith catalysts need to last ten times longer than a slurry catalyst to present a practical economic alternative. Other comparisons of slurry reactors and monoliths have been reported by Cybulski *et al.* (1999) and Boger *et al.* (2003).

An alternative monolith configuration suggested by the Delft group (Albers *et al.*, 1998; Hoek *et al.*, 2004) is the Novel Monolithic Stirred Reactor. This design attempts to combine the advantages of the slurry reactor, but with a fixed catalyst in the form of monoliths. The illustration, shown in Figure 3.4, consists of an arrangement of monolithic blocks in a stirrer-like configuration. The power requirement for stirring is relatively low compared to the agitation of slurry reactors due to the permeability of the monolith structure, and a large throughput can be realised through the catalyst channels. This design could be operated in batch or CSTR mode. They have only operated lab-scale prototypes in liquid-phase mode, but with the addition of a sparger to create a bubbly froth, two-phase reactions should also be possible.

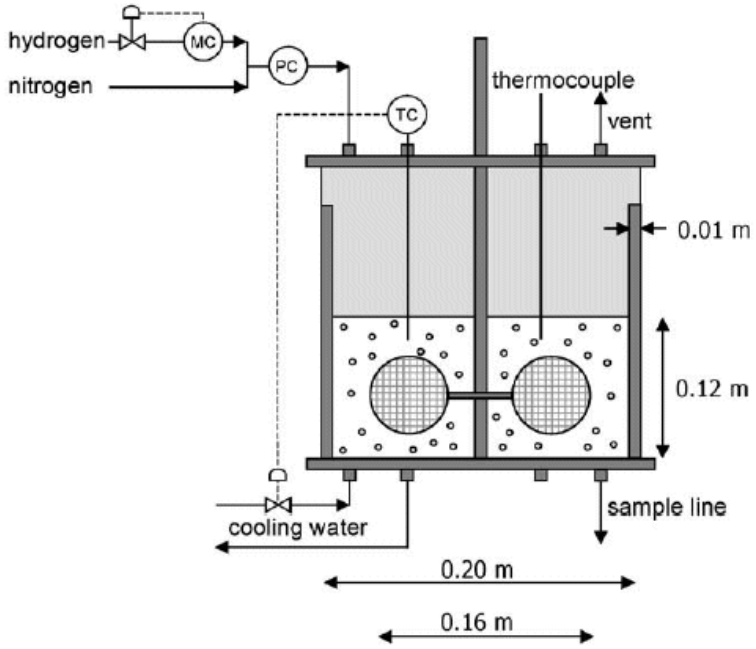


Figure 3.4 – Schematic drawing of the Novel Monolithic Stirred Reactor. Reproduced from Hoek *et al.* (2004).

## 3.4.2. Trickle bed reactors

Comparing trickle bed and monolith reactors, it is found that trickle bed reactors are sensitive to flow maldistributions that can lead to stagnant regions, hot-spot formation and runaway behaviour. They also can have a wide range of residence times due to short-cut routes which form through the reactor, and the presence of stagnant regions. In contrast, monolith reactors tend to have far sharper RTDs due to parallel flow and operation in the Taylor flow regime where the gas bubbles act as pistons separating the liquid slugs and preventing back-mixing (Edvinsson and Cybulski, 1995; Thulasidas *et al.*, 1999). Taylor flow is covered in detail in Section 3.7. The pressure drop across a monolith reactor is around 20-100 times less than for an equivalent randomly packed bed (Satterfield and Ozel, 1977; Liu *et al.*, 2002). Experiments have shown that monoliths can be 50% more productive per reactor volume than trickle beds (Nijhuis *et al.*, 2001), even though the porosity of the monoliths is apparently very high. The activity of the catalyst fixed in the trickle bed was actually higher, indicating the superiority of the monolith structure over the trickle bed in terms of performance. There is room for improvement of the catalyst activity on the support by better deposition, which would further increase the effectiveness of the

monoliths. The thin coating of catalyst particles gives a short diffusion length, and the high mass transfer rates possible allow near-intrinsic reaction rates to be achieved. Also, the catalyst is found to be better utilised in the monolith reactor; around 40 times more efficiently than in a trickle bed (Kapteijn *et al.*, 2001), reducing precious metal costs significantly. A comparison of the performance of a trickle bed and monolith reactor is presented in Figure 3.5. Stankiewicz (2001) reported that for saturated, single phase flow the monolith reactor required a reactor volume between 6 and 53 times less than for a conventional packed bed reactor.

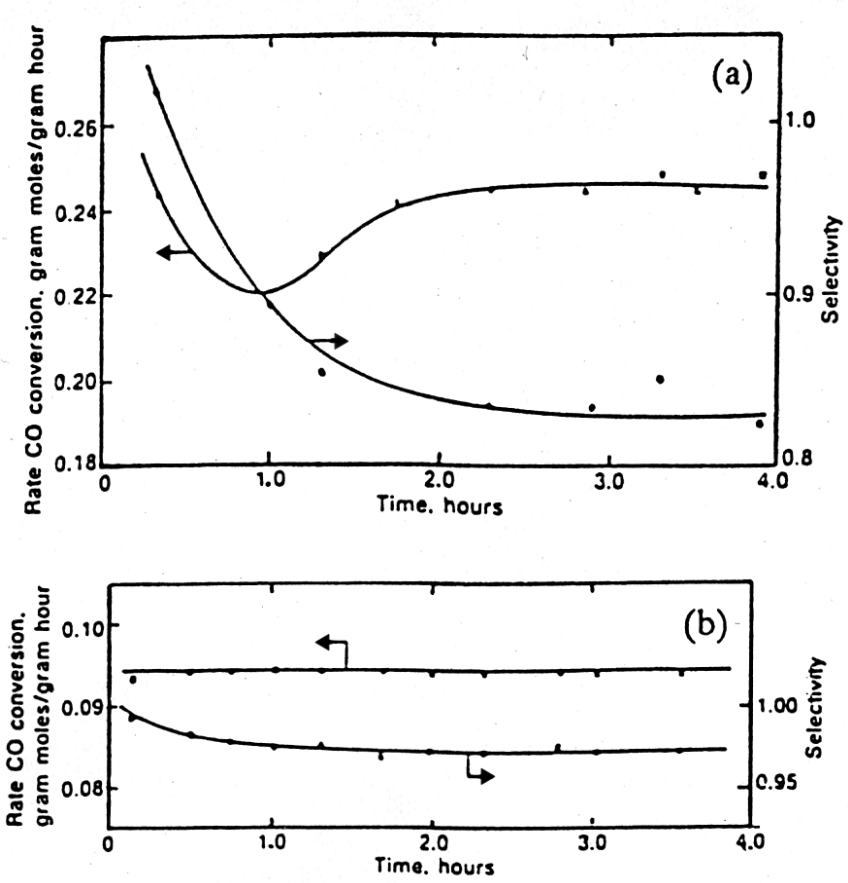


Figure 3.5 – Methanation conversions and selectivities (a) packed bed, (b) monolith. Reproduced from Tucci and Thomson (1979).

## 3.5. Industrial application of monolith reactors

The commercial uptake of monolith reactors is a relatively recent development occurring over the last 10-15 years. They have found application in several areas of industry despite their novelty. Two-phase flow in monolith reactors is used for

petrochemical, organic and inorganic synthesis. Fine chemicals, biotechnology and electrochemistry are rapidly growing sectors. This section will cover some of these applications in more detail.

Typical applications of monolith reactors include steam reforming, methanation, hydrogenation and oxidation processes, however the only currently commercially significant, industrial-scale, three phase reaction involving monolith reactors is found in the anthraquinone hydrogenation process. This is a main step in the auto-oxidation process for hydrogen peroxide manufacture, and is practised by Eka-Nobel Chemicals on the scale of over 200 kt/yr (Edvinsson and Cybulski, 1995; Cybulski *et al.*, 1999; Albers *et al.*, 2001). It is worth noting that this process demands very high selectivity of over 99%. The composition of the oil can vary considerably and so a batch process is usually the method adopted. The slurry reactor has been the traditional choice (Edvardsson and Irandoust, 1994), but monolith reactors allow fast and easy adaptation to new process requirements, and avoid the complex and expensive recovery of the catalyst.

The remainder of this section will discuss some specific literature on lab- or pilot-scale reactors which compare reactions in monoliths with other reactor types. These are of potential interest to industry, though not yet implemented on a production-scale.

For example, Irandoust and Andersson (1988), illustrate the performance difference between a slurry and monolith reactor for the hydrogenation of 2-ethyl-hexenal under similar conditions, shown in Figure 3.6. The graphs appear similar, but the timescales are very different, which clearly illustrates the superior mass transfer properties of the monolith reactor. They acknowledge that part of the lower activity in the slurry reactor may be due to deactivation of the catalyst during grinding.

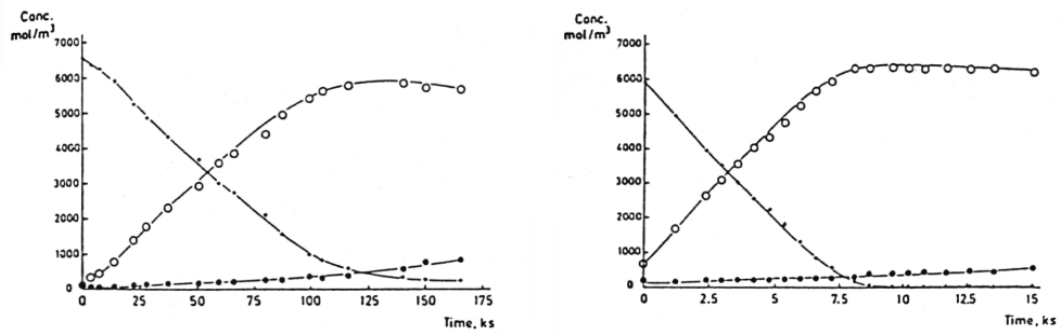


Figure 3.6 – A comparison of hydrogenation of 2-ethyl-hexenal in (a) slurry and (b) monolith reactors. Hydrogen pressure 9.8 bar; temperature 160°C; (+) 2-ethyl-hexenal; (○) 2-ethyl-hexanal; (●) 2-ethylhexanol. Reproduced from Irandoust and Andersson (1988).

Cybulski *et al.* (1999) compared the performance of a monolith reactor to a mechanically agitated slurry reactor for the hydrogenation of 3-hydroxypropanal to 1,3-propanediol. They found that the productivity of a monolith reactor was generally much higher, while the energy needed for mixing / circulation was much lower. They also point out that the slurry reactor becomes preferable in the case of fast catalyst deactivation, since it is far simpler to replace catalyst material in the case of the slurry reactor.

Winterbottom *et al.* (2003b) compared the hydrogenation of 2-butyne-1,4-diol in a monolith reactor to that in a slurry reactor and a stirred tank reactor. They found that the rate of reaction was greater for the monolith reactor, and that the selectivity was higher for the monolith and slurry reactor (0.980-0.993) than for the stirred tank (0.90-0.95) at a conversion approaching 100%. They concluded that the monolith has similar performance to the slurry reactor, but with the added advantage of easier catalyst handling and separation. This group later also compared residence time distribution and selectivity in a monolith reactor and trickle bed (Kulkarni *et al.*, 2005). They found that the RTDs for the two reactors were relatively similar, however the mixing characteristics of the monolith reactor gave better selectivities, shown in Figure 3.7.

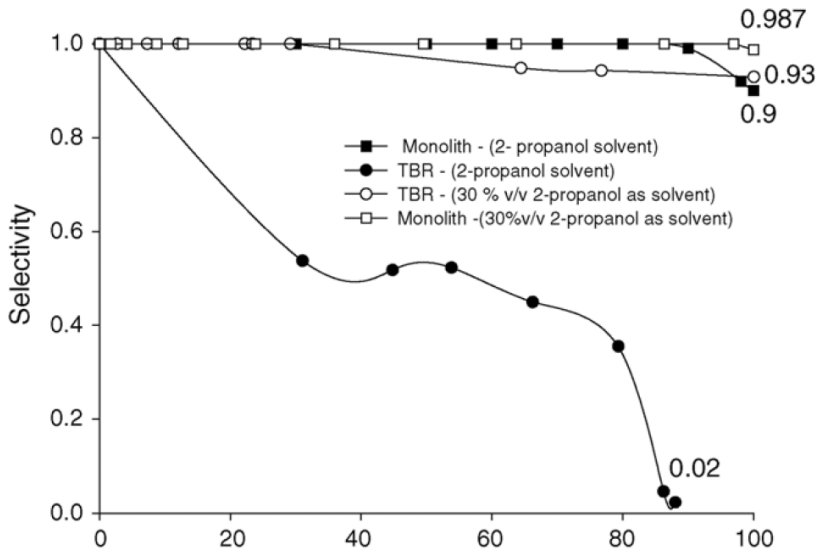


Figure 3.7 – Selectivity vs. conversion of 2-butyne-1,4-diol in TBR and monolith reactor, with carrying solvent concentrations. Reproduced from Kulkarni *et al.* (2005).

Other examples of comparing monoliths and packed beds includes work by Katziantoniou and Andersson (1984) who studied the hydrogenation of nitrobenzoic acid in a palladium-coated monolith. The effectiveness factor for the monolith ranged from 0.081 to 0.115, while for 5 mm spheres in a trickle bed, it amounted to 0.021 to 0.024, mainly attributed to shorter diffusion lengths in the monolith. Similar results were obtained by Soni and Crynes (1980), who studied hydrodesulphurisation and hydrodenitrification in monolith and trickle bed reactors. The monolith catalyst activity was higher for that of the pellets, and they attributed this to the differences in average pore radii and intra-particular diffusion lengths – for monoliths these were 8.0 nm and 0.114 mm versus 3.3 nm and 1 mm for pellets. Effectiveness factors for the monolith and pellets were found to be 0.94 and 0.216 respectively.

Liu *et al.* (2002) compared a monolith reactor (made of iron oxide), a packed bed reactor and a packed bed in radial flow configuration for the dehydrogenation of ethylbenzene to styrene. For clarification, packed beds may be operated in radial flow where the reaction mixture flows into a central channel and then outwards through an annulus of catalyst packing. This gives far lower pressure drops than for conventional packed bed reactors, within the range of monolith reactors. However, these types of reactor have low space-utilisation and there are problems with the mechanical support of the catalyst particles under high-stress operation. The authors found that the

pressure drop across the monolith reactor was around 20 times less than that for a conventional packed bed of 1/8 inch catalyst beads. They also claim that a monolith reactor will outperform a radial flow reactor (packed with catalyst beads) by 10% for the same reactor volume, and recommend retro-fitting existing reactors with their new design. Since monolith reactors are more compact, this allows the inclusion of additional catalyst volume or additional heat transfer equipment. Around 30-50% of reactor space can be freed in this way, as shown in Figure 3.8. This is especially useful in highly exothermic or endothermic reactions since inter-stage heat exchangers can be built in, improving reaction kinetics and hence achieving better conversion.

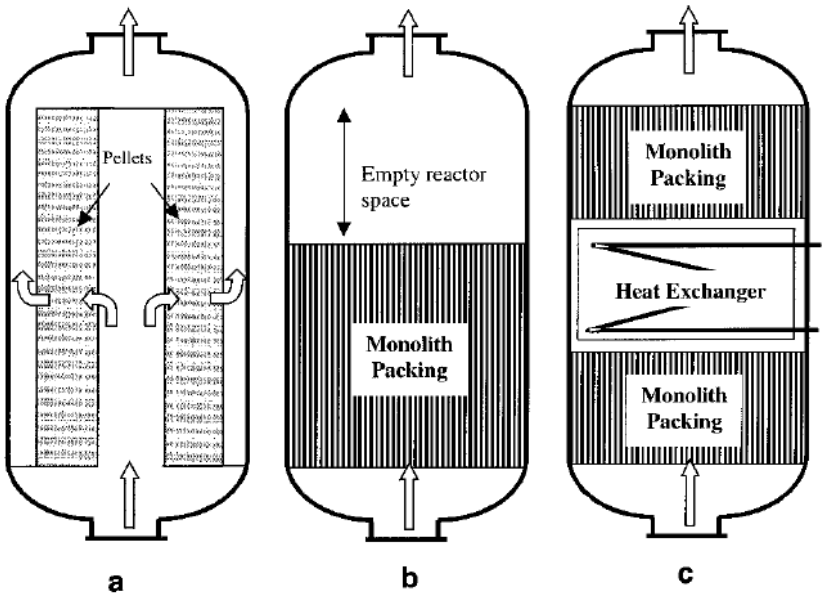


Figure 3.8 – Illustration of (a) a radial flow reactor; (b) a monolith reactor of equivalent performance; and (c) an illustration of how this empty space might be utilised. Reproduced from Liu *et al.* (2002).

The reverse process (hydrogenation of styrene) was studied by Nijhuis *et al.* (2003) in monolith and trickle beds. They found that the monolith gave a three times higher volumetric productivity with much less catalyst. They estimate that the concentration of hydrogen at the catalyst surface is 34 and 2.2 mol/m$^3$ for the monolith and trickle bed respectively, and this correlates to the much improved catalyst utilisation. They also point out that gum formation within the monolith reactor is significantly lower due to much better hydrogen mass transfer.

Other studies of reactions in monolith reactors have been reported in the literature. Examples include: nitrate reductions (for water treatment) by Warna *et al.* (1994), the hydrogenation of phenylacetylene by Vergunst *et al.* (2001), the hydrogenation of edible oils by Edvardsson and Irandoust (1994) and Boger *et al.* (2003), methanation by Tucci and Thomson (1979), reactive stripping (counter-current mode) by Schildhauer *et al.* (2005), alcylation of aromatics by Beers *et al.* (2001a), Fischer-Tropsch synthesis by Schanke *et al.* (2001) and Kapteijn *et al.* (2005), glucose oxidation by Kawakami *et al.* (1989), and the synthesis of vinyl acetate by Nicolau *et al.* (1998).

In summary, the literature claims that monoliths obtain higher effectiveness factors, rates of reaction, selectivities and better catalyst utilisation than traditional reactors.

## 3.6. Fluid distribution

The method of fluid distribution to the monolith reactor needs careful thought, since the low pressure drop across the monolith reactor is a potential problem – unlike packed beds, monoliths do not damp non-uniformity of flow which may enter at the inlet to the reactor, and so this heterogeneity is propagated along the length of the reactor. Of course, the opposite is also true – a good distribution will not be spoilt while propagating down the reactor. Therefore it is important to equalise the flow across the cross section before the entrance to the monolith using a reliable flow distributor. The question of fluid distribution is a crucial one, though it seems to have been neglected by many authors partly because the quantification of the quality of the distribution is far from straightforward, and most observations are based on indirect effects such as extent of reaction and residence time distributions. Several authors have used the criterion of a "repeatable pressure drop" to quantify the distribution. They propose that if, by a small adjustment of the packing or distributor, the pressure drop is not reproducible, then the system is not well-distributed. However, it is far more desirable to have a direct measurement of the dynamics inside the reactor itself, as this will ultimately lead to better understanding of the processes involved. It will be shown in Chapter 4 that MRI and other tomographic methods can provide a superior measure of maldistribution in a non-invasive and quantitative manner, and Chapter 10 will show the optimisation of distributor design by MRI measurements.

The following two sections of this chapter will outline some fluid distribution methods and outline their advantages and applications.

### 3.6.1. Traditional distribution methods

To ensure an even gas-liquid distribution, the two-phase flow to the reactor is generally created by either a showerhead in the case of liquid distribution, or a fritted glass plate for bubbly flow. Characteristics of different porous and fritted glass plate designs are investigated in Akita and Yoshida (1974), with reference to the resultant bubble sizes and mass transfer coefficients. The simplest solution to the problem of achieving homogeneous fluid distribution has proved to be to spray the liquid in small drops over the surface of the monolith (Satterfield and Ozel, 1977). The drops must be much smaller than the monolith channels, and in spraying them the monolith becomes wetted. The liquid begins to flow downward in a thin film on the channel wall, and as more liquid accumulates a meniscus is formed, and then a liquid plug. Sieve plates have also been used; these use larger drops, and are arranged such that the distance between the holes is of the same order as the channels. The drops break up when they hit the surface and form slugs. Irandoust *et al.* (1989) used a sieve plate for distribution, with holes somewhat smaller than the monolith channels and found they gave good, consistent performance. A low pressure drop across the perforated plates is desirable to minimise the required liquid pumping energy. However, below a critical limit (100-200 mm of water) they found that the liquid becomes poorly distributed, resulting in a much less fine dispersion. For the same porosity of a plate, smaller holes will give better results, though very small holes are liable to plug. The advantage of a sieve plate is that it requires a lower pressure drop, however both methods give little control over bubble and slug lengths. Optimisation of these requires better understanding and control over parameters such as gas and liquid flow rates, and the monolith channel density (higher density cells will slow the descent of the fluid).

Lebens *et al.* (1999a) report that, using a spray distributor and internally-finned monolith, flow maldistribution was around 10% for average to high flow rates, but increased to 25% at lower flow rates. However, they do not discuss how they took

these measurements, or quantified the maldistribution, though it is likely to be similar to other methods used at the Department of Reactor and Catalysis Engineering at Delft University. These methods, along with others, are discussed later in Section 7.5.1.

An alternative method to showerhead and spray distributors is to use pulsed flow to generate slugs and bubbles (Irandoust and Andersson, 1988). Flooding the monolith with liquid for a short time gives distribution to all channels. The linear velocity in the monolith is very high, in the order of 0.2-0.5 m/s, and since the slug lengths should be on the order of 1-2 cm, the frequency of pulsing must be around 10-50 Hz. This method does not provide a very even distribution unless the liquid can be added sufficiently fast and uniformly over the monolith.

### 3.6.2. Novel distributors

A further variation is to use a bubble sparger setup for gas distribution in a high liquid holdup system, where gas is bubbled up through a glass frit, and passes over the monolith co-currently with the liquid phase. A sparger used in this work is pictured in Figure 3.9.

Figure 3.9 – A photograph of the bubble sparger used in this work, consisting of a sintered glass frit of diameter 20 mm.

An extension of this configuration is the monolith foam reactor design, which also utilises a porous glass frit in upflow configuration, though with a high gas holdup configuration (Crynes *et al.*, 1995; Thulasidas *et al.*, 1995b). Gas and liquid are introduced through the glass frit and a stable bubbly foam is produced, which then moves upwards through the monolith catalyst. A large pressure drop (around 10 bar) is applied over the glass frit to produce this bubbly foam. The advantage of operating in a foam configuration with high gas holdup is that the gas and liquid are intimately mixed, with the liquid forming a very thin film surrounding the gas bubbles, aiding mass transfer. Characteristics of the foam can be controlled by varying the gas and liquid flow rates, and the pore size in the distribution plate. Hence, the development of slug-flow behaviour can be controlled, but optimal conditions must be found empirically. The foam reactor obtains very good mass transfer properties, however the residence time distribution of the liquid phase corresponded to a stirred tank reactor, indicating a high degree of undesirable back-mixing within the foam (Patrick *et al.*, 1995). With foam reactors, it has been shown that flow can proceed in any direction; up (Crynes *et al.*, 1995), down and horizontally (Stankiewicz and Moulijn, 2000).

Air Products and Chemicals Inc. have implemented a liquid-motive gas ejector (Broekhuis *et al.*, 2001). The ejector uses liquid flow to entrain and compress recycled hydrogen. When properly engineered, the ejector also serves as an effective gas-liquid contactor, fully pre-saturating the liquid before it enters the reactor, and producing a fine dispersion of gas bubbles in the liquid, which results in excellent gas-liquid distribution to the reactor. Because the ejector acts as a gas compressor, the gas-liquid dispersion can be delivered to the reactor inlet at a pressure greater than the outlet. The efficiency of the ejector is relatively low i.e. only a small fraction of the pressure energy of the liquid is converted into the pressure energy of the gas, however, the additional energy dissipation ensures that the gas and liquid form a fine dispersion. While this consumes more energy than gravity-driven flow, it greatly expands the range of practical operating conditions, and moreover allows higher density monoliths (with greater surface area) to be used. This ejector system is similar in principle to the turbulent froth reactor (Section 8.2.4) and CDC reactor (Winterbottom *et al.*, 2000). These ejector reactors also operate in downflow with a high liquid holdup. Gas and liquid are introduced through an orifice, which creates a

turbulent froth, and this entrained, bubbly froth then passes down over the monolith. The CDC reactor (Figure 3.10a) was designed by the Department of Chemical Engineering at Birmingham, and the turbulent froth reactor (Figure 3.10b) is a modification developed by the Department of Chemical Engineering at the University of Cambridge.

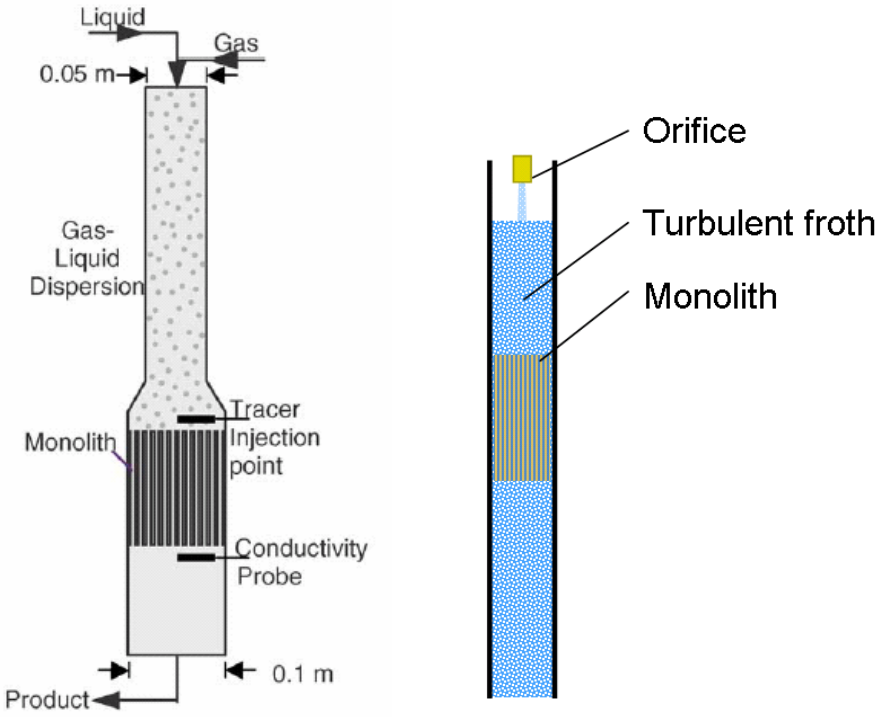


Figure 3.10 – Schematic of the setups for monolith two-phase flow configurations; (a) The CDC reactor (reproduced from Kulkarni *et al.*, 2005); and (b) the turbulent froth reactor (Section 8.2.4) developed in the Department of Chemical Engineering at the University of Cambridge.

In regards to the more traditional methods (showerhead / spray distributors), there is in general an optimum nozzle position, and at this position the local velocities are even over the monolith cross section (Heibel *et al.*, 2001). At low superficial velocities, maldistributions are easily formed, however at higher, more commercially relevant throughputs an even distribution is easily formed (Nijhuis *et al.*, 2001).

In this work, a selection of different distribution methods have been implemented in order to compare their relative performance. Bubble spargers, showerheads, spray

distributors and froth reactors have all been used, and the results of MRI investigations are summarised in Chapters 8 to 12.

## 3.7. Flow phenomena in monolith reactors

To fully understand the characteristics of two-phase flow in an array of channels such as the monolith reactor, we must first understand the flow regimes which we may encounter within a single channel. There is considerable literature devoted to the study of individual capillaries under single- and two-phase flow. This section will outline the flow regimes we may expect, highlight the advantages of Taylor flow, then summarise some of the mass transfer effects under varying operation conditions, and conclude with a summary of RTD investigations of monolith reactors.

### 3.7.1. Flow regimes

Some researchers have investigated the use of liquid saturated with gas, as this behaves as a single-phase and avoids the complication of trying to control two-phase flow behaviour, greatly simplifying the design of a reactor system (Vaarkamp *et al.*, 2001). However, there are distinct advantages in operating under a two-phase flow regime, as will be shown later. The choice of flow regime has a large effect on the mass-transfer characteristics and hence the performance of the monolith reactor. This section will outline the advantages of choosing Taylor flow (slug flow) as the flow regime of choice in a monolith reactor, and report mass transfer studies to support this choice.

Much work has been done to try to create a flow-regime map to predict two-phase flow patterns. The main flow regimes are illustrated in Figure 3.11.

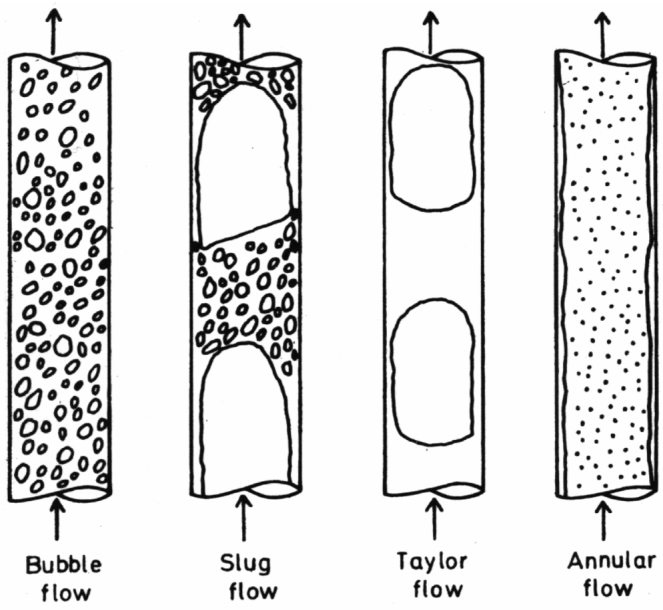


Figure 3.11 – Illustration of flow regimes for two-phase flow, with increasing gas holdup from left to right. Reproduced from Irandoust and Andersson (1989).

There has been relatively little work done in capillaries where surface tension forces become significant (Table 3.3), compared to the volumes of literature available for larger diameter pipes. However, efforts to relate large and small scale phenomena are generally progressing well. A comprehensive summary has been written by Triplett *et al.* (1999).

Table 3.3 – Literature references for selected flow transition maps

| Author | Geometry | Pipe diameter | Orientation |
|---|---|---|---|
| Suo and Griffith (1964) | Circular | Dimensionless capillaries | Horizontal |
| Barnea *et al.* (1983) | Circular | 4 - 12 mm | Horizontal and vertical |
| Barajas and Panton (1993) | Circular | 1.6 mm | Horizontal |
| Fukano and Kariyasaki (1993) | Circular | 1.0 - 4.9 mm | Horizontal and vertical |
| Zhang and Giot (1995) (counter-current flow maps) | Circular | various | Vertical |
| Mishima and Hibiki (1996) | Circular | 1 - 4 mm | Vertical |
| Triplett *et al.* (1999) | Circular and Triangular | 1.1 - 1.49 mm | Horizontal |
| Zhao and Bi (2001) | Triangular | $d_h$ = 0.9 - 2.9 mm | Vertical |

Taylor flow was first described by Taylor (1961) and describes the streamlines which occur around a bubble passing along a fluid-filled tube, illustrated in Figure 3.12. The development of laminar parabolic flow in the liquid slugs is prevented by the presence of bubbles. The liquid phase shows internal re-circulation, and because of this radial mass transfer is increased. Near-plug flow is obtained as gas and liquid have sharp residence time distributions. An illustration of the plug flow behaviour is shown in Figure 3.13 for a bundle of such capillaries. Note that the RTD is not as sharp as might be expected, due to the variance in behaviour of the bundle of capillaries which results in overall spreading of the RTD compared to that of a single capillary.

Figure 3.12 – Illustration of Taylor flow showing liquid recirculation within liquid slugs. Reproduced from Irandoust and Andersson (1989).

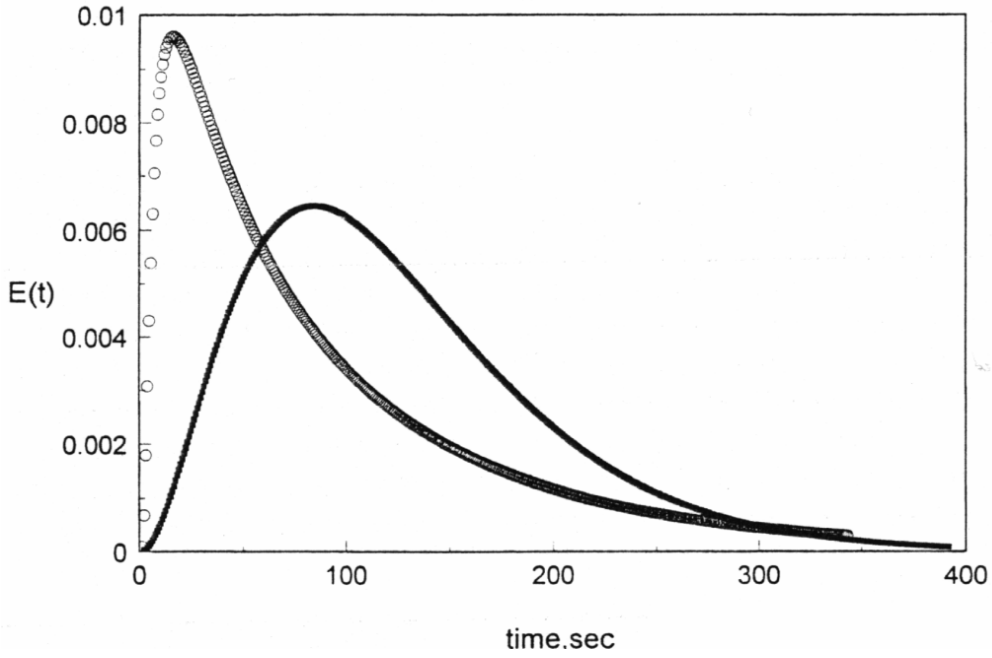


Figure 3.13 – RTD curves from a tracer study of flow within a capillary bundle. (■) liquid flow only, (○) liquid and gas under Taylor flow. Reproduced from Thulasidas *et al.* (1999).

In very short slugs, the liquid flow is not fully developed – the close proximity of bubble-ends restricts the development of the streamlines. Thulasidas *et al.* (1997) used Particle Imaging Velocimetry (PIV) to report that systems with liquid slugs longer than 1.5 times the diameter of the capillary show clear indications of fully developed Poiseuille flow between adjacent bubbles. Their measurement of characteristic recirculation times indicate that a typical PIV particle within a slug will move from one end to the other end of the slug during the time the slug takes to travel a distance of twice its length. This increased to three times its length for a square capillary, indicating a subtle difference in flow patterns. These observations indicate that in shorter slugs, fluid will come into contact with the wall more often, and improve mass transfer. So in summary, the liquid slug must be above a certain length for useful circulation to occur, but too long a slug will decrease the total amount of mixing; a short slug will complete more recirculation cycles than a long slug, resulting in increased radial transfer rates.

Visual measurements of capillaries of diameter similar to monolith channels indicate that the slug lengths obtained are around 2-5 times the diameter of the channel, depending on the gas and liquid flow rates (Heiszwolf *et al.*, 2001; Kreutzer *et al.*, 2001).

Thulasidas *et al.* (1995a) found a correlation between the ratio of bubble length to unit cell length as a function of the ratio of liquid flow to total liquid flow, shown in Figure 3.14 and Figure 3.15. It should be possible to verify some of these results experimentally using MRI, and an analysis of slug lengths is presented later in Chapter 9 for comparison. The ratio of liquid to total flow rate can be estimated for a monolith as a whole, and hence we can attempt to infer a value for $\beta$, which relates the slug and bubble lengths, which defines the mass transfer characteristics of the system. The authors found that this relationship was also valid (on average) for capillary bundles.

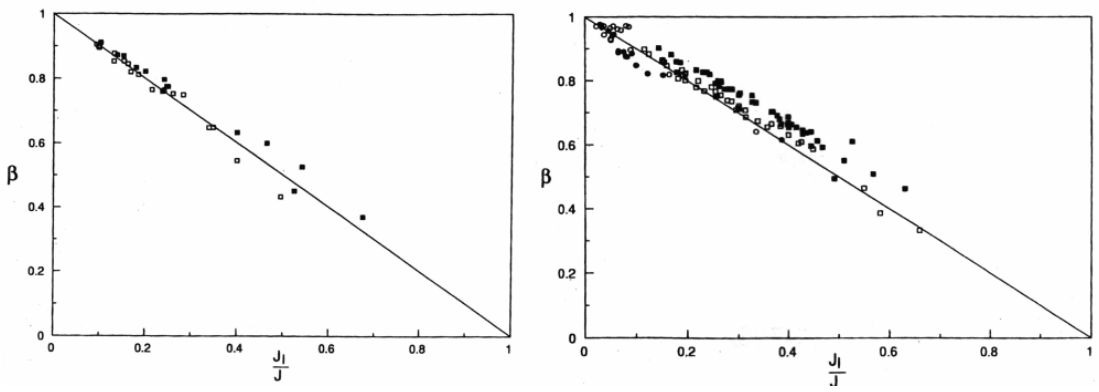


Figure 3.14 – Ratio of bubble length to unit cell length as a function of ratio of liquid flow to total flow rate for (a) circular and (b) square capillaries. Data for (■) 100 centistoke silicone oil and (□) 50 centistoke silicone oil, both in bubble-train flow. Reproduced from Thulasidas *et al.* (1995a).

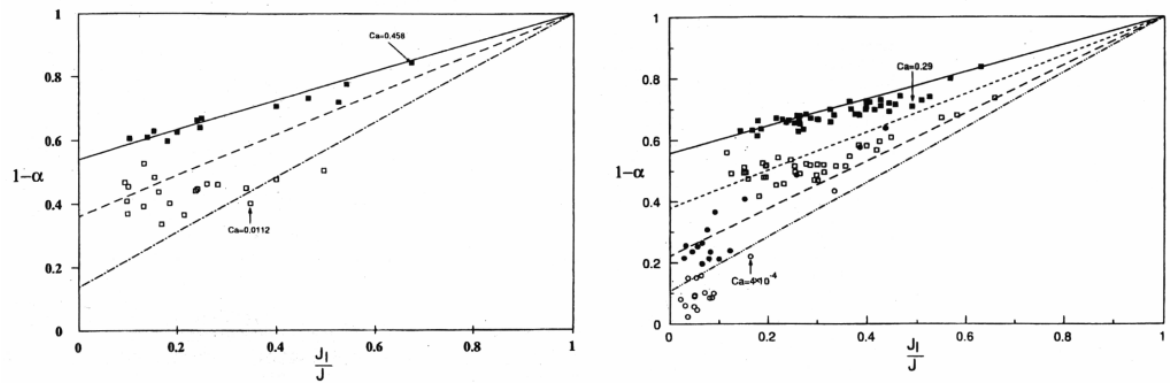


Figure 3.15 – Volume fraction of liquid as a function of ratio of liquid flow rate to total flow rate for (a) circular and (b) square capillaries. Data for (■) 100 centistoke silicone oil and (□) 50 centistoke silicone oil, both in bubble-train flow. The lines are drawn at constant Ca (top Ca = 1.0; middle Ca = 0.10; bottom Ca = 0.01). Reproduced from Thulasidas *et al.* (1995a).

Heiszwolf *et al.* (2001), Kreutzer *et al.* (2005c) and Yawalkar *et al.* (2005), claim that the liquid slug lengths can be estimated from a simple pressure drop measurement. The justification behind this bold statement originates from work done showing that the liquid slug length is only a function of the hold-up and does not depend on the gas or liquid velocities. For slugs which do not exceed 10 times the channel diameter, they claim that the pressure drop can be correlated directly to the holdup and hence the average slug length. If this works, then from a simple and non-intrusive method the mass transfer rates can be modelled using the predicted slug length. They also

note that the experimental slug length is very much dependent on the type of fluid distributor used, as well as the choice of flow regime.

### 3.7.2. Mass transfer in Taylor flow

For three-phase reactions, the gas phase must diffuse through the liquid to reach the solid catalyst. The rate of diffusion is generally small compared to the rate of reaction, and so mass transfer rates are important. In Taylor flow, gas bubbles move through the monolith channels filling nearly the entire width of the channel, and gas is separated from the catalyst by only a very thin liquid film, giving a high mass-transfer of gas to the catalyst. This liquid film is continuously regenerated by the circulating fluid and remains stable despite its very small thickness, which is in the range of 30-200 μm (Fukano and Kariyasaki, 1993). The liquid film thickness has been found by many people to be independent of the length of Taylor bubbles and liquid slugs for a given system at a given flow rate (Irandoust and Andersson, 1989). However, for very small Ca[(1)] of less than 0.01, Schwartz *et al.* (1986) found that the film thickness depended on the bubble length as well as the Ca. It was also found that the direction of flow had little effect on the measured values of the liquid film thickness.

In fact, the presence of Taylor flow gives mass transfer rates between the bulk fluid and the wall of an order of magnitude larger than for single phase liquid flow (Horvath *et al.*, 1973; Hatziantoniou and Andersson, 1982; Kapteijn *et al.*, 2001), while the increase in friction is much less. Heat transfer rates up to 2.5 times that for pure liquids have been reported (Oliver and Wright, 1964; Oliver and Young Hoon, 1968). Horvath *et al.* (1973) reported that heat transfer was seen to increase with Reynolds number, but to *rapidly* increase with a decrease in liquid slug length. This has led to the use of inert gases to improve liquid phase reactions in monoliths by inducing bubble-flow, or the addition of a second liquid phase, to enhance selectivity or extract a product (Thulasidas *et al.*, 1999).

---

[1] Capillary number, defined as $Ca = \mu U / \sigma$, where $\mu$ is the viscosity of the liquid and $\sigma$ is the liquid-air surface tension. For comparison, the capillary number in a monolith channel is typically in the region of 0.001-0.01 (based on $\mu$ = 0.001002 $Nsm^{-2}$ and $\sigma$ = 0.0728 $Nm^{-1}$ for water at 293 K).

Taylor flow is exhibited in physiological systems, such as for blood flow through small capillaries. Red blood cells, separated by plasma, are easily deformable under stress and in small blood vessels behave like bubbles. Blood plasma under these conditions is observed to undergo circulating motion forward along the capillary axis and back along the wall. This motion was found to be very effective on nutrient distribution and reaching gas equilibrium in blood plasma (Prothero and Burton, 1961). Similarly, the efficiency of other devices such as tubular reactors with absorbing walls for liquid chromatography or tubular dialysers may be improved by the introduction of Taylor flow with an inert gas phase.

### 3.7.3. Mass transfer: literature review

This section will summarise some of the literature available regarding the measurement of mass transfer performance in various systems of interest. It is hoped to validate some of these findings and reports of flow phenomena using MRI measurement techniques.

It has been observed that the gas-liquid mass transfer parameter increases with decreasing liquid slug length and increases with the velocity of the channel flow up to around 0.5 m/s (Heiszwolf *et al.*, 2001; Machado *et al.*, 2005). Therefore, they recommend operating monoliths in relatively rapid flow-through mode. Practically, this means that monoliths will need to be operated in a flow loop, with a relatively low conversion per pass. Radial flow was found to increase with Reynolds number, and rapidly increased with a decrease in slug length. Researchers agree that the optimal operation of monolith reactors will be determined from intrinsic reaction kinetics and the length and sequence of gas bubbles and liquid slugs.

In downflow, the flow due to gravity subtracts from the total flow into the film, thereby reducing the final film thickness, and *vice versa* for upflow. This effect is less pronounced at high fluid velocities since the inertial effects will dominate the viscous effects. According to Bretherton (1961) the liquid film in a vertical tube will be thicker than that in a horizontal tube in Taylor flow due to gravity effects. Thulasidas *et al.* (1995a) found that for horizontal flows and Ca $\geq$ 0.3, where viscous

forces are dominant, the diameters of bubbles are larger than found in up- or down-flow, corroborating Bretherton's results. For low Ca, the horizontal and downward flow diameters are almost identical. The thickness of the liquid film for upward flow was found to be larger than for downward flow. For tubes larger than a few centimetres in size, inertial and gravity forces play a larger role, and the effect of capillarity is very small.

The Taylor flow regime allows gas and liquid to contact the walls separately (in succession), and this avoids the free mixing of gas and liquid, which often results in the gas being dislodged by the liquid e.g. in trickle beds, the liquid phase tends to take the path of least resistance, forming a layer over certain regions of the pellets and not others. This means that regions of the bed become liquid-free, and where liquid is present gas must diffuse through a thick liquid film. The alternating gas / liquid flow and thin liquid films of the Taylor regime avoids such occurrences.

Heiszwolf *et al.* (2001) and Nijhuis *et al.* (2003) have studied flow stability in the Taylor flow regime, and found that for the flow to be stable, the liquid fraction in the channels should be in the region of 20-80%. This gives a wide operating window, although it is not as flexible as a trickle bed reactor in which the flow rates can almost be varied independently while still maintaining the trickle flow regime. They also recommend superficial flow velocities of around 10 cm/s for downflow. These high velocities are necessary otherwise gravity pulls the liquid faster than the desired flow rate, causing maldistribution.

Irandoust and Andersson (1988) found that the activity of their monolith was actually greater than that of the crushed catalyst[(2)] (which they assign to deactivation during the grinding process). They also calculate a mass transfer rate which is more than three times the observed reaction rate for their hydrogenation, putting the performance of the reactor well within the kinetically-limited regime.

[2] It is reported in the literature that commercial catalyst pellets show around 50% of the activity of crushed catalyst particles; so for a pellet to show higher activity than the crushed material is unusual.

Bercic and Pintar (1997) measured the gas-liquid mass transfer in round capillaries by physical absorption of methane in water. They were able to vary the length of liquid slugs and gas bubbles independently. Remarkably, the gas-liquid mass transfer coefficient is not a function of the bubble length or channel diameter. It is however a strong function of the liquid slug length and the fluid velocity. In the absence of reaction at the wall, the layer close to the wall remains saturated with gas, and no gas is transferred as the bubble passes by. Therefore, they are essentially only measuring the axial mass transfer from the caps of the bubble to the liquid slug. It is unlikely that this is the dominating factor under reaction conditions, where mass transfer in the film region is the critical factor. The strong relationship between liquid slug length and mass transfer means that for a system which contains a wide range of slug lengths, we cannot simply average the behaviour. Hence, it is desirable to have a well-behaved system with well-defined, optimum flow. Irandoust *et al.* (1992) also studied mass transfer of Taylor flow in circular capillaries (of 1.5 and 2.2 mm i.d.) and were able to independently vary gas bubble lengths between 3.4 and 29.1 mm, and liquid slug lengths between 2.9 and 67.0 mm. Velocities were in the range 0.092-0.56 m/s. They report that the gas-liquid mass transfer is mainly determined by the liquid-phase diffusion coefficient and the gas-liquid contact area. They report that the flow rate and liquid slug lengths had a minor influence, whereas other authors claim a stronger link. The gas contact area is affected only by the channel diameter and gas hold-up and is easy to calculate. The results from both groups are for upflow in capillaries of comparable sizes. Bercic and Pintar (1997) found that the length of liquid slugs affected the mass transfer much more than the length of gas bubbles.

An additional method of improving the mass transfer in multi-phase reactors is to utilise low-frequency vibrations. Vandu *et al.* (2004) show that the application of frequency 60 Hz and wavelength 0.5 mm significantly improved mass transfer in the monolith and bubble columns under investigation by around 70%. Additionally, they improved the gas-liquid distribution within the monolith channels. However, it is unclear how practical this technique would be to implement on a large scale in industry, and what the cost-benefit analysis of the extra equipment and energy costs would reveal.

Kreutzer *et al.* (2001) report that the variation of holdup had little impact on the mass transfer characteristics of their system, and construct a model that predicts a linear increase in mass transfer with monolith cell density.

It should be noted that the presence of impurities can affect the internal recirculation and mass transfer since they will accumulate at the gas-liquid interface and make the liquid behave as a rigid film (Irandoust *et al.*, 1992). This reduces the effect of recirculation due to Taylor flow, so for dirty fluids the mass transfer may be much less than anticipated. Care should be taken in experiments to keep the fluids pure (or consistently dirty, if this is of interest). Even aerated fluid or rough monolith walls can affect the interface properties, and change the predicted flow behaviour from that of idealised lab experiments.

### 3.7.4. Residence time distributions

The observation of residence time distributions (RTDs) can give valuable insight into the operation and behaviour of reactors. While an indirect measurement, RTDs can give an indication of the degree of mixing or plug-flow behaviour, and the spread of residence times experienced by fluid packets. Some findings from the literature are now presented.

Liquid residence time decreases as liquid or gas flows are increased, however the Peclet number remains approximately constant (Patrick *et al.*, 1995), increasing slightly with gas flow rate. The Peclet number is a measure of the degree of mixing, and can be expressed as the ratio of transport by convection to that of dispersion. Using their experimental setup of a froth reactor in upflow, with a porous glass frit creating the stable froth, the RTD obtained has a sharp peak followed by a long tail, indicating axial mixing in the liquid phase. They show that the RTD curve for upflow is approximately equal to 1.15 CSTRs in series. This has been explained by the presence of falling liquid films as the gas rises, and hence fluid passes from one slug to the next. This is more evident in square capillaries where there is always a continuous liquid phase due to surface tension drawing liquid into the corners, and so

there is significant downflow at the corners. This effect is less noticeable in co-current downflow, with near plug-flow achieved.

The monolith channels can become predominantly filled as the liquid flow rate becomes a significant fraction of the total flow rate, and in these channels reaction rate will be greatly reduced by lack of gaseous phase reactant and also by the drop in mass transfer rate associated with leaving the Taylor flow regime. Upflow operation is generally considered to be unstable due to the small viscous forces and the large density difference of bubbles and liquid slugs. Hence, channels with high gas holdup will accelerate and there will be slow flow or even downflow in channels of high liquid holdup. Patrick *et al.* (1995) observed that as more channels fill, the head will increase over the reactor until some breakthrough occurs, and a cycle is observed. These cycles should increase in frequency as the gas flow rate increases, and hence the irregularities are averaged out. Pinto *et al.* (1997) report that bubbles which are larger than a monolith channel will only enter if there is co-current liquid upflow within that channel. This leads to the creation of stagnant channels in upflow. Thulasidas *et al.* (1995b) reported that in an array of square capillaries there were a wide range of behaviours observed. Some capillaries were very active, while others showed little or no flow. The length of the bubbles and slugs were also found to vary within a wide range.

Pinto *et al.* (1997) also report that the velocities in bubble train flow through a capillary are constant, except for the moment when a new bubble enters the capillary – then there is around a 200 ms drop in velocity due to a pressure build-up as the bubble is pulled in. It is hoped that some of the faster MRI techniques may be able to study fluctuations on this timescale, and this application is discussed later in Section 11.4.2.

From the aforementioned observations, it follows that the design of monoliths cannot simply be scaled up from observations of a single capillary, as the phenomena underlying the flow distribution are more complex when multiple channels are considered. The following section covers some of the necessary considerations of moving from single- to multiple-channel systems, including criteria for stable flow.

## 3.8. Multiple channel systems

If the pressure within a channel is dominated by friction, then an increase in flow rate is accompanied by an increase in pressure drop. As a result, the dynamic response to a perturbation of the velocity will be to dampen out any fluctuations. If other contributions to the pressure drop, such as the hydrostatic contribution, become important, then an increase in flow rate may cause a decrease in pressure drop and the system becomes inherently unstable, leading to oscillatory behaviour. This behaviour is described by Grolman *et al.* (1996), Reinecke *et al.* (1999) and Kreutzer *et al.* (2005a). These authors report the following criterion for a multiple-channel system to be hydrodynamically stable, namely that the response to an increase in flow rate within a channel should be an increase in the resistance to flow within that channel, as described by Eq. 3.1.

$$\frac{\partial}{\partial u}\left(\frac{\partial P}{\partial L}\right) > 0 \qquad \textbf{(Eq. 3. 1)}$$

Kreutzer *et al.* (2005a) then used this criterion to produce a flow map of the stable operating regime for a monolith reactor. They confirm that upflow is always slightly unstable, and that downflow becomes unstable at low flow rates. Winterbottom *et al.* (2003a) demonstrated that downflow operation gave a narrow RTD (indicating minimal back-mixing) in comparison to single-phase laminar flow, and also corresponded to a higher selectivity, highlighting the importance of plug flow behaviour for reaction control. Edvinsson *et al.* (2001) showed that only in downflow could an even distribution of liquid over the monolith be achieved. At low velocities, downflow became unstable, and in upflow a good distribution could not be achieved due to channelling, in agreement with the predictions above.

In an ideal world, the many parallel channels would behave independently[3] (and there is some evidence that this is the case at high flow rates). At lower flow rates,

[3] As a side note, it is generally assumed that the channels of the monolith have no radial flow between channels. However, depending on the porosity of the substrate, interaction between adjacent channels may be possible. Reactors which have significantly permeable walls are often termed "membrane reactors". The permeable walls allow the transfer of certain chemicals out of the reactor – this can aid chemical conversion in the case of limiting equilibria, though the cost of membrane reactors is high and there are significant engineering and materials selection problems, so it is likely they will only find a place in specialised chemical production, such as pharmaceuticals. Also the stability and resistance to fouling of membranes is important to ensure a long lifetime. They are not discussed further here.

instabilities and maldistribution occur due to subtle differences between channels, small viscous forces, and the large difference in density between the bubbles and slugs. Literature is available for lab systems which mimic two-phase split-flow in parallel channels, primarily associated with heat-exchange boilers found in the nuclear industry. For example, Iloeje *et al.* (1982) use three parallel channels of internal diameter 23.6 mm. They observed both co-current flow directions, as well as counter-current operation, and claim that the mode of operation is history dependent. Of course, the applicability of such research to the far smaller channels of the monolith is questionable, but it is useful nonetheless to have some background information on the behaviour of connected parallel flows. A more complex system was investigated by Stark and Manga (2000), analysing the motion of long bubbles in a network of tubes. They used this as a model of porous media, but their approach may also be applicable to monoliths. Qu and Mudawar (2003) studied the pressure fluctuations for two-phase flow in 21 parallel microchannels. They observed hydrodynamic instabilities, which were a result of fluid holdup variations and feedback interaction between channels.

The choice of upflow or downflow (or even counter-current flow) is another significant design parameter, and the following two sections will discuss some of the advantages and disadvantages associated with each. The stability of flow in each of these orientations is discussed, drawing on conclusions from the literature.

### 3.8.1. Counter-current flow

Most reports in the literature have investigated co-current downflow, though counter-current operation is possible in the wavy film regime. Many commercially important reactions benefit from counter-current operation, where we have equilibrium-limited reactions or strong product-inhibition e.g. desulphurisation, esterification, hydrocracking, *etc*. Counter-current operation overcomes the equilibrium composition, and leads to nearly complete conversion of reactants in these cases. Successful counter-current monolith operation has been aided by the use of internally finned channels (Lebens *et al.*, 1998; Ellenberger and Krishna, 1999; Lebens *et al.*, 1999c; Beers *et al.*, 2001b; Heibel *et al.*, 2001; Nijhuis *et al.*, 2002). These help

stabilise the wavy liquid film and prevent flooding, while keeping a large surface of catalyst to overcome the problems of the low mass transfer rate in the laminar liquid film. Much higher liquid velocities are achievable (up to 0.06 m/s in a 25 cpsi monolith) than for unfinned tubes, and this is important for industrial-scale throughputs. Counter-current flow is strongly limited by the behaviour of the liquid exiting at the bottom end of the channels, and so the design of dedicated inlet and outlet systems for the reactors becomes important. Outlet construction to guide the liquid away from the exit will allow good drainage, and help prevent flooding. A bevelled edge, or an angular cut of monolith aids drainage of fluid, and helps prevent bridging of the outlet, leading to slug formation and flooding (Heibel *et al.*, 2002). A more complex approach involves the insertion of a capillary into the outlet such that the gas is ensured to enter without creating liquid slugs. This approach is very effective, though awkward to align accurately for all channels (Cybulski and Moulijn, 1998). It has been found that monolith reactors can be designed to operate at flooding limits within the range of other commercial structured packings, and allow operation under commercially viable conditions.

Even with perfect drainage, flooding is eventually reached by the formation of growing, unstable waves in the liquid film which bridge the channel. This effect is much stronger for liquids with high surface tension like water. Wallis' flooding correlation (Wallis, 1961) is found to work well for both finned and unfinned tubes. In such wavy-film counter-current reactors, it is noted that the values for mass transfer are considerably higher in the upper than in the lower part, attributed to the developing concentration and velocity profiles, and the initiation of waves in the falling film (Lebens *et al.*, 1999c). There is a smooth transition to the flooding point, with a more gradual increase in pressure drop. Lebens *et al.* (1998) attributed this to the behaviour in the onset of flooding in the bottom section of the monolith, which breaks down over the length of the monolith and does not propagate to the top.

### 3.8.2. Upflow versus downflow

The main reason for persisting with the upflow configuration despite its apparently inherent instability issues, is that it gives greater control over the superficial velocity

of the fluid and hence control over residence time. The design of the foam reactor (Section 3.6.2) helps stable performance, though the RTD obtained is still that of a CSTR rather than plug flow.

Achieving uniform liquid and gas distribution is easier in downflow than upflow, and has the advantage that the gravity-driven flow results in essentially constant pressure, allowing spontaneous internal gas recirculation, provided that there is an open connection between the top and bottom of the reactor (Edvinsson and Cybulski, 1995). Irandoust *et al.* (1989) found that downflow reactors could be scaled up very simply, and production capacity was nearly constant, indicating no scaleup effects, provided the liquid distribution was carefully controlled to ensure a good distribution. This has been reported elsewhere in the literature, showing that since there is no interaction between the channels, the behaviour of the monolith is independent of scale. Machado *et al.* (2005) also reported little productivity difference in scale-up, although they mention that their catalyst deactivated more rapidly at the larger pilot-plant scale.

Kawakami *et al.* (1989) found that the gas-liquid mass transfer in their monolith bio-reactor was 2-5 times higher and the liquid-solid mass transfer was 4-5 times higher in upflow than downflow configuration. They attribute this to a larger gas-liquid interfacial area in upflow. Kolb and Cerro (1993) suggest that for square channels in upflow, there is always some degree of downward flow in the corners, giving perfect mixing for upflow (and plug flow for downflow). This would lead to greater mass transfer within channels in upflow, in agreement with the results of Kawakami. However, Irandoust *et al.* (1989) performed experiments on larger diameter monoliths and found that the production capacity of their monolith was 30-50% *lower* in upflow than in downflow. They attribute this to the uneven distribution of gas and liquid for upflow configuration, which was worst in their largest reactor.

## 3.9. Optimisation

Since we have a plethora of design options, we require methods for selection of optimal choices. Krishna and Sie (1994) give a general overview of multiphase reactor design, covering three scales of design; catalyst design, injection and dispersion strategies, and choice of hydrodynamic flow regime.

The design of a monolith reactor has many design variables – materials, cell density and geometry, washcoat thickness, catalyst loading, flow direction, fluid distribution method, gas and liquid flow rates, *etc*. Many of the parameters are inter-related and can have a significant effect on performance. Several studies have addressed the optimisation of these parameters, including Vergunst *et al.* (2001) and Roy *et al.* (2004b). The influence of flow pattern, flow distribution, pressure drop and mass transfer has been investigated in the upflow configuration (Irandoust *et al.*, 1989), showing clearly that slug flow gives the highest production rate with very small scale up effects. The scaling up of reactors in downflow is very straightforward, showing the same behaviour as a single capillary, and reaction rates per unit volume remain almost constant.

With respect to pressure drop, as the liquid phase enters the top of the monolith, it accelerates / decelerates until it reaches the velocity at which the hydrostatic pressure drop balances the frictional force. If the liquid load is less than this limit, gas will be entrained as well. For example, if the liquid flows under gravity at, say, 0.50 m/s, by lowering the liquid flow rate, gas will be entrained to make the flow 0.50 m/s, thereby ensuring that slug flow is formed in all the channels. The velocity at which the forces balance is approximately given by Eq. 3.2.

$$U_l = \frac{\rho_l \, \mathrm{g} \, d_l^{\,2}}{32 \mu_l} \qquad \textbf{(Eq. 3. 2)}$$

Of course, if the flow is gravity- rather than pressure-driven, we have little control over the velocity within the monolith. One solution is to tailor the channel density, since $U_l \propto d_l^2$, and so decreasing the diameter by a factor of $\sqrt{10}$ will decrease the velocity by a factor of 10. The maximum cell density available is 1600 cpsi, which

limits our flexibility. In general, some degree of recirculation will still be necessary to achieve the required conversion.

Optima may be found empirically for cell density and catalyst coating thickness, though the case of side reactions can complicate the design process – diffusion limitations may become important, and the washcoat may need to be so thin that the reactor size becomes unfeasibly large (Vergunst *et al.*, 2001). High cell densities are generally desirable because they give high geometric surface area and volume fraction of catalyst (for constant coating thickness). Increasing the washcoat thickness will initially increase performance of a monolith reactor, but there typically exists a maximum, beyond which performance decreases. This is due to the reduction in surface area associated with coating a thicker layer on the inside of the channels, the associated pressure drop of a narrower channel, and the decrease in selectivity associated with a thicker catalyst coating.

The rate of chemical reaction is also an important factor. For example, with slow reactions, it is the amount of catalyst present in the reactor that becomes conversion-limiting. For moderate reactions, diffusion can become limiting when the catalytic layer is too thick, and for fast reactions diffusion limitation and external mass transfer problems may occur.

Hence, we have a large number of degrees of freedom to consider in the design of a monolith system. Modelling of flow behaviour in such systems will help to reduce the number of empirical measurements to be performed in the design process. The state of knowledge regarding the modelling of monolith reactors is discussed in the following section.

## 3.10.Modelling

The modelling of monolith reactors is generally considered to be more straightforward than modelling of random packed beds, due to their well-defined structure. However, the environment outside the monolith (e.g. inlet fluid distribution, outlet configuration) affects its performance to a great extent. Since the

combustion of various pollutants in exhaust gas was the earliest application for monoliths, extensive models for these processes were created and gas phase flow is well understood and characterised. Laminar flow, turbulent flow, heat and mass transfer effects, and reaction kinetics have all been addressed thoroughly. A comprehensive review of gas phase modelling in monoliths is covered in Cybulski and Moulijn (1998), and further reading includes the work of Beeckman and Hegedus (1991), Hayes and Kolaczkowski (1994), Jahn *et al.* (1997), Balakotaiah *et al.* (2000), Gupta and Balakotaiah (2001), Canu and Vecchi (2002) and Bhattacharya *et al.* (2004).

Modelling of two-phase gas / liquid flows in monoliths is more complex, and the literature tends to rely on empirical or semi-empirical correlations. Most existing correlations for monolith design come from limited numbers of experiments, and they should not be used outside the investigated operating range. Surface tension, inertia, gravitation and viscous effects are all important, and for some Newtonian fluids the Navier-Stokes equations can be used. The effects of compressibility and capillary action may be important, depending on the system of interest. Other important parameters include bubble shape, film thickness, fluid distribution, wall roughness, pressure drop, mass transfer (through the liquid film and the bubble ends), axial dispersion, heat transfer (in non-adiabatic cases) and reaction kinetics.

Much work has been done on numerical solutions for flow in capillaries, including Bretherton (1961), Suo and Griffith (1964), de Moraes *et al.* (1978), Bendiksen (1985), Ratulowski and Chang (1989), Thulasidas *et al.* (1995b), Edvinsson and Irandoust (1996), Bercic (2001), Woehl and Cerro (2001), Hazel and Heil (2002) and Kreutzer *et al.* (2005b). They use a variety of methods, from potential-flow to finite element modelling, computational fluid mechanics and others. In general, flows at low Ca ($Ca < 1$) are simpler to model and agree well with experimental data. As Ca increases, the solutions become more complex, but in general slug flow in a capillary is very well documented and understood. The development of a similarly robust model for a monolith process (whether theoretical or empirical) will allow the design to be optimised, both in terms of performance and economics. Some examples for further reading can be found in Cybulski and Moulijn (1998).

Adiabatic conditions are a reasonable assumption for all but the most endothermic / endothermic processes since the presence of a liquid phase absorbs most of the energy with very little temperature rise, and so isothermal plug flow can generally be assumed. This is in contrast to similar gas phase reactions, where the same cannot be said due to the much lower specific heat capacity of the gas.

The film thickness will depend on the shape of the channel, and also on the wall roughness and the liquid parameters. Square capillaries have been addressed by Ratulowski and Chang (1989), Kolb and Cerro (1993), Thulasidas *et al.* (1995a) and Thulasidas *et al.* (1995b), taking into account the non-axisymmetric properties of a bubble within a square channel, where the liquid film is thicker in the corners of the channel, and hence a two-dimensional model is required. The film thickness is a critical parameter, since it directly determines the rate of mass transfer from the gas to the catalyst.

A model for two-phase flow in a monolith was developed by Hatziantoniou and Andersson (1984). They took into account pressure drop, mass transfer, and a hydrogenation reaction. Cybulski *et al.* (1993) modelled two phase flow with mass transfer and reaction for methanol synthesis. The mass transfer and pressure drop characteristics have been modelled by Edvinsson and Cybulski (1995). Vergunst *et al.* (2001) have developed a model that takes into account kinetics, diffusion and mass transfer.

## 3.11. Conclusions

This chapter has discussed in detail the literature regarding monolith reactors, including their material construction, catalyst loading, geometrical parameters, mass transfer characteristics (in relation to traditional reactor designs), and fluid distribution methods. The flow regime of major interest is Taylor flow due to its superior mass transfer properties, but both downflow and upflow have attractive qualities so the decision is not so clear cut. Both types will be studied as a part of this work, as well as systems with different methods of fluid distribution, and different ranges of superficial velocities (from 1-30 cm/s). Numerous parameters have been

reported from the literature, many of which can be verified with fast MRI techniques (as will be shown in Chapters 8-12).

Chapter 4 will discuss specific tomographic techniques for studying monolith reactors, which will give valuable additional information over and above the traditional techniques. MRI will be the focus of this work, and its application to porous systems is discussed in detail.

## 3.12.Nomenclature

| **Symbols** | **Definition** |
|---|---|
| Ca | Capillary number |
| $d_h$ | Hydraulic diameter |
| g | Acceleration due to gravity |
| $L$ | Length |
| $P$ | Pressure |
| Re,$Re_P$ | Reynolds number (based on the pipe diameter), Reynolds number based on the particle diameter |
| $u,U$ | Superficial velocity |
| $\beta$ | Ratio of bubble length to slug length |
| $\mu$ | Viscosity |
| $\rho$ | Material density |
| $\sigma$ | Surface tension |

| **Abbreviations** | **Definition** |
|---|---|
| CDC | Co-current downflow contactor |
| cpsi | Cells per square inch |
| MRI | Magnetic resonance imaging |
| PIV | Particle imaging velocimetry |
| RTD | Residence time distribution |
| STP | Standard temperature and pressure |
| VOC | Volatile organic compounds |

## 3.13.References


Akita, K. and F. Yoshida (1974). "Bubble Size, Interfacial Area, and Liquid-Phase Mass-Transfer Coefficient in Bubble Columns." Industrial & Engineering Chemistry Process Design and Development **13**(1): 84-91.

Albers, R. E., M. Nystrom*, et al.* (2001). "Development of a monolith-based process for H2O2-production: from idea to large-scale implementation." Catalysis Today **69**(1-4): 247-252.

Albers, R. K. E., M. J. J. Houterman*, et al.* (1998). "Novel monolithic stirred reactor." Aiche Journal **44**(11): 2459-2464.

Bakker, J. J. W., M. T. Kreutzer*, et al.* (2005). "Hydrodynamic properties of a novel 'open wall' monolith reactor." Catalysis Today **105**(3-4): 385-390.

Balakotaiah, V., N. Gupta, *et al.* (2000). "A simplified model for analyzing catalytic reactions in short monoliths." Chemical Engineering Science **55**(22): 5367-5383.

Barajas, A. M. and R. L. Panton (1993). "The Effects of Contact-Angle on 2-Phase Flow in Capillary Tubes." International Journal of Multiphase Flow **19**(2): 337-346.

Barnea, D., Y. Luninski, *et al.* (1983). "Flow Pattern in Horizontal and Vertical 2 Phase Flow in Small Diameter Pipes." Canadian Journal of Chemical Engineering **61**(5): 617-620.

Beeckman, J. W. and L. L. Hegedus (1991). "Design of Monolith Catalysts for Power-Plant Nox Emission Control." Industrial & Engineering Chemistry Research **30**(5): 969-978.

Beers, A. E. W., T. A. Nijhuis, *et al.* (2001a). "Zeolite coated structures for the acylation of aromatics." Microporous and Mesoporous Materials **48**(1-3): 279-284.

Beers, A. E. W., R. A. Spruijt, *et al.* (2001b). "Esterification in a structured catalytic reactor with counter-current water removal." Catalysis Today **66**(2-4): 175-181.

Bendiksen, K. H. (1985). "On the Motion of Long Bubbles in Vertical Tubes." International Journal of Multiphase Flow **11**(6): 797-812.

Bercic, G. and A. Pintar (1997). "The role of gas bubbles and liquid slug lengths on mass transport in the Taylor flow through capillaries." Chemical Engineering Science **52**(21-22): 3709-3719.

Bercic, G. (2001). "Influence of operating conditions on the observed reaction rate in the single channel monolith reactor." Catalysis Today **69**(1-4): 147-152.

Bhattacharya, M., M. P. Harold, *et al.* (2004). "Mass-transfer coefficients in washcoated monoliths." Aiche Journal **50**(11): 2939-2955.

Boger, T., S. Roy, *et al.* (2003). "A monolith loop reactor as an attractive alternative to slurry reactors." Catalysis Today **79**(1-4): 441-451.

Bretherton, F. P. (1961). "The Motion of Long Bubbles in Tubes." Journal of Fluid Mechanics **14**: 81-96.

Broekhuis, R. R., R. M. Machado, *et al.* (2001). "The ejector-driven monolith loop reactor - experiments and modeling." Catalysis Today **69**(1-4): 87-93.

Canu, P. and S. Vecchi (2002). "CFD simulation of reactive flows: Catalytic combustion in a monolith." Aiche Journal **48**(12): 2921-2935.

Crynes, L. L., R. L. Cerro, *et al.* (1995). "Monolith Froth Reactor - Development of a Novel 3-Phase Catalytic-System." Aiche Journal **41**(2): 337-345.

Cybulski, A., R. Edvinsson, *et al.* (1993). "Liquid-Phase Methanol Synthesis - Modeling of a Monolithic Reactor." Chemical Engineering Science **48**(20): 3463-3478.

Cybulski, A. and J. A. Moulijn (1994). "Monoliths in Heterogeneous Catalysis." Catalysis Reviews-Science and Engineering **36**(2): 179-270.

Cybulski, A. and J. Moulijn (1998). Structured Catalysts and Reactors, Marcel Dekker.

Cybulski, A., A. Stankiewicz, *et al.* (1999). "Monolithic reactors for fine chemicals industries: A comparative analysis of a monolithic reactor and a mechanically agitated slurry reactor." Chemical Engineering Science **54**(13-14): 2351-2358.

Dautzenberg, F. M. and M. Mukherjee (2001). "Process intensification using multifunctional reactors." Chemical Engineering Science **56**(2): 251-267.

de Lathouder, K. M., J. Bakker, *et al.* (2004). "Structured reactors for enzyme immobilization: advantages of tuning the wall morphology." Chemical Engineering Science **59**(22-23): 5027-5033.

de Moraes, F. F., J. F. Davidson, *et al.* (1978). "The Motion of a Large Gas Bubble Rising through Liquid Flowing in a Tube." Journal of Fluid Mechanics **89**(3): 497-514.

Edvardsson, J. and S. Irandoust (1994). "Reactors for Hydrogenation of Edible Oils." Journal of the American Oil Chemists Society **71**(3): 235-242.

Edvinsson, R., R. E. Albers, *et al.* (2001). "Development of a Monolith-based Process for H2O2 Production: From Idea to Large-scale Implementation." Catalysis Today **69**(1-4): 247-252.

Edvinsson, R. K. and A. Cybulski (1995). "A Comparison between the Monolithic Reactor and the Trickle-Bed Reactor for Liquid-Phase Hydrogenations." Catalysis Today **24**(1-2): 173-179.

Edvinsson, R. K. and S. Irandoust (1996). "Finite-element analysis of Taylor flow." Aiche Journal **42**(7): 1815-1823.

Ellenberger, J. and R. Krishna (1999). "Counter-current operation of structured catalytically packed distillation columns: pressure drop, holdup and mixing." Chemical Engineering Science **54**(10): 1339-1345.

Flytzani-Stephanopoulos, M. and G. E. Voecks (1980). "Autothermal reforming of n-tetradecane and benzene solutions of napthalene on pellet catalysts, and steam reforming of n-hexane on pellet and monolithic catalyst beds." Final Report to DOE **DE-AI03-78ET-11326**.

Fukano, T. and A. Kariyasaki (1993). "Characteristics of Gas-Liquid 2-Phase Flow in a Capillary-Tube." Nuclear Engineering and Design **141**(1-2): 59-68.

Garcia-Bordeje, E., F. Kapteijn, *et al.* (2002). "Preparation and characterisation of carbon-coated monoliths for catalyst supports." Carbon **40**(7): 1079-1088.

Grolman, E., R. Edvinsson, *et al.* (1996). Hydrodynamic Instabilities in Gas-Liquid Monolithic Reactors. New York, ASME Press.

Gupta, N. and V. Balakotaiah (2001). "Heat and mass transfer coefficients in catalytic monoliths." Chemical Engineering Science **56**(16): 4771-4786.

Hatziantoniou, V. and B. Andersson (1982). "Solid-Liquid Mass-Transfer in Segmented Gas-Liquid Flow through a Capillary." Industrial & Engineering Chemistry Fundamentals **21**(4): 451-456.

Hatziantoniou, V. and B. Andersson (1984). "The Segmented 2-Phase Flow Monolithic Catalyst Reactor - an Alternative for Liquid-Phase Hydrogenations." Industrial & Engineering Chemistry Fundamentals **23**(1): 82-88.

Hawthorn, R. D. (1974). "Afterburner catlaysis - effects of heat and mass transfer between gas and catalyst surface." AICHE Symp. Ser. **70**(137): 428-438.

Hayes, R. E. and S. T. Kolaczkowski (1994). "Mass and Heat-Transfer Effects in Catalytic Monolith Reactors." Chemical Engineering Science **49**(21): 3587-3599.

Hazel, A. L. and M. Heil (2002). "The steady propagation of a semi-infinite bubble into a tube of elliptical or rectangular cross-section." Journal of Fluid Mechanics **470**: 91-114.

Heibel, A. K., J. J. Heiszwolf, *et al.* (2001). "Influence of channel geometry on hydrodynamics and mass transfer in the monolith film flow reactor." Catalysis Today **69**(1-4): 153-163.

Heibel, A. K., F. Kapteijn, *et al.* (2002). "Flooding performance of square channel monolith structures." Industrial & Engineering Chemistry Research **41**(26): 6759-6771.

Heiszwolf, J. J., M. T. Kreutzer, *et al.* (2001). "Gas-liquid mass transfer of aqueous Taylor flow in monoliths." Catalysis Today **69**(1-4): 51-55.

Hoek, I., T. A. Nijhuis, *et al.* (2004). "Performance of the monolithic stirrer reactor: applicability in multi-phase processes." Chemical Engineering Science **59**(22-23): 4975-4981.

Horvath, C., B. A. Solomon, *et al.* (1973). "Measurement of Radial Transport in Slug Flow Using Enzyme Tubes." Industrial & Engineering Chemistry Fundamentals **12**(4): 431-439.

Iloeje, O. C., J. A. Kervinen, *et al.* (1982). "Flow Split Relationships in 2-Phase Parallel Channel Flows." Journal of Fluids Engineering-Transactions of the Asme **104**(4): 455-462.

Irandoust, S. and B. Andersson (1988). "Mass-Transfer and Liquid-Phase Reactions in a Segmented 2-Phase Flow Monolithic Catalyst Reactor." Chemical Engineering Science **43**(8): 1983-1988.

Irandoust, S. and B. Andersson (1989). "Liquid-Film in Taylor Flow through a Capillary." Industrial & Engineering Chemistry Research **28**(11): 1684-1688.

Irandoust, S., B. Andersson, *et al.* (1989). "Scaling up of a Monolithic Catalyst Reactor with 2-Phase Flow." Industrial & Engineering Chemistry Research **28**(10): 1489-1493.

Irandoust, S., S. Ertle, *et al.* (1992). "Gas-Liquid Mass-Transfer in Taylor Flow through a Capillary." Canadian Journal of Chemical Engineering **70**(1): 115-119.

Jahn, R., D. Snita, *et al.* (1997). "3-D modeling of monolith reactors." Catalysis Today **38**(1): 39-46.

Kapteijn, F., T. A. Nijhuis, *et al.* (2001). "New non-traditional multiphase catalytic reactors based on monolithic structures." Catalysis Today **66**(2-4): 133-144.

Kapteijn, F., R. M. de Deugd, *et al.* (2005). "Fischer-Tropsch synthesis using monolithic catalysts." Catalysis Today **105**(3-4): 350-356.

Kawakami, K., K. Kawasaki, *et al.* (1989). "Performance of a Honeycomb Monolith Bioreactor in a Gas-Liquid Solid 3-Phase System." Industrial & Engineering Chemistry Research **28**(4): 394-400.

Kolb, W. B. and R. L. Cerro (1991). "Coating the inside of a Capillary of Square Cross-Section." Chemical Engineering Science **46**(9): 2181-2195.

Kolb, W. B. and R. L. Cerro (1993). "Film Flow in the Space between a Circular Bubble and a Square Tube." Journal of Colloid and Interface Science **159**(2): 302-311.

Kreutzer, M. T., P. Du, *et al.* (2001). "Mass transfer characteristics of three-phase monolith reactors." Chemical Engineering Science **56**(21-22): 6015-6023.

Kreutzer, M. T., J. J. W. Bakker, *et al.* (2005a). "Scaling-up multiphase monolith reactors: Linking residence time distribution and feed maldistribution." Industrial & Engineering Chemistry Research **44**(14): 4898-4913.

Kreutzer, M. T., F. Kapteijn, *et al.* (2005b). "Multiphase monolith reactors: Chemical reaction engineering of segmented flow in microchannels." Chemical Engineering Science **60**(22): 5895-5916.

Kreutzer, M. T., M. G. van der Eijnden, *et al.* (2005c). "The pressure drop experiment to determine slug lengths in multiphase monoliths." Catalysis Today **105**(3-4): 667-672.

Krishna, R. and S. T. Sie (1994). "Strategies for Multiphase Reactor Selection." Chemical Engineering Science **49**(24A): 4029-4065.

Kulkarni, R., R. Natividad, *et al.* (2005). "A comparative study of residence time distribution and selectivity in a monolith CDC reactor and a trickle bed reactor." Catalysis Today **105**(3-4): 455-463.

Lebens, P. J. M., R. K. Edvinsson, *et al.* (1998). "Effect of entrance and exit geometry on pressure drop and flooding limits in a single channel of an internally finned monolith." Industrial & Engineering Chemistry Research **37**(9): 3722-3730.

Lebens, P. J. M., J. J. Heiszwolf, *et al.* (1999a). "Gas-liquid mass transfer in an internally finned monolith operated countercurrently in the film flow regime." Chemical Engineering Science **54**(21): 5119-5125.

Lebens, P. J. M., F. Kapteijn, *et al.* (1999b). "Potentials of internally finned monoliths as a packing for multifunctional reactors." Chemical Engineering Science **54**(10): 1359-1365.

Lebens, P. J. M., M. M. Stork, *et al.* (1999c). "Hydrodynamics and mass transfer issues in a countercurrent gas-liquid internally finned monolith reactor." Chemical Engineering Science **54**(13-14): 2381-2389.

Liu, W., W. P. Addiego, *et al.* (2002). "Monolith reactor for the dehydrogenation of ethylbenzene to styrene." Industrial & Engineering Chemistry Research **41**(13): 3131-3138.

Liu, W. and S. Roy (2004). "Effect of channel shape on gas/liquid catalytic reaction performance in structured catalyst/reactor." Chemical Engineering Science **59**(22-23): 4927-4939.

Machado, R. M., R. R. Broekhuis, *et al.* (2005). "Applying monolith reactors for hydrogenations in the production of specialty chemicals - process and economic considerations." Catalysis Today **105**(3-4): 305-317.

Mishima, K. and T. Hibiki (1996). "Some characteristics of air-water two-phase flow in small diameter vertical tubes." International Journal of Multiphase Flow **22**(4): 703-712.

Nicolau, I., P. M. Colling, *et al.* (1998). US Patent 5 705 679 and 5 854 171.

Nijhuis, T. A., M. T. Kreutzer, *et al.* (2001). "Monolithic catalysts as efficient three-phase reactors." Chemical Engineering Science **56**(3): 823-829.

Nijhuis, T. A., A. E. W. Beers, *et al.* (2002). "Water removal by reactive stripping for a solid-acid catalyzed esterification in a monolithic reactor." Chemical Engineering Science **57**(9): 1627-1632.

Nijhuis, T. A., F. M. Dautzenberg, *et al.* (2003). "Modeling of monolithic and trickle-bed reactors for the hydrogenation of styrene." Chemical Engineering Science **58**(7): 1113-1124.

Oliver, D. R. and S. J. Wright (1964). "Pressure Drop and Heat Transfer in Gas-Liquid Slug FLow in Horizontal Pipes." British Chemical Engineering **9**: 590-596.

Oliver, D. R. and A. Young Hoon (1968). "Two-Phase Non-Newtonian Flow - Heat Transfer." Trans. Instn. Chem. Engrs. **46**: 116-122.

Patrick, R. H., T. Klindera, *et al.* (1995). "Residence Time Distribution in 3-Phase Monolith Reactor." Aiche Journal **41**(3): 649-657.

Perez-Cadenas, A. F., M. M. P. Zieverink, *et al.* (2005). "High performance monolithic catalysts for hydrogenation reactions." Catalysis Today **105**(3-4): 623-628.

Pfefferle, L. D. and W. C. Pfefferle (1987). "Catalysis in Combustion." Catalysis Reviews-Science and Engineering **29**(2-3): 219-267.

Pinto, M. D. T., M. A. Abraham, *et al.* (1997). "How do bubbles enter a capillary?" Chemical Engineering Science **52**(11): 1685-&.

Prothero, J. and A. C. Burton (1961). "The Physics of Blood Flow in Capillaries - The Nature of the Motion." Biophysical Journal **1**: 565-575.

Qu, W. L. and I. Mudawar (2003). "Measurement and prediction of pressure drop in two-phase micro-channel heat sinks." International Journal of Heat and Mass Transfer **46**(15): 2737-2753.

Ratulowski, J. and H. C. Chang (1989). "Transport of Gas-Bubbles in Capillaries." Physics of Fluids a-Fluid Dynamics **1**(10): 1642-1655.

Reinecke, N. and D. Mewes (1999). "Oscillatory Transient Two-phase Flows in Single Channels with Reference to Monolithic Catalyst Supports." International Journal of Multiphase Flow **25**(6-7): 1373-1393.

Roy, S., T. Bauer*, et al.* (2004a). "Monoliths as multiphase reactors: A review." Aiche Journal **50**(11): 2918-2938.

Roy, S., A. K. Heibel*, et al.* (2004b). "Design of monolithic catalysts for multiphase reactions." Chemical Engineering Science **59**(5): 957-966.

Satterfield, C. N. and F. Ozel (1977). "Some Characteristics of 2-Phase Flow in Monolithic Catalyst Structures." Industrial & Engineering Chemistry Fundamentals **16**(1): 61-67.

Schanke, D., E. Bergene*, et al.* (2001). US Patent 6 211 255.

Schildhauer, T. J., F. Kapteijn*, et al.* (2005). "Reactive stripping in pilot scale monolith reactors - application to esterification." Chemical Engineering and Processing **44**(6): 695-699.

Schwartz, L. W., H. M. Princen*, et al.* (1986). "On the Motion of Bubbles in Capillary Tubes." Journal of Fluid Mechanics **172**: 259-275.

Soni, D. S. and B. L. Crynes (1980). "Comparison of the Hydrodesulfurization and Hydrodenitrogenation Activities of Monolith Alumina Impregnated with Co and Mo and Nalcomo 474 Catalysts." Abstracts of Papers of the American Chemical Society **179**(MAR): 25-FUEL.

Stankiewicz, A. and J. Moulijn (2000). "Process Intensification: Transforming Chemical Engineering." CHemical Engineering Progress **96**(1): 22-34.

Stankiewicz, A. (2001). "Process intensification in in-line monolithic reactor." Chemical Engineering Science **56**(2): 359-364.

Stark, J. and M. Manga (2000). "The motion of long bubbles in a network of tubes." Transport in Porous Media **40**(2): 201-218.

Suo, M. and P. Griffith (1964). "Two-Phase Flow in Capillary Tubes." Trans. ASME, Journal of Basic Engineering: 576-582.

Taylor, G. I. (1961). "Deposotion of a Viscous Fluid on the Wall of a Tube." Journal of Fluid Mechanics **10**: 161-165.

Thulasidas, T. C., M. A. Abraham*, et al.* (1995a). "Bubble-Train Flow in Capillaries of Circular and Square Cross-Section." Chemical Engineering Science **50**(2): 183-199.

Thulasidas, T. C., R. L. Cerro*, et al.* (1995b). "The Monolith Froth Reactor - Residence Time Modeling and Analysis." Chemical Engineering Research & Design **73**(A3): 314-319.

Thulasidas, T. C., M. A. Abraham*, et al.* (1997). "Flow patterns in liquid slugs during bubble-train flow inside capillaries." Chemical Engineering Science **52**(17): 2947-2962.

Thulasidas, T. C., M. A. Abraham*, et al.* (1999). "Dispersion during bubble-train flow in capillaries." Chemical Engineering Science **54**(1): 61-76.

Triplett, K. A., S. M. Ghiaasiaan*, et al.* (1999). "Gas-liquid two-phase flow in microchannels - Part I: two-phase flow patterns." International Journal of Multiphase Flow **25**(3): 377-394.

Tucci, E. R. and W. J. Thomson (1979). "Monolith Catalyst Favored for Methanation." Hydrocarbon Processing **58**(2): 123-126.

Vaarkamp, M., W. Dijkstra*, et al.* (2001). "Hurdles and solutions for reactions between gas and liquid in a monolithic reactor." Catalysis Today **69**(1-4): 131-135.

Vandu, C. O., J. Ellenberger*, et al.* (2004). "Hydrodynamics and mass transfer in an upflow monolith loop reactor: influence of vibration excitement." Chemical Engineering Science **59**(22-23): 4999-5008.

Vergunst, T., F. Kapteijn*, et al.* (2001). "Optimization of geometric properties of a monolithic catalyst for the selective hydrogenation of phenylacetylene." Industrial & Engineering Chemistry Research **40**(13): 2801-2809.

Wallis, G. B. (1961). "Flooding Velocities for Air and Water in Vertical Tubes." AEEW - R123 (AEE Winfrith Report).

Warna, J., I. Turunen*, et al.* (1994). "Kinetics of nitrate reduction in monolith reactor." Chemical Engineering Science **49**(24B): 5763-5773.

Williams, J. L. (2001). "Monolith structures, materials, properties and uses." Catalysis Today **69**(1-4): 3-9.

Winterbottom, J. M., Z. Khan*, et al.* (2000). "A comparison of triglyceride oil hydrogenation in a downflow bubble column using slurry and fixed bed catalysts." Journal of Chemical Technology and Biotechnology **75**(11): 1015-1025.

Winterbottom, J. M., H. Marwan*, et al.* (2003a). "Selectivity, Hydrodynamics and Solvent Effects in a Monolithic Cocurrent Downflow (CDC) Reactor." Canadian Journal of Chemical Engineering **81**: 838-845.

Winterbottom, J. M., H. Marwan*, et al.* (2003b). "The palladium catalysed hydrogenation of 2-butyne-1,4-diol in a monolith bubble column reactor." Catalysis Today **79**(1-4): 391-399.

Woehl, P. and R. L. Cerro (2001). "Pressure drop in monolith reactors." Catalysis Today **69**(1-4): 171-174.

Yawalkar, A. A., R. Sood*, et al.* (2005). "Axial mixing in monolith reactors: Effect of channel size." Industrial & Engineering Chemistry Research **44**(7): 2046-2057.

Zhang, J. and M. Giot (1995). Phenomenological Modeling of Flow Regime Map in Vertical Gas-Liquid Countercurrent Flows. Pisa.

Zhao, T. S. and Q. C. Bi (2001). "Co-current air-water two-phase flow patterns in vertical triangular microchannels." International Journal of Multiphase Flow **27**(5): 765-782.

# Chapter 4 – Tomographic Methods

Tomographic measurement techniques allow the visualisation of opaque 3D systems, such as monolith reactors. This chapter begins with a brief literature review of current tomographic techniques, and is followed by a section addressing some of the issues related to applying nuclear magnetic resonance (NMR) and magnetic resonance imaging (MRI) both to porous media in general, and specifically to the monolith reactor itself.

**CHAPTER 4 – TOMOGRAPHIC METHODS**



## 4.1. Tomographic methods

So far, the previous chapter has only discussed traditional measurement techniques such as pressure drop measurements, residence time distributions (RTDs) and reactor conversion. However, these are all indirect measurements of the reactor system and yield limited information on the processes occurring within the reactor itself, and so reactor design is limited by poor understanding of the physical processes. Within the

last 25 years tomographic methods have been sufficiently developed to the stage that they can provide a wealth of information about fluid distribution, velocities, frequencies of oscillations, phase fractions, chemical conversion / selectivity, and flow / conversion heterogeneities. This allows their effects on reactor performance to be much better understood, and in a truly non-intrusive manner. This chapter will briefly cover the main tomographic methods available, and highlight some of their advantages and limitations regarding the probing of monolith reactors. Most tomographic techniques allow the measurement of gas / liquid properties of the reactor such as phase fractions and spatial distributions. The different techniques, such as x-ray, electrical conductance and MRI tomography have various advantages and disadvantages, most notably in terms of sensitivity to different material properties, as well as the temporal and spatial resolutions possible. These are summarised in Table 4.1. Each of these methods will now be addressed in turn.

Table 4.1 – Summary of typical spatial and temporal resolution of tomographic techniques.

| Technique | Spatial resolution | Temporal resolution |
|---|---|---|
| Catchpots | ~ 10% of diameter | seconds |
| Conductivity probes | 1D, mm | seconds |
| Neutron / x-ray / gamma ray | 100 μm | minutes |
| Optical | mm | seconds |
| Particle imaging velocimetry | mm | seconds |
| Capacitance tomography | 5 - 10% of diameter | milliseconds |
| NMR / MRI | mm | minutes |

### 4.1.1. Catchpots

Perhaps not a tomographic method, catchpots still deserve a mention because of their long history in traditional chemical engineering. By arranging a series of collection devices at the exit of a reactor, a map of liquid holdup and reactor conversion may be compiled as illustrated in Wang *et al.* (1998) and Boyer *et al.* (2002). Dynamically-controlled catchpots can also measure the *rate* at which liquid flows into each catchpot, and therefore give a crude velocity distribution map. Typically the spatial

resolution obtained is very limited, but it can give an indication of gross heterogeneity in reactor performance, or the reproducibility of fluid distribution.

However, care must be taken in the use of catchpots – as will be shown in Section 4.1.1.1 – the presence of a catchpot below the monolith reactor affects the hydrodynamics. This is because the pressure drop across the monolith reactor is very low, and the presence of a collection device alters the pressure drop sufficiently to affect performance. Therefore catchpots are not a non-invasive measuring technique.

#### 4.1.1.1. A catchpot experiment

This section will cover a brief diversion highlighting an experiment performed as part of this work. A catchpot device was constructed for use with the monolith reactor, and is shown in Figure 4.1a. A brass collection device consisting of 16 tubes with counter-sunk entrances to aid fluid entry, with each tube connecting to a plastic pipe. Figure 4.1b shows the catchpot device in action below the monolith reactor, with 16 collection devices attached by equal lengths of plastic piping.

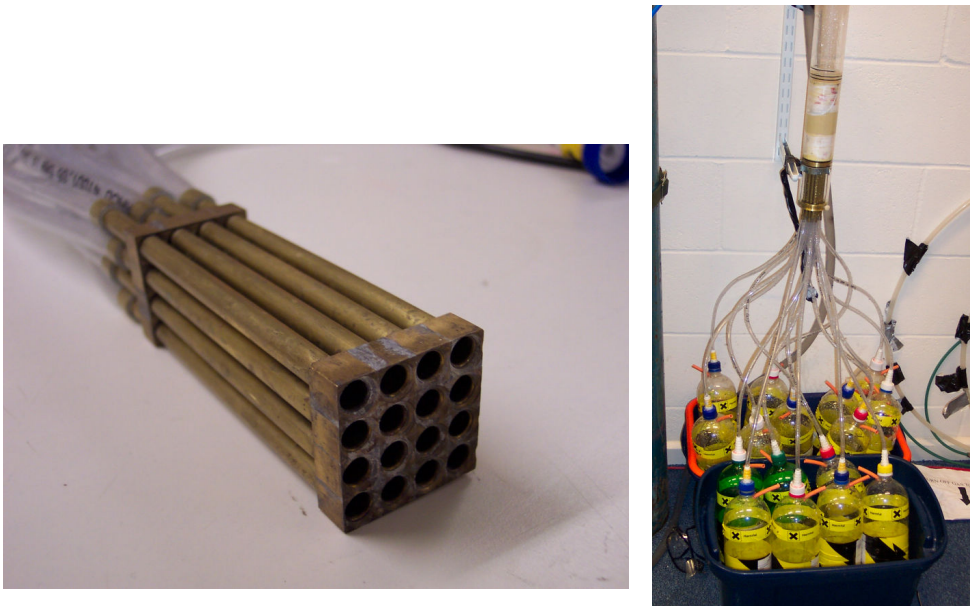

(a) (b)

Figure 4.1 – The catchpot device; (a) the fluid separation system, which is a series of brass conduits, each of which connects with plastic piping; and (b) the catchpot setup with the device located below the monolith reactor.

Figure 4.2 shows four experiments using the catchpot device on a 300 cpsi monolith, with the intention of probing the repeatability of the technique. The flow was provided by a showerhead distributor at a rate of 150 l/hr (setup is described Section 8.2.2), and the flow was stopped and re-started between experiments. It can be seen that the catchpot method is relatively repeatable, with only minor variations.

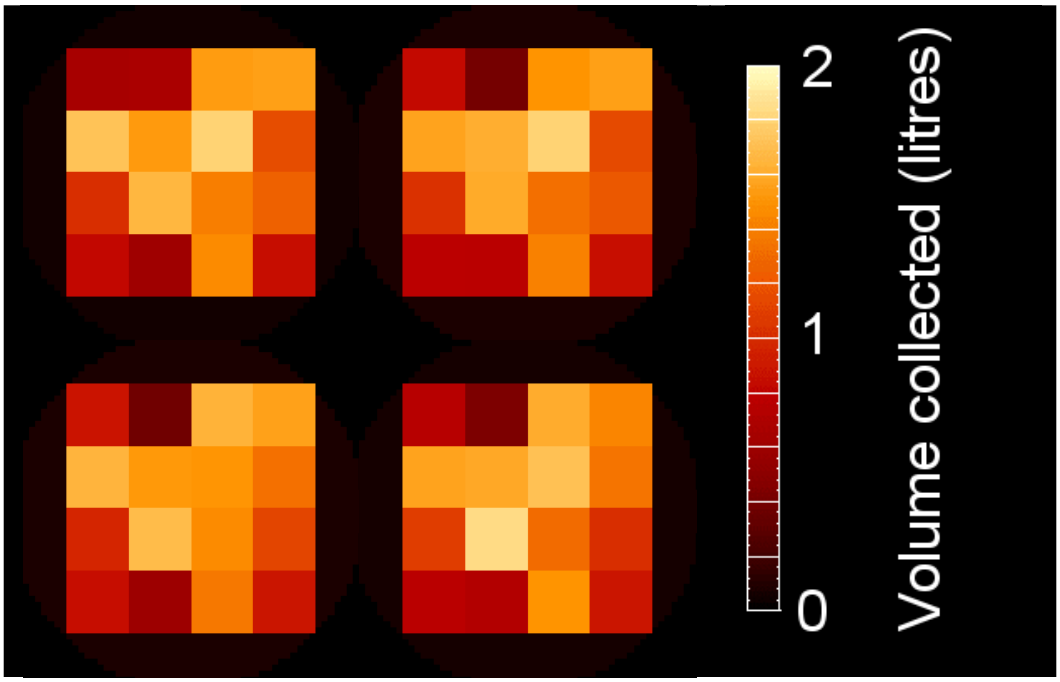


Figure 4.2 – Catchpot flow-distribution measurements. The four images show four separate measurements of distribution in 16 regions for 150 l/hr water by a showerhead distributor. The experiments were performed over 9 minutes, and the volume of water collected during this time is indicated by the relative intensities.

Even though the repeatability of the measurement has been established, the presence of the catchpot device is still an invasive measurement as is shown in the MRI images of Figure 4.3. These are acquired with the RARE pulse sequence (Section 2.3.1) and they show fluid as high intensity. It can be seen that the fluid distribution is far less homogeneous in the case with the catchpots present. This shows that the pressure drop imposed by the catchpot device is significant in relation to the low pressure drop across the monolith reactor. Hence, it has been shown that while the catchpot might still be an appropriate device for other reactors, in regards to the monolith reactor it is a consistent but flawed measurement technique.

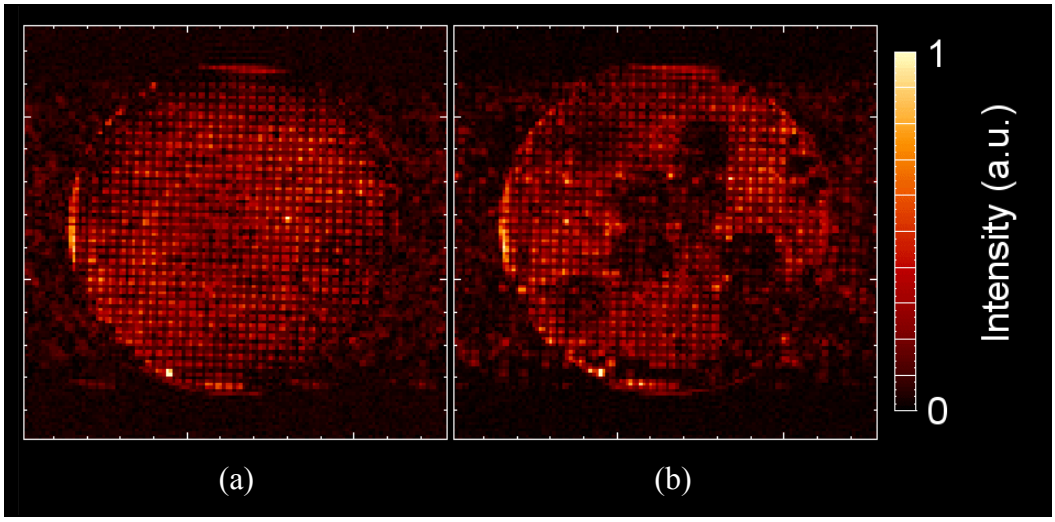


Figure 4.3 – Images of flow through the monolith at 110 l/hr water, using a showerhead distributor; (a) with no catchpot device present; and (b) with the catchpot present below the monolith section.

### 4.1.2. Conductivity probes

The next tomographic technique for consideration is the conductivity probe. The method of pushing probes into a reactor system to measure local conductivity is an invasive technique, but it is well established in the literature and if carefully implemented should not affect the flow significantly. If the liquid in the reactor is electrically conductive we may quantify the holdup between two probes, or we may detect the passing of liquid slugs and bubbles by the difference in local conductivity. By measuring the gas and liquid slug lengths in this manner, we may estimate the mass transfer properties of the system. (Prasser *et al.*, 1998; Schmitz and Mewes, 2000; Heiszwolf *et al.*, 2001; Boyer *et al.*, 2002).

Similar techniques have been used to obtain temperature measurements in monolith stacks (Albers, 2005). They show that it is possible to estimate reactor conversion *in-situ* and even detect differences in activity between catalyst sections.

### 4.1.3. Neutron / x-ray / gamma ray absorption

These radiation-attenuation techniques give 2D images, which record the density of material along the radiation path. By rotating either the sample or the radiation source

through 360°, we may obtain a series of images which can then be processed via various back-projection algorithms to build up a complete 3D representation of the sample. The resolution of these techniques is only limited by their wavelength and is often below 100 μm, depending on the sample size. The sample rotation is slow however, on the order of minutes, and is not suitable for dynamic systems. There are also issues with radiation safety and cost. Some examples of studies of packed beds using x-ray tomography are presented in Nakumara *et al.* (1991), Kantzas (1994) and Toye *et al.* (1995), showing the measurement of holdup and fluid distribution.

Roy and Al-Dahhan (2005) used gamma-ray tomography to study the time-averaged fluid distribution in a two-phase monolith reactor. The holdup measured with this technique was accurate to within 3.5% of measurements obtained by stopping the flow and measuring the quantity of fluid within the reactor volume. Another example of work with two-phase flows is reported in Reinecke *et al.* (1998).

### 4.1.4. Optical methods

These techniques require optically transparent materials, and are best suited to observations of single capillaries – in bundles of capillaries, it is difficult to resolve the individual flow patterns. However, studies of single capillaries can give information on bubble shape, slug lengths, *etc*., and this data can then feed into modelling more complex systems such as monoliths.

However, it has been shown by Haakana *et al.* (2004) that it is possible to estimate the performance of individual capillaries within a bundle. The authors sliced open a monolith and replaced a row of the cordierite channels with a row of glass capillaries, which extended past the end of the monolith section. Thereby it was possible to visually observe the flow patterns within a row of channels. However, this is definitely an invasive method and it is unclear whether the hydrodynamics would really be unaffected such an intrusion.

### 4.1.5. Particle imaging velocimetry

Particle Imaging Velocimetry (PIV) is a flow visualisation technique capable of acquiring 2D planar images and is a well-tested method for visualising streamlines in single phase flow. Seed particles are introduced into the fluid and a plane is illuminated with laser light. The motion of the particles is then tracked by cross-correlation processing (Section 7.5.2.6). Velocity measurements may be calculated by comparing subsequent frames from a video. This technique gives time averaged data, and so cannot identify small velocity fluctuations. For two-phase flow, PIV requires the refractive index matching of the phases since it is an optical technique. A review is given in Santiago *et al.* (1998).

### 4.1.6. Capacitance tomography

This technique consists of a number of electrodes placed around the sample, and the conductivity between pairs of electrodes is measured rapidly in sequence. This non-invasive technique typically gives a very high temporal resolution but low spatial resolution. The interpretation of the images is difficult since the fluid distribution within the sample will affect the path of the conductivity measurements, and the attenuation is non-linear. Incorrect reconstruction will lead to distorted images and there is an issue developing reliable reconstruction algorithms. It is however very cheap and can give a rapid qualitative measurement, perhaps quantitative if the system is well understood *a priori*. Applications include probing fluid distribution in packed beds, monolith reactors, fluidised beds and pipe flow. In industry, capacitance tomography is often used for real-time process control due to its relatively low cost and rapid imaging. Examples of applications of capacitance tomography include work by Dickin *et al.* (1992), Wang *et al.* (1995), Reinecke *et al.* (1996), Reinecke and Mewes (1996), Reinecke and Mewes (1997), Reinecke *et al.* (1998) and Mewes *et al.* (1999).

### 4.1.7. NMR / MRI

NMR and MRI have been widely developed in the field of medical imaging, but these techniques are becoming of use in investigating chemical engineering applications. They have been applied to packed bed reactors and more recently monolith reactors.

MRI has a good trade-off between the advantages of x-ray and conductance tomography, with relatively fast imaging times but reasonable resolution. The development of new pulse sequences and imaging techniques has pushed the temporal and spatial resolution far beyond what was possible twenty years ago, and it is now arguably the most useful of the tomographic techniques in terms of the information that can be obtained. The resolution of MRI is sufficient to identify individual monolith channels, even at high cell densities, while new imaging techniques allow the investigation of dynamic systems such as two-phase flow without significant time-averaging. Moreover, since it is a spectroscopic technique it is also possible to obtain *in-situ*, chemically-specific information which is not possible by other methods.

In summary, MRI has several advantages over other tomographic / visualisation techniques, such as non-invasiveness, lack of directional preference, the ability to quantitatively measure various parameters and its ability to work in opaque systems. The main downside is the cost and size of equipment. It seems unlikely that MRI will ever become a standard process control technique, although the recent emergence of hand-held devices holds promise. Portable NMR probes are already used in oil well drilling operations to probe the properties of the surrounding rock.

NMR and MRI are well established in the study of porous media. Examples include hydrodynamic studies, velocity measurements, diffusion measurements and pore sizing in trickle bed reactors, packed beds, porous rock and more. General literature on MRI of chemical engineering reactors and porous media include Callaghan (1991), Gladden (1993), Gladden and Alexander (1996), Watson and Chang (1997), de Swiet *et al.* (1998), Chang and Watson (1999), Gladden (1999), Barrie (2000), Krapitchev *et al.* (2002), Mantle and Sederman (2003), Gladden *et al.* (2004), and Koptyug *et al.* (2005). Another excellent source of literature are the proceedings of the conference series “International Conference on Magnetic Resonance in Porous Media”, which are published in the journal Magnetic Resonance Imaging. A brief summary will be included here to illustrate the scope and applicability of MRI to a range of porous media.

### 4.1.8. Application of NMR to porous systems

Packed bed reactors have been extensively studied by MRI, including velocity visualisation and structure / flow correlations (Sederman *et al.*, 1998). These observations have been validated against lattice-Boltzmann simulations (Mantle *et al.*, 2001). MRI measurement of fluid distribution and wetting efficiency have been reported by Sederman and Gladden (2001). Johns *et al.* (2000), showed that local transitions in packed beds could be identified, and correlated with a critical local Reynolds number. Anadon *et al.* (2005), later showed that these hydrodynamic fluctuations could be observed in real time within trickle beds, which grow with flow rate and cause a transition to the pulsing flow regime. Sains *et al.* (2005) have shown an example of using ultra-fast imaging techniques to probe unsteady flow in a narrow packed bed, where traditional MRI velocity techniques would fail. Yuen *et al.* (2002) and Koptyug *et al.* (2005) have showed the quantitative measurement of spatially-resolved chemical conversion within packed beds.

The three main research groups studying monolith reactors with MRI are those at the Reactor and Catalysis Engineering group at Delft University, the International Tomography Centre in Novosibirsk, and the Department of Chemical Engineering at the University of Cambridge. A summary of their published literature is outlined in this section.

Koptyug *et al.* (2000) and Koptyug *et al.* (2001) have used MRI to obtain velocity images of single-phase steady-state gas flow and liquid flow through a 25 cpsi monolith catalyst support (Figure 4.4a). Typical image resolution was 400 μm × 400 μm with acquisition times between 20 and 40 minutes. Koptyug *et al.* (2005) studied spin-density data of monolith and packed bed reactors, and reduced the acquisition times from 4-5 minutes to around 10 seconds by doping with a small amount of paramagnetic manganese to reduce the $T_1$ of the system (Figure 4.4b). They were able to visualise pellet temperature oscillations under reaction conditions; the reaction is autocatalytic and tends to generate oscillations when coupled with mass transfer limitations, a behaviour they were able to model. They also observed liquid distribution within the channels of a 400 cpsi monolith reactor, but this was operated in steady-state rather than dynamic Taylor flow.

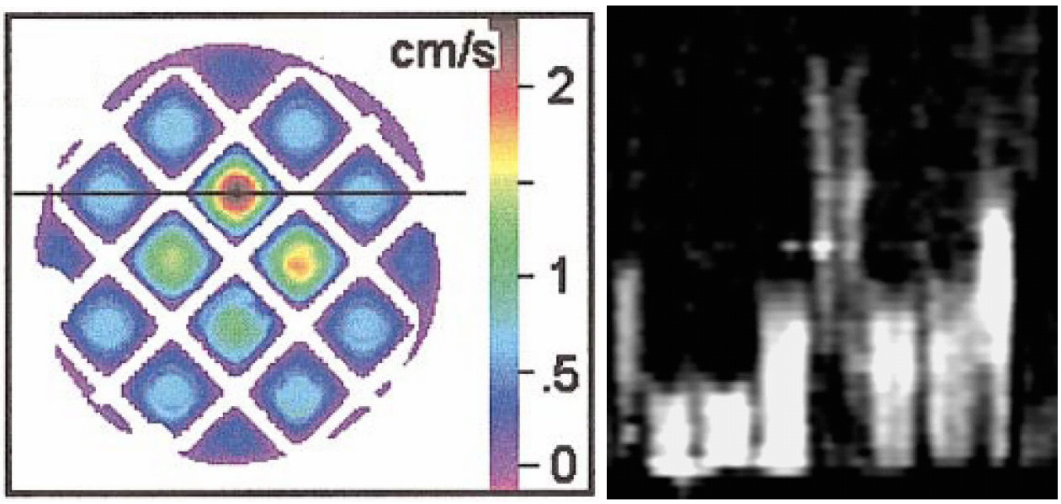


(a) (b)

Figure 4.4 – (a) Two-dimensional velocity map of water flowing through a 25 cpsi monolith at a flow rate of 76 ml/min and a maximum flow velocity of 2.4 cm/s (Koptyug *et al.*, 2000), (b) distribution of the liquid phase in a row of channels in a 400 cpsi monolith at a octane flow rate of 36 ml/s (Koptyug *et al.*, 2005).

Heibel *et al.* (2001) investigated steady-state film-flow in a 25 cpsi monolith. The resolution was 390 μm × 390 μm, with a data acquisition time of 17 minutes (Figure 4.5). They were able to visualise the shape of the film in the corners of the channels and used this data to model the flow fields within the channels which gave RTDs in good agreement with experimental results.

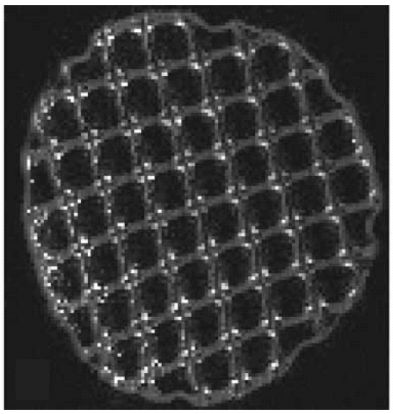

Figure 4.5 – Spin-density image of a 25 cpsi monolith in the film-flow regime, with a superficial liquid flow rate of 2.25 cm/s. Reproduced from Heibel *et al.* (2001).

Yuen (2003) investigated an esterification within a catalytically active 400 cpsi monolith reactor using MRI. The length of the reactor was 3.5 cm, and was operated

in single-phase flow at 1 ml/min, giving a residence time of approx 40 minutes and a steady-state conversion of around 5%. This low conversion is most likely due to carrying out the reaction at room temperature. The conclusions were that the flow rate and conversion in most channels were very similar, to within 5% of the mean flow (Figure 4.6) and 20% of the mean conversion.

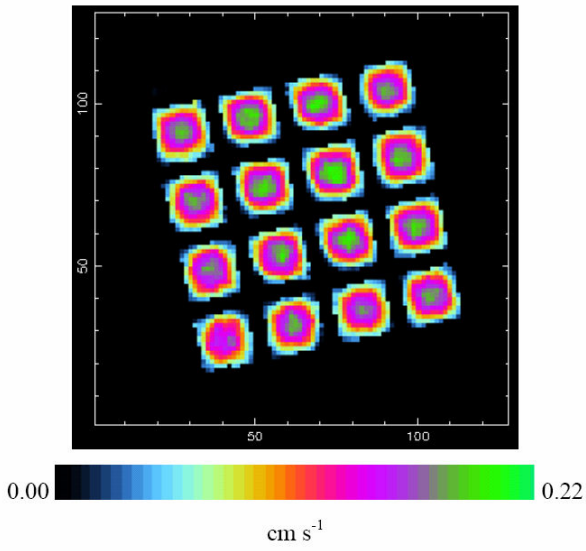


Figure 4.6 – Velocity map for single-phase liquid flow through a 400 cpsi monolith at a flow rate of 1 ml/min. Reproduced from Yuen (2003).

Recent developments in ultra-fast MRI techniques allow the capturing of transient behaviour in gas-liquid flow. For example, Mantle *et al.* (2002) and Sederman *et al.* (2003) acquired images of dynamic Taylor flow within a monoliths of density 200-400 cpsi, undergoing gas upflow through stagnant water. They used a resolution of 391 μm × 798 μm, and an acquisition time of 160 ms. The SEMI-RARE pulse sequence (Section 5.2) was sufficiently rapid to visualise two-phase Taylor flow without motion artefacts (Figure 4.7). They showed that it was possible to measure bubble-size distributions (Figure 4.8a), phase volume fractions and bubble-velocity distributions (Figure 4.8b) by measuring the displacement between successive images. By acquiring four images, each 160 ms apart, they state that the bubble velocity varied between these pairs of images, and therefore is not a constant over the 640 ms observation time. This shows the need for rapid techniques to resolve instantaneous velocities rather than time-averaged data. They also observe channels in reverse flow with respect to the superficial flow velocity, which is in agreement with the literature

on upflow configuration (Section 3.7). They state no correlation between bubble size and bubble velocity. They observe that the velocity distributions show a bimodal distribution for larger channels sizes, and a unimodal distribution for smaller channels (Figure 4.9).

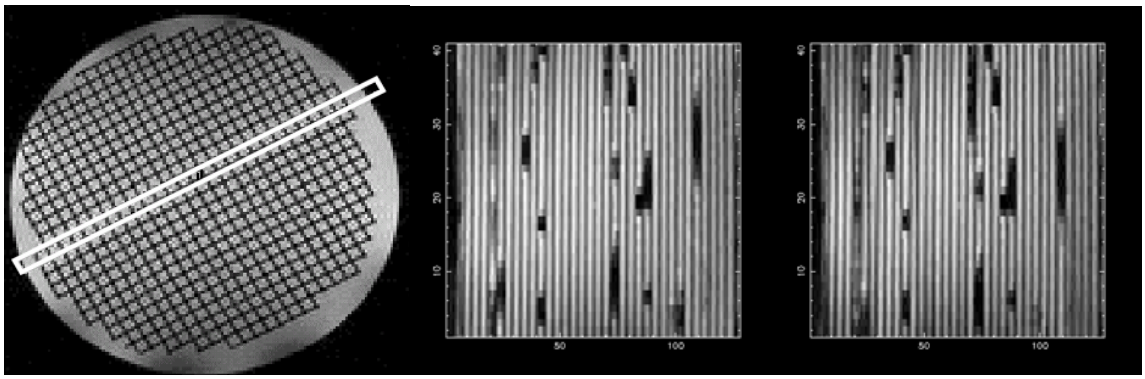

(a) (b) (c)

Figure 4.7 – An example of SEMI-RARE applied to two-phase flow within a 200 cpsi ceramic monolith; (a) shows a cross-section of the water-saturated monolith; (b) a slice through a single row of channels under two-phase flow conditions; (c) is the second image in the SEMI-RARE acquisition. Reproduced from Mantle *et al.* (2002).

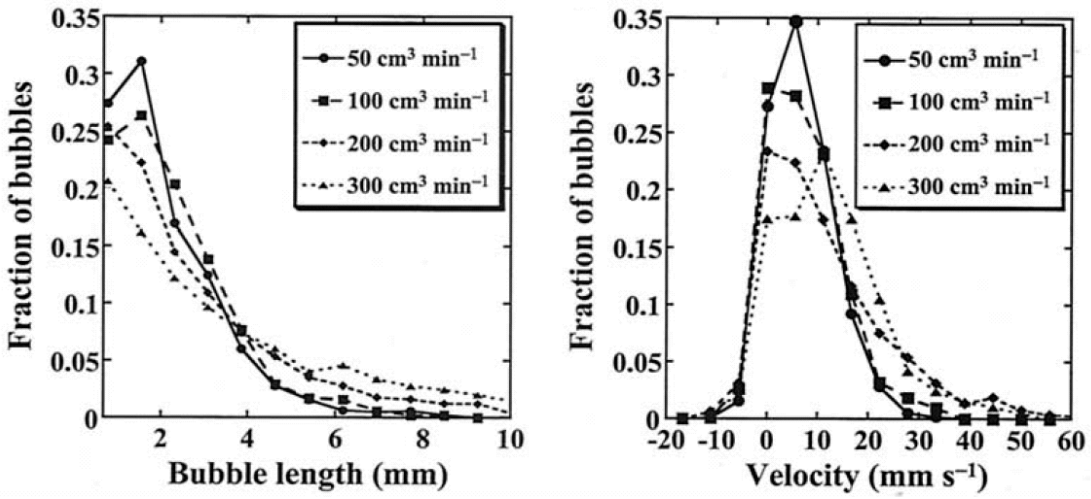


(a) (b)

Figure 4.8 – Distributions of (a) bubble length, and (b) bubble velocity, for a 400 cpsi monolith obtained from processing SEMI-RARE data. Reproduced from Sederman *et al.* (2003).

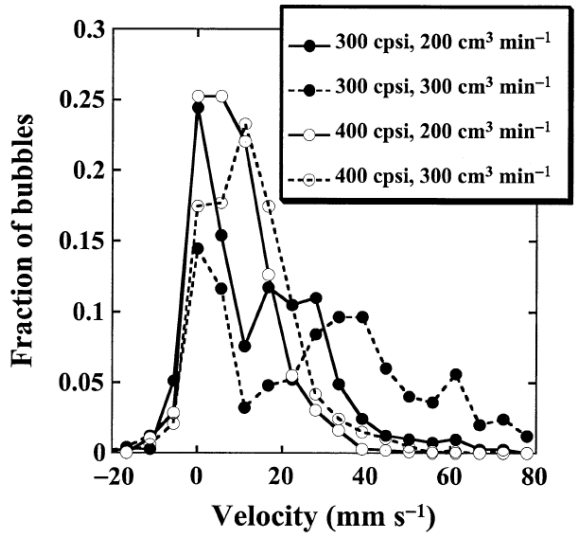


Figure 4.9 – Bubble velocity distribution for monoliths of different cell densities. Reproduced from Sederman *et al.* (2003).

In summary, it has been shown that the bubble and slug profiles within the channels may be observed non-invasively, and velocity distributions may be measured for the monolith using time-of-flight techniques. The dependence of these characteristics has been probed as a function of gas/liquid velocity and the channel size. Under upflow operation, some channels are observed to be stagnant and even reverse direction. The measurement of bubble and slug lengths, velocity distributions, holdup and correlations between these parameters will lead to better understanding and operation of these reactors.

It is clear that MRI is a very powerful technique for characterising porous media such as structured reactors, and in the next section some of the challenges in working with MRI of porous media will be outlined.

## 4.2. Challenges in studying porous media with MRI

A porous medium may be characterised as a medium containing a complex system of internal surfaces and phase boundaries. Examples include sandstones, zeolites, catalyst particles and monoliths, although monoliths contain a much higher degree of ordering and uniformity of pore length scale than many materials which are

considered “porous”. The following sections will address some of the complications arising in MRI observations of such systems.

### 4.2.1. NMR relaxometry

NMR relaxation times of fluids within porous media are shorter than in corresponding bulk fluids. This is due to the interaction between the fluid and the solid surface, which increases the effect of $T_2^*$ relaxation. This is another challenge to performing experiments since signal-to-noise will typically be far worse than for bulk fluids. Acquisition times need to be kept short, and segmented scanning is often used (Section 2.3).

### 4.2.2. Magnetic susceptibility

As was mentioned in Section 2.1.8, the presence of multiple gas/liquid/solid interfaces has the effect of increasing local magnetic susceptibilities. This is related to the broadening of the spectral line-width, and can also result in distortion of the image. This distortion is due to local magnetic field gradients altering the Lamor frequencies, and since we typically frequency-encode our images the apparent location of the spins will be incorrect. The line-broadening limits the resolution we can achieve, and distortions are observed at interfaces. This distortion is more apparent in the presence of small applied gradients ($\mathbf{B_1}$).

For example, with the EPI sequence we may have 2 kHz per point in the read direction, but only 25 Hz in the phase direction. A local inhomogeneity in the static magnetic field of more than 25 Hz will give rise to a distortion in the phase-encode direction of more than one pixel. Careful shimming of the field coils can generally remove the worst of this effect, but care is required and the results may not always be quantitative in porous media. Further information is available in Peters and Bowtell (1999). The extent of the magnetic susceptibility distortions in monolith systems is investigated in detail in the following section.

### 4.2.2.1. A magnetic susceptibility experiment

This section will cover a brief diversion highlighting an experiment performed as part of this work. To investigate how extensive the magnetic susceptibility effects within the monolith would be, a series of experiments were performed. These were performed by using a standard 2D spin-warp imaging technique (Section 2.2.7) to excite a slice through the sample under different imaging gradient strengths, and then observing the extent of any distortions occurring at air/liquid/solid interfaces. This is illustrated in Figure 4.10, showing data taken for a full and half-full beaker of water.

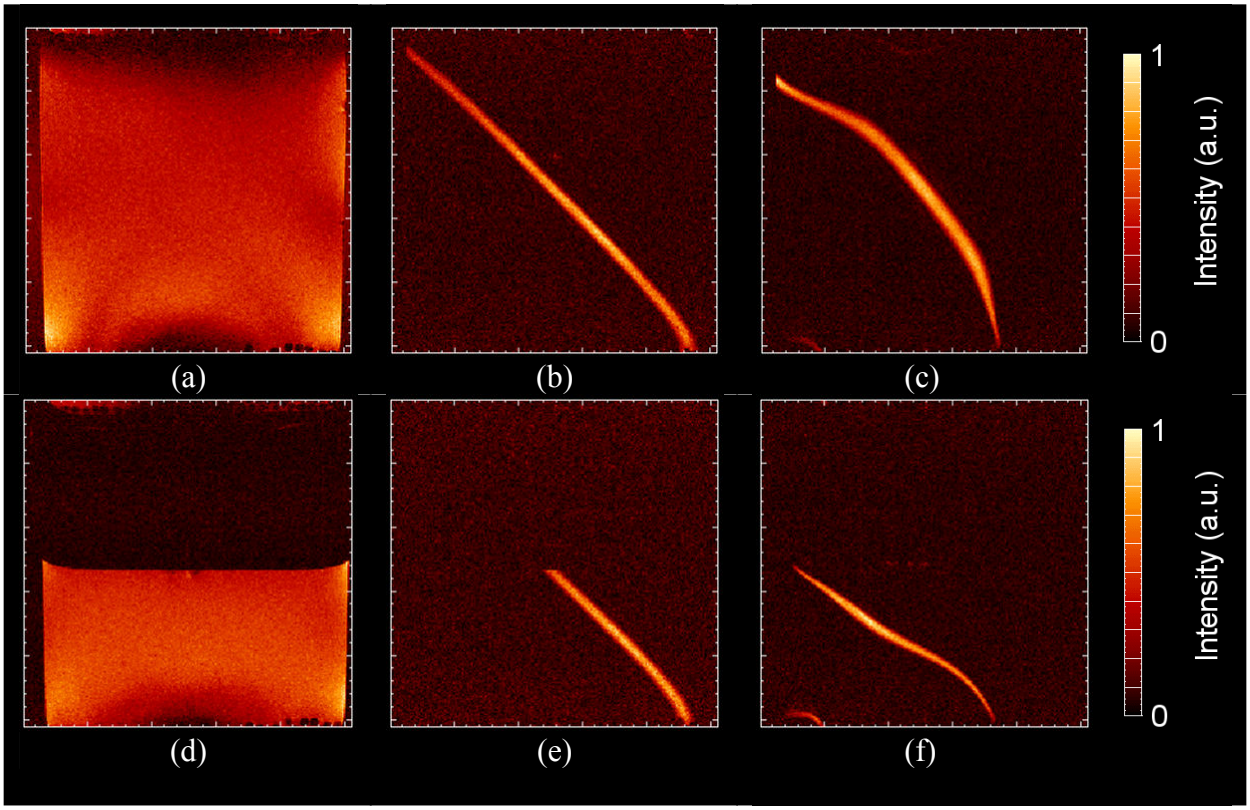


Figure 4.10 – Magnetic susceptibility investigation, showing data taken for full and half-full beakers of water; (a) is a 2D vertical cross-section through a beaker of water, which fills the field of view; (b) shows an excited slice through (a) with a high gradient strength (0.91 $Gcm^{-1}$); (c) shows the same excitation with a lower gradient strength (0.069 $Gcm^{-1}$); (d)-(f) are corresponding images to (a)-(c) for a half-filled beaker of water.

From the data of Figure 4.10, we see that more distortion is apparent at lower gradient strengths. Figure 4.10(b) and (e) are acquired with a high gradient strength (similar to the RARE sequence), which show little distortion; Figure 4.10(c) and (f) show the

same acquisition with a lower gradient strength (similar to the EPI sequence), showing greater distortion, particularly where there is an air-water interface.

The same magnetic susceptibility investigation was then performed on a monolith section, and the results are shown in Figure 4.11. Again, we observe distortion at low gradient strengths due to the presence of the air-water interface. It should be noted that (d) shows the fluid within the monolith channels is drawn upwards due to capillary forces, and is actually higher than the water in the beaker surrounding the monolith.

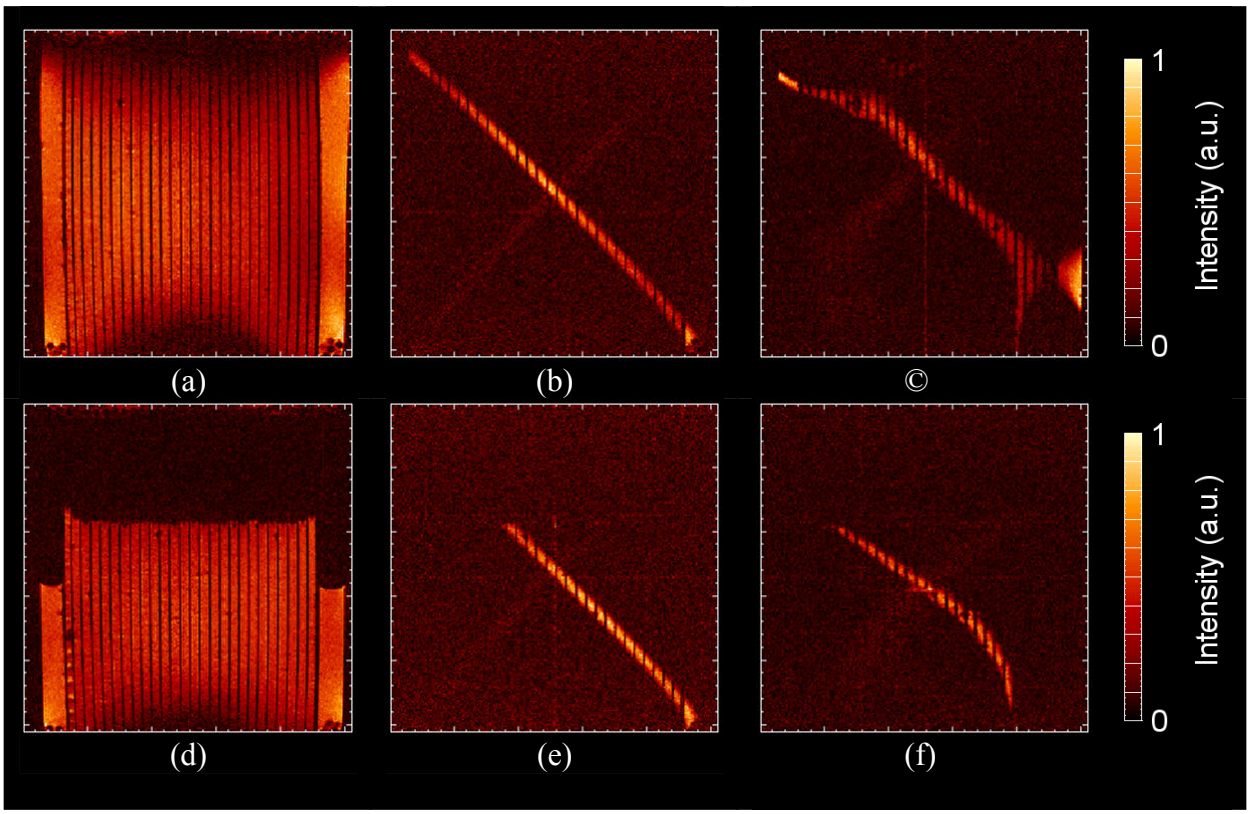


Figure 4.11 – Magnetic susceptibility investigation of monolith section, which was stood in a (a)-(c) full and (d)-(f) half-full beaker of water. Experimental conditions are the same as for Figure 4.10.

An x-y investigation on the monolith was then performed, and this was achieved by blocking the base of the monolith with silicone sealant (see Figure 4.12). This meant that some channels became liquid-filled and others did not. The results are shown in Figure 4.13, and we can see that there is little visible effect of magnetic susceptibility in this orientation.

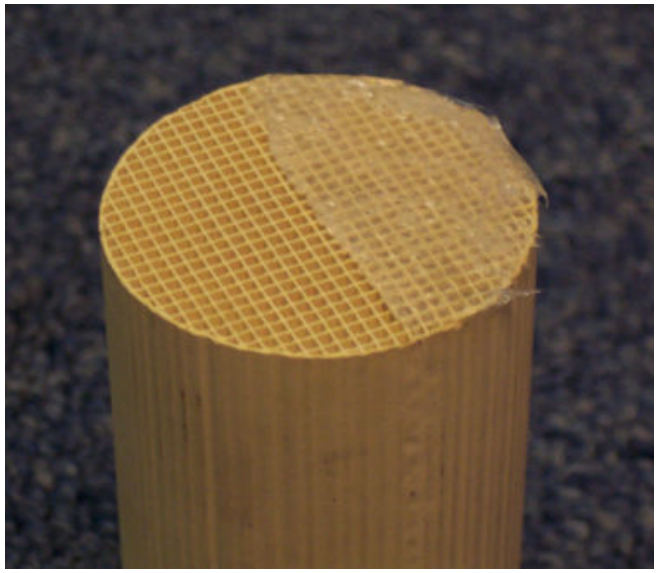

Figure 4.12 – Photograph of the monolith section with several channels blocked with silicone sealant.

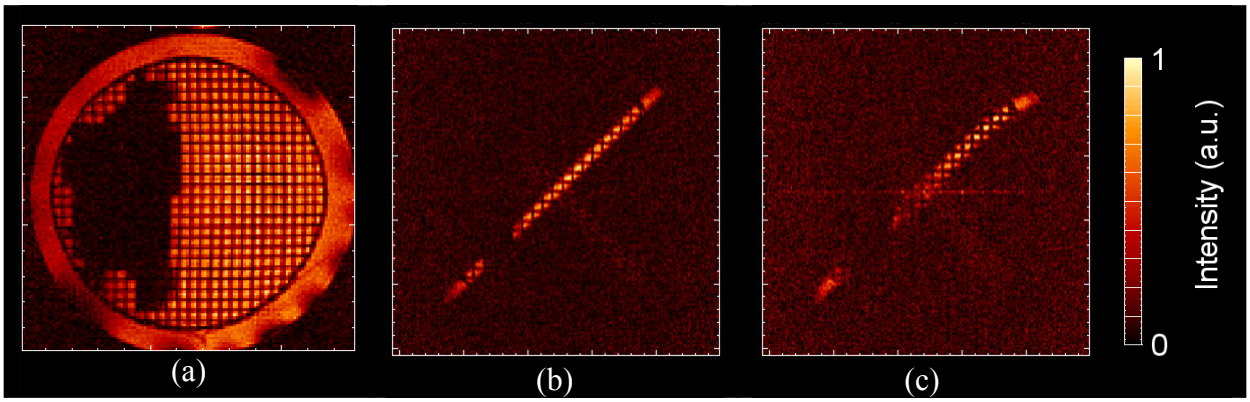


Figure 4.13 – Magnetic susceptibility investigation of the monolith cross-section; (a) x-y cross-sectional image through a half-filled monolith; (b) a line excitation through the monolith taken with high strength imaging gradients (0.91 $Gcm^{-1}$); and (c) the same slice with lower strength gradients (0.069 $Gcm^{-1}$).

The excitation slice taken at low gradient strength (c) is seen to be slightly distorted, although the extent is not as severe as in the previous vertical images. The question as to whether the magnetic susceptibility of the system is acceptable or not depends on the strength of the phase blip in the EPI / GERVAIS[(1)] sequences. This is calculated from the phase rewind gradient which moves us from the centre of **k**-space to the edge of the **k** raster; we divide this by half of the number of phase steps to get

[1] GERVAIS is a novel MRI technique, which is a modification of the EPI sequence. It therefore behaves in a similar manner as regards image artefacts. GERVAIS is described in detail in Section 5.6.

the phase gradient blip strength. This gives a phase gradient strength of 0.13 $Gcm^{-1}$ (for an EPI image with 48 phase steps) which is slightly higher than the low-gradient experiments shown above. Hence, any distortions will be less severe than those seen above, and should be acceptable, particularly in the x-y direction where there is relatively little distortion. This is the orientation we are typically most interested in.

### 4.2.3. Imaging of porous systems

In multi-phase systems, different phases are distinguished by either exploiting the differences in their NMR relaxation properties, differences in NMR signal strength, or with chemical shift imaging, differences in chemical composition.

Contrasting agents (dopants) can be deployed to further enhance the magnetic resonance images of dynamic systems. Deuterium ($^2H$), which is invisible in $^1H$ NMR, is often used, though it is prohibitively expensive. We may add heavy metal elements, such as gadolinium or dysprosium – when dissolved as a salt in the fluid, they can drastically alter the $T_1$ and $T_2$ properties of the fluid. This then allows the tailoring of the pulse sequence – for example, if we reduce the $T_1$ of water from around five seconds to 100 ms, we may repeat the excitation / acquisition process far more rapidly and acquire more data in less time.

### 4.2.4. Measuring velocity

MRI as a tool for visualising fluid flow was reviewed by Fukushima (1999). The main motivation for flow visualisation in porous media comes from the oil and chemicals industry, where the main interest is to determine optimum configuration of porous reactors such as trickle beds and monolith reactors. Early work was performed using "time-of-flight" techniques, which relied on tagging a slice and using the evolution in the position of the signal intensity to measure velocity. However, this was prone to imaging artefacts and more reliable measurements are obtained from the phase-encoding technique (Section 2.4.1). There are two ways of extracting velocity information from the data:

- calculate the phase shift induced in each pixel (relative to a zero-flow image)
- Fourier transform the data with respect to both **k**- and **q**-space.

It has been shown that in complex heterogeneous porous media where each pixel contains many pores, the **q**-space method (Section 2.4.2) yields more accurate velocity measurements (Chang and Watson, 1999). This in, in part, due to partial volume effects – if we visualise a voxel which contains spins moving at different velocities, the overall phase-shift measurement for that voxel will be an average of the spins which it contains. However, it is not accurate to average the real and imaginary parts separately, such as happens with an MRI experiment.

Using these techniques, significant heterogeneities may be visualised in an apparently well-packed bed; Sederman *et al.* (1998) observed 8% of the pores carried 40% of the volumetric flow. Structural correlations showed a link between flow velocity within a pore space and both the local Reynolds number and coordination number.

Magnetic resonance velocity measurement has provided an excellent way of verifying computational fluid dynamic (CFD) models. For example, velocities in a packed bed of glass beads were predicted from Lattice-Boltzmann numerical simulations. The results were compared with MRI and a quantitative agreement between experiment and simulation was found (Manz *et al.*, 1999; El-Bachir, 2005).

Many groups have experimented with flow-encoded EPI sequences (Kose, 1991; Guilfoyle and Mansfield, 1992; Sederman *et al.*, 2004). Waggoner and Fukushima (1996) have tackled the problems of measuring very slow flows and as well as flows in the presence of background magnetic field gradients.

## 4.3. Conclusion

NMR has been a powerful tool in the study of porous media over the past forty years. Many techniques have been developed and provide invaluable insights into the understanding of porous media, despite the issues associated with imaging complex porous structures. Previous work on MRI of monolith reactors has shown some of the potential of the technique, and this work aims to further develop the state of knowledge to produce novel tools and analysis methods for characterising monolith reactors using MRI.

The effect of magnetic susceptibility has been shown experimentally to have an effect on the imaging of monolith reactors (particularly with the EPI and GERVAIS sequences), however it is not so significant that the data will necessarily lose their quantitativeness.

The following chapter will describe some fast MRI techniques developed in the Department of Chemical Engineering at the University of Cambridge, and also some novel techniques developed specifically as a part of this work.

## 4.4. Nomenclature

| **Symbols** | **Definition** |
|---|---|
| $\mathbf{B_1}$ | RF magnetic field |
| **k** | Reciprocal space vector |
| **q** | Dynamic reciprocal space vector |
| $T_1$ | Spin-lattice relaxation constant |
| $T_2^*$ | Apparent spin-spin relaxation constant |

| **Abbreviations** | **Definition** |
|---|---|
| CFD | Computational fluid dynamics |
| cpsi | Cells per square inch |
| EPI | Echo-planar imaging |
| GERVAIS | Gradient echo rapid velocity and acceleration imaging sequence |
| MRI | Magnetic resonance imaging |
| NMR | Nuclear magnetic resonance |
| PIV | Particle imaging velocimetry |
| RARE | Rapid acquisition with relaxation enhancement |
| RF | Radio frequency |
| RTD | Residence time distribution |
| SEMI-RARE | Single-excitation, multiple image RARE |

## 4.5. References


Albers, R. E. (2005). "Characterization of the performance of an industrial monolith reactor by accurate mapping of temperature differences." Catalysis Today **105**(3-4): 391-395.

Anadon, L. D., M. H. M. Lim*, et al.* (2005). "Hydrodynamics in two-phase flow within porous media." Magnetic Resonance Imaging **23**(2): 291-294.

Barrie, P. J. (2000). Characterization of porous media using NMR methods. Annual Reports on Nmr Spectroscopy, Vol 41. **41:** 265-316.

Boyer, C., A. M. Duquenne*, et al.* (2002). "Measuring techniques in gas-liquid and gas-liquid-solid reactors." Chemical Engineering Science **57**(16): 3185-3215.

Callaghan, P. T. (1991). Principles of Nuclear Magnetic Resonance Microscopy. Oxford, Clarendon Press.

Chang, C. T. P. and A. T. Watson (1999). "NMR imaging of flow velocity in porous media." Aiche Journal **45**(3): 437-444.

de Swiet, T. M., M. Tomaselli*, et al.* (1998). "In situ NMR analysis of fluids contained in sedimentary rock." Journal of Magnetic Resonance **133**(2): 385-387.

Dickin, F. J., B. S. Hoyle*, et al.* (1992). "Tomographic Imaging of Industrial-Process Equipment - Techniques and Applications." Iee Proceedings-G Circuits Devices and Systems **139**(1): 72-82.

El-Bachir, M. S. (2005). MRI Study of Flow in Narrow Packed Beds. Dept of Chemical Engineering. Cambridge, University of Cambridge**:** 286.

Fukushima, E. (1999). "Nuclear magnetic resonance as a tool to study flow." Annual Review of Fluid Mechanics **31**: 95-123.

Gladden, L. F. (1993). "Nuclear-Magnetic-Resonance Studies of Porous-Media." Chemical Engineering Research & Design **71**(A6): 657-674.

Gladden, L. F. and P. Alexander (1996). "Applications of nuclear magnetic resonance imaging in process engineering." Measurement Science & Technology **7**(3): 423-435.

Gladden, L. F. (1999). "Applications of in situ magnetic resonance techniques in chemical reaction engineering." Topics in Catalysis **8**(1-2): 87-95.

Gladden, L. F., C. Buckley*, et al.* (2004). "'Looking into' chemical products and processes." Current Applied Physics **4**(2-4): 93-97.

Guilfoyle, D. N. and P. Mansfield (1992). "Fluid Flow Measurement in Porous Media by Echo-Planar Imaging." Journal of Magnetic Resonance **97**: 342-358.

Haakana, T., E. Kolehmainen*, et al.* (2004). "The development of monolith reactors: general strategy with a case study." Chemical Engineering Science **59**(22-23): 5629-5635.

Heibel, A. K., T. W. J. Scheenen*, et al.* (2001). "Gas and liquid phase distribution and their effect on reactor performance in the monolith film flow reactor." Chemical Engineering Science **56**(21-22): 5935-5944.

Heiszwolf, J. J., M. T. Kreutzer*, et al.* (2001). "Gas-liquid mass transfer of aqueous Taylor flow in monoliths." Catalysis Today **69**(1-4): 51-55.

Johns, M. L., A. J. Sederman*, et al.* (2000). "Local transitions in flow phenomena through packed beds identified by MRI." Aiche Journal **46**(11): 2151-2161.

Kantzas, A. (1994). "Computation of Holdups in Fluidized and Trickle Beds by Computer-Assisted Tomography." Aiche Journal **40**: 1254-1261.

Khrapitchev, A. A., S. Stapf*, et al.* (2002). "NMR visualization of displacement correlations for flow in porous media." Physical Review E **66**(5).

Koptyug, I. V., S. A. Altobelli*, et al.* (2000). "Thermally polarized H-1 NMR microimaging studies of liquid and gas flow in monolithic catalysts." Journal of Magnetic Resonance **147**(1): 36-42.

Koptyug, I. V., L. Y. Ilyina*, et al.* (2001). "Liquid and gas flow and related phenomena in monolithic catalysts studied by H-1 NMR microimaging." Catalysis Today **69**(1-4): 385-392.

Koptyug, I. V., A. A. Lysova*, et al.* (2005). "In situ MRI of the structure and function of multiphase catalytic reactors." Catalysis Today **105**(3-4): 464-468.

Kose, K. (1991). "One-Shot Velocity Mapping Using Multiple Spin-Echo Epi and Its Application to Turbulent-Flow." Journal of Magnetic Resonance **92**(3): 631-635.

Mantle, M. D., A. J. Sederman, *et al.* (2001). "Single- and two-phase flow in fixed-bed reactors: MRI flow visualisation and lattice-Boltzmann simulations." Chemical Engineering Science **56**(2): 523-529.

Mantle, M. D., A. J. Sederman, *et al.* (2002). "Dynamic MRI visualization of two-phase flow in a ceramic monolith." Aiche Journal **48**(4): 909-912.

Mantle, M. D. and A. J. Sederman (2003). "Dynamic MRI in chemical process and reaction engineering." Progress in Nuclear Magnetic Resonance Spectroscopy **43**(1-2): 3-60.

Manz, B., L. F. Gladden, *et al.* (1999). "Flow and Dispersion in Porous Media: Lattice-Boltzmann and NMR Studies." Aiche Journal **45**: 1845-1854.

Mewes, D., T. Loser, *et al.* (1999). "Modelling of two-phase flow in packings and monoliths." Chemical Engineering Science **54**(21): 4729-4747.

Nakumara, M., H. Mori, *et al.* (1991). Visualisation of Liquid Flow in Trickle Bed Reactor by X-Ray Computer Tomography Method. Tsukuba.

Peters, A. M. and R. Bowtell (1999). "Resolution in high field echo planar microscopy." Journal of Magnetic Resonance **137**(1): 196-205.

Prasser, H. M., A. Bottger, *et al.* (1998). "A new electrode-mesh tomograph for gas-liquid flows." Flow Measurement and Instrumentation **9**(2): 111-119.

Reinecke, N., M. Boddem, *et al.* (1996). "Tomographic measurement of the relative phase fractions in binary flow of fluid phases." Chemie Ingenieur Technik **68**(11): 1404-1412.

Reinecke, N. and D. Mewes (1996). "Tomographic imaging of trickle-bed reactors." Chemical Engineering Science **51**(10): 2131-&.

Reinecke, N. and D. Mewes (1997). "Investigation of the two phase flow in trickle-bed reactors using capacitance tomography." Chemical Engineering Science **52**(13): 2111-2127.

Reinecke, N., G. Petritsch, *et al.* (1998). "Tomographic measurement techniques - Visualization of multiphase flows." Chemical Engineering & Technology **21**(1): 7-18.

Roy, S. and M. Al-Dahhan (2005). "Flow distribution characteristics of a gas-liquid monolith reactor." Catalysis Today **105**(3-4): 396-400.

Sains, M. C., M. S. El-Bachir, *et al.* (2005). "Rapid imaging of fluid flow patterns in a narrow packed bed using MRI." Magnetic Resonance Imaging **23**(2): 391-393.

Santiago, J. G., S. T. Wereley, *et al.* (1998). "A particle image velocimetry system for microfluidics." Experiments in Fluids **25**(4): 316-319.

Schmitz, D. and D. Mewes (2000). "Tomographic imaging of transient multiphase flow in bubble columns." Chemical Engineering Journal **77**(1-2): 99-104.

Sederman, A. J., M. L. Johns, *et al.* (1998). "Structure-flow correlations in packed beds." Chemical Engineering Science **53**(12): 2117-2128.

Sederman, A. J. and L. F. Gladden (2001). "Magnetic resonance imaging as a quantitative probe of gas-liquid distribution and wetting efficiency in trickle-bed reactors." Chemical Engineering Science **56**(8): 2615-2628.

Sederman, A. J., M. D. Mantle, *et al.* (2003). "Quantitative 'real-time' imaging of multi-phase flow in ceramic monoliths." Magnetic Resonance Imaging **21**(3-4): 359-361.

Sederman, A. J., M. D. Mantle, *et al.* (2004). "MRI technique for measurement of velocity vectors, acceleration, and autocorrelation functions in turbulent flow." Journal of Magnetic Resonance **166**(2): 182-189.

Toye, D., P. Marchot, *et al.* (1995). "Analysis of Liquid Flow Distribution in Trickling Flow Reactor using Computer-assisted X-ray tomography." Chemical Engineering Research & Design **73**: 258-262.

Waggoner, R. and E. Fukushima (1996). "Velocity Distribution of Slow Fluid Flows in Bentheimer Sandstones." Magnetic Resonance Imaging **14**: 1085-1091.

Wang, S. J., T. Dyakowski, *et al.* (1995). "Real-Time Capacitance Imaging of Bubble Formation at the Distributor of a Fluidized-Bed." Chemical Engineering Journal and the Biochemical Engineering Journal **56**(3): 95-100.

Wang, Y. F., Z. S. Mao, *et al.* (1998). "Scale and variance of radial liquid maldistribution in trickle beds." Chemical Engineering Science **53**(6): 1153-1162.

Watson, A. T. and C. T. P. Chang (1997). "Characterizing porous media with NMR methods." Progress in Nuclear Magnetic Resonance Spectroscopy **31**: 343-386.

Yuen, E. H. L., A. J. Sederman, *et al.* (2002). "In situ magnetic resonance visualisation of the spatial variation of catalytic conversion within a fixed-bed reactor." Applied Catalysis a-General **232**(1-2): 29-38.

Yuen, E. H. L. (2003). Aspects of Nuclear Magnetic Resonance Studies of Transport and Reaction in Packed Bed Reactors. Dept of Chemical Engineering. Cambridge, University of Cambridge**:** 257.

# Chapter 5 – Novel MRI techniques

MRI is well established, and is an ideal analytical tool for many chemical engineering applications. The fundamentals of NMR and some traditional imaging sequences have been covered in Chapter 2. This chapter will highlight some of the recent progress made within the field of ultra-fast velocity imaging techniques, both as a part of this research and by others in the Department of Chemical Engineering at the University of Cambridge. The development and theory behind the creation of the pulse techniques will be summarised, followed by quantification of the pulse sequences.

**CHAPTER 5 – NOVEL MRI TECHNIQUES**



## 5.1. Introduction

With MRI there is always a strong drive towards faster, more efficient pulse techniques in order to reduce imaging times and allow the study of dynamic processes, rather than acquiring time-averaged data. Five novel MRI pulse sequences for measuring velocities will be introduced in this section, along with illustrations of their application, quantification and limitations.

- **SEMI-RARE** – This is a modification of the RARE pulse sequence, developed within this group by Sederman *et al.* (2003), and allows a sequence of images to be acquired from a single excitation at increasing time from the excitation. Velocities may then be measured by comparing the movement of features between the successive images; it is in essence a time-of-flight (TOF) velocity technique. Typical imaging times are of the order of 150 ms per image (128×64 pixels). The number of images that can be acquired from the initial excitation is limited by the relaxation properties of the system, and also by diffusional attenuation.

- **Line-excitation** – this is a novel, RARE-based, TOF technique developed as part of this work, and works by tagging a line of fluid and then observing the spreading of the tagged spins. The rapid acquisition speed and applicability to dynamic systems make it a useful imaging technique. It will be shown in Chapter 14 that propagator-type information may be extracted from this data for an entire row of monolith channels in 150 ms.

- **COTTAGE** – this is a novel modification of the line-excitation method, developed as part of this work. 1D profiles are acquired, and by sacrificing spatial resolution in this way we may acquire far higher temporal resolution. Phase-encoding is used to measure the velocity directly. This technique allows the observation of dynamic fluctuations in velocity, on the order of around 3 ms.

- **MULTI-SLICE COTTAGE** – this is a further modification of the phase-encoded COTTAGE technique, where rather than refocusing the same 1D profile repeatedly, we scan across the cross-section of our imaging region, thereby building up a 2D velocity image in around 80 ms, at the expense of temporal resolution.

- **GERVAIS** – this is a phase-encoded velocity and acceleration sequence based on EPI, developed within this group by Sederman *et al.* (2004). It allows the acquisition of a velocity component in 20 ms, and all three velocity

components in 60 ms. It will be shown that there are problems in applying this sequence to porous media, and also for two-phase flow. However, it is still the fastest 2D velocity imaging sequence available and is of considerable use in dynamic systems, particularly turbulent flows.

This chapter will introduce the pulse sequences, and illustrate their quantitative nature in simple systems such as laminar / turbulent flow in a pipe. The following chapter will then focus on the application of these novel techniques to the monolith reactor, and highlight the extent of their applicability and the quality of data we may extract in each case.

Where stable laminar flow is required, the criterion of Christiansen and Lemmon (1965) was used. The authors performed experiments on a Plexiglass pipe with i.d. of 40 mm and report that the length after which the centreline velocity deviates by less than 1% from the parabolic form is given by Eq. 5.1.

$$\frac{L_{99\%}}{D} = 0.056\,\mathrm{Re} \qquad \textbf{(Eq. 5.1)}$$

We can apply this criterion to the flow configurations used in this work. For example, if we have a tube of diameter 29 mm, and a settling length of 60 tube diameters, then the maximum superficial velocity for steady flow is 3.7 cm/s (88 l/hr). In a tube of diameter 49 mm, and a settling length of 20 tube diameters, the maximum superficial velocity for steady flow is 7.3 mm/s (49 l/hr).

## 5.2. SEMI-RARE

The SEMI-RARE pulse sequence (Single Excitation, Multiple Image RARE) is based on the RARE sequence (Section 2.3.1) and was developed by Sederman *et al.* (2003). It has the same advantages as RARE, which allows the traversing of **k**-space in one shot rather than line-by-line, and is a robust acquisition method which is relatively insensitive to precise shimming and systems characterised by significant magnetic susceptibility variations (Section 4.2.3). The SEMI-RARE sequence acquires a RARE image as normal, but then uses the remaining magnetisation to then re-scan over **k**-space and thereby acquire successive images from a single excitation. This is

advantageous since it allows the rapid collection of images; a new excitation and acquisition could not be performed this rapidly since the system must be allowed to relax between excitations. The SEMI-RARE pulse sequence is shown in Figure 5.1.

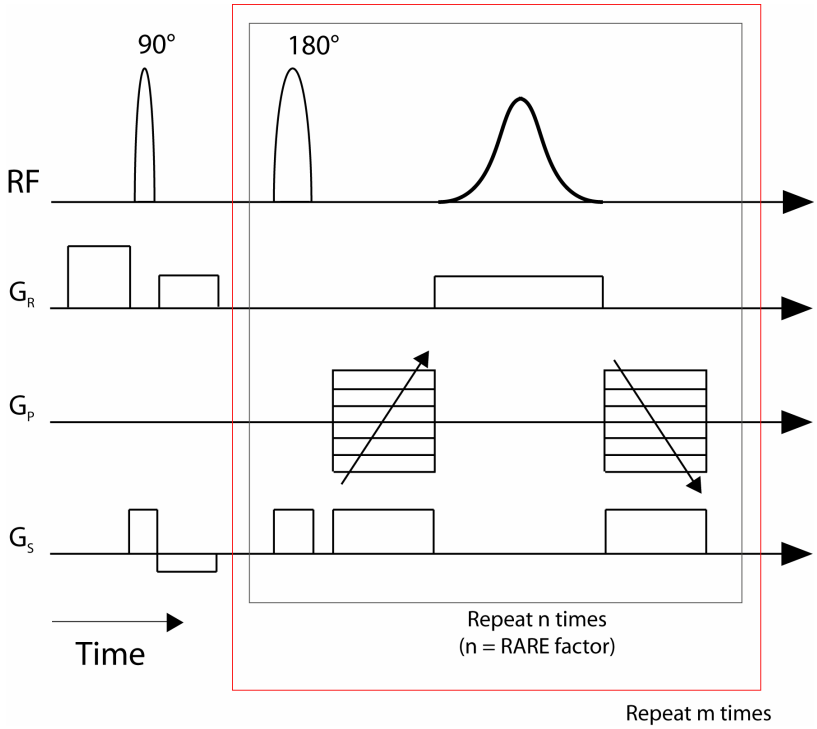


Figure 5.1 – Illustration of the SEMI-RARE pulse sequence. It is identical to the RARE pulse sequence (Section 2.3.1), except we repeat the echo loop to obtain *m* images, until magnetisation is lost.

An example of a series of images acquired using the SEMI-RARE sequence is presented in Figure 5.2, showing a series of 120 images acquired from a single excitation. Signal-to-noise decreases with each image in the SEMI-RARE sequence, according to the $T_2$ relaxation properties of each liquid and the diffusive attenuation of the system. It can be seen that water has the longest $T_2$, followed by octanol and glycerol with the shortest. By tracking the decrease in intensity of each voxel, $T_2$ curves may be fitted to estimate the relaxation characteristics.

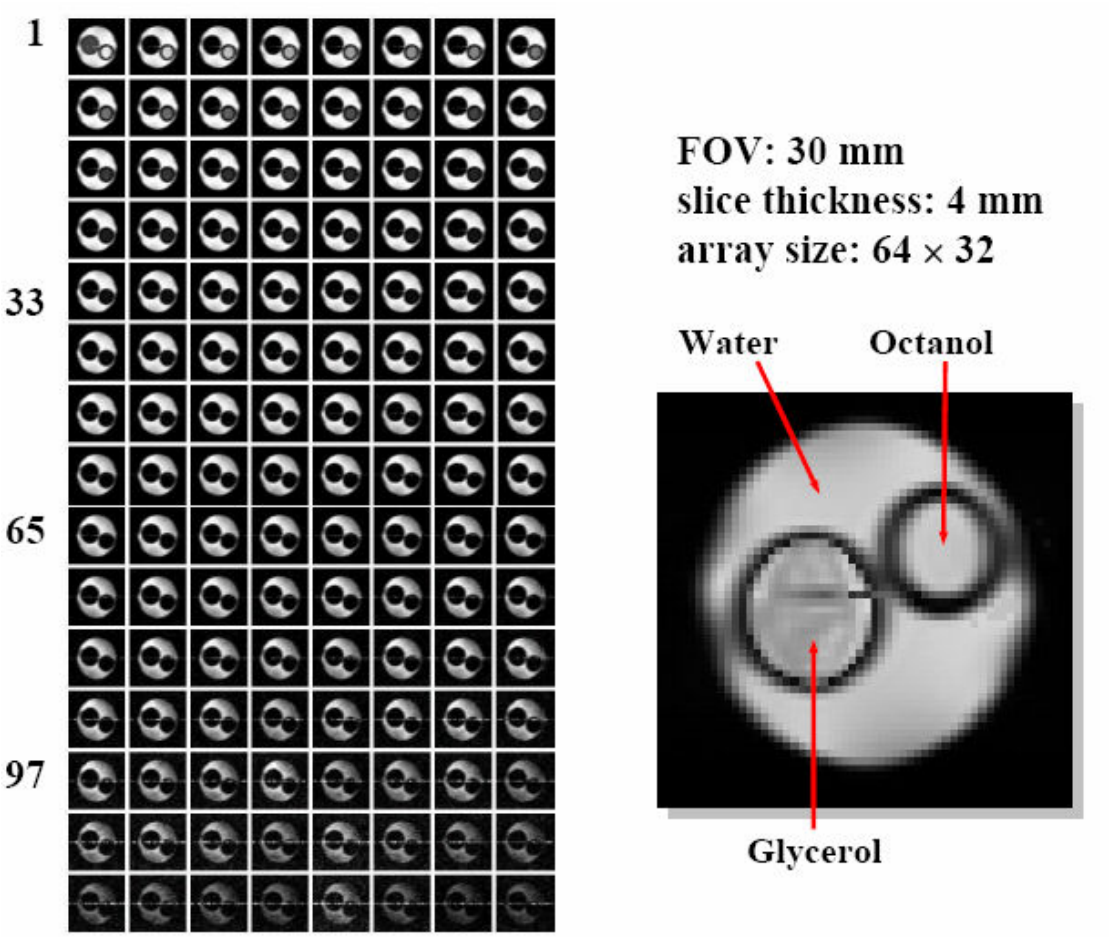


(a) (b)

Figure 5.2 – An example of a SEMI-RARE acquisition, showing 120 images acquired from a single acquisition. The phantom of water contains two test tubes of octanol and glycerol. The different $T_2$ relaxation properties of the substances can be seen in the rates of decay of the relative intensities within the series of images. Reproduced from Sederman *et al.* (2003).

In less ideal systems than this example, or for images of higher resolution (longer imaging times), fewer images can be acquired before magnetisation is lost. For example, in the monolith system (Figure 5.3) we typically only acquire four images of suitable signal strength, although this is sufficient to allow detailed analysis procedures. It is possible to infer velocities by measuring the movement of details within the successive images; it is in essence a time-of-flight velocity technique. Typical imaging times are of the order of 150 ms per image. The SEMI-RARE sequence therefore allows the visualisation of a dynamic system such as the monolith reactor, while retaining the robustness of the RARE method in application to porous media.

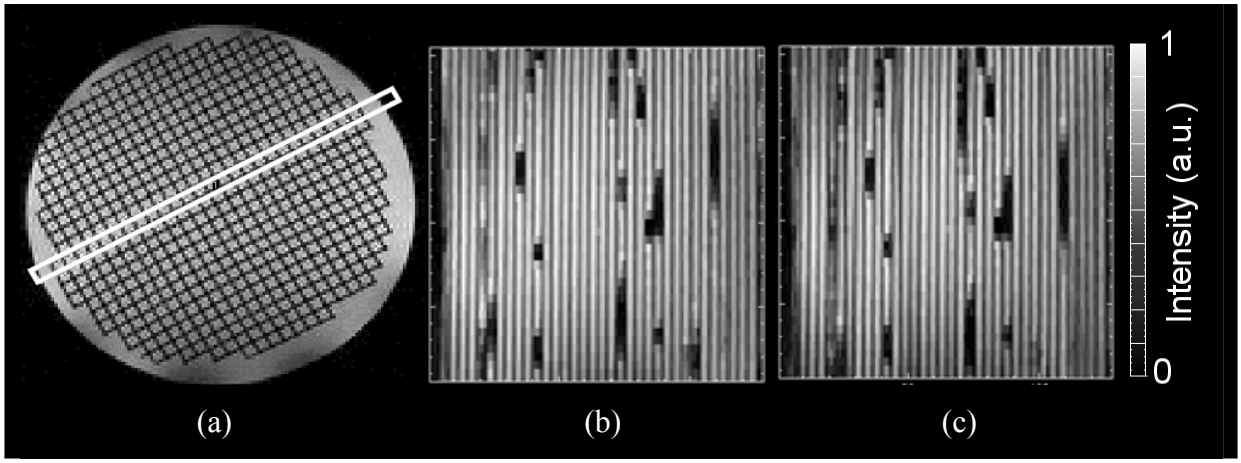


Figure 5.3 – An example of SEMI-RARE applied to two-phase flow within a ceramic monolith; (a) shows a cross-section of the water-saturated monolith, showing the position of the channels shown in (b) and (c); (b) a slice through a single row of channels under two-phase flow conditions, showing the presence of Taylor flow; (c) is the second image in the SEMI-RARE acquisition, which allows us to track slug interfaces and therefore measure local velocities. Reproduced from Mantle *et al.* (2002).

The SEMI-RARE tracking method is inherently quantitative, since the displacement of the liquid / gas interfaces are proportional to the fluid velocity in each channel. Further information may be gleaned from these data, such as bubble / slug lengths and frequencies, holdup measurements, and distribution homogeneity. These aspects will be covered in detail in Chapters 8-11.

The disadvantage of the SEMI-RARE method is that it does not provide any velocity-tracking information for the case of single-phase flow, since there exist no interfaces or features to track[1]. For this reason, the "line-excitation" modification was introduced, and is described in the following section.

## 5.3. Line-excitation

The proposed line-excitation method is also based upon the RARE sequence. The pulse sequence was implemented in this work by exciting a slice of fluid within the reactor in the transverse plane, and then refocusing the spins in the direction of flow, such that only a single line of magnetisation is refocused, as illustrated in Figure 5.4.

[1] There is however the possibility of examining the inflow of non-excited spins into the imaging volume, and thereby estimating the velocity of the fluid. This is discussed further in Section 11.2.1.

We acquire a time-averaged image of the displacement of the tagged spins over time, similar in nature to the DANTE-tagging technique (Morris and Freeman, 1978). The pulse sequence for the line-excitation method is presented in Figure 5.5.

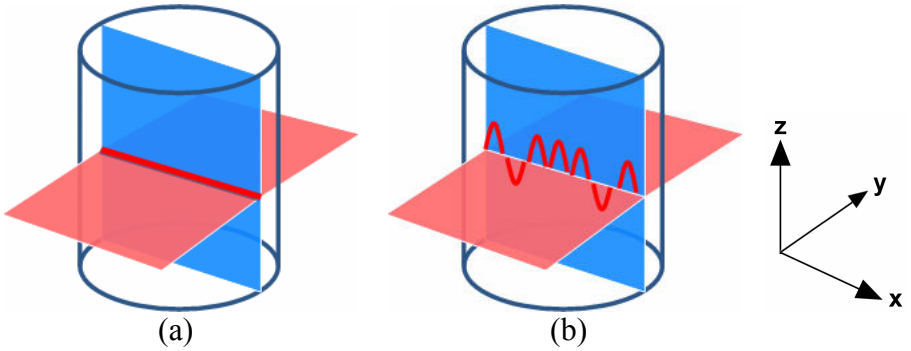


Figure 5.4 – Schematic of the orientation of the excitation plane (red) and read plane (blue), (a) with zero flow, a single line is refocused; (b) with flow, the line is displaced according to the axial flow velocity.

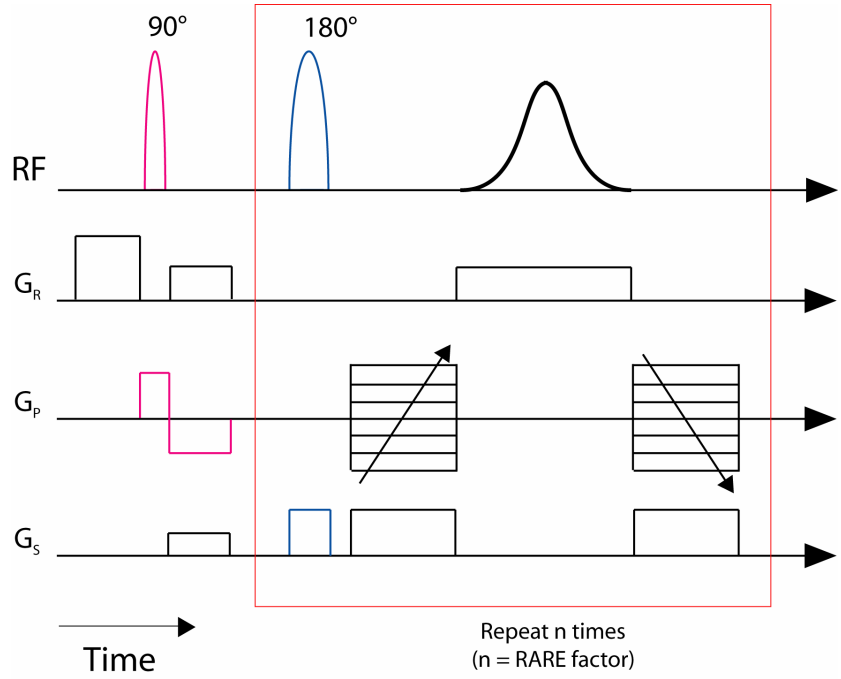


Figure 5.5 – Pulse sequence diagram of the SEMI-RARE line-excitation sequence. It is based upon the RARE sequence (Section 2.3.1) but the orientation of the initial excitation pulse (red) has been switched such that it is orthogonal to the refocusing pulse (blue), thus giving only a line of refocused magnetisation within our 2D image.

The modification of the RARE pulse sequence to obtain the line-excitation sequence is relatively straightforward – the gradient pulse associated with the 90° RF excitation pulse is swapped from the slice-direction ($x$) to the phase-direction ($z$), along with its

refocusing gradient. This changes the orientation of the excitation plane relative to the read plane, and the intersection is seen as a thin line which is displaced by movement over the acquisition time.

The line-excitation method achieves a rapid fingerprint of instantaneous flow, and this time-of-flight technique can be translated directly into velocity information for systems with confined geometries and where the flow is laminar, i.e. where we have uni-axial flow with negligible *x*- and *y*-velocities. This is therefore applicable to flow in a pipe or within the channels of a monolith, since fluid elements are constrained primarily to axial flow with relatively insignificant radial mixing (Section 3.4.2), and so the displacement is indeed proportional to the velocity of each fluid element. It will be shown later in Chapter 14 that even propagator information can be inferred from the degree of spreading of the intensity[(2)].

This rapid imaging sequence has many advantages over traditional phase-encoded methods for measuring displacement / velocity information. Phase-encoded methods suffer from partial volume effects in materials of dense geometry, giving incorrect phase-averaging between different velocities which may exist within the same voxel[(3)]. In comparison, the line-excitation method measures the actual position of the tagged spins, and therefore does not suffer from problems with phase-averaging; the intensity acquired in a voxel is a direct measure of the number of tagged spins in each voxel (with some caveats regarding relaxation effects), and it is rapid enough at 150 ms not to suffer from significant time-averaging.

An example of the line-excitation technique applied to laminar flow in a pipe is shown in Figure 5.6, showing four images acquired from a single excitation. The parabolic profile is clearly observed.

[2] However, it will also be shown that the propagator information is not complete since we lose intensity from spins in regions of high shear, in short-$T_2$ environments (near the walls) and due to diffusive attenuation. This means that typically, we lose signal from slow-moving fluid near the wall, which skews the propagator measurement. However, in this example where we are simply looking at the movement of spins, rather than focussing on their particular distribution, we can still extract an accurate displacement.

[3] Incorrect phase averaging occurs since the real parts of the signal are averaged separate to the imaginary parts of the signal. This is not necessarily a true average of the phase angle.

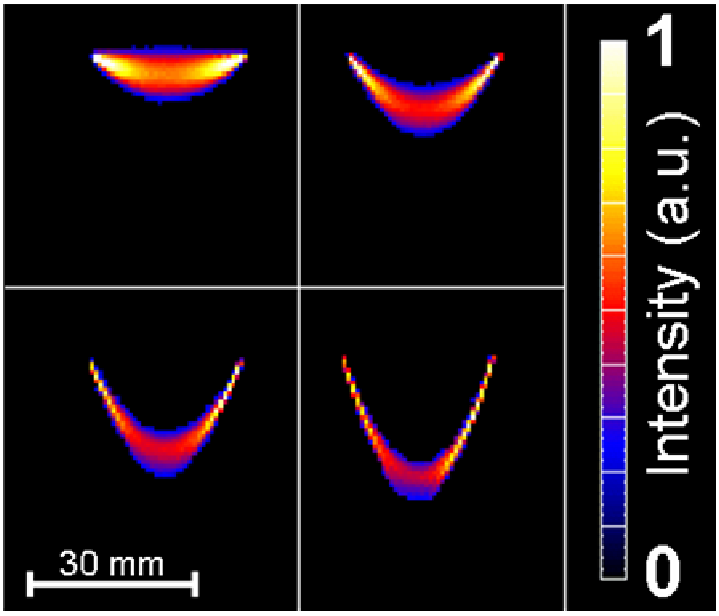


Figure 5.6 – Illustration of the SEMI-RARE time-of-flight technique applied to laminar flow in a pipe. The flow rate of water was 60 l/hr, and the pipe i.d. was 30 mm. The slice thickness is 3 mm and was positioned at the centreline of the pipe. The four images are acquired 156 ms apart, after a single excitation pulse.

The following example shows the quantification of the SEMI-RARE time-of-flight technique. The data of Figure 5.6 were analysed, and velocities were extracted by measuring the position of the maximum intensity at each $x$-coordinate; this displacement was then converted to a velocity by knowledge of the echo time of the SEMI-RARE data, as shown in Eq. 5.2.

$$v = \frac{s}{t_e} \qquad \textbf{(Eq. 5.2)}$$

where $v$ is the superficial velocity of the spins, and $s$ is the measured displacement at the observed echo time, $t_e$. As can be seen in Figure 5.7a, the intensity spread follows the analytical profile very closely, particularly at lower flow rates. The technique generally gives slightly lower measurements than expected at each flow rate, and this is within the errors associated with measuring the volumetric flow rate in the system at less than 10%. Figure 5.7b shows a comparison of the "true" flow rate (as measured manually, with a measuring cylinder and stopwatch) versus the measured MRI flow rate. It is seen that at higher flow rates (where laminar flow breaks down) the technique is no longer quantitative since the assumption of axial flow with no radial mixing no longer holds.

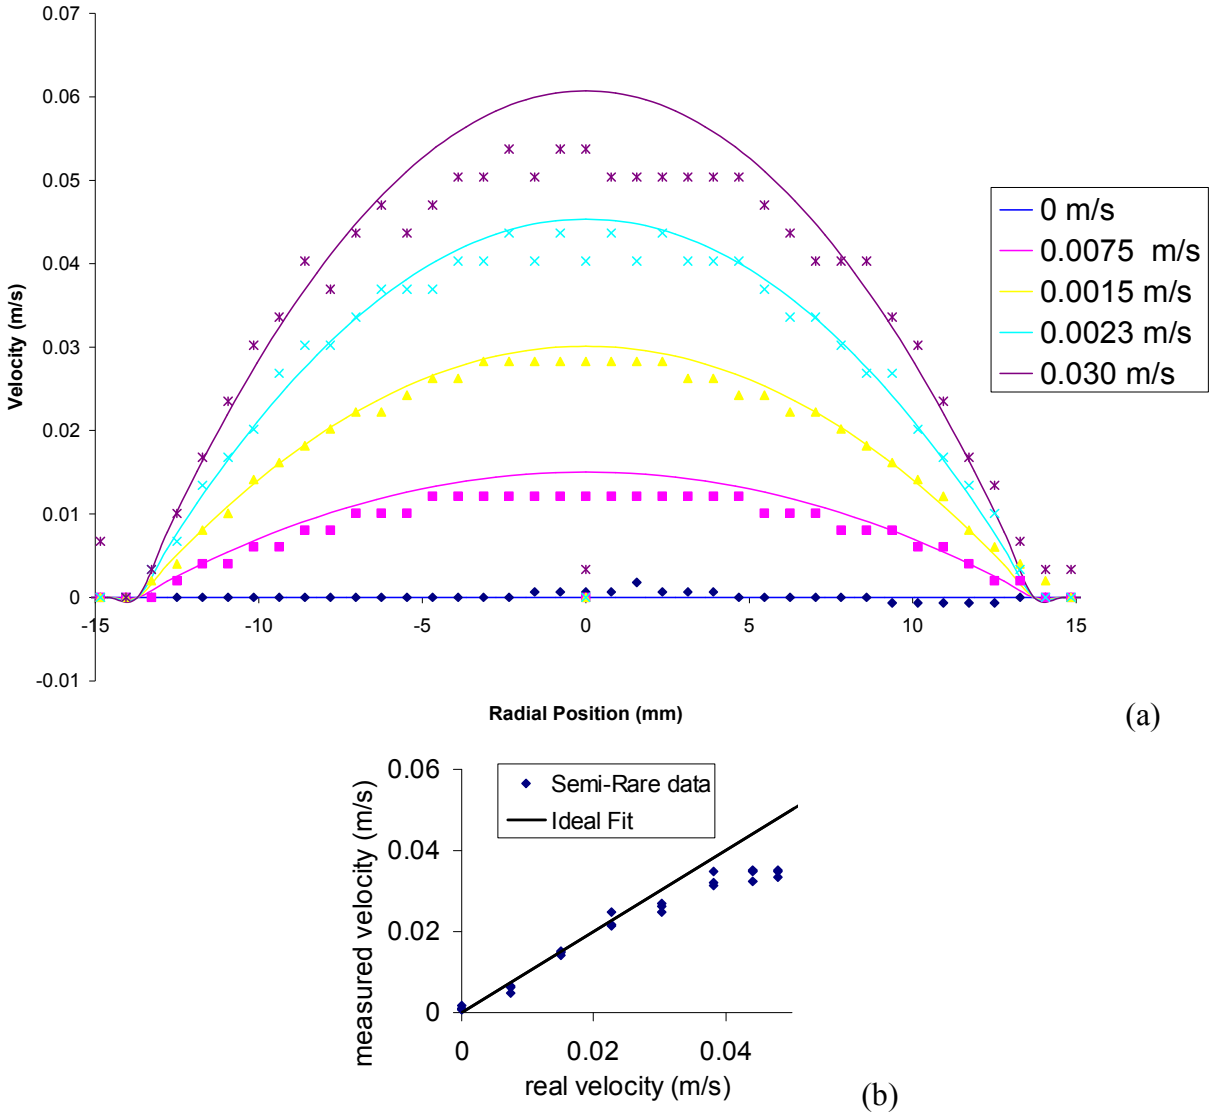


Figure 5.7 – Quantification of the SEMI-RARE time-of-flight technique; (a) a series of SEMI-RARE profiles plotted against the analytical solution at increasing superficial velocity, and (b) a plot of measured superficial velocity versus the actual superficial velocity through the system.

The combination of the line-excitation method with the SEMI-RARE technique allows the acquisition of consecutive images from one excitation, which in turn allows a greater range of velocities to be studied.

## 5.4. COTTAGE

The novel COTTAGE (COntinuous Translation Tracking with Alternating Gradient Encode) sequence will be described in this section. It was developed by modifying the line-excitation sequence such that we keep the same excitation and refocusing pulses as used to read a narrow line, but remove the ramping phase gradient and simply replace it with a user-defined, constant phase gradient[4] – see the pulse sequence diagram in Figure 5.8.

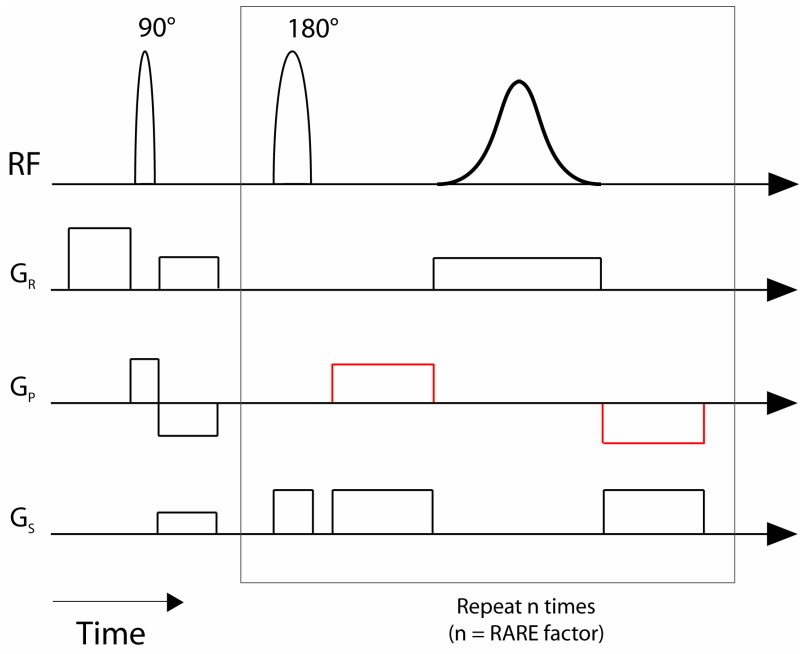


Figure 5.8 – Pulse sequence diagram for the COTTAGE technique. It is similar to the line-excitation sequence (Figure 5.5) but has a constant, rather than ramped, phase gradient (red).

We will therefore acquire the same line of **k**-space during each echo of the RARE sequence. Each line, when Fourier transformed, will correspond to a projection of the system, and by examining the phase-encoding of the signal in each line, we will obtain information on the average displacement of the fluid in that pixel at the time of the echo. The echo train will therefore give us a series of 1D profiles showing the evolution of the phase-angle over the acquisition time. In short, we have sacrificed information on the spread of the displacement for enhanced time resolution.

[4] In fact any user-defined phase gradient may be used. For example, it may be advisable to have a stronger phase gradient during the earlier echo trains to give better resolution of phase angle, but to gradually decrease the phase of the signal at later times to prevent phase-wrapping. However, for simplicity and ease of post-processing, a constant phase gradient was used in this work.

COTTAGE can detect changes in dynamic systems on the order of 3 ms (one echo in the RARE sequence). In a similar way to other phase-encoded velocity imaging sequences, we must acquire and subtract a zero flow dataset from the phase angles obtained.

The phase is imparted to the tagged spins by the first phase gradient in the loop, and the read gradient is then applied, measuring the properties of the system. The second phase gradient serves to remove the imparted phase from the spins ready for the next loop. By comparing the zero-flow data to the measurement of the phase we acquire, we may calculate the displacement of the spins at that particular time. The measured phase is related to the displacement of the spins by Eq. 5.3.

$$s = \frac{\phi}{\gamma \delta G} \qquad \textbf{(Eq. 5.3)}$$

where $\phi$ is the measured phase angle, $\gamma$ is the gyromagnetic ratio, $\delta$ and $G$ are the duration and strength of the phase-encoding gradient. An example of a COTTAGE measurement showing the evolution of displacement with time is shown in Figure 5.9, for laminar flow in a 10 mm i.d. pipe. In this image, the vertical axis corresponds to time, and the horizontal direction corresponds to profiles of the average displacement at that time. Parabolic profiles are seen in all the horizontal profiles as expected (Figure 5.10), and it grows linearly with time (Figure 5.11).

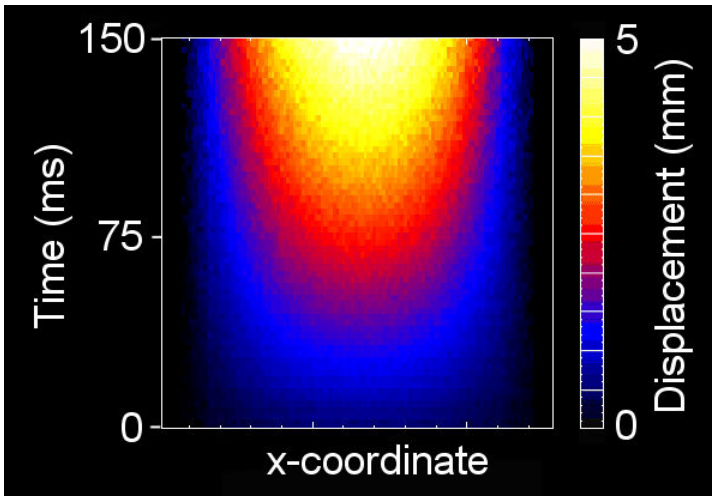


Figure 5.9 – A COTTAGE measurement of laminar flow in a pipe of diameter 10 mm at a flow rate of 4.6 l/hr. The data is read from the central vertical slice, the x-axis corresponds to the x-distance, and the vertical axis shows the profiles acquired in time. The intensity in the image corresponds to the measured displacement.

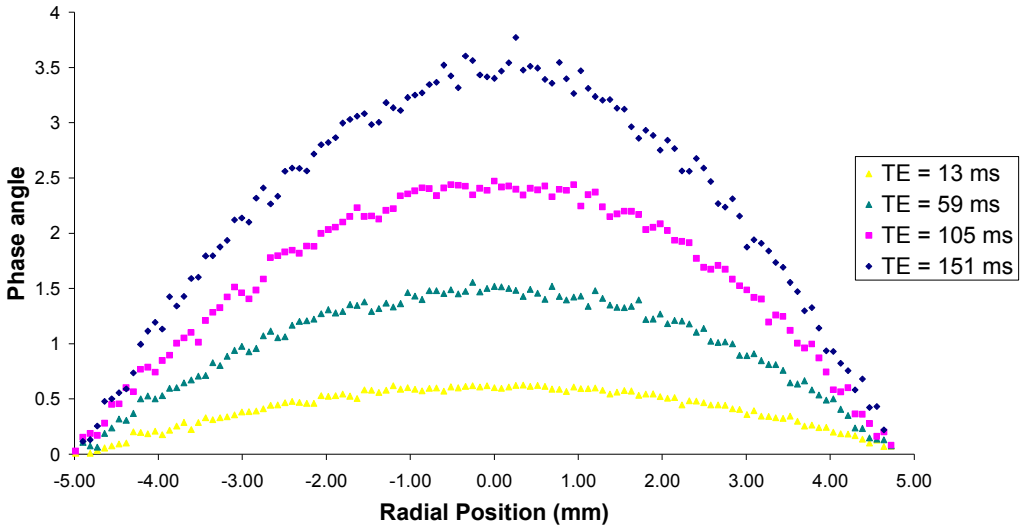


Figure 5.10 – Profiles taken from Figure 5.9 at increasing echo time, showing the parabolic nature of the measurement.

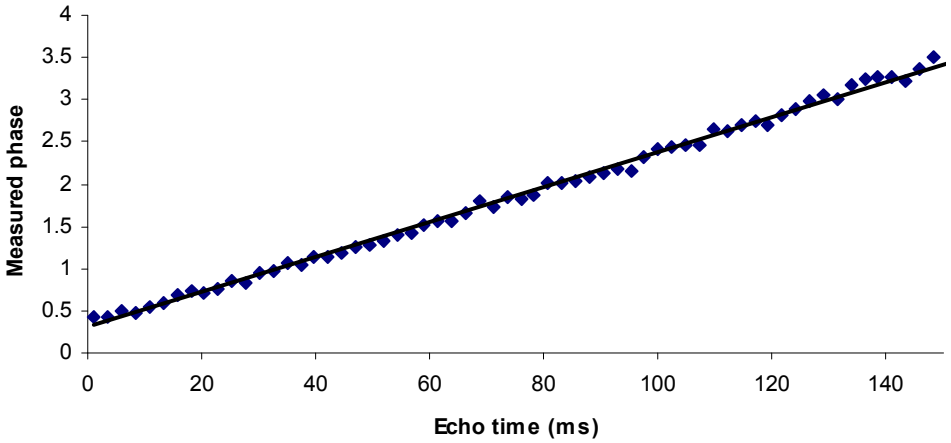


Figure 5.11 – A profile extracted from the data of Figure 5.9 (taken in the vertical direction), showing the linear increase of measured phase with time. The phase measurement is directly proportional to the displacement, by Eq. 5.3.

It can be seen from Figure 5.11 that the increase in displacement is indeed linear with time, as expected, however there is an initial phase offset. It is unclear where this offset comes from, but it needs to be corrected for before quantitative results may be obtained. Once we have corrected for this apparent phase offset, the data of Figure 5.9 (which showed displacement at increasing time) can be converted into a velocity measurement by dividing through by the appropriate echo time, as is shown in Figure 5.12a. It can be seen that the velocity increases very slightly with time, rather than being constant as expected for laminar flow – this is due to the assumption that during the phase encoding / decoding part of the pulse sequence, the fluid does not move. If

there is motion within this time, the phase imparted by the first phase gradient will not be fully removed by the second phase gradient. Hence, we will retain an increasing residual phase, which will be interpreted as an increased displacement. For example, if we observe the centreline velocity of Figure 5.12a, the measured velocity ranges between 1.64 cm/s for the 1st echo and 1.87 cm/s for the 64th echo. This effect is therefore significant (a total of 14% over 64 echoes), and needs to be corrected for. The additional phase which is measured is proportional to the velocity of the spins (described by Eq. 5.4), and is therefore proportional to the phase measured by the pulse sequence. By calculating the relative contribution to this from the residual phase, we can calculate that we must reduce the value of the measured velocity by a factor shown in Eq. 5.5.

$$\phi_{additional} = \gamma \delta \Delta G v \qquad \textbf{(Eq. 5.4)}$$

$$V_{TRUE} = V_{MEASURED} - \left(0.0022 \times V_{MEASURED} \times NE\right) \qquad \textbf{(Eq. 5.5)}$$

where *NE* is the echo number. The result of applying this correction factor is shown in Figure 5.12b.

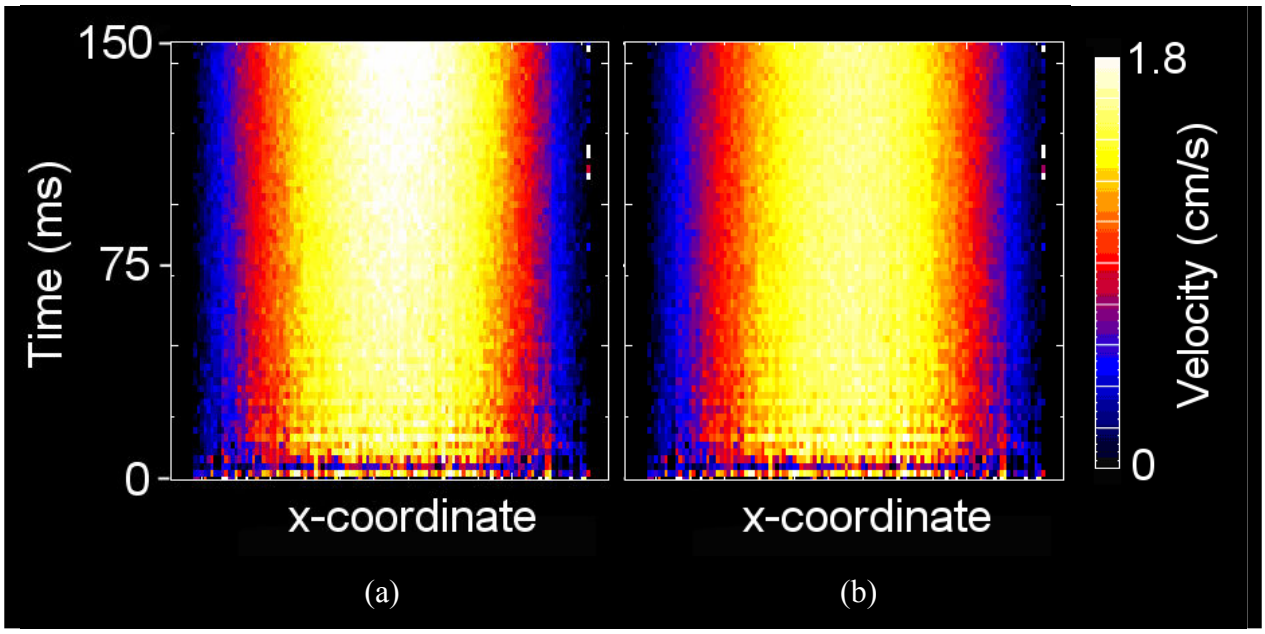


Figure 5.12 – (a) The data of Figure 5.9 is shown corrected for the time of each echo train, thereby obtaining a series of velocity profiles; (b) shows the data of (a) after correcting for the systematic increase in measured velocity with echo time, which occur due to assumptions in the pulse sequence design.

The strength of the phase gradient in the COTTAGE sequence must be chosen based on a number of factors. A low value of phase gradient will retain a good signal (since we will be acquiring data near the centre of **k**-space, where signal is strongest). However, if the fluid flow rate is small we may have low resolution of the phase-angle. Having a large phase gradient would give a larger phase angle to measure and therefore have less uncertainty in the measurement, but would also dephase the signal further and give noisier data. In practice, the optimum lies somewhere in between, and is found experimentally. Typically, one might consider taking two measurements; one with low and one with high gradients. This would allow us to resolve both high flow rates and low flow rates within the same system. The effect of altering the gradient strength is illustrated in Figure 5.13 – the images are taken at the same flow rate, but with low and high phase gradients respectively. Figure 5.13b has larger values of phase-angle, and the profiles at low echo times are considerably less noisy than for Figure 5.13a. However, we have massive foldover at long times where the fluid has displaced further.

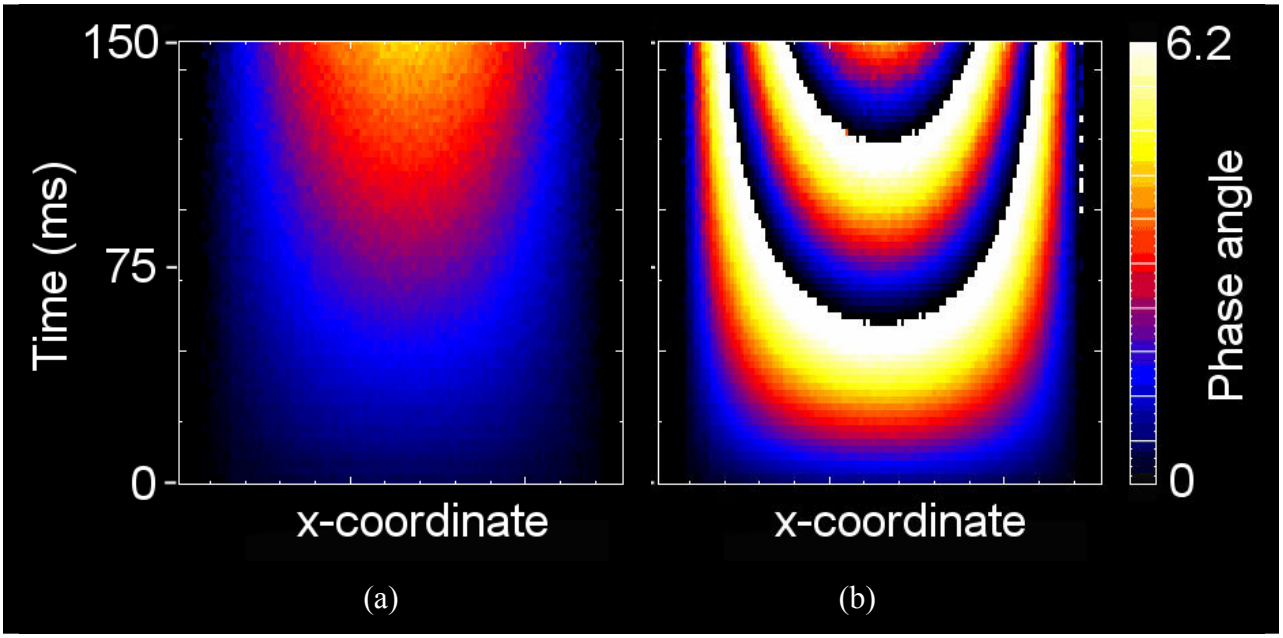


Figure 5.13 – An illustration of the COTTAGE sequence with (a) low-strength phase gradient (0.22 G/cm), and (b) higher strength phase gradient (1.77 G/cm), both at a flow rate of 3.5 l/hr.

The quantification of the COTTAGE technique was performed by selecting profiles from the data at a range of flow rates, and comparing this to the analytical solution for laminar flow in a pipe. The results are shown in Figure 5.14. It can be seen that the profiles are in good agreement, though the velocities are very slightly lower than expected.

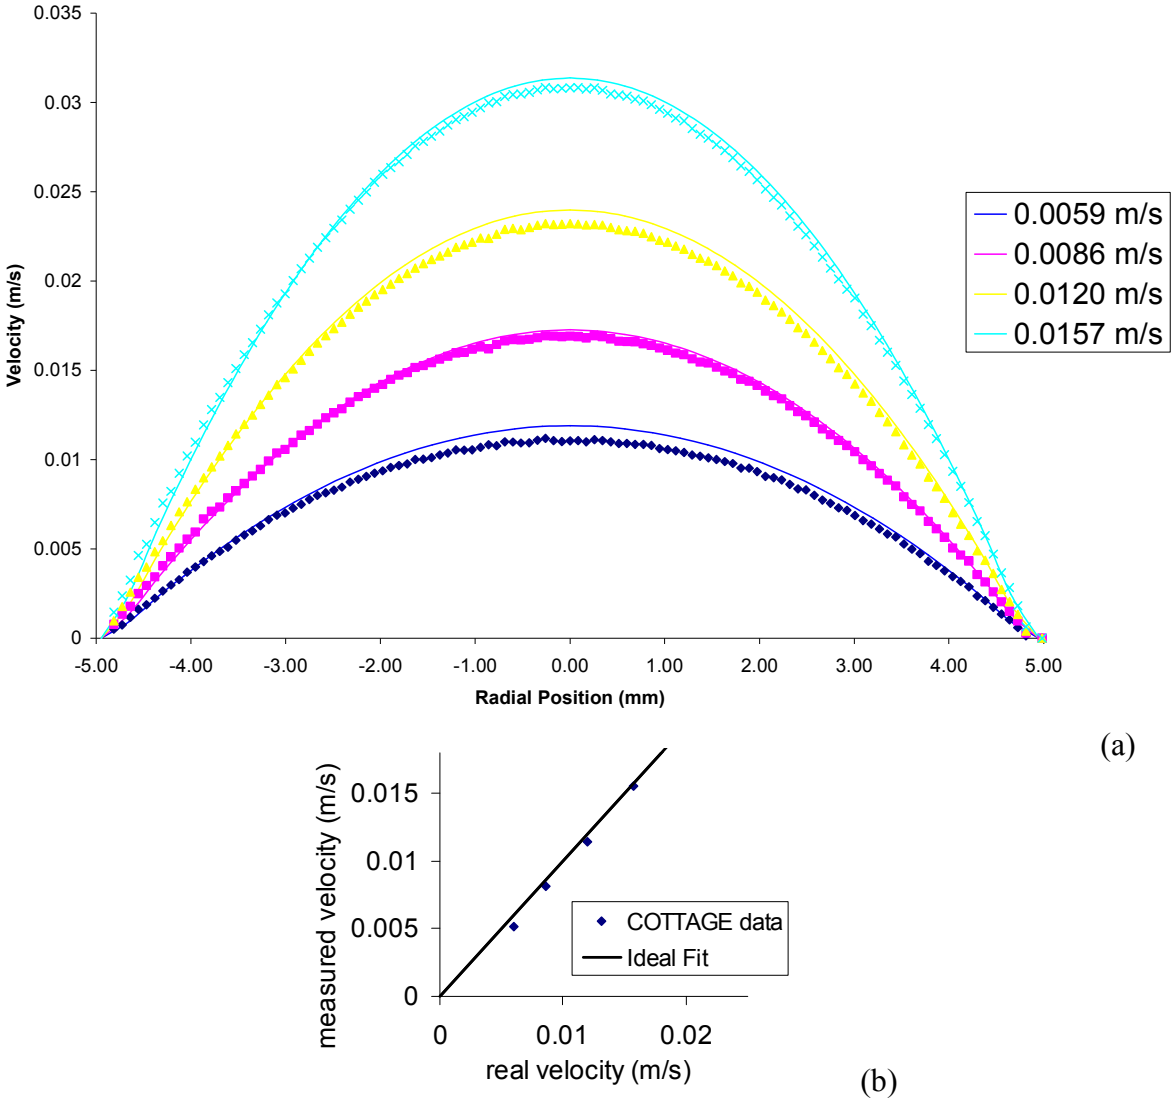


Figure 5.14 – Quantification of the COTTAGE technique; (a) a series of COTTAGE profiles plotted against the analytical solution for laminar flow in a pipe, and (b) a plot of measured total flow rate versus the actual flow rate through the system.

The COTTAGE technique can be applied to very high flow rates, the only limiting case being where the fluid has sufficient *x*- or *y*-velocity that it leaves the central refocusing slice, losing magnetisation. In practice, velocities of up to 30 cm/s have

been studied with this method, including turbulent flow. An example of a COTTAGE velocity experiment on turbulent pipe flow at a Reynolds number of 2500 is shown in Figure 5.15. It can be seen that there is a flattened profile (compared to that of laminar flow), that the flow is not steady with time, and that divergent fluctuations can be observed (such as the dip in axial velocity towards the top left of the image).

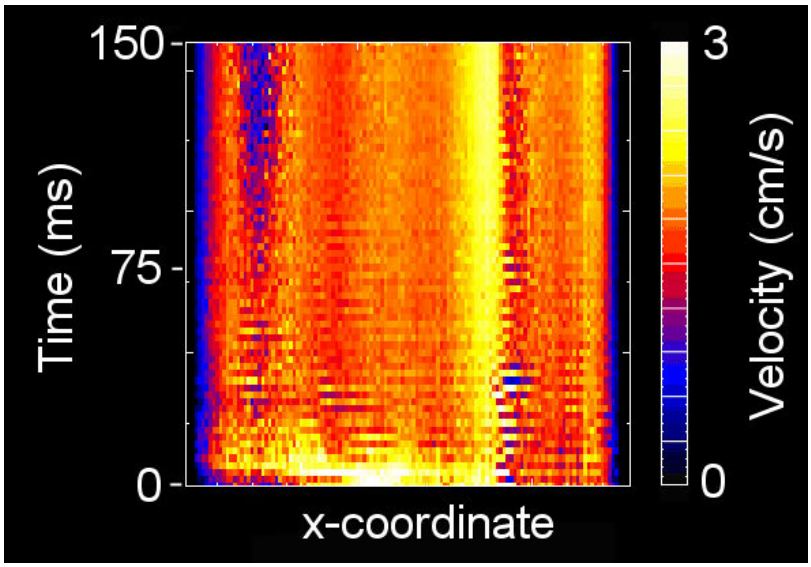


Figure 5.15 – An example of COTTAGE applied to turbulent flow in a 5 cm pipe at a Reynolds number of 2500, where the intensity corresponds to axial velocity.

## 5.5. Multi-slice cottage

The phase-encoded COTTAGE sequence was then further modified, with the purpose of acquiring 2D velocity maps. Instead of acquiring a series of 1D profiles from the same line of spins over time, this method moves the slice position across the imaging volume during the acquisition process, thereby scanning a 2D image by a series of 1D profiles (see Figure 5.16). When the refocusing slice has scanned across the cross-section, we may return to the beginning and begin another pass. This can continue (in the manner of the SEMI-RARE sequence) until the magnetisation has fully decayed. This technique is capable of measuring 2D axial-velocity maps in 80 ms, and four successive velocity maps over 320 ms (limited by the signal-to-noise of the system, which is dependent on the flow and relaxation properties of the system).

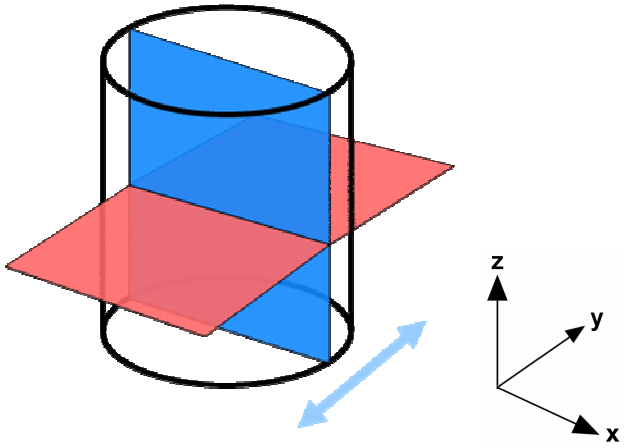


Figure 5.16 – Schematic of the orientation of the excitation plane (red) and read plane (blue). The arrow shows the movement of the read plane, corresponding to a multi-slice scanning setup.

To implement this sequence, we require that during the phase-encoding and read steps the only signal in the transverse plane is that of fluid within the current slice of interest. Therefore, we must store the remainder of the signal in the longitudinal plane. This is normally done with a stimulated echo (Section 2.1.7), which stores the magnetisation and phase information in *z*-storage and prevents dephasing. The pulse sequence is shown in Figure 5.17. An initial 90° soft pulse is used to excite the spins within a slice in the transverse plane, then a dephasing time is waited, and then a second 90° is used to put all the spins from the excitation slice into *z*-storage. This halts further dephasing of the spins. Then, we extract a row of spins of interest by applying a 90° pulse to bring these spins back into the transverse plane. The dephasing time is waited, and a stimulated echo is formed and acquired. Then the dephasing time is waited again and the spins are returned to *z*-storage by another 90° soft pulse. We then repeat the sequence for every slice in the cross-section, extracting, reading and storing the spins in turn.

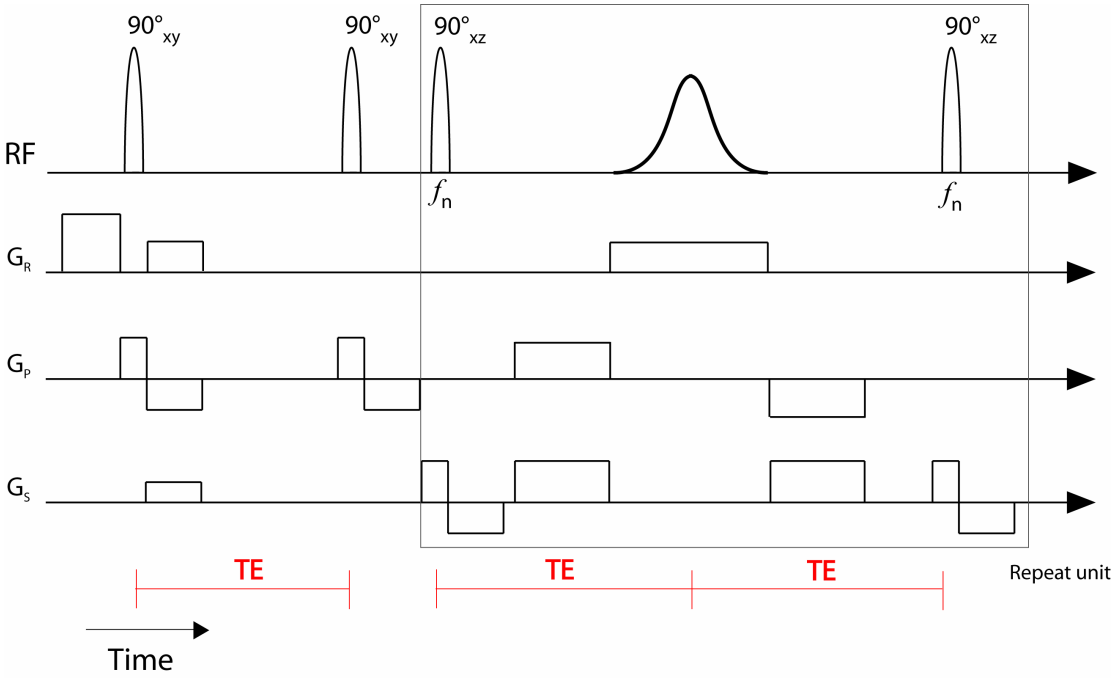


Figure 5.17 – Pulse sequence diagram of the multi-slice COTTAGE technique. The time associated with the stimulated echo is shown in red. Spins of interest are removed from *z*-storage with a 90° pulse with variable frequency offset, $f_n$. This offset is changed to move the position of the slice over the acquisition process.

To improve the pulse sequence, a novel method was implemented – that of using a double-pulse which would *simultaneously* return spins to *z*-storage for the current slice while retrieving those of the next slice. This elegant (and faster) sequence was implemented by creating a double-sinc3 pulse from two single sinc3 pulses[5] of bandwidth 9180 Hz, and an offset between the two pulses of 9180 Hz, as shown in Figure 5.18. An illustration of the double-sinc pulse scanning method is shown in Figure 5.19.

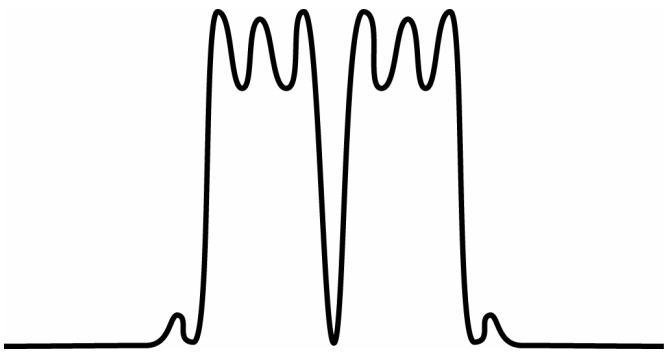

Figure 5.18 – Illustration of the double-sinc3 pulse, created by the convolution of two sinc pulses with different frequency offsets.

[5] See Appendix 6.9 for a description of the sinc3 pulse shape.

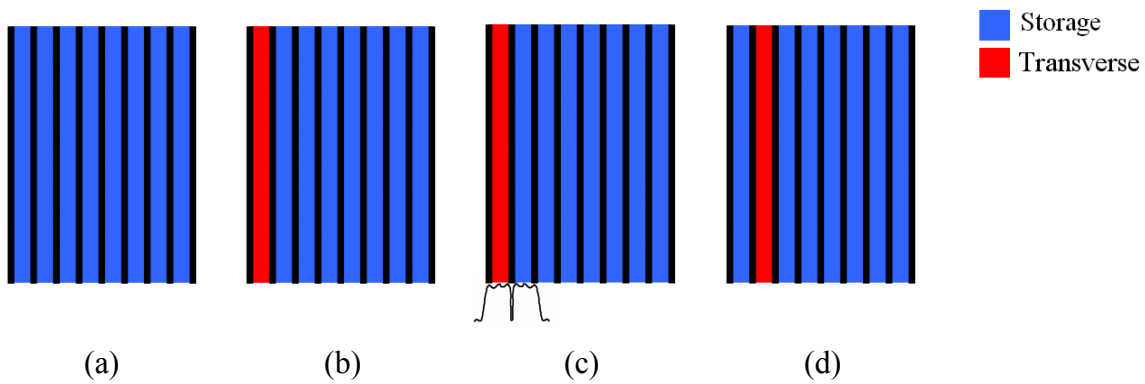


Figure 5.19 – Schematic of the double-sinc scanning method (a) the experiment begins with the excitation of a slice in the transverse plane, and these spins are all placed into *z*-storage; (b) then the first vertical slice is refocused (with a standard sinc pulse); (c) the double-sinc pulse is then applied to simultaneously place the excited spins back into storage, and to bring the spins of the next slice into the read plane; (d) this gives us the situation shown here, and repeated application of the double-sinc pulse will move the excitation slice across the cross-section.

The improved MS-COTTAGE sequence, utilising this double pulse, is shown in Figure 5.20.

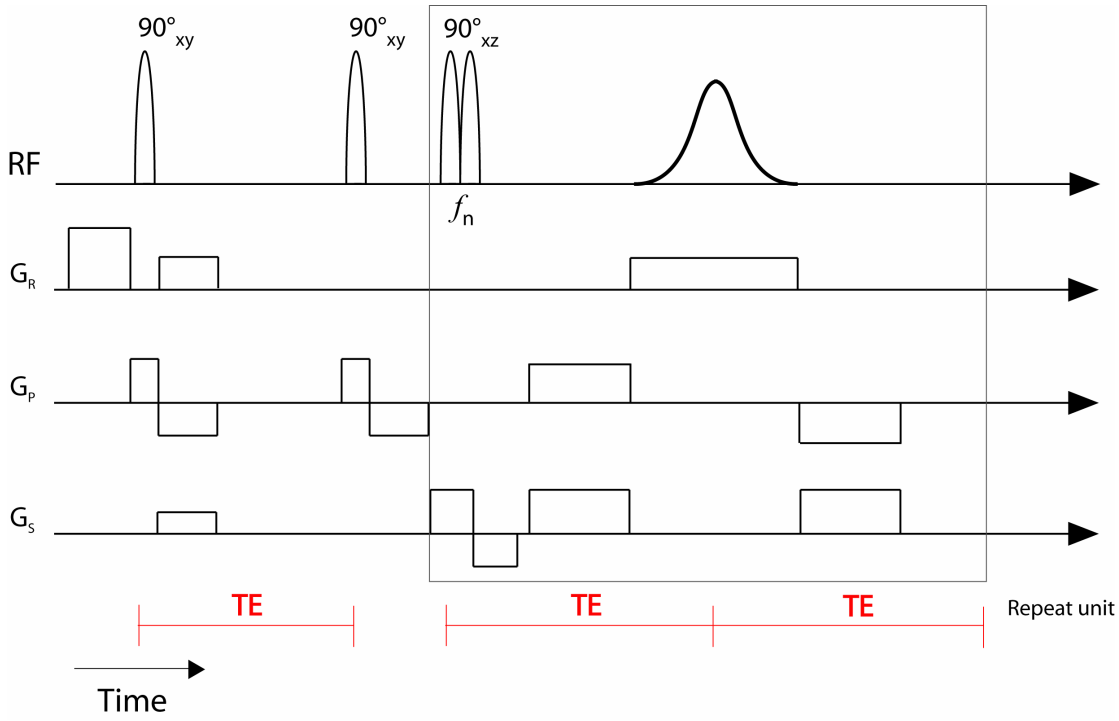


Figure 5.20 – Pulse sequence diagram of the improved MS-COTTAGE technique. The double-sinc pulse is used to replace the pair of sinc pulses in the original pulse sequence.

The multi-slice sequence is very sensitive to the powers of the excitation pulse and the refocusing double pulse. If the pulse angles are not 90°, signal will be lost due to the strong dephasing homospoil gradients in the sequence which are used to remove unwanted signal. Typically, we want the excitation pulse to be strong to give us as much signal as possible to measure, but if the refocusing double pulse is not exactly 90°, then the signal will decay rapidly as we scan across the monolith. The pulse strength is set manually to maximise the observed signal strength. It should be noted that increasing the excitation slice thickness is not necessarily beneficial, since if it is too thick then the phase of the spins may exceed a full revolution and give misleading averaged values.

As was mentioned in Section 2.3.1, the RARE pulse sequence (which COTTAGE is based upon) combines a spin echo and a stimulated echo which add constructively. However, in the COTTAGE sequence we are only interested in the stimulated echo; the stimulated echo corresponds to the echo of the spins we are currently manipulating, whereas the spin echo corresponds to the echo of spins from other slices. The solution to this was to purposefully dephase the contribution from the spin echoes – this leads to an overall decrease in the signal-to-noise of the experiment, but ensures that the signal acquired originates from the spins of interest.

To investigate the quantitative nature of this technique, the usual system of laminar flow in a pipe (5 cm i.d.) was again utilised. The resulting data are shown in Figure 5.17. The parabolic nature of the flow can clearly be seen, and the successive scans of the pipe measure the same velocity (though at lower signal-to-noise). The four scans are acquired in 309 ms. The field of view is 55 mm, and 31 slices are taken across the scanning direction. This gives a resolution of 430 μm × 1700 μm. The excitation and slice thicknesses were both 1.7 mm.

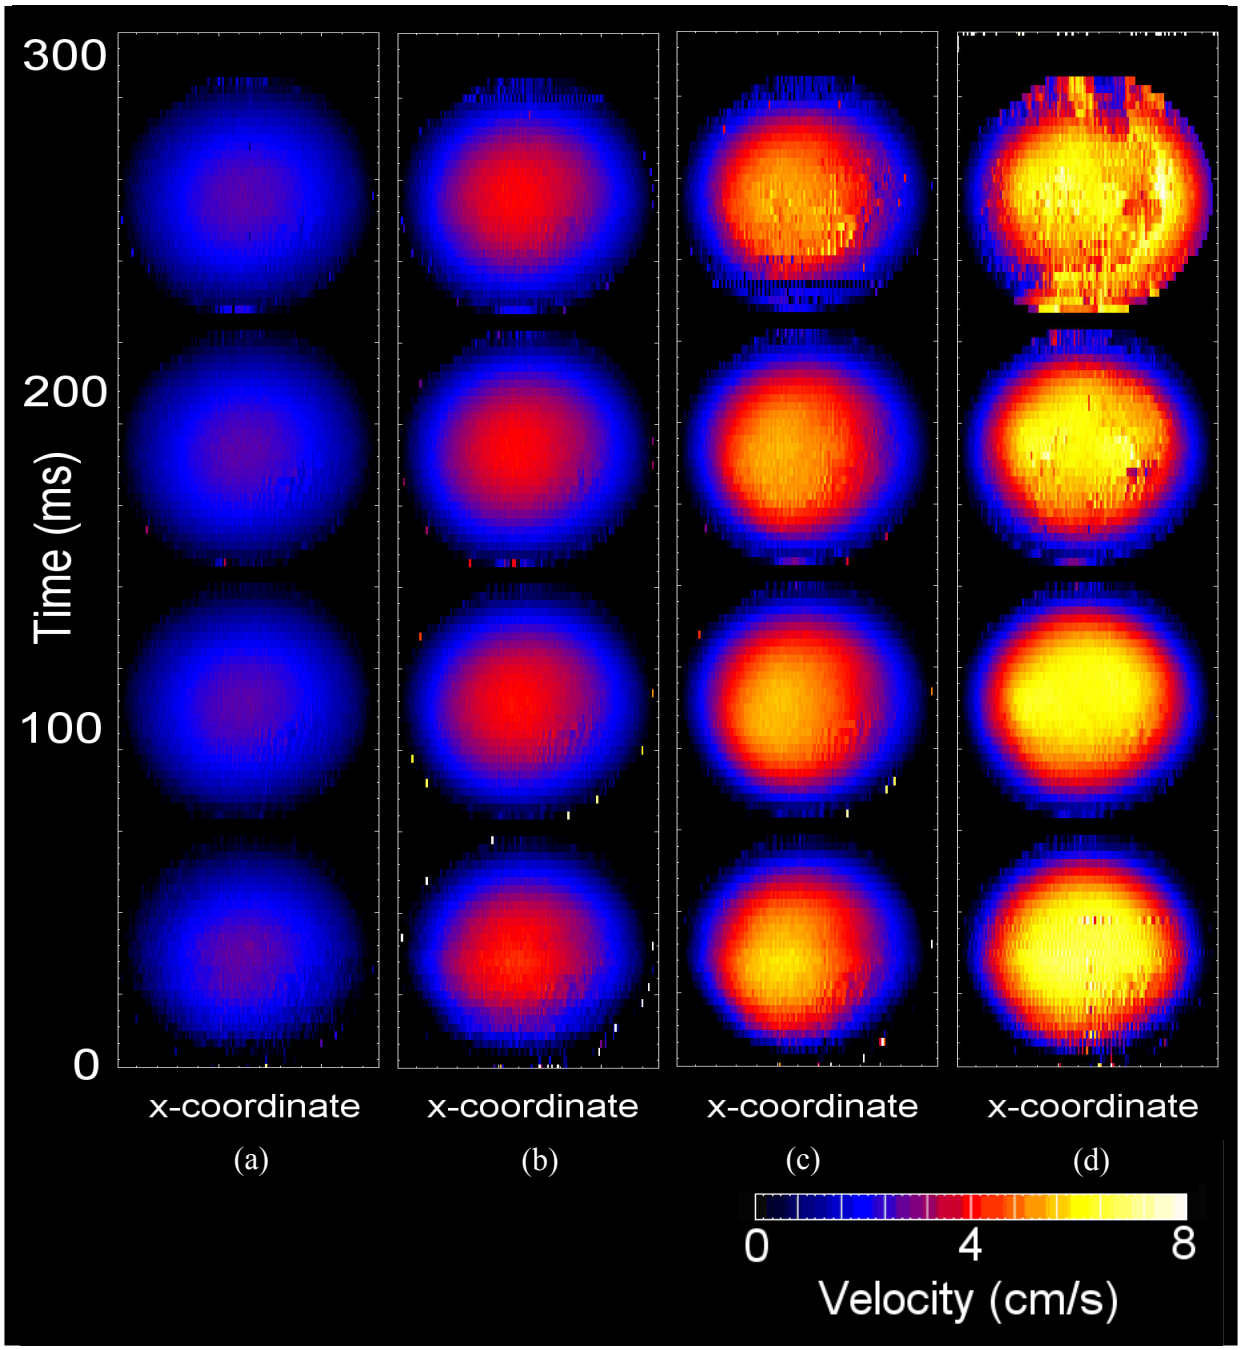


Figure 5.21 – Example of the multi-slice COTTAGE technique applied to laminar flow in a 5 cm pipe. The intensity corresponds to the velocity, and acquisition time is increasing in the y-axis. Flow rates are (a) 33.5 l/hr, (b) 50.0 l/hr, (c) 68.7 l/hr and (d) 84.8 l/hr.

Some extracted profiles from the above data are shown in Figure 5.22a, illustrating the quantitative nature of the data. The sequence shows some problems at higher flow rates where the laminar flow becomes unsteady. These non-quantitative profiles are believed to be due to incorrect phase-averaging occurring – in this particular experiment, data were signal-averaged *before* Fourier transforming, but this is not

appropriate for non-steady state flow. Subsequent data acquisitions were acquired and Fourier transformed, and *then* averaged together.

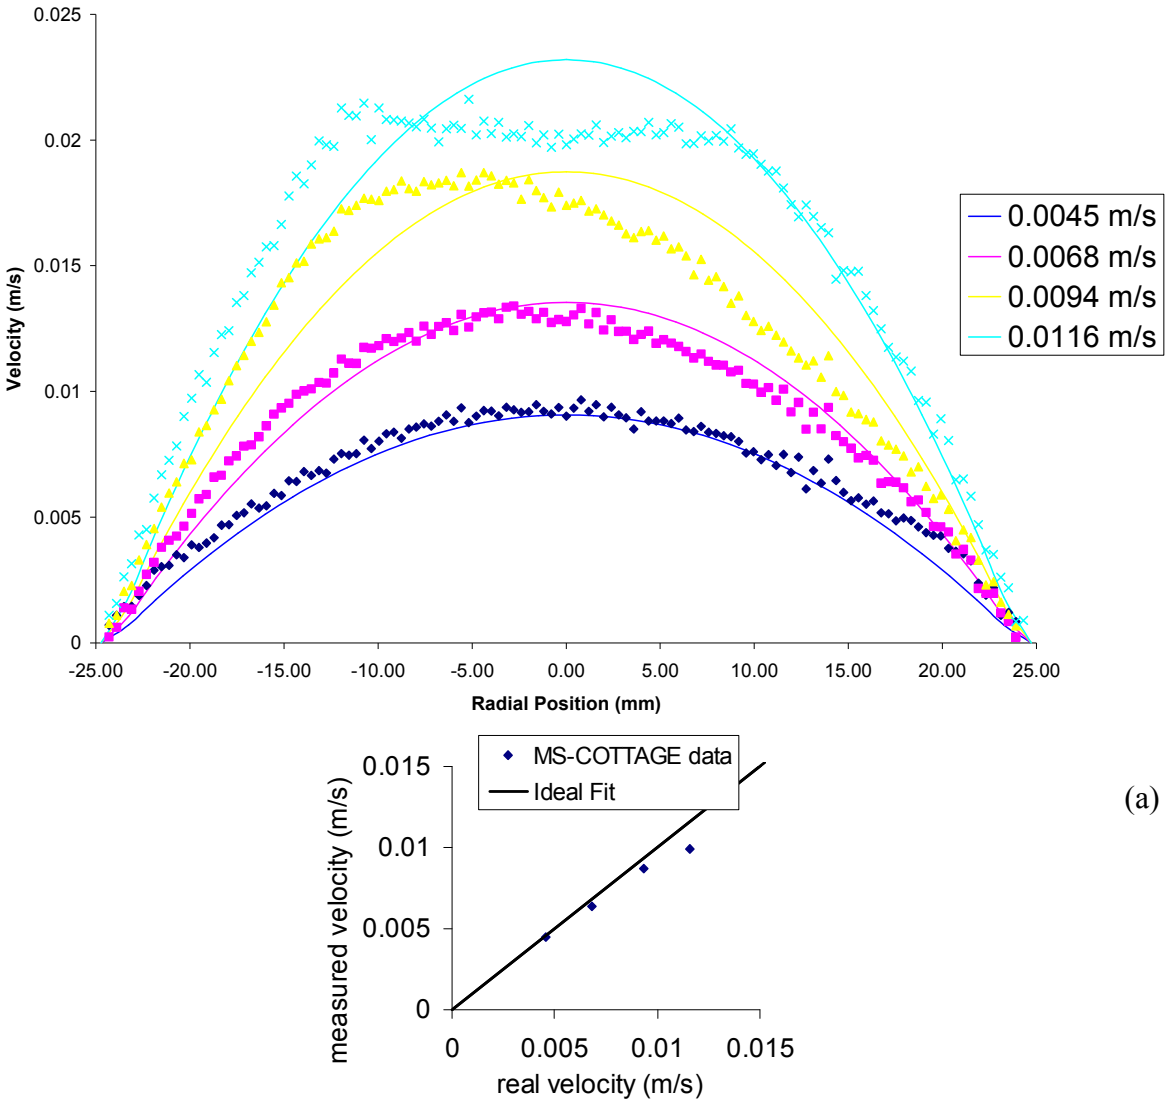


Figure 5.22 – Quantification of the MS-COTTAGE technique; (a) a series of MS-COTTAGE profiles plotted against the analytical solution for laminar flow in a pipe, and (b) a plot of measured total flow rate versus the actual flow rate through the system.

The next section will deal with studying unsteady flow using the MS-COTTAGE technique, and here the total flow across the pipe will be considered, not simply a slice through the profile. Some examples of MS-COTTAGE applied to unsteady flow are shown in Figure 5.23 at increasing Reynolds numbers. The transition from laminar parabolic flow to unsteady flow can be clearly seen. Figure 5.24 shows a series of acquisitions at a Reynolds number of 1660, and the onset of fingering and other turbulent features is observed.

### 5.5.1. An application of MS-COTTAGE

By using the same pipe flow setup, and increasing the flow rate past the steady limit (Eq. 5.1), we can investigate the applicability of MS-COTTAGE to unsteady flows. Data acquired at a range of increasing Reynolds numbers are shown in Figure 5.23. The non-parabolic nature of the flow can be seen, with clear velocity fluctuations at higher flow rates. MS-COTTAGE data are all acquired in 78 ms. The field of view is 55 mm × 55 mm, and 31 slices are taken across the scanning direction. This gives a resolution of 430 μm × 1700 μm. The excitation and slice thicknesses were both 1.7 mm. Data were acquired up to flow rates of Re 3300. Automatic phase-unwrapping (Section 7.7) was used where necessary to remove fold-over artefacts.

Above Re 1700, flow rates were quite unsteady, and above Re 2500 significant manual correction of the phase data was necessary. Moreover, to avoid signal-loss associated with sideways motion (which causes the dephasing of spins), the resolution of the acquisitions was lowered above Re 2500 from 31 slices to 15 slices – this halves the acquisition time, and reduces the extent of signal-loss due to sideways motion and allows data to be acquired up to Re 3300.

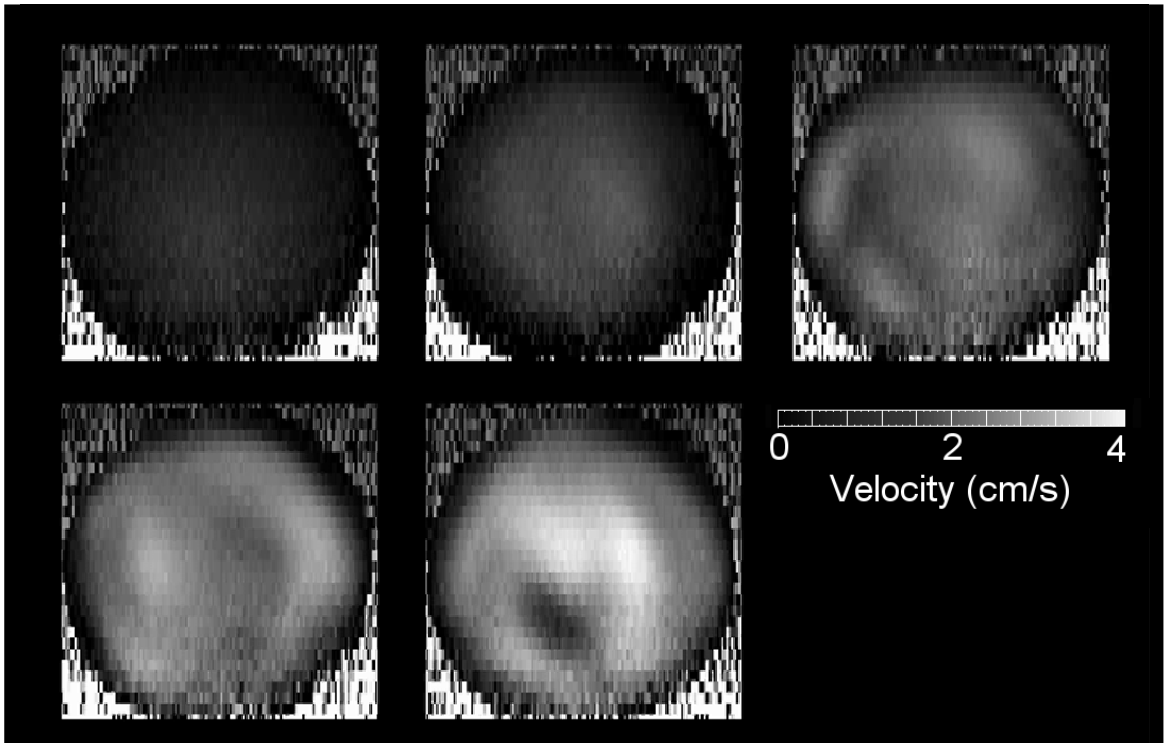


Figure 5.23 – Data from the MS-COTTAGE technique applied to unsteady laminar flow in a pipe at Reynolds numbers of (a) 250, (b) 500, (c) 760, (d) 980, and (e) 1140.

Figure 5.24 shows successive acquisitions at a Reynolds number (Re) of 1664. At this flow rate, evidence of fingering and the onset of turbulence can be observed.

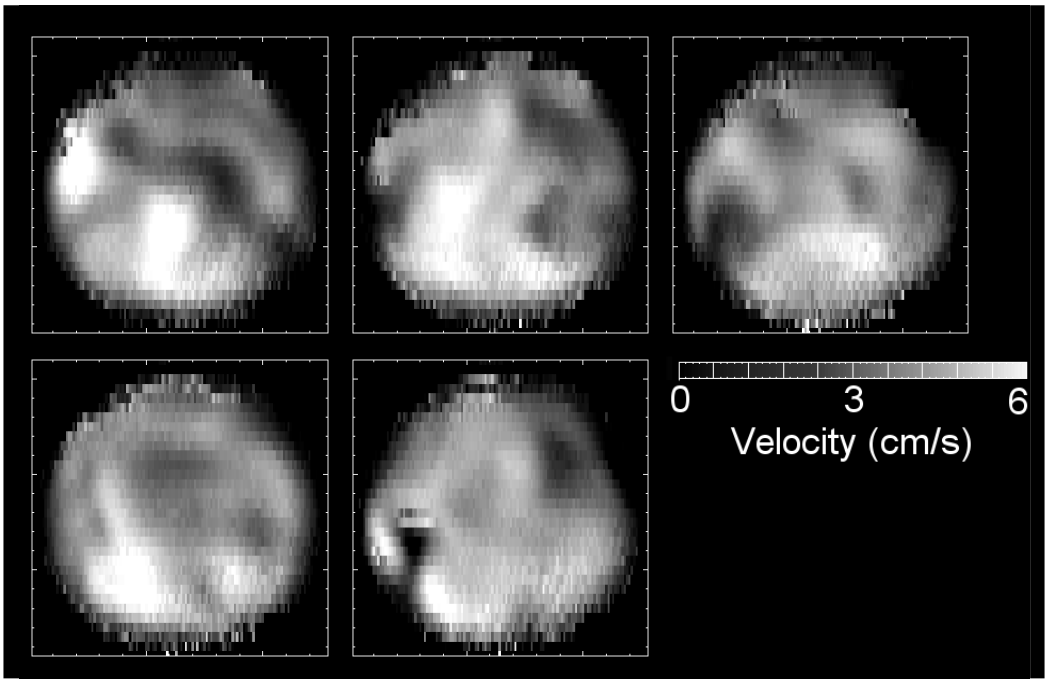


Figure 5.24 – A selection of MS-COTTAGE images at a Reynolds number of 1664. The beginnings of turbulent vortices and fingering may be observed.

By measuring the total flow rate over the cross section, we may compare the measured superficial velocity to the true superficial velocity, as shown in Figure 5.25. It is seen that the quantitative nature of the technique is excellent, and is still robust at Reynolds numbers up to 3300. This shows that the technique is quantitative even in the turbulent flow regime.

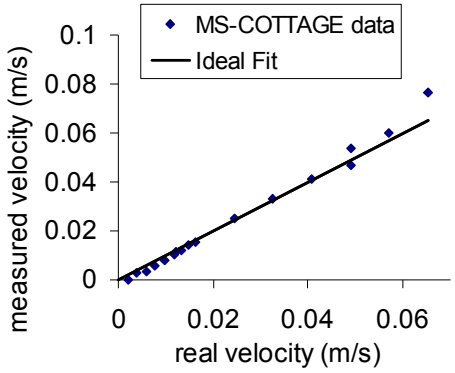


Figure 5.25 – Illustration of the quantitative nature of the MS-COTTAGE technique applied to unsteady flow, showing a plot of the MS-COTTAGE superficial velocity versus the actual superficial velocity through the system.

## 5.5.2. Discussion of the MS-COTTAGE technique

There are some limitations of the MS-COTTAGE technique which need to be discussed, and they will now be addressed, with some examples for illustration of the problems encountered.

The MS-COTTAGE technique begins to have problems when the *x*- and *y*-velocities become significant, since the assumption of axial flow no longer holds (leading to incorrect measurement of velocity, as the velocity is no longer proportional to the *z*-displacement). This lateral mixing has two effects – signal is lost due to fluid moving to a position where the double-sinc pulse no longer correctly retrieves / stores the spins. These spins are therefore dephased, and signal is lost. Figure 5.26 shows an example of the velocity and corresponding intensity images at a flow rate of Re 1664, and it can be seen that the regions of largest velocity heterogeneity map onto the regions of greatest signal loss. This effect becomes more significant at higher flow rates.

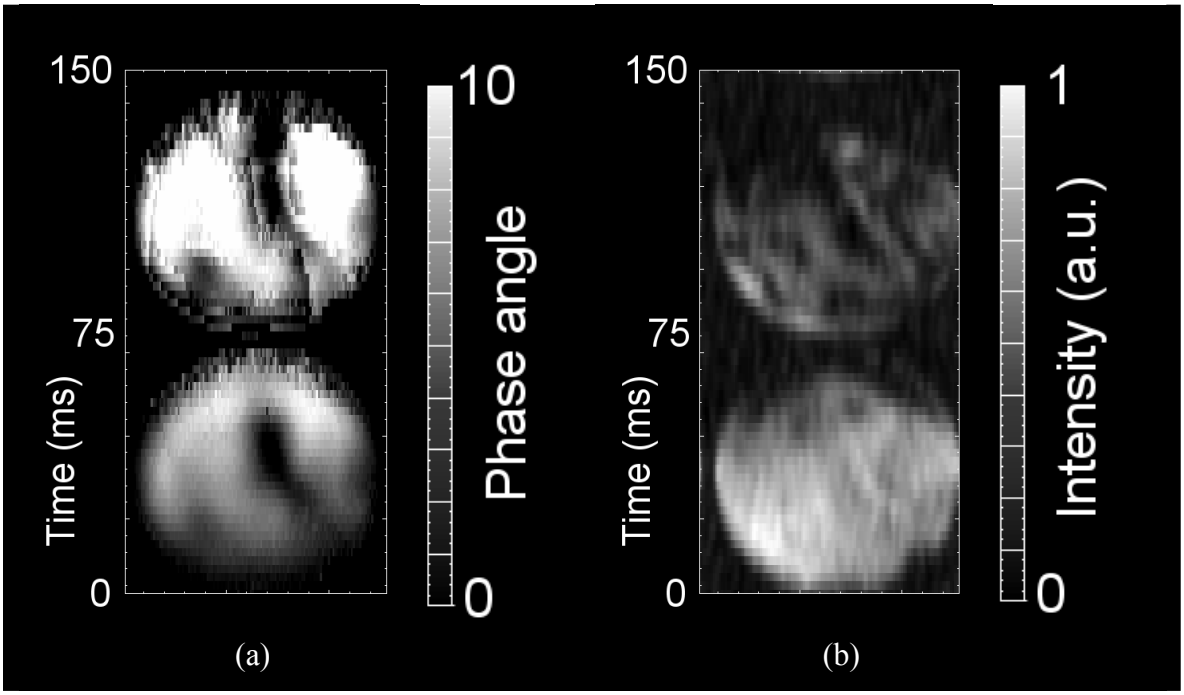


Figure 5.26 – Illustration of loss of intensity due to sideways motion (a) shows the velocity image at a Reynolds number of 1660, and (b) is the corresponding intensity image.

The second artefact effect we observe is that incorrect phase averaging occurs; since spins move significantly over the acquisition time, they will be mixing with spins of different phases. Where two spins with a phase value of, say, 5 and 7 find themselves within the same voxel, the result would be expected to be an average of 6. However, since the spin of phase 7 will in fact wrap around to $7 - 2\pi$, the resulting average is in fact $6 - \pi$. Hence, it is possible to observe phase-wrap artefacts within an image, whose value does not differ by $2\pi$ but by $\pi$. This is illustrated in Figure 5.27. The data can still be accurately phase corrected by adding $\pi$ to the corresponding region.

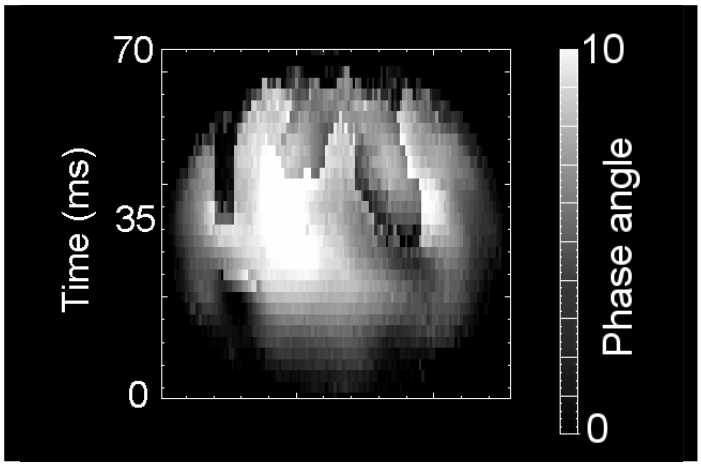


Figure 5.27 – Illustration of a phase-averaging artefact. The Reynolds number here is 2080, and there is significant sideways motion at this velocity. The fingering region contains values close to the $2\pi$ foldover limit, and they differ from their neighbouring voxels by a value of $\pi$, not a value of $2\pi$ as with a normal phase-wrap.

However, there are very few existing MRI techniques which are able to study detail in such high-velocity, turbulent systems (the most notable of which is GERVAIS, described in Section 5.6), and so MS-COTTAGE constitutes a significant contribution to the field, despite its limitations.

## 5.6. GERVAIS

GERVAIS (Gradient Echo Rapid Velocity and Acceleration Imaging Sequence) is a modification of the EPI sequence, developed by Sederman *et al.* (2004), and is shown in Figure 5.28. It includes additional phase-encoding gradients which are designed to obtain 3D velocity and acceleration images. The indicated region is repeated for each

velocity or acceleration component required. The pairs of phase-encoding gradients are placed in different directions during each repeat unit – for example, we might cycle the encoding gradients X, Y, Z and the three echo trains we acquire would correspond to the X, Y and Z velocity components. If we cycle them X, Y, Z, Z then we will obtain three velocity components, and the fourth echo train will be an acceleration image. GERVAIS has the advantages over RARE of removing the need for a phase-rewinding step, thereby saving on imaging time, and also removing the train of 180° RF pulses which create unwanted echoes. It does however retain all the same imaging artefacts as EPI, but is very fast and so is useful for dynamic velocity imaging. It has been shown to be able to study turbulent pipe flows up to Re 6000. Velocities are calculated by the method described in Eq. 2.26 (Section 2.4.1).

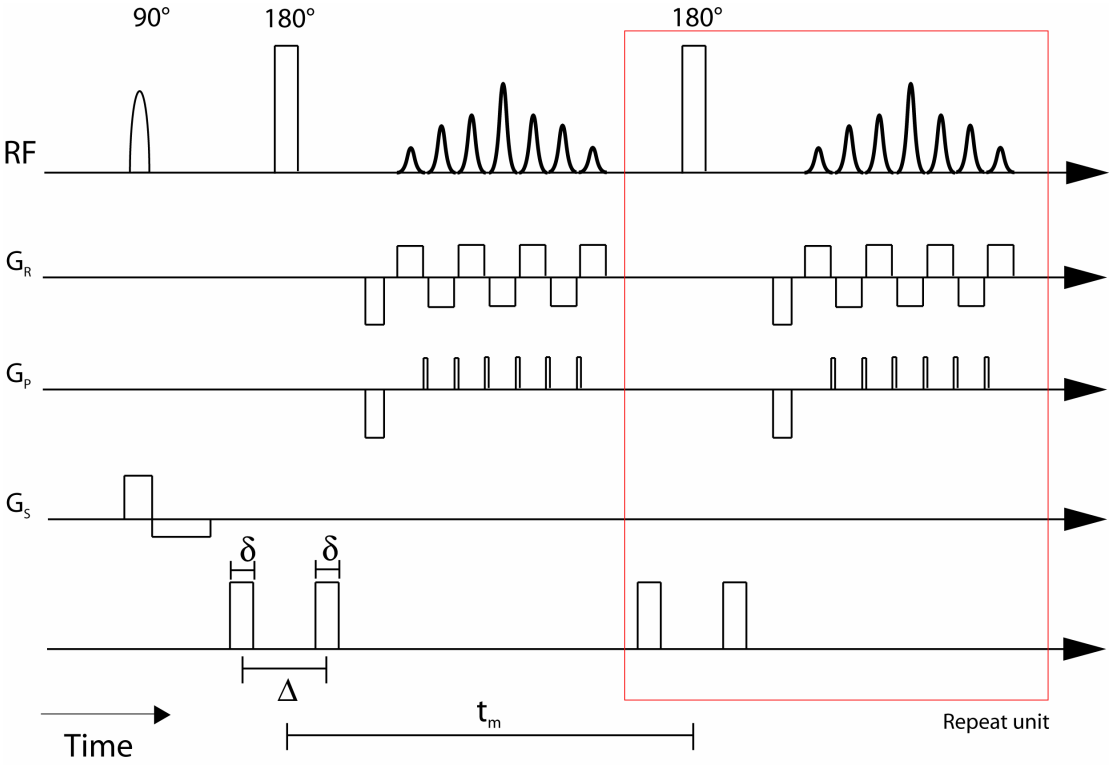


Figure 5.28 – The GERVAIS pulse sequence diagram. It is based upon the EPI sequence (Section 2.3.2) but with additional phase-encoding gradients, and the sequence is repeated for each velocity component.

GERVAIS data were acquired for flow in a 5 cm i.d. pipe at a range of flow rates up to Re 3300. The field of view was 7 cm, giving a resolution of 1094 μm × 1094 μm.

The *z*-velocity components were acquired in 20 ms and a slice thickness of 1.7 mm. A sample data set is shown in Figure 5.29 for a flow of 80 l/hr.

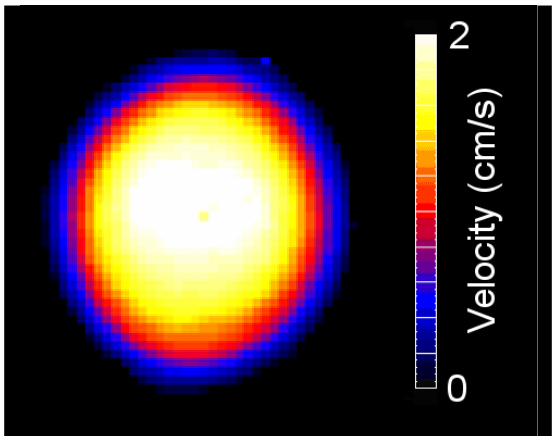


Figure 5.29 – Example of the GERVAIS technique applied to laminar flow in a pipe, at 80 l/hr. The data were acquired in 20 ms.

Profiles through the centreline of the data at increasing flow rate are shown in Figure 5.30, illustrating good agreement with the analytical solution for laminar pipe flow. Further quantification is shown in Figure 5.31, which shows the measured flow rate through the pipe cross section as measured with MRI, compared to the real flow rate through the system.

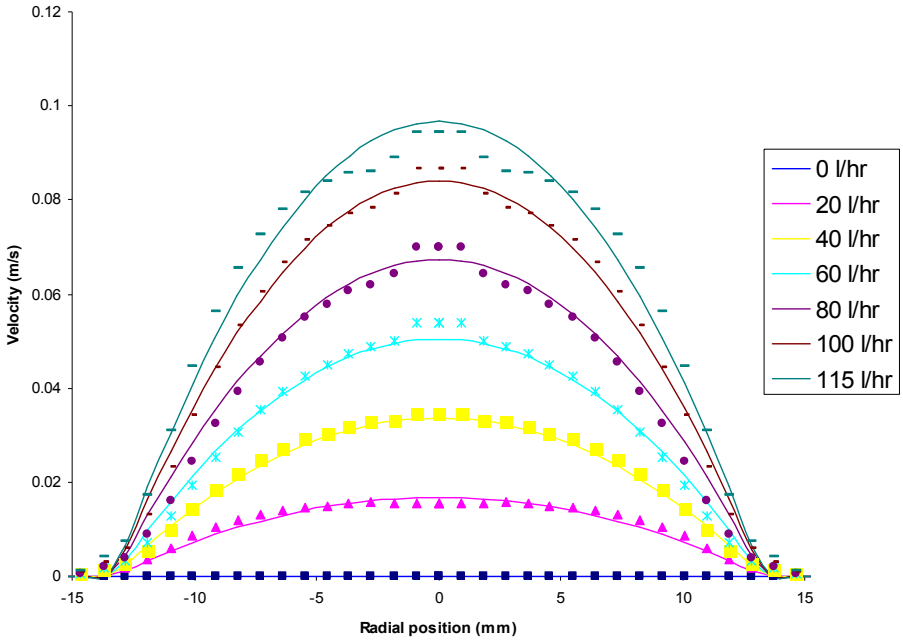


Figure 5.30 – Comparison of the analytical solution for the centreline velocity, with that measured by the GERVAIS pulse sequence.

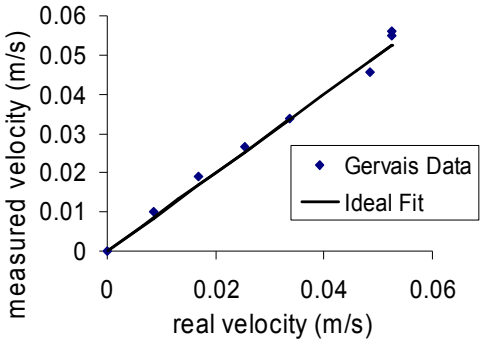


Figure 5.31 – Graph illustrating the quantitative nature of the GERVAIS sequences, showing the comparison of measured velocity versus the true superficial velocity.

GERVAIS is capable of measuring the dynamic flow fields present in turbulence and has been applied to turbulent flow to give the visualisation of turbulent eddies with MRI. Some example data acquired by Dr Sains (in the Department of Chemical Engineering at the University of Cambridge) are shown in Figure 5.32 for *x*-*y* velocity components, and Figure 5.33 for *x*, *y* and *z* components. Turbulence is an important phenomenon in chemical engineering, as well as other disciplines. Empirical descriptions exist but accurate theories elude us, partly because until relatively recently it has not been possible to image the flow fields in detail.

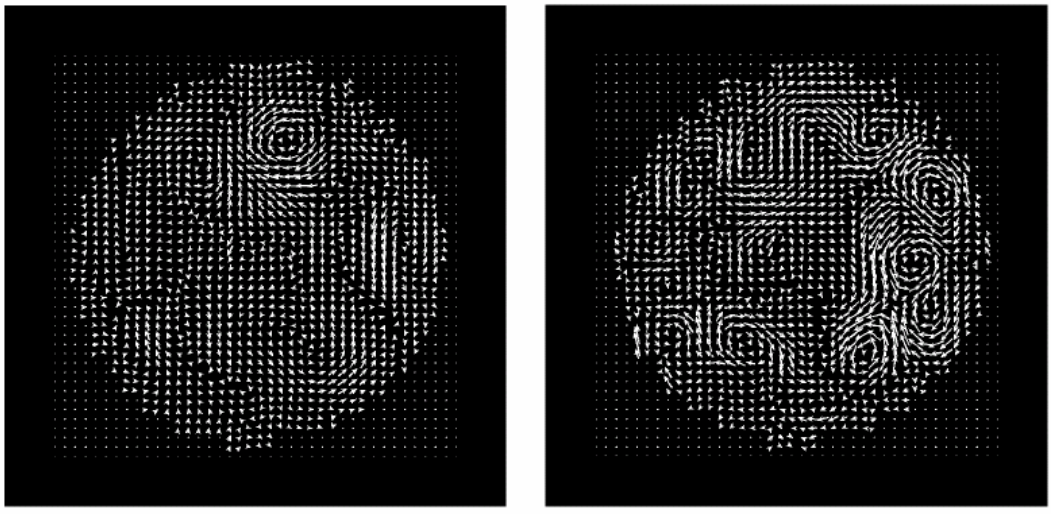


Figure 5.32 – Example of GERVAIS data, showing x-y vectors from a turbulent imaging experiment of water in a pipe. The presence of vortices can be clearly seen.

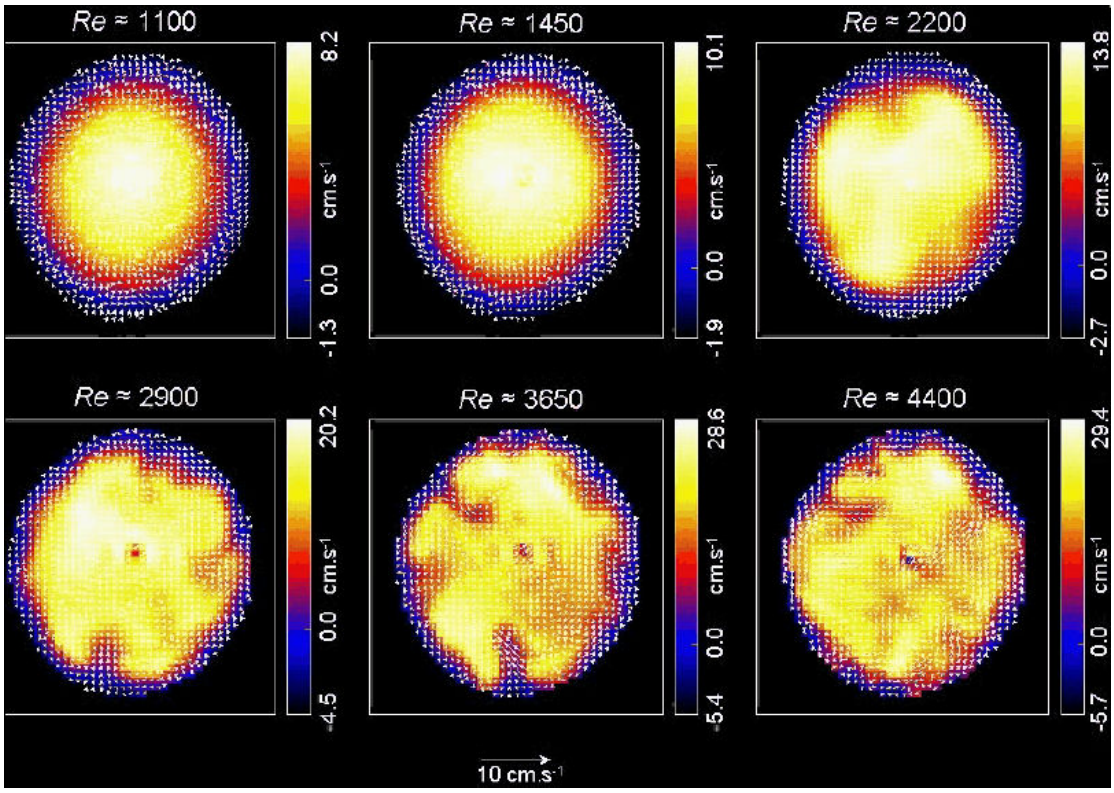


Figure 5.33 – A selection of 3-component velocity images. The *x*-*y* vectors are shown by arrows, and the *z*-velocity is denoted by the colour scale. The velocity of the water was increased, showing the transition from laminar flow, through unsteady flow, to fully turbulent flow at high Reynolds numbers.

GERVAIS has the advantage that it can capture all three orthogonal velocity components, rather than just the axial velocity component of MS-COTTAGE. It also retains very good signal-to-noise even when studying the turbulent flow regime. However, the MS-COTTAGE technique is considerably more robust to susceptibility effects, and is expected to be quantitative in the monolith system, whereas GERVAIS cannot be quantified in such systems (as will be shown in Section 6.6).

Since the GERVAIS and MS-COTTAGE techniques are both quantitative, two-dimensional, rapid, velocity-measurement methods, the following section will show a brief comparison between the data acquired with these two methods.

## 5.7. A comparison of GERVAIS and MS-COTTAGE

This section will apply an image processing technique from Chapter 7 to compare the results obtained with the MS-COTTAGE and GERVAIS velocity imaging techniques. We have already established that the two MRI techniques are quantitative; we have shown that the average superficial flow rate is consistent with the true flow rate, and that the central profile of the velocity map corresponds quantitatively with laminar flow in a pipe. In this section, the data are to be compared to identify whether they contain similar features within the images (degree of heterogeneity, instability, size of features). It is expected that the trends should be similar for both types of data, though some differences are expected due to the different natures of the pulse sequences, and their different acquisitions times – 20 ms for GERVAIS and 80 ms for MS-COTTAGE – therefore the degree of time-averaging of the features will vary between the data types.

The data to be analysed were acquired on the same system (water flowing within a 5 cm i.d. pipe), at a range of flow rates between Reynolds numbers of zero and 3300. One hundred repetitions were acquired at each flow rate in order to obtain statistical information on the fluctuations in the data. The MS-COTTAGE data was acquired at higher resolution than the GERVAIS data, so for a fair comparison the MS-COTTAGE data was interpolated to match the lower resolution. All the other MRI parameters were kept constant between the two techniques. It is not expected that a sharp transition between laminar and turbulent flow would be observed, primarily due to the fact that insufficient entry length was available for true laminar flow to fully develop.

In order to quantify the features of the texture we observe, the method of Section 7.5.2.5 was implemented to identify the characteristic texture length present within an image. For this method, the custom autocorrelation function with the subtractive kernel was used, and the results of the analysis are shown in Figure 5.34. It can be seen that the trend in each case is that the characteristic radius identified decreases with Reynolds number, which would be expected as we transition from laminar flow, to unsteady flow, and finally to turbulent flow where the flow contains more local fluctuations and hence the characteristic features become smaller. The

results of the two data sets agree well, with the MS-COTTAGE being slightly lower in all cases. However, the quantitative nature of the autocorrelation function has not been investigated, so the actual numbers are not necessarily comparable. The trends are quite clear in both cases.

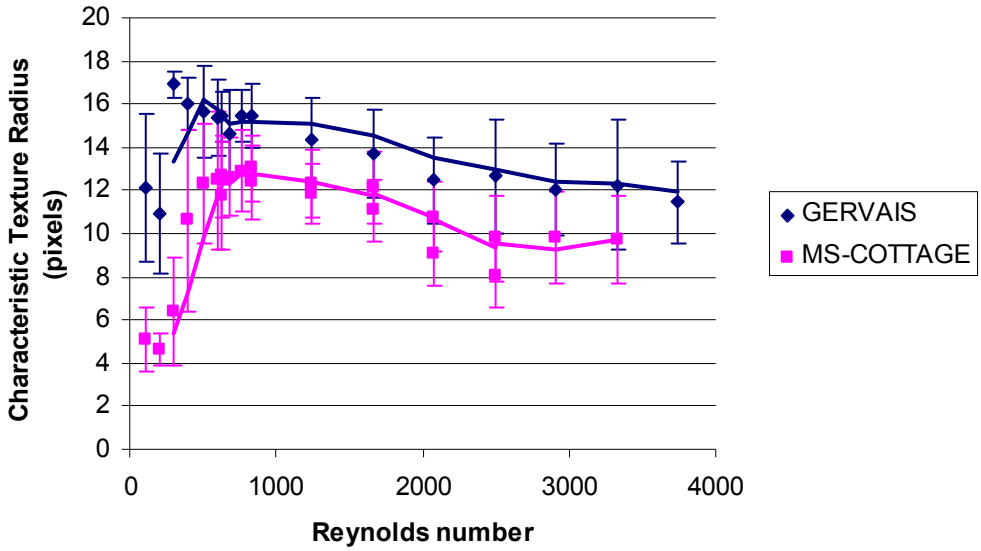


Figure 5.34 – Comparison of the results of the characteristic texture length analysis applied to the GERVAIS and MS-COTTAGE data. The data points are the average of 100 velocity maps, and the standard deviation of these measurements are shown as bars to give an indication of the range of variation present.

## 5.8. Conclusions

This chapter has introduced several novel MRI pulse sequences. SEMI-RARE and GERVAIS were developed by Dr Sederman in the Department of Chemical Engineering at the University of Cambridge, and the line-excitation, COTTAGE and MS-COTTAGE sequences were developed as a part of this work. Each of the techniques were introduced and their principles described in turn, and all have been shown to be quantitative in nature regarding flow in a pipe.

COTTAGE is a 1D velocity technique with rapid time resolution (2 ms), and MS-COTTAGE extends this to 2D velocity imaging in around 80 ms, which is a significant advance in MRI techniques. It is not as fast as GERVAIS, but has

advantages such as better resistance to magnetic susceptibility artefacts, and so should hopefully be of more use in porous systems, such as the monolith reactor.

Having exhausted the application of these novel MRI techniques to flow in a pipe, the following chapter will illustrate their application to the monolith system in single phase flow, and show the extent to which we may quantify the novel MRI techniques within this porous system.

## 5.9. Nomenclature

| **Symbols** | **Definition** |
|---|---|
| $D$ | Diameter |
| $f_n$ | Variable frequency offset of a soft pulse |
| G | Amplitude of phase gradient |
| **k** | Reciprocal space vector |
| $L_{99\%}$ | Entrance length at which laminar flow deviates by less than 1% from a parabolic profile |
| $m, n$ | Repetition counters |
| $NE$ | Echo number |
| Re | Reynolds number |
| $s$ | Displacement |
| $t_e$, $TE$ | Echo time |
| $T_2$ | Spin-spin relaxation constant |
| $v$, $V$ | Velocity |
| | |
| $\Delta$ | Time between phase gradient pairs |
| $\phi$ | Phase angle |
| $\delta$ | Duration of phase gradient |
| $\gamma$ | Gyromagnetic ratio |

| **Abbreviations** | **Definition** |
|---|---|
| COTTAGE | Continuous translation tracking with alternating gradient encode |
| cpsi | Cells per square inch |
| EPI | Echo-planar imaging |
| FOV | Field of view |
| GERVAIS | Gradient echo rapid velocity and acceleration imaging sequence |
| MRI | Magnetic resonance imaging |
| MS-COTTAGE | Multi-slice COTTAGE |
| NMR | Nuclear magnetic resonance |
| RARE | Rapid acquisition with relaxation enhancement |
| RF | Radio frequency |
| SEMI-RARE | Single excitation, multiple image RARE |
| TOF | Time of flight |

# 5.10.References

Christiansen, E. B. and H. E. Lemmon (1965). "Entrance region flow." Aiche Journal **11**(6): 995-999.

Mantle, M. D., A. J. Sederman*, et al.* (2002). "Dynamic MRI visualization of two-phase flow in a ceramic monolith." Aiche Journal **48**(4): 909-912.

Morris, G. A. and R. Freeman (1978). "Selective Excitation in Fourier-Transform Nuclear Magnetic-Resonance." Journal of Magnetic Resonance **29**(3): 433-462.

Sederman, A. J., M. D. Mantle*, et al.* (2003). "Single excitation multiple image RARE (SEMI-RARE): ultra-fast imaging of static and flowing systems." Journal of Magnetic Resonance **161**(1): 15-24.

Sederman, A. J., M. D. Mantle*, et al.* (2004). "MRI technique for measurement of velocity vectors, acceleration, and autocorrelation functions in turbulent flow." Journal of Magnetic Resonance **166**(2): 182-189.

# Chapter 6 – Single-phase Flow in a Monolith

This chapter will apply the novel MRI techniques described in Chapter 5 to single-phase flow in a monolith. The quantitative nature of the techniques will be investigated in this more complex, porous environment. Results will be presented and discussed.

**CHAPTER 6 – SINGLE-PHASE FLOW IN A MONOLITH**



## 6.1. Introduction

This chapter will investigate the performance of the novel MRI techniques when applied to single-phase flow in a monolith. This is a necessary step before attempting to measure two-phase flow, as it will show whether the techniques are still quantitative when applied to porous media before adding the further complication of trying to measure dynamic, two-phase systems.

## 6.2. Flow configuration

A flow loop consisting of a Perspex column, of 50 mm i.d., containing a cylindrical, square-channel monolith was positioned within the vertical bore of the magnet, as shown in Figure 6.1. A monolith of diameter 42 mm and length 0.15 m rated at 200 channels per square inch (cpsi) was used, with each channel having square cross-

sections of side 1.7 mm (Figure 6.2). The monolith was held in place with a rubber 'O'-ring and a Teflon wrapping which ensured that no fluid bypassed the monolith at the sides.

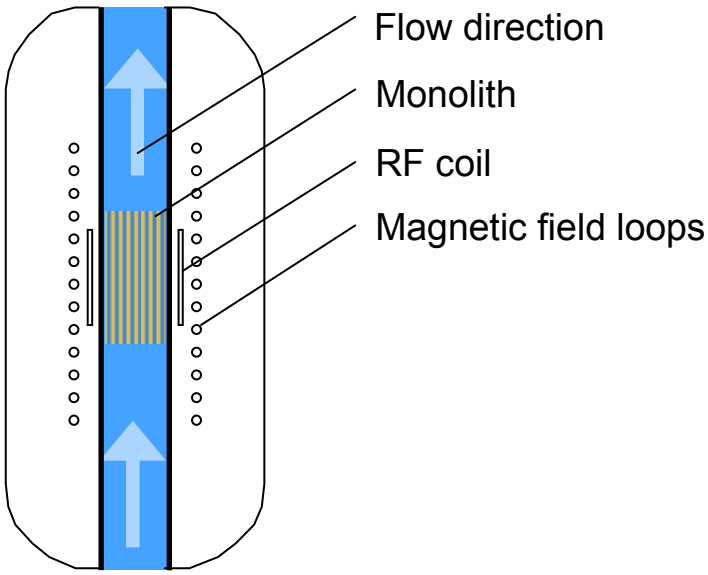


Figure 6.1 – Schematic of the setup for single-phase liquid flow through the monolith.

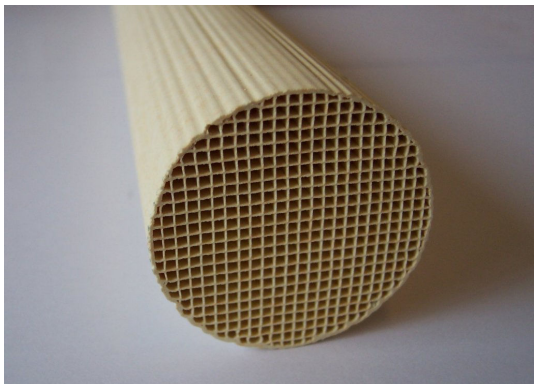

Figure 6.2 – Photograph of the 200 cpsi monolith section used in this work.

For the case of single-phase flow, the flow direction was chosen to be upwards for two reasons; (i) to use gravity to even out velocity fluctuations, and (ii) to ensure that there were no trapped gas cavities within the flow loop, since air would be pushed upwards and out of the experimental rig. This choice is further justified in Appendix 6.8.

Each of the novel MRI sequences will now be applied to the monolith in turn, to investigate the extent to which the pulse sequences are still quantitative in such a porous system.

## 6.3. Line-excitation

To quantify the line-excitation method (Section 5.3) in the monolith system, data were acquired for single-phase flow in a 200 cpsi monolith. The field of view was 35 mm in the direction of flow and 51 mm orthogonally, with resolutions of 547 μm and 398 μm respectively. The image slice thickness was taken as 1.7 mm, corresponding to the channel thicknesses. Sinc3 pulses were used for the excitation and refocusing pulses (see Appendix 6.9). The acquisition time was 156 ms. Acquisitions were improved significantly by homospoiling to remove unwanted signal. This reduced ghosts and phase artefacts, and aided the detection of signal in otherwise noisy images. Homospoils were used in both the read and slice directions, and it was found that there was an optimum strength – too low, and they would not quell unwanted signal; too high and they could begin to act as velocity weighting gradients, and also reduce the strength of the signal acquired. The direction of superficial flow was chosen to be in the phase direction so as to minimise flow artefacts in the images, and the read gradient was chosen to be in the *x*- direction since the read direction is not flow-compensated.

Data from a slice through the central row of channels in a 200 cpsi monolith were acquired at a range of liquid flow rates from 0-115 l/hr, and an example dataset is shown in Figure 6.3a. At the edges of the data, the stagnant zone can be observed; this region was not included in the analysis but was a useful reference feature for measuring the displacement of moving spins. The average displacement at each radial position was then calculated by finding the position of maximum intensity (which should correspond to the mean displacement at half of the acquisition time). Performing this analysis leads to the plot of Figure 6.3b.

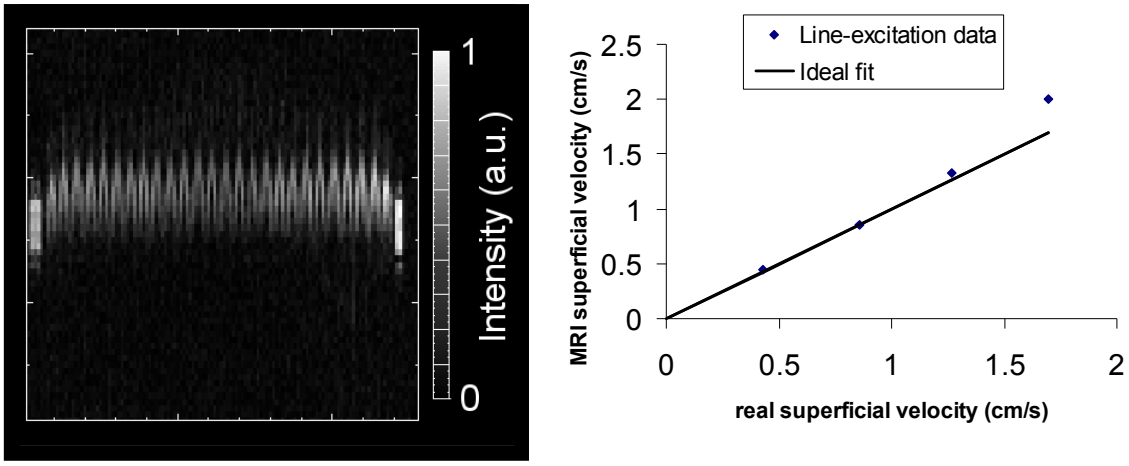


(a) (b)

Figure 6.3 – Quantification of the line-excitation method; (a) shows a sample experimental result for single-phase flow at 115 l/hr through a 200 cpsi monolith; (b) shows the measured superficial fluid velocity versus the real superficial velocity through the system.

It can be seen that the average superficial velocity as measured by this method compares very well with the true superficial velocity, as measured by the liquid flow supplied to the system. There is a deviation at higher flow rates, whereby we measure a larger velocity than is actually present. It is believed that this is due to increased velocity attenuation near the walls of the monolith, which loses signal from these slower moving spins and hence gives an artificially high average measure of the displacement.

## 6.4. COTTAGE

This section will discuss the application of the COTTAGE pulse sequence (Section 5.4) to the monolith reactor configurations. A test section was constructed to aid the development of this novel MRI sequence and allow high-resolution testing of COTTAGE within the monolith reactor. A small section of monolith was cut of length 10 cm and was placed within a 20 mm flow cell (Figure 6.4). Some glass ballotini were added at the entrance to the flow cell to dampen the flow and ensure a good distribution. However, this may also explain the flow maldistribution seen in Figure 6.5b, since there may be a ballotini obstruction in particular channels.

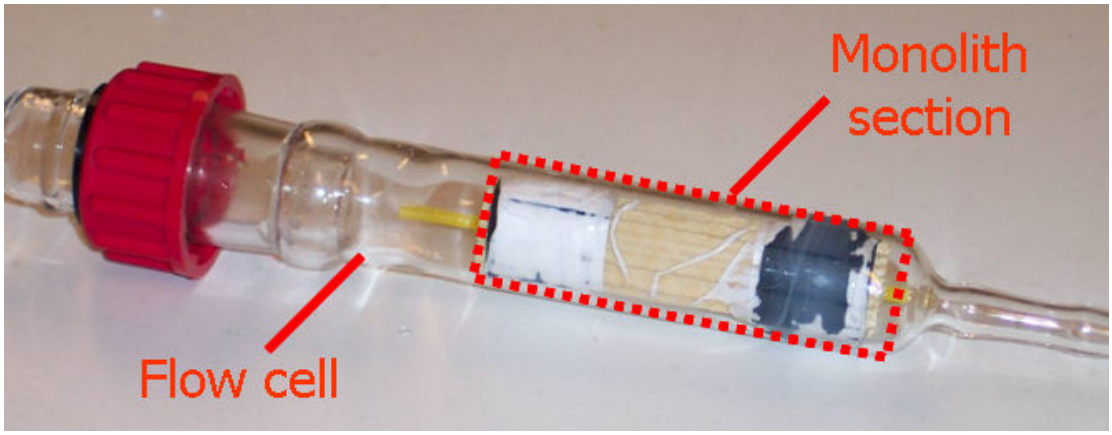


Figure 6.4 – Photograph of the monolith flow cell used in the initial COTTAGE experiments.

A cross-sectional MRI image of the liquid-saturated monolith section is shown in Figure 6.5a. This image has a resolution of 98 μm × 98 μm and a field-of-view of 25 × 25 mm. The slice thickness was 5 mm. A traditional spin-echo phase-encoded velocity image (Section 2.4.1) is shown in Figure 6.5b, for a single-phase liquid flow of 136 cc/min. The resolution is 195 μm × 195 μm and the field-of-view was 25 × 25 mm. The slice thickness was 5 mm, and the total acquisition time was 4.5 minutes. Some fluid by-passing around the side of the monolith is seen.

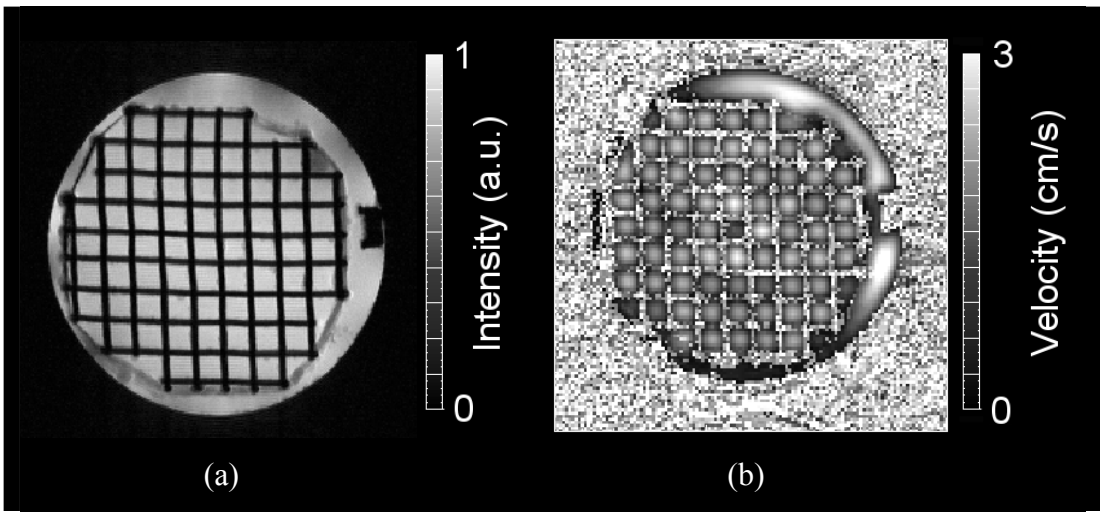


Figure 6.5 – (a) A water-saturated, zero-flow image of the monolith test section (200 cpsi) within a 20 mm diameter glass flow cell; (b) a traditional spin-echo, phase-encoded velocity experiment for single-phase flow at 136 cc/min. The acquisition time was 4.5 minutes.

High-resolution COTTAGE data were acquired for a slice through the 200 cpsi monolith. The resolution in the *x*-direction was 195 μm and the field of view was 25 mm. The slice thickness was 1.7 mm in both the excitation and refocusing planes, and the pulses were sinc3 pulses. 100 profiles were acquired, taken 2.8 ms apart, corresponding to measurements over 280 ms. Eight averages were used to improve the signal-to-noise ratio, which was possible because the system was in steady-state. The first 55 profiles at a single-phase flow rate of 136 cc/min are shown in Figure 6.6, illustrating the increase in the measured displacement with time.

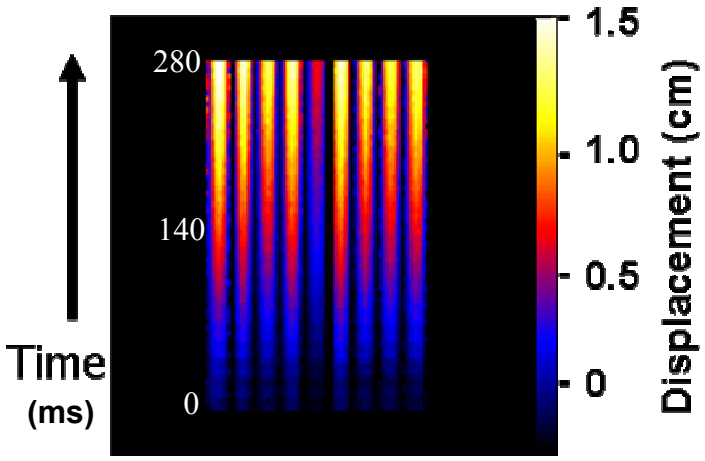


Figure 6.6 – A dataset showing COTTAGE data acquired under single-phase flow in the high-resolution monolith system. Flow rate is 136 cc/min. Total acquisition time is 280 ms.

Extracting profiles from the data of Figure 6.6 at increasing time steps, gives the plots shown in Figure 6.7. These clearly show the parabolic profiles within each channel, and how the parabolas grow with time. Extracting profiles at different x-coordinates gives the plots shown in Figure 6.8. These plots show the linear increase of displacement with time, and this can therefore be converted to a velocity. This illustrates the power of this technique; whereas for the line-excitation data we must acquire a 2D image over 150 ms, the COTTAGE technique allows us to acquire **k**-space a line at a time, and we need only use as many lines as are of suitable signal strength. For example, the first 10 lines of the dataset are acquired over the first 30 ms, and ten points are more than enough to extract a velocity. This technique should therefore be able to perform at higher velocities and in systems where the line-excitation technique fails from poor signal strength.

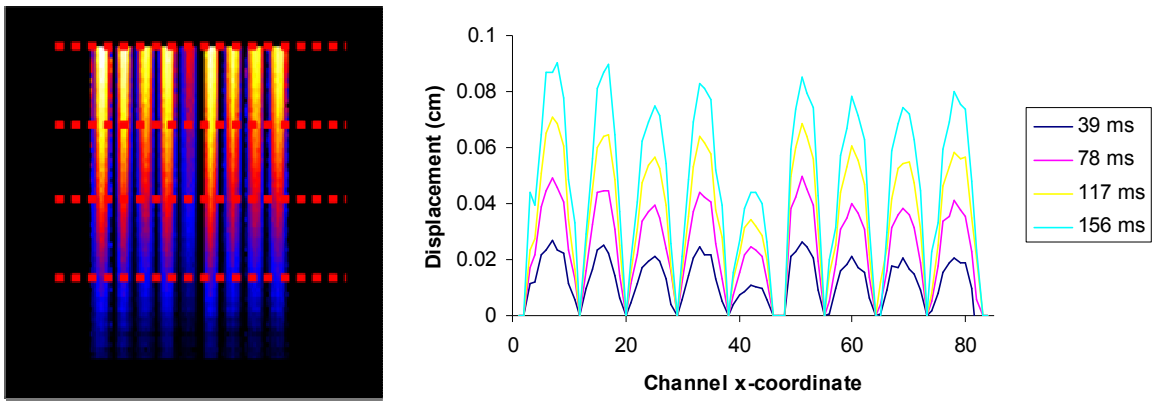


Figure 6.7 – (a) Four selected profiles from Figure 6.6 are shown in (b) corresponding to 39, 78, 117 and 156 ms after excitation. This illustrates the evolution of the profiles within each channel, as well as the parabolic nature within each channel.

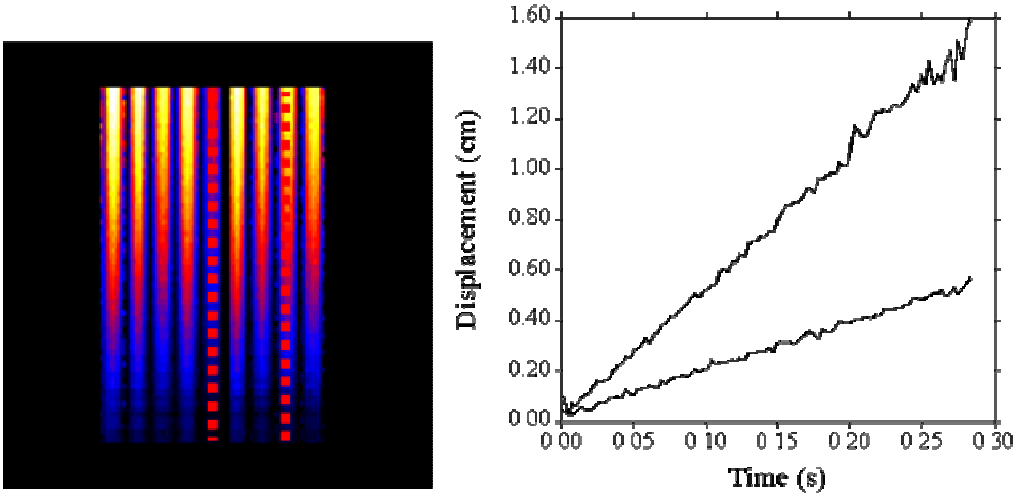


Figure 6.8 – (a) Two selected profiles from Figure 6.6 are shown in (b) corresponding to the centre line of two separate monolith channels. This illustrates the linear increase of displacement with time for single channels, which allows the measure of an average velocity.

The quantification of these velocities are shown in Figure 6.9. However, as we have seen in Figure 6.5b there is fluid by-pass around the edge of the monolith, and so the velocities from one row of channels are not necessarily representative of the overall superficial velocity. The linear increase is apparent, and this is a good indication of quantification. A 2D or multi-slice method is required to quantify the velocity in this system, and this will be addressed by the multi-slice COTTAGE sequence in the following section.

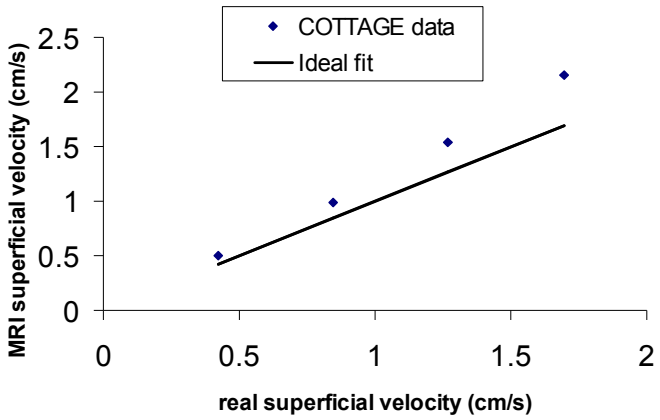


Figure 6.9 – Quantification of the COTTAGE pulse sequence. The measured average velocity by MRI is plotted against the true, measured superficial velocity for the system.

## 6.5. Multi-slice COTTAGE

The multi-slice adaptation of the COTTAGE sequence (Section 5.5) was applied to single-phase flow in the monolith, and the results will now be discussed. The first two scans of the cross-section were in general of sufficient signal-to-noise to produce reliable velocity plots (Figure 6.10a), however at the higher flow rates even the second image was badly degraded. Only the first images after the excitation were used in the analysis

Data were acquired on the flow system of Section 6.2, using a 200 cpsi monolith in the 5 cm i.d. pipe. A range of flow rates between 0 and 170 l/hr were used. The MRI data have a field of view of 55 mm, giving a resolution of 430 μm × 1770 μm. The echo time between each of the thirty-one 1D scans was 3.3 ms. The entire 2D cross section is acquired in 102 ms. Sinc3 pulses were used for both the excitation and refocusing pulses, and both were calibrated to 1.7 mm thickness.

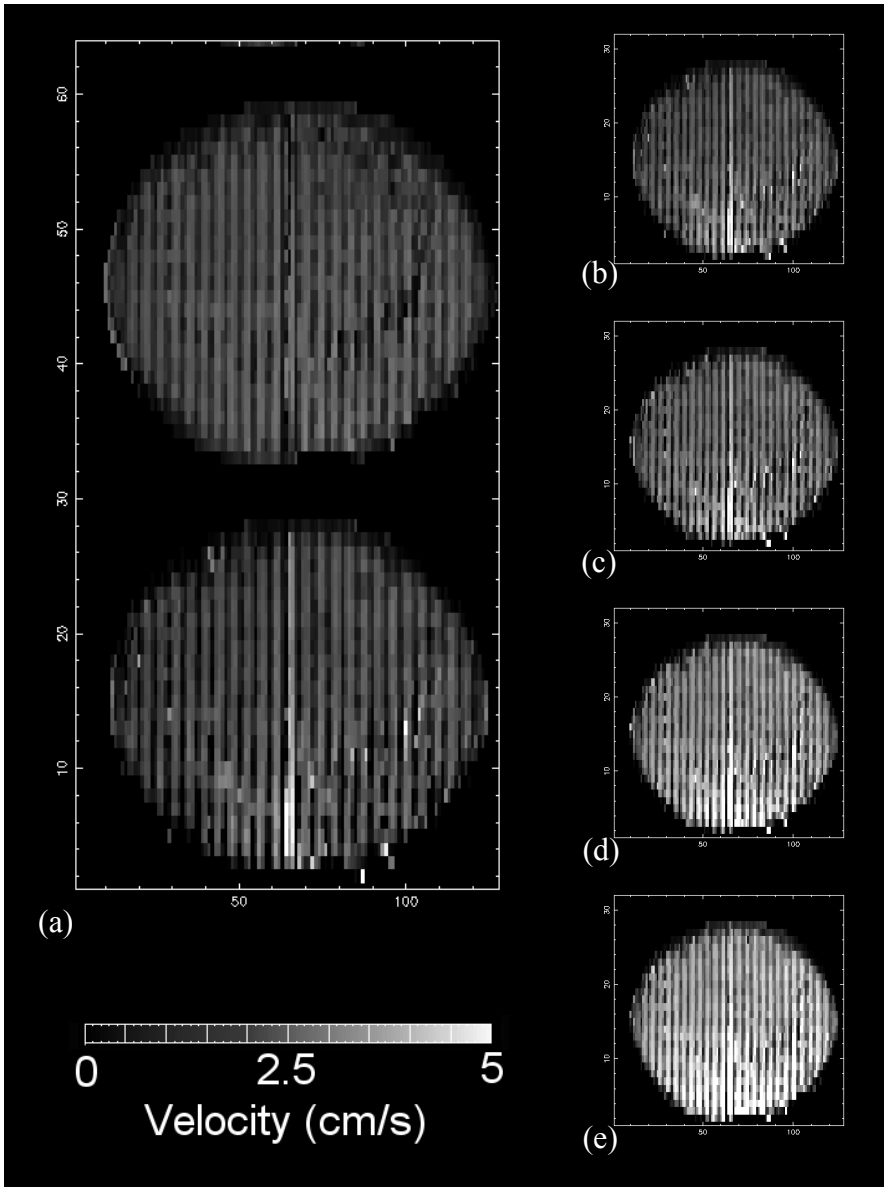


Figure 6.10 – Sample data acquired using the MS-COTTAGE technique on single-phase flow in the monolith reactor; (a) shows the first two scans from an acquisition at a flow rate of 24 cc/s; (b-e) show the first scan from experiments at flow rates of 24 cc/s, 33 cc/s, 41 cc/s and 47 cc/s.

If we examine the quantitative nature of the measurements, we obtain the plot of Figure 6.11. Here we can see that the average superficial velocity measured by the MS-COTTAGE technique compares very well with the true measured superficial velocity. This is important because we have obtained a 2D quantitative velocity measurement in a porous media in only 102 ms. This rapid measurement should allow us to probe dynamic systems, and this is discussed further in Chapter 11.

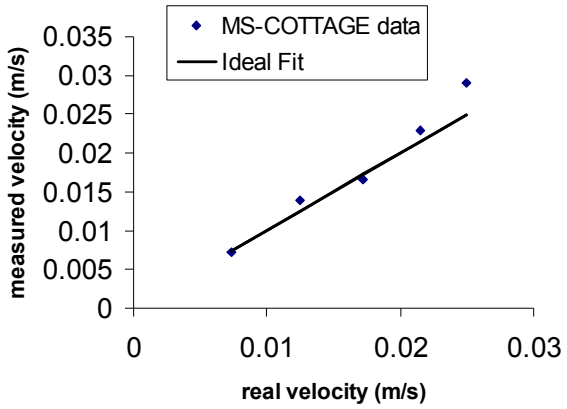


Figure 6.11 – Quantification of the MS-COTTAGE method for single-phase flow through a 200 cpsi monolith. The graph shows the measured superficial fluid velocity versus the real superficial velocity through the system.

## 6.6. GERVAIS

The GERVAIS pulse sequence (Section 5.6) was applied to the single-phase flow configuration (Section 6.2). Data were acquired with a field of view of 70 mm × 70 mm, giving a resolution of 1094 μm × 1094 μm. The *z*-velocity component was acquired in 20 ms. Results of the quantification are shown in Figure 6.12. It can be seen that the measured velocities from the GERVAIS sequence are at all flow rates slightly higher than the true flow rates. This is believed to be due to the loss of signal near the walls of the monolith, where slow-moving spins are lost. GERVAIS is also more sensitive to magnetic susceptibility errors, and this may contribute to the effect observed.

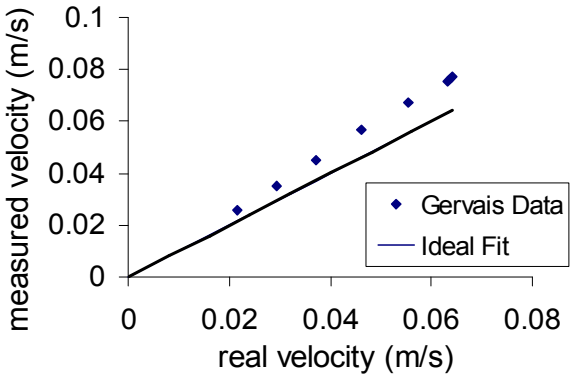


Figure 6.12 – Quantification of the GERVAIS sequence for single-phase flow in the monolith reactor. The graph shows the measured superficial fluid velocity versus the real superficial velocity through the system.

Some velocity distributions for single-phase flow at increasing flow rate are shown in Figure 6.13, and these are consistent with the excepted narrow distribution, which shifts and broadens at higher flow rates.

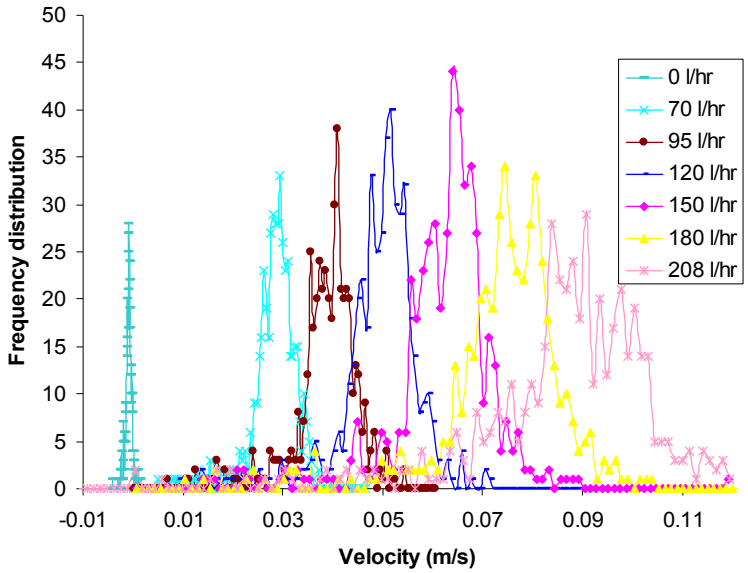


Figure 6.13 – Velocity distributions obtained at increasing flow rate, as measured by the GERVAIS sequence.

Even if the GERVAIS technique is not fully quantitative, the data are acquired extremely rapidly, and so is still a very useful *qualitative* technique. For example, we may utilise it to observe flow fluctuations or maldistribution on a short timescale where other techniques are too slow.

## 6.7. Conclusions

This chapter has applied the novel MRI techniques of Chapter 5 to single-phase flow in the monolith reactor. Each of the techniques has been shown to be qualitative, apart from COTTAGE and GERVAIS. This was because COTTAGE is a 1D technique, and could not measure the total flow rate through the monolith for comparison with the measured volumetric flow rate. GERVAIS encounters a systematic error in measuring velocities in the monolith, which is probably caused by magnetic susceptibility errors.

## 6.8. Appendix - bubbles

This appendix will cover an observation of single-phase flow in a pipe. It highlights the need to operate single-phase flow in an upright configuration, and also highlights the sensitivity of the monolith reactor to small pressure fluctuations.

A series of experiments were performed to investigate the sensitivity of the monolith reactor to local pressure changes. The configuration of Section 6.2 was used, but with *downwards* single-phase flow. Bubbles of different sizes were introduced into the flow. In general, they were unable to rise through the monolith since the liquid flow opposed their motion, and they sat on the underside of the monolith section, as indicated in Figure 6.14. A large bubble can be seen on the bottom surface of the monolith. A refractive index mis-match makes it tricky to observe the base of the monolith when filled with water, but the bubble can still be clearly seen.

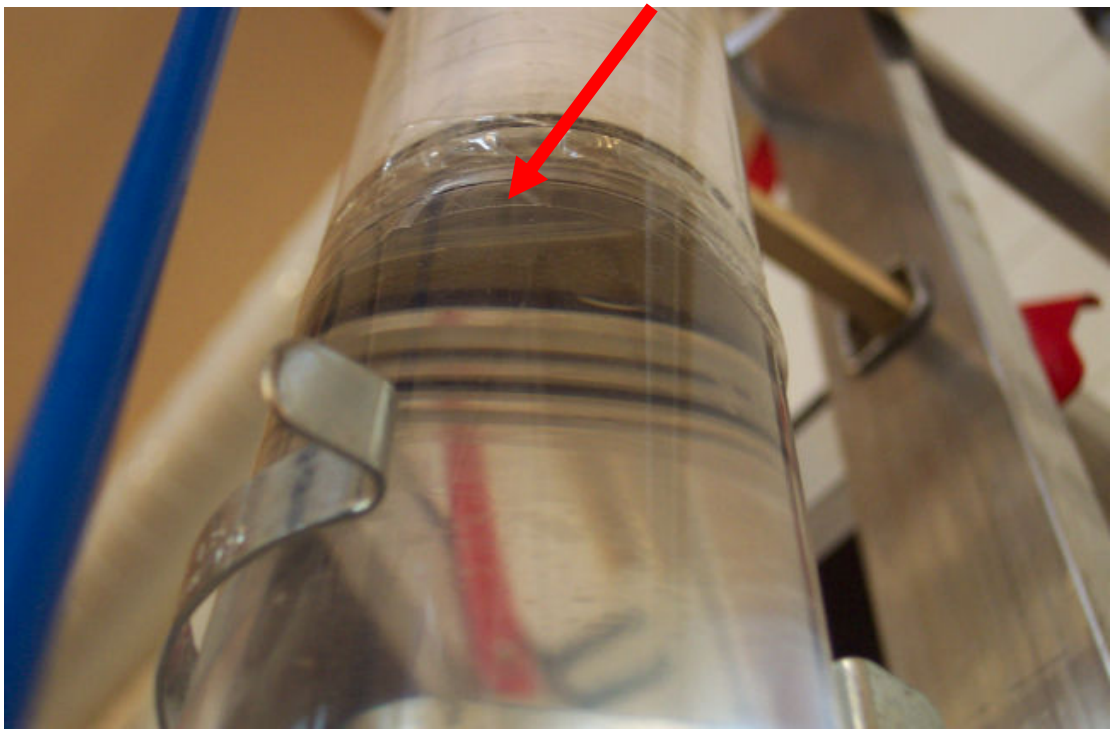

Figure 6.14 – Photograph of a bubble lodged under the monolith reactor, whose upwards movement is resisted by the presence of liquid downflow.

A series of GERVAIS velocity data with the parameters of Section 6.6 were taken to investigate the effect of such bubbles on the velocity distribution. The results at a flow rate of 45 l/hr and a range of different bubbles sizes are shown in Figure 6.15. They clearly show that the presence of an air bubble creates a large enough pressure drop to cause the flow in the channels above it to stop, and this therefore increases the velocity through the other channels. This illustrates the sensitivity of the monolith to pressure fluctuations – only a small increase in pressure is required to stop flow in a channel, and in practice this could be caused by fouling deposits, or even a fluid with high-surface tension moving through the channels.

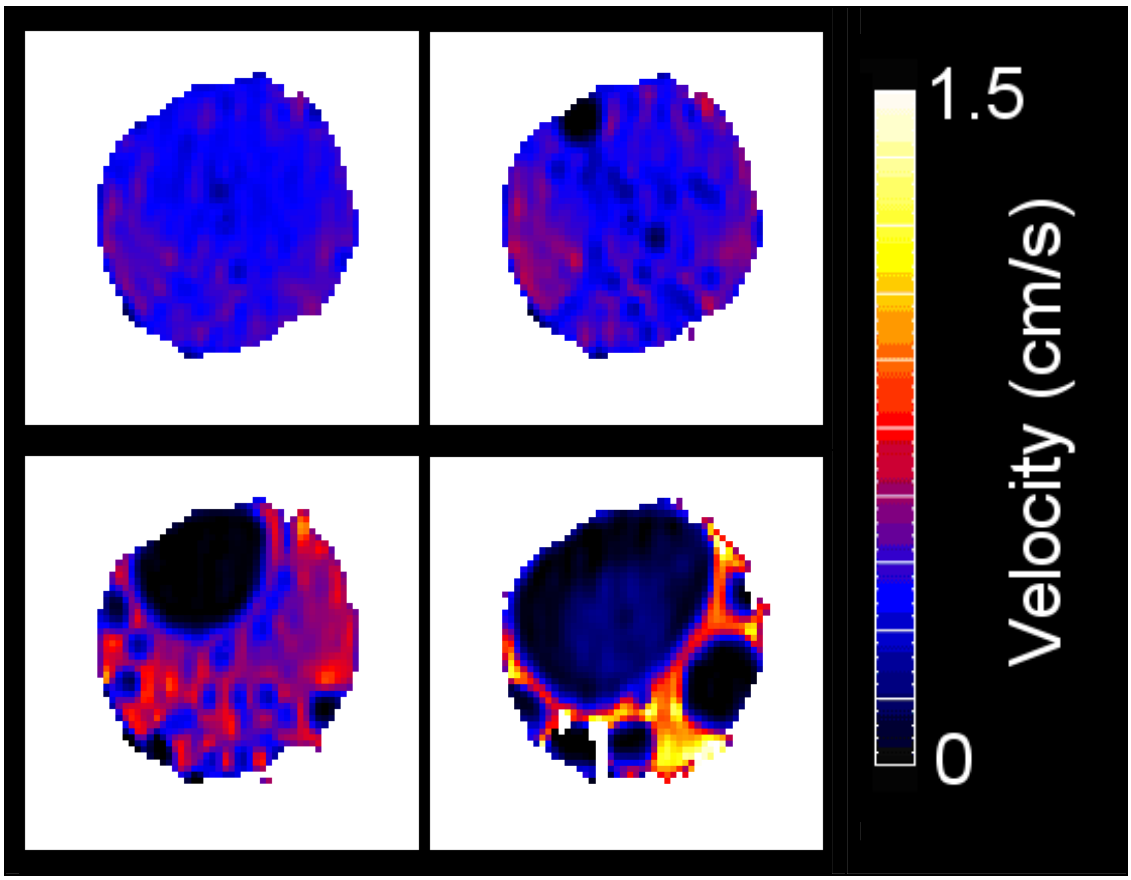


Figure 6.15 – GERVAIS velocity data acquired at a flow rate of 45 l/hr, and with a range of different sizes of bubble present at the base of the monolith.

## 6.9. Appendix – slice selection

The choice of excitation pulse shape is critical, since it defines the spatial distribution of spins that we will excite and refocus. There are many pulse shapes available, but by far the most common are the Gaussian, sinc3 and sinc5 pulses, illustrated in Figure 6.16.

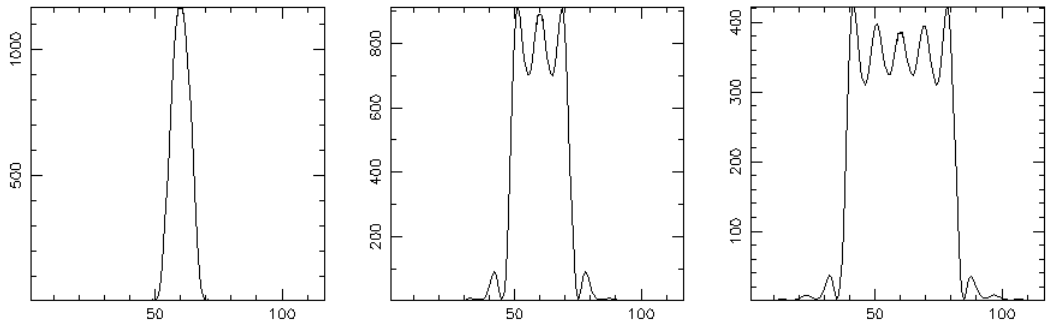


Figure 6.16 – Example of three excitation pulse shapes; (a) Gaussian, (b) sinc3, and (c) sinc5. The widths of the pulses shown here are arbitrary; the bandwidth is tailored to match the chosen slice thickness for the experiment.

Calibration of the pulse sequences to the monolith system gave the following parameters for a 180° pulse; Gaussian - 1.1 dB, bandwidth of 1950 Hz; sinc3 - 6 dB, bandwidth of 5500 Hz; sinc5 - 6 dB, bandwidth of 9765 Hz. The bandwidth of the Gaussian is measured as the width at the half-height, for the sinc pulses it is measured as the bandwidth between the two side-lobes.

It can therefore be seen that, if we choose the Gaussian to excite a slice there will be significant excitation beyond the edges of that region – all be it at low intensity – whereas the sinc pulses have well-defined, near-vertical edges and our excited fluid will be almost exclusively contained within the region of interest.

The sinc3 pulse was chosen in preference over the Gaussian for the excitation and refocusing pulses, due to its ability to excite a sharp, narrow slice. This will result in only obtaining signal from spins within the channel of interest. However, the sinc gives a poorer signal to noise, and the pulse length must be longer to excite very narrow slices due to its much larger spectral width. For example, we could use a 512 μs gauss pulse to excite a 1.2 mm slice, but the sinc3 pulse had to be extended to

1085 μs to achieve the same result. This is the reason for not using the sinc5 pulse; even though it is a better approximation to a rectangular pulse, the spectral width is larger again and the increase in pulse time means a longer image acquisition, longer dephasing time, and hence a poorer signal to noise ratio. The sinc3 was judged to be the optimum excitation shape for our purposes.

In Figure 6.17, we see the extent of the excitation from the sinc3 pulse, highlighting its ability to only excite fluid within the targeted row of channel, and very little outwith it. Figures are shown for both 200 and 400 cpsi monoliths. Data are taken with a FOV of 50 mm × 50 mm, a resolution of 195 μm × 195 μm, and slice thickness of 5 mm. The sinc3 width is 1.7 mm for the 200 cpsi monolith, and 1.2 mm for the 400 cpsi monolith.

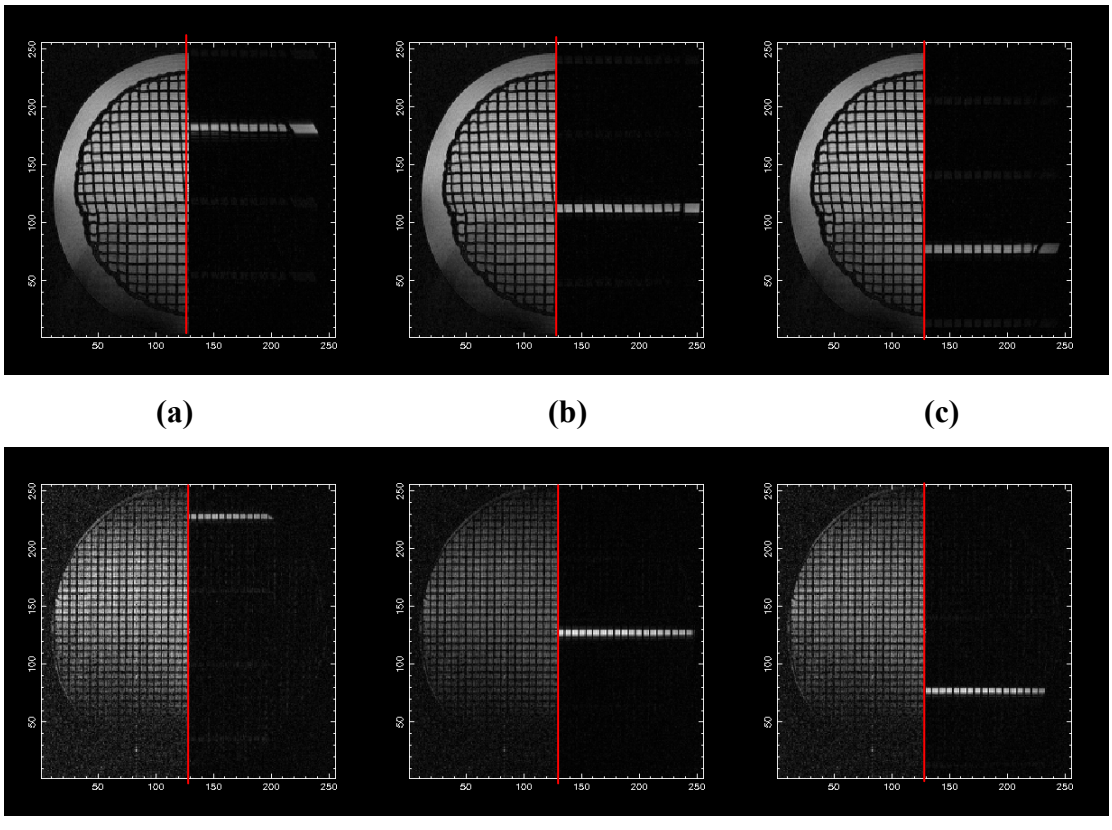


Figure 6.17 – Sample images showing the positioning of the sinc3 refocussing pulse; the left side of the image is a standard 2D cross-section of the monolith, the right side shows only the region excited by the vertical refocusing pulse. Figures (a-c) are for the 200 cpsi monolith, and (d-f) are for the 400 cpsi monolith.

## 6.10. Nomenclature

| Symbols | Definition |
|---|---|
| $D$ | Diameter |
| $f_n$ | Variable frequency offset of a soft pulse |
| G | Amplitude of phase gradient |
| **k** | Reciprocal space vector |
| $L_{99\%}$ | Entrance length at which laminar flow deviates by less than 1% from a parabolic profile |
| $m, n$ | Repetition counters |
| $NE$ | Echo number |
| Re | Reynolds number |
| $s$ | Displacement |
| $t_e$, $TE$ | Echo time |
| $T_2$ | Spin-spin relaxation constant |
| $v$, $V$ | Velocity |
| $\Delta$ | Time between phase gradient pairs |
| $\phi$ | Phase angle |
| $\delta$ | Duration of phase gradient |
| $\gamma$ | Gyromagnetic ratio |

| Abbreviations | Definition |
|---|---|
| COTTAGE | Continuous translation tracking with alternating gradient encode |
| cpsi | Cells per square inch |
| EPI | Echo-planar imaging |
| FOV | Field of view |
| GERVAIS | Gradient echo rapid velocity and acceleration imaging sequence |
| MRI | Magnetic resonance imaging |
| MS-COTTAGE | Multi-slice COTTAGE |
| NMR | Nuclear magnetic resonance |
| RARE | Rapid acquisition with relaxation enhancement |
| RF | Radio frequency |
| SEMI-RARE | Single excitation, multiple image RARE |
| TOF | Time of flight |

# Chapter 7 – Image analysis techniques

Data obtained through MRI is generally in the form of 2D images, which then must be interpreted to extract quantitative information. This leads naturally to the development of image analysis techniques. This chapter presents a review of the theory behind the techniques used in this work, and outlines the development and implementation of novel analysis methods, with examples of application to real MRI data where appropriate to highlight their usefulness and applicability. These techniques are later applied to data from real monolith systems in Chapters 9-11.

**CHAPTER 7 – IMAGE ANALYSIS TECHNIQUES**

## 7.1. Introduction

Humans are primarily visual creatures, relying on our eyes for 99% or more of the information we receive about our surroundings (Russ, 2002). Human vision is primarily qualitative and comparative rather than quantitative or absolute. This is shown in the existence of optical illusions where the brain is tricked into the mis-interpretation of information. Image processing is used for two different purposes – improving the visual appearance of images to the human eye, and preparing images for quantitative measurement of the features and structures present. Often we may be able to instinctively see trends or patterns in data, which we must then attempt to quantify in a more rigorous manner. Most procedures focus on enhancing features in the image in order to make them easier to detect and process. In this chapter many images are presented in false colour – even though the originals are simple grey-scale – to ease interpretation by the human eye, since we may distinguish more levels of colour than levels of greyscale.

MRI experiments typically result in the acquisition of many images, which must then be analysed. Hence we have the need for quantitative, robust image treatment and analysis techniques. The techniques presented in this chapter will focus on the extraction of numbers and trends from data via image enhancement and feature quantification techniques. These techniques will be applied to some real reactor systems for illustration, in particular those of narrow packed beds, turbulent flow in a pipe, and profiles of slug flow in monolith channels. First we must cover a few basic image processing techniques (Section 7.2) which are illustrated with sample data. This is then followed by a literature review of background levelling techniques (Section 7.3) and a short section on algorithm optimisation and convergence (Section 7.4). Then a literature review of texture processing techniques for distinguishing between data of different distribution homogeneities (Section 7.5), followed by the development and implementation of these techniques to real experimental data. Then a review of some feature processing techniques for extracting specific information from characteristic features within an image (Section 7.6), again with development and application to real MRI data. The chapter concludes with a section on phase-unwrapping techniques specifically for correcting MRI velocity data (Section 7.7).

The techniques reviewed and developed here are done so with the aim of application to our system of interest, the monolith reactor. These techniques are later applied to real monolith MRI data in Chapters 9-11.

## 7.2. Basic image processing

First, we must cover some basic image processing techniques. This section will cover binary gating, partial volume effects and smoothing techniques in turn.

### 7.2.1. Binary gating

When acquiring MRI data from a slice, signal is received from the entire volume; coherent signal from the nuclei of the imaged phase, and noise from the remainder of the imaged volume. Binary gating is an important step in the quantification of MRI data, since we identify which regions contain valid data, and hence prevent our calculations being swamped by noise data. The binary gating procedure is used in the creation of mask images from data, where we assign regions which contain valid data a value of unity, with a value of zero elsewhere. This allows us to identify regions of interest and blank out others. This is illustrated in Figure 7.1 with data acquired by El-Bachir (2005) in the N = 2 packed bed (described in Section 7.5.2.3). Figure 7.1a shows the spin-density image of flow through the bed. The data has an isotropic field of view of 55 mm, a resolution of 215 μm × 215 μm and was acquired with a spin-warp imaging sequence (Section 2.2.7). The contrast visible within the image is caused by flow attenuation, since the fluid is moving with a Reynolds number of 216 (based on the particle diameter). Figure 7.1b shows the corresponding histogram. It can be seen that the histogram is mainly bimodal, with a peak at low signal intensity (below $8\times10^4$), and another at high signal intensity (between $1.8\times10^5$ and $3\times10^5$). The former is associated with noise, and the latter with coherent water signal. By choosing a gating level just higher than that associated with the noise population, the image is binary gated. In the example, Figure 7.1c shows the data after binary gating at a level of $8\times10^4$. Gating levels for some datasets are easily chosen, however it is often difficult to distinguish noise pixels from data pixels based solely on their relative intensities. For such data, a good rule of thumb is that the gating for the noise

level is best estimated by the following method – by selecting a region which clearly contains only noise pixels, and calculating five times the standard deviation of this region, we obtain a reasonable estimate for the background noise level. This method was first proposed by Sederman (1998) and has been shown to be a reasonable estimate for many real data situations, and is considerably more robust than estimating the level manually.

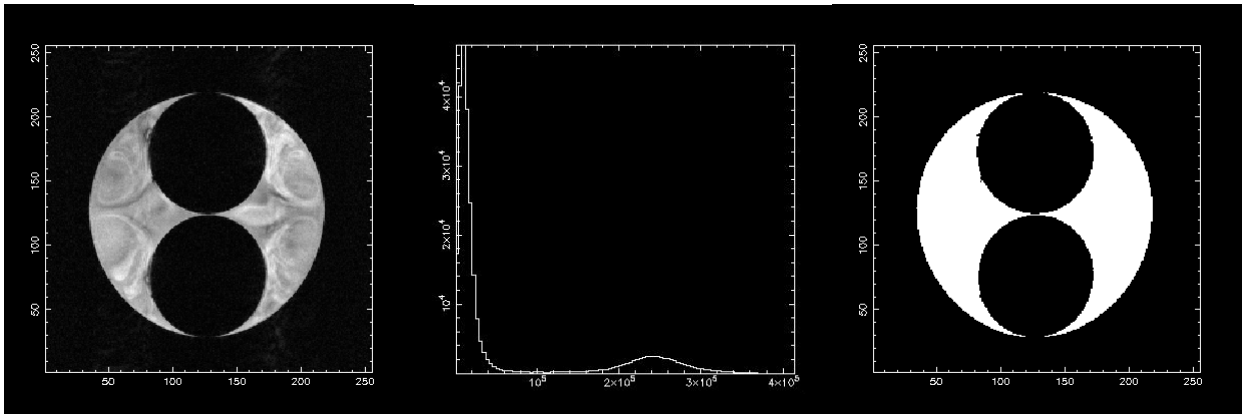


(a) (b) (c)

Figure 7.1 – The binary gating process. (a) a 2D spin-density image of water flowing through a packed bed of N = 2; (b) a histogram of the signal intensity in (a); (c) the mask obtained after binary gating at a level of $8\times10^4$.

## 7.2.2. Partial volume effects

The MRI data are discretised into voxels, typically in a rectilinear grid. Voxels located at the interface of the imaged phase and another phase may be partially filled by both, and in this case the voxel will take an intermediate value of signal intensity. Depending on the gating level, the voxel may be identified as signal or otherwise. This is termed as the partial volume effect, and often means that a range of possible threshold values exist within an image. There should be two overriding requirements of a gated image to help in identifying a suitable result; firstly the level must be such that the noise regions are nulled out; and secondly that the total area of the mask is equal to the real area of the apparatus – this greatly aids calculation precision, for example when converting from velocities to superficial flow rates we need accurate measurements of the area involved. Partial volume effects are also of specific concern to velocity imaging techniques, and this is covered in Section 2.4.

## 7.2.3. Smoothing functions

In the process of binary gating, it is often seen that images will contain lone pixels caused by noise fluctuations close to the gating level. This is illustrated in Figure 7.2, where the data of Figure 7.1a have been corrupted by adding synthetic noise. The effects of noise on binary gating are seen in Figure 7.2b.

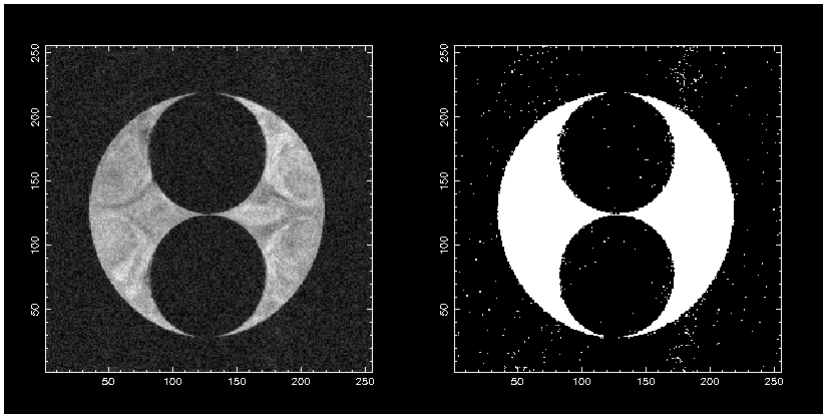


(a) (b)

Figure 7.2 – The effect of noise on binary gating; (a) the data of Figure 7.1a is shown with the addition of synthetic noise; (b) shows the result of binary gating as before, this should be compared with Figure 7.1c.

In the past such artefacts have been removed by successive "erosion" and "dilation" steps – in the erosion step, pixels on the outer edge of an interface are removed, and then in the dilation step the new interface is expanded back out. For most of the image, the effect of the two processes should cancel out and no change should be observed. However, lone noise pixels will be erased in the erosion step and hence will not be present for the dilation step, and so are totally removed. This process is illustrated in Figure 7.3, and the result is a reasonable approximation to the noiseless case of Figure 7.1c. However, these two steps are not completely reversible and fine detail is lost in some regions, so a more robust method was sought.

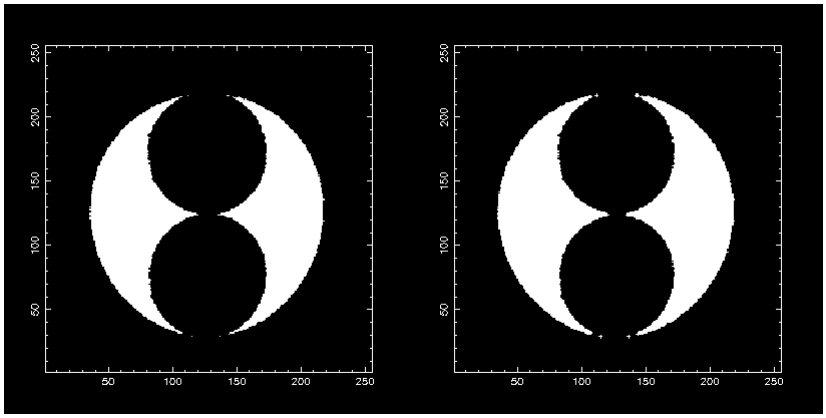


(a) (b)

Figure 7.3 – The effect of (a) an erosion step, and (b) a subsequent dilation step on the data of Figure 7.2b. The result is similar to the ideal case of Figure 7.1c.

In this work, a smoothing method was implemented, whereby the images were Fourier transformed, a smoothing function was applied, and then Fourier transformed back[1]. This has the effect of removing high-frequency, short scale features from the image, thereby blurring it. This method has the advantage that we are not classifying pixels by intensity alone, but rather by removing details corresponding to short-scale features, which are most likely to be noise-related. We effectively remove the noise variations in the image while retaining information on the main features of the image. Hence, we can now binary gate robustly. An example of applying smoothing before gating is shown in Figure 7.4. Smoothing in the Fourier domain is covered in more detail in Section 13.6.1.

In most cases a Gaussian smoothing function is chosen for simplicity and because of its suitability for image processing applications. The Gaussian filter produces the least edge-blurring for a given amount of noise reduction (Russ, 2002), but it cannot eliminate the blurring altogether. The strength of the smoothing can be tailored by adjusting the width of the function, and we may vary this separately in each dimension, only smoothing in one direction if required.

[1] In much the same manner as in NMR processing, whereby raw data is multiplied by a function such as sine or q-sine to suppress signal at the edge of **k**-space (associated with noise) while retaining signal near the centre. The result is to remove noise from an image, at the expense of some loss of detail (i.e. the blurring of fine structure).

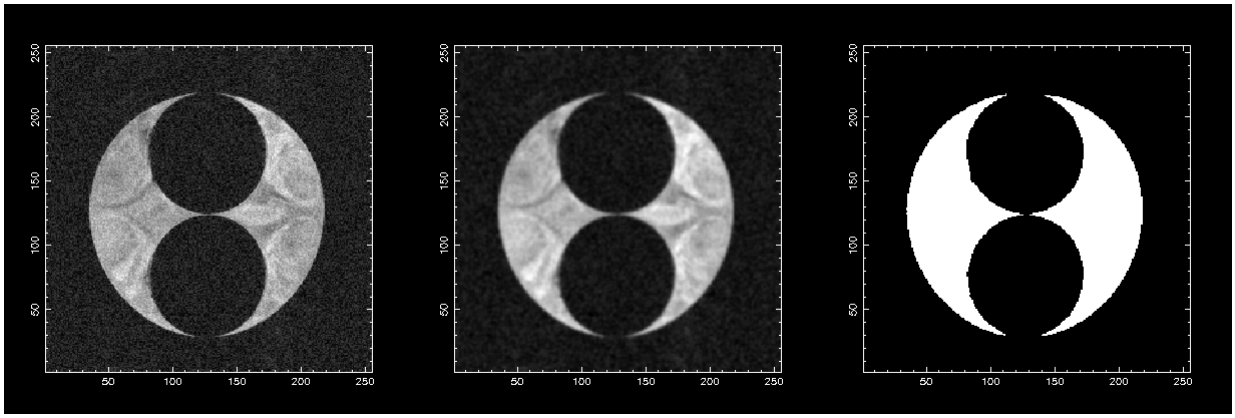


(a) (b) (c)

Figure 7.4 – Illustration of the binary gating process applied to noisy data which has undergone smoothing; (a) shows the raw, noisy data; (b) shows the smoothed data; (c) shows the result of binary gating of (b).

A single smoothing pass is generally sufficient to remove most of the noise-related errors. By adjusting our choice of smoothing function and its parameters, we can control the degree of smoothing. Alternatives to the Gaussian function are sine and q-sine (sine squared), both of which are found in NMR processing options (see Figure 7.5).

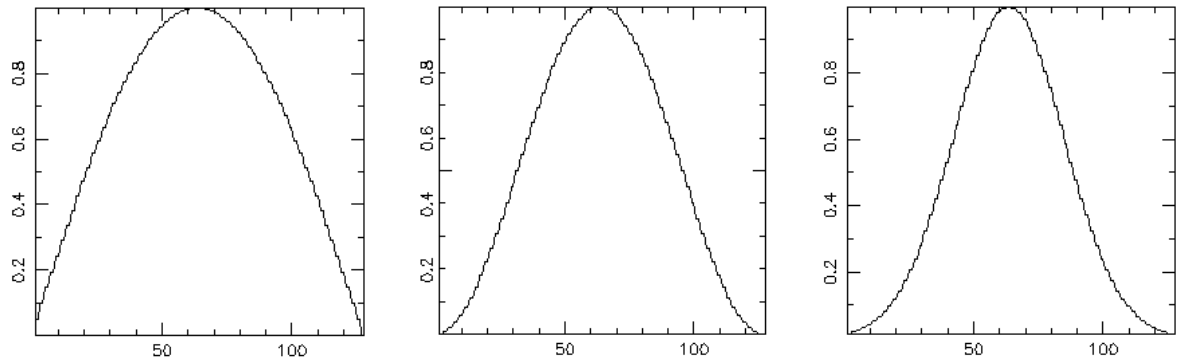


(a) (b) (c)

Figure 7.5 – Examples of three basic smoothing functions; (a) sine, (b) q-sine (sine squared), and (c) Gaussian functions.

Other, more sophisticated smoothing functions are available, for example Savitsky-Golay smoothing (Savitsky and Golay, 1964) which attempts to compute a smooth background level by fitting to the local data. This is a low-pass filter, similar to a

moving average. The coefficients in the model are chosen to preserve higher moments in the data, thus reducing the distortion of essential features such as peak heights and line widths in a spectrum, while the efficiency of the suppression of random noise is effectively unchanged. An example of this function is given in Figure 7.6.

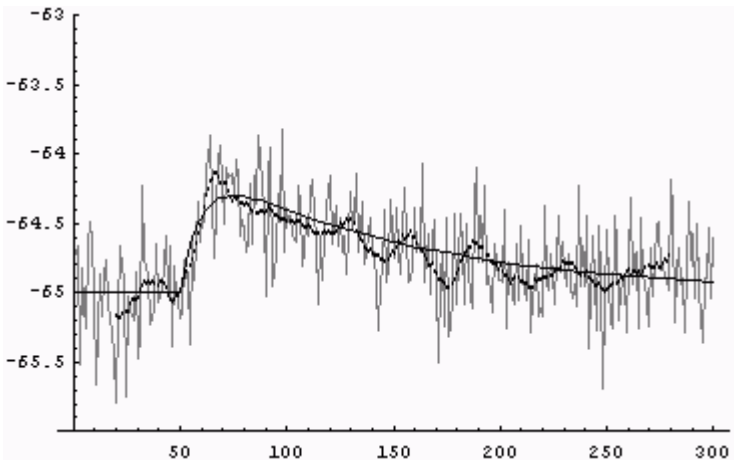


Figure 7.6 – Example of Savitsky-Golay filtering to remove noise fluctuations. The raw data is shown (—), along with smoothed (▬) and further smoothed (▬) results. Comparison to the model curve is now far more straightforward. Reproduced from the documentation for Wolfram Research's Matlab software.

Some methods smooth images while attempting to retain edge sharpness by having non-zero values at the edge of **k**-space. An example of this is the Laplacian of Gaussian (LoG) function, as shown in Figure 7.7. The LoG function combines a Gaussian filter to blur the image, and a Laplacian filter to enhance edges. More information on such functions can be found in Wirth (2002b).

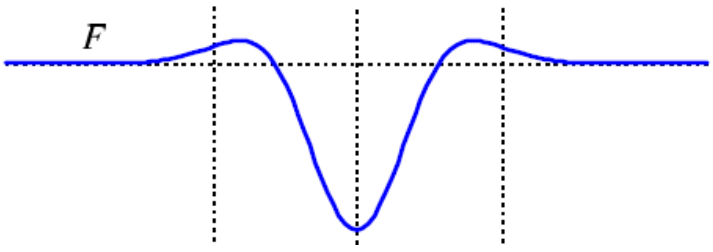


Figure 7.7 – Laplacian of Gaussian (LoG) smoothing function, which has non-zero values at the edge of **k**-space.

## 7.3. Background levelling

A further problem with binary gating is that the sensitivity of the RF coil can vary significantly across the imaging volume – for example, we may see variations in intensity even within a phantom of water despite having a constant spin density. These effects are more pronounced in MRI images when we have susceptibility variations within the coil, resulting in images with uneven local sensitivities. The signal may fall off towards the edges of the coil, or there may be localised regions of low intensity, also termed "bites". If the images are to be binary gated in some way, then this poses a major problem since the correct gating level will not be consistent across the whole image. One solution is to take an image of a phantom of water in the coil to measure the inherent characteristics of the RF coil. This may not always be possible – for example in clinical MRI it is not always practical to do this – and it may be too time-consuming to perform the acquisition. It also will not correct for any distortion of the magnetic field lines caused by the properties of the sample itself.

There exist several methods of varying complexity to estimate and correct for these variations in an MRI image. Polynomial fitting and unsharp masking will be covered. Other more sophisticated statistical methods exist, and one will be shown later.

### 7.3.1. Polynomial fitting

Polynomial fittings attempt to follow large-scale trends in the data intensity by a best-fit method. If we are interested in applying this method, we may select a number of points from the background part of an image, and perform a least-squares fitting of a function that approximates the background. Then the original image can be subtracted or divided by this function. Extending this further, we could even fit two envelopes to the image, one for the background level and another for the intensity level and normalise the image between this range (Figure 7.8). We may have problems finding similar features throughout the image, or if the features of interest vary in intensity relative to each other. It is relatively difficult to extend to 3 dimensions, since we would be fitting curves to what is effectively a 4D surface (intensity variations over a 3D volume).

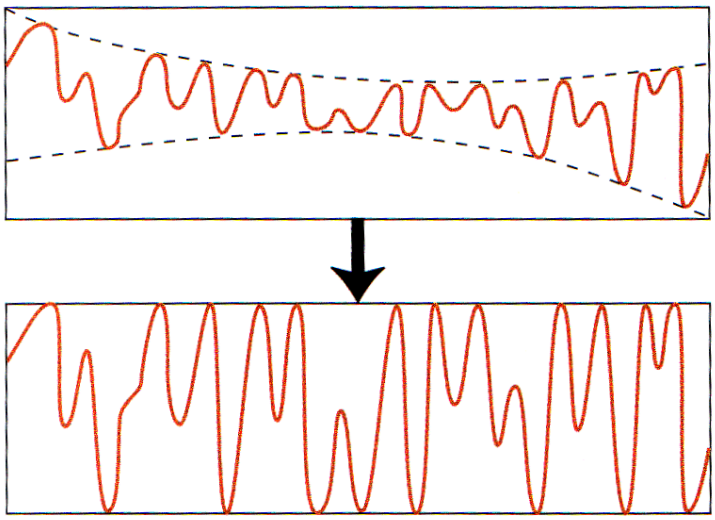

Figure 7.8 – Illustration of a polynomial fitting method for background levelling. Reproduced from Russ (2002).

### 7.3.2. Unsharp masking

The unsharp masking operation is of use when the background level varies more abruptly than can be fitted to simple polynomial functions. The assumptions in this method are that the features of interest are limited in size, and on a smaller length scale than the background features.

One of the most useful applications of background levelling is in astrophysics, where we encounter the problem that the dynamic range of intensity from the night sky is very wide – we have regions of dim light from distant galaxies and very bright points from stars i.e. we have a large range of contrast levels. If the receiver gain of the observation device is adjusted such that we capture both the dim and bright regions, then the galaxies will be unobservable since their intensities are near the lower end of the data, there will be little information in the centre of the histogram and stars will be at the very high end. This is an inefficient utilisation of the histogram, and one of the techniques which is employed is to divide through by an average intensity map, in a technique known as "unsharp masking". This is done by blurring the image and dividing the original image through by the blurred version. This has the effect of levelling the contrast levels through reducing the intensity of local regions by their

average intensity. This allows both fine detail in galaxy structures and bright stars to be observed within the same image, and is illustrated in Figure 7.9.

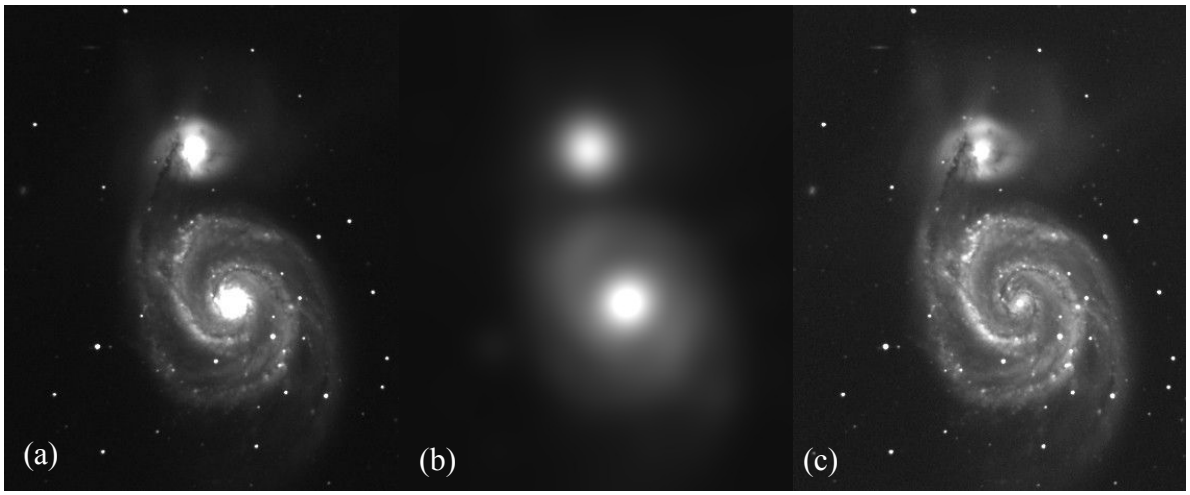


Figure 7.9 – Illustration of detail enhancement by better utilisation of the image histogram, in application to an astrophysics image of a galaxy; (a) the starting image; (b) a blurred image of (a) to obtain the local brightness; (c) the final image, produced by using (b) as an opacity (or exposure) mask for (a). The details are far easier to discern in this image. Reproduced from www.astrosurf.com/buil/iris/tutorial1/doc6_us.htm

### 7.3.3. TINA

We will now briefly cover a sophisticated implementation of background levelling, which utilises complex statistical methods and sophisticated algorithms specifically designed for MRI image correction. TINA stands for "TINA Is No Acronym" (or alternatively, "There Is No Alternative"), and is the name of an open-source development project designed for machine-vision and medical image processing, amongst other uses (http://www.tina-vision.net).

Vokurka *et al.* (1999) and Siu (2004) present a non-parametric approach for correcting non-uniformity, based on smooth shifts in intensity within homogeneous materials. This is implemented in the TINA package and consists of a sophisticated algorithm for coil inhomogeneity correction. It calculates estimates for additive image noise, local intensity gradients and smooth local derivatives, and computes an overall gain map which is a measure of the coil sensitivity and is used to correct the original data. Some example results of this package are shown in Figure 7.10 and

Figure 7.11. This technique has not been implemented as a part of this work, but would be of use in systems where the background sensitivity is not known *a priori*.

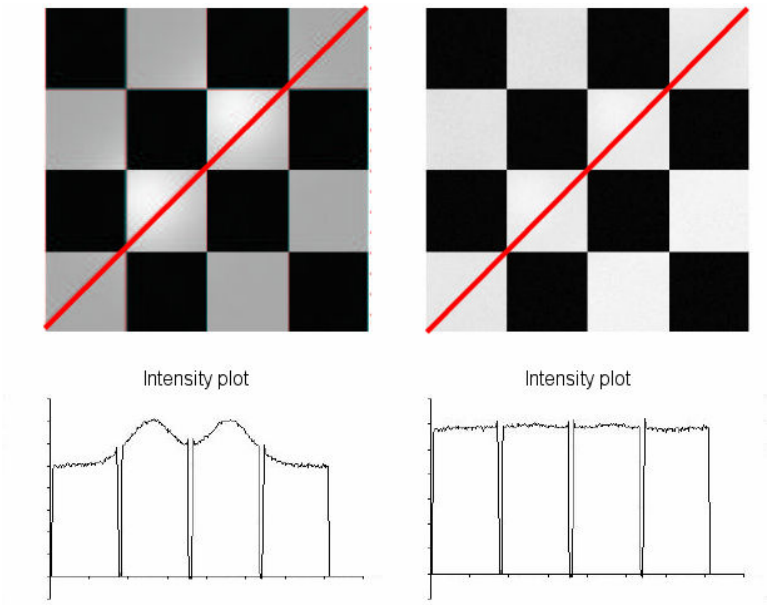


Figure 7.10 – Illustration of background levelling technique on synthetic data, showing the removal of background sensitivity inhomogeneities from a chequerboard image. Reproduced from Vokurka *et al.* (1999).

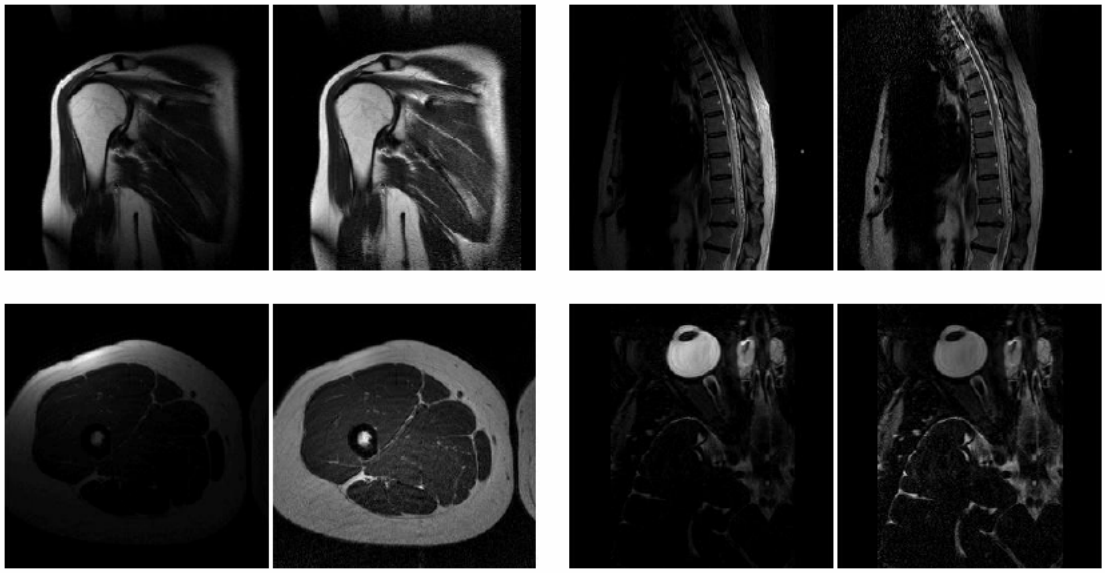

Figure 7.11 – Background levelling technique, applied to real MRI data, showing a significant improvement in image quality. Raw data are shown on the left, and corrected data on the right. Reproduced from Vokurka *et al.* (1999).

## 7.4. Algorithm verification and optimisation

In image processing we often must choose parameters for the processing algorithms. It would be of great use to have a utility to assess the quality of the final image result to ensure that algorithms are performing satisfactorily and that we have chosen optimum parameters. An information measure typically used in image quality control is the entropy measurement, as defined by Shannon (1948). This is calculated in the manner shown in Eq. 7.1.

$$H = -\sum_{i}^{N} P_i\left(\{g_a\}\right) \log P_i\left(\{g_a\}\right) \qquad \textbf{(Eq. 7.1)}$$

where the subset of pixel grey-scale values are given by Bauer *et al.* (2005). (This definition is analogous to the entropy measure outlined for texture analysis in Section 7.5.3). The value of $H$ is non-negative and takes its maximum value when the distribution of grey-scale values is flat, and is often regarded as a measure of the "peakiness" of the image histogram. This can be used as an iterative measure to find the optimum processing parameters, with the aim of minimising $H$. The assumption is that RF field inhomogeneities will introduce "information" into the histogram and therefore appear as an increase in the value of $H$. However, Eq. 7.1 is not scale invariant i.e. it cannot be used as a measure of absolute information content between two images, but only as a relative measure for a single image under different processing conditions. Thacker *et al.* (2002) propose an improvement to this information measure which removes these disadvantages and extends it to two-dimensions, shown in Eq. 7.2.

$$B = \iint \sqrt{\frac{P(g_1, g_2)}{g_1 g_2}} \mathrm{d}g_1 \mathrm{d}g_2 \qquad \textbf{(Eq. 7.2)}$$

This is a likelihood measure with strong scale invariance and weak polynomial invariant characteristics, which they illustrate in their publication. This measure has not been implemented as a part of this work, but it is included here for completeness. It is a useful measure for the case of iterative algorithms (such as iterative deconvolution in Section 13.4) to identify when an iteration has completed. Many iterative algorithms, if not well-posed, can overshoot the true result and must be halted. This was one of the problems encountered in this work in implementing the Simulated Annealing method of Section 13.4.5 and Section 13.6.5.

## 7.5. Texture processing

The aim of this section is to quantify the homogeneity of the data within an image such that we may distinguish between a poor fluid distribution and a good one. The structural property of the distribution may be thought of as a "texture", in that we are interested in the pattern of local distribution or the structure present within the image. Texture occurs over a region rather than a point and has qualities such as periodicity and scale; it may be described in terms of direction, coarseness, contrast and so on. We will now cover some texture analysis and processing techniques, namely standard deviations, autocorrelation functions and the GLCM (Grey Level Co-occurrence Matrix), along with some real-world examples of their applications. These procedures have application in quantifying textures and for feature identification.

### 7.5.1. Methods for quantification of distribution in the literature

Most methods for quantifying the distribution of intensity within an image rely on standard deviations, but give no feeling for the degree of clustering or structure observed in many images.

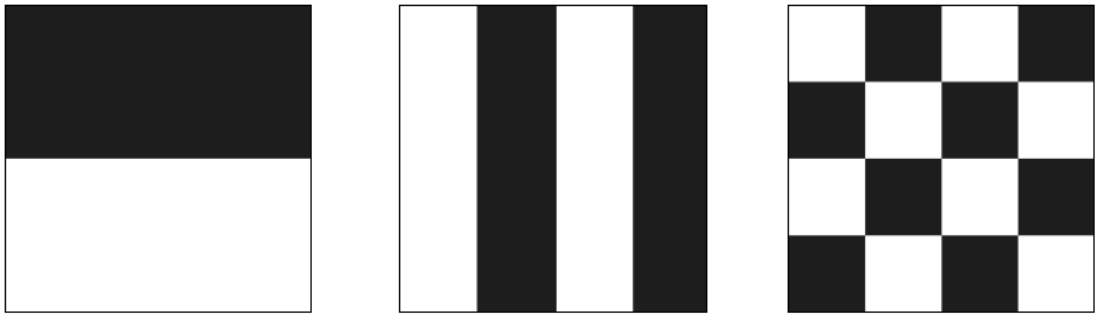

(a) (b) (c)

Figure 7.12 – Three different images with identical intensity distribution and identical standard deviations, but very different spatial distributions.

As can be seen in Figure 7.12, a measurement of the intensity distribution or standard deviation is inadequate to quantify the quality of the distribution, since we do not take into account the relative spatial distribution of intensities. Figure 7.12c consists of a

considerably more homogeneous distribution than Figure 7.12a, and we require a method to quantify this result which is readily apparent from a visual inspection.

There are many maldistribution factors referred to in the literature, but these for the most part are simply variations on the theme of a standard distribution. Some examples include the equations shown in Eq. 7.3 – Eq. 7.7.

$$\mathrm{mf} = \sqrt{\sum_{i=1}^{r} \frac{A_i}{A_0}\left(1 - \frac{|V_i|}{V_0}\right)^2} \quad \text{(Jiang } et\ al.\text{, 1999)} \qquad \textbf{(Eq. 7.3)}$$

$$\Phi_i = \sqrt{\frac{\sum_{j=1}^{4}\left[\left(T_{i,j} - T_0\right) - \Delta T_i\right]^2}{(\Delta T_i)^2}} \quad \text{(Albers, 2005)} \qquad \textbf{(Eq. 7.4)}$$

$$\mathrm{Mf} = \sqrt{\frac{1}{n}\sum_{n}\left(\frac{w_i - w_0}{w_0}\right)^2} \quad \text{(Darakchiev and Dodev, 2002)} \qquad \textbf{(Eq. 7.5)}$$

$$\omega_{L,i} = \left|\frac{u_{L,i} - \bar{u}_L}{\bar{u}_L}\right|,\ \omega_L = \frac{\iint_{A_{mono}} \omega_{L,i}\,\mathrm{dA}}{A_{mono}}, \text{ (Heibel } et\ al.\text{, 2001a)} \qquad \textbf{(Eq. 7.6)}$$

$$M_h = \frac{1}{A}\iint\left(\frac{u_{loc} - u_{avg}}{u_{avg}}\right)^2 dA \quad \text{(Edwards } et\ al.\text{, 1999; Heibel } et\ al.\text{, 2001b)} \qquad \textbf{(Eq. 7.7)}$$

All of the above equations from the literature are based upon standard deviations though their nomenclature and formation differ slightly[(2)]. None are able to differentiate between the data of Figure 7.12.

Wang *et al.* (1998) used a method based on the power spectrum density function (PSDF) to investigate maldistribution in trickle bed reactors. This is mathematically related to the autocorrelation function (Section 7.5.2), and calculated by taking the square of the magnitude of the Fourier transform of the signal, as shown in Eq. 7.8 (Russ, 2002).

$$\Phi(\omega) = \left|\frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty} f(t)e^{-i\omega t}dt\right|^2 = F(\omega)F^*(\omega) \qquad \textbf{(Eq. 7.8)}$$

[2] Note that the individual nomenclatures have not been explained here – the important information is the similarity of the equations, and their dependence on a measure of standard deviation. For more information, please see the specific source publications.

The PSDF describes the general frequency composition of random rata and can therefore be used to analyse the fluctuation of flow rate across the bed cross-section. The peaks in the PSDF represent the wavelengths of the dominant fluctuations of liquid flow rate in the radial dimension. Its applications include estimating texture lengths and, possibly, motion blurring. In their paper, Wang *et al.* (1998) showed that this method could detect that increasing liquid flow rate resulted in more uniform liquid distribution. They also report that the large-scale liquid distribution was not affected greatly by the gas flow rate. This technique could be applied to the monolith reactor or other systems, but has not been included as a part of this work. It was found that the autocorrelation function could provide all these advantages, and this technique is described in the following section.

## 7.5.2. Autocorrelation functions

This section will outline the formulation of a standard autocorrelation function, and then illustrate the creation of a custom function to overcome some of its limitations. This new function will then be validated by processing some real data following a method published in the literature, and finally an extension to this method is proposed to measure characteristic length scales within data images.

### 7.5.2.1. The autocorrelation function

An alternative method to quantify distribution uniformity is to use an autocorrelation function, such as shown in Eq. 7.9.

$$G(u,v) = \iint g(x,y)g*(x+u,y+v)dxdy \qquad \textbf{(Eq. 7.9)}$$

The autocorrelation of an image with itself gives a measure of the length scales of the structure within the image, i.e. lots of fine, short-scale structure would indicate a homogeneous distribution, whereas long-scale structure would indicate a heterogeneous, clustered distribution. This basis of the autocorrelation function analysis involves convoluting the image with itself, and the function effectively scans through each pixel of the image and compares it to those at all distances from it. The resulting autocorrelation image then contains the information of these pixel-pairs,

indicating typical length scales, clustering, orientation directions and other structural features within the original image.

### 7.5.2.2. Experimental application of the autocorrelation function

Some MRI images of a 400 cpsi monolith reactor are shown in Figure 7.13. These are x-y images taken with a field of view of 51 mm and a resolution of 797 μm × 797 μm. The slice thickness is 4 cm and the data was acquired in 156 ms using the RARE pulse sequence. Flow is provided with a showerhead distributor (19 holes of 1 mm at a hexagonal pitch, described in Section 8.2.2) at a range of liquid and gas flow rates. The high intensity is due the presence of water, and fast moving regions appear darker since they move through too fast to be imaged (even with the large slice thickness). This poses a problem since air also appears as low-intensity. But despite these issues with the interpretation of intensity, this simple imaging technique effectively shows that we may obtain a measure of the *quality* of the distribution in each case. The images acquired in Figure 7.13 are clearly very different from each other, and show the variation in distribution obtained with changes in flow rate. Autocorrelation images of these data are shown in Figure 7.14.

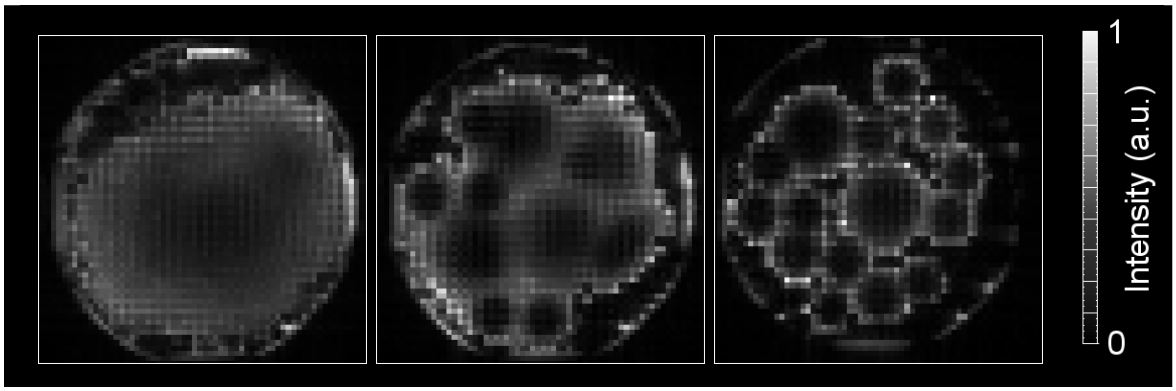


(a) (b) (c)

Figure 7.13 – A summary of the effect seen on increasing the liquid flow rate in the system. (a) 80 l/hr water, 200 cc/min air, (b)100 l/hr water, 200 cc/min air, (c)120 l/hr water, 420 cc/min air.

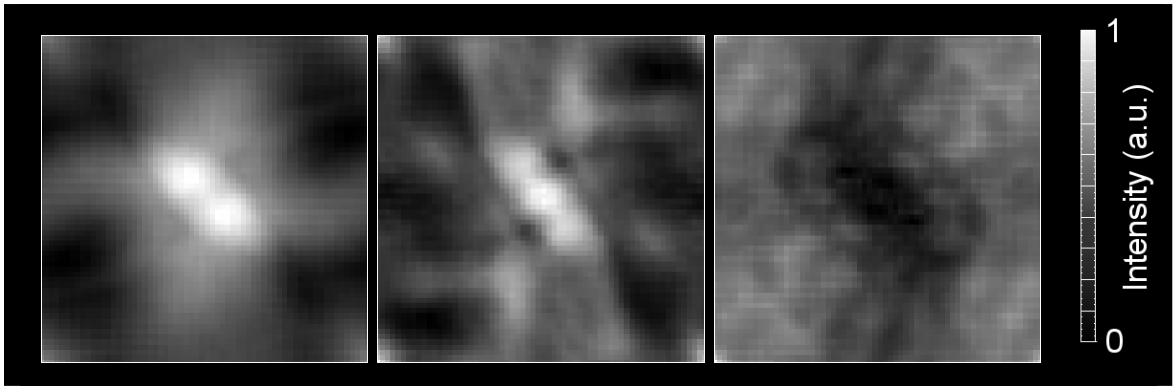


(a) (b) (c)

Figure 7.14 – The corresponding autocorrelation images of the data in Figure 7.13.

The information shown in the autocorrelation images is interpreted in the following manner; the centre of the image corresponds to the origin, and the intensity at any particular point away from the origin indicates the similarity between pixel pairs which are that displacement and angle away from each other. So, a region of high intensity near the origin indicates self-similarity at a short length scale, whereas high intensity further away indicates longer-range similarity. If a distribution is homogeneous, then it will correlate more strongly with itself and there will be a bright region near the centre, as is seen in Figure 7.13a and Figure 7.14a. These autocorrelation images require quantification in order that they may be compared directly and we may achieve a method for ranking the homogeneity of a distribution.

A problem encountered in autocorrelation analyses is that of reactor geometry – the reactors investigated typically have an outer wall, packing elements, or some other sort of geometry (cylindrical geometry in the case of the monolith reactor). Ordinarily we would only wish to consider data from within the reactor itself and neglect surrounding data; however use of the above function results in an autocorrelation image which wraps around the geometry – if the distance under investigation goes beyond the edge of the image, the extra distance wraps around to the other side of the image. Hence the image is treated as a continuous space, rather one with finite geometry. This error can be accounted for (or at least the effects reduced) by dividing through by the autocorrelated image of a binary gated mask image of the data (which therefore contains only contributions from geometry, not data). This was the method used by El-Bachir (2005) and the procedure is outlined in the following section.

## 7.5.2.3. Data analysis method of El-Bachir (2005)

We will now briefly examine a direct application of the autocorrelation function to a real reactor setup. This analysis was originally performed by El-Bachir (2005). The N = 2 reactor (Figure 7.15) is a narrow packed bed and is of interest to steam reforming in particular, but is also a useful case study to develop techniques for application to more complex flow systems. Some sample velocity images are shown in Figure 7.16. The flow is at first glance very homogeneous along the length of the bed, but through careful stability analysis of the data the author showed that there are subtle differences which may be highlighted, allowing the quantification of entrance and exit effects along the length of the reactor.

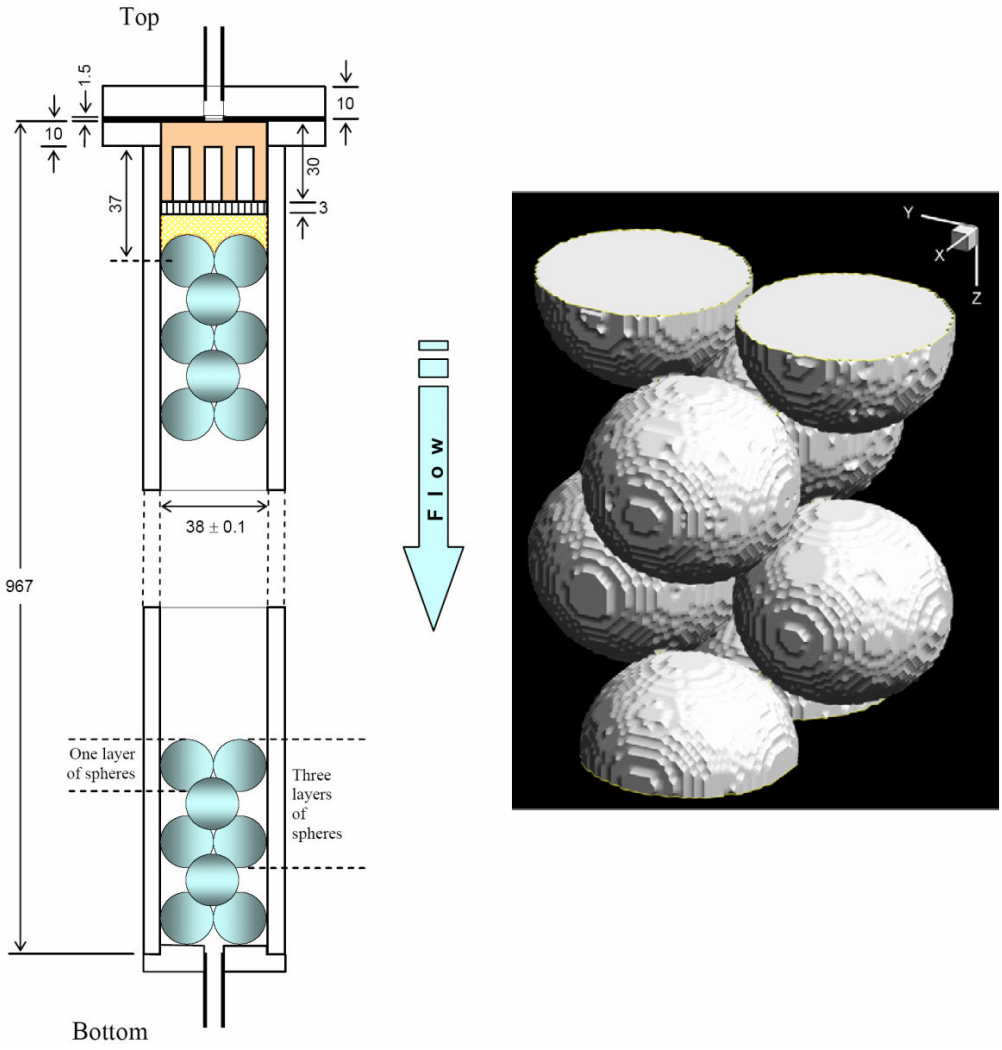


(a) (b)

Figure 7.15 – (a) Schematic of the N = 2 reactor, consisting of 69 layers of plastic spheres inside a Perspex tube; (b) a 3D MRI scan of two layers of packing, illustrating the stacking periodicity.

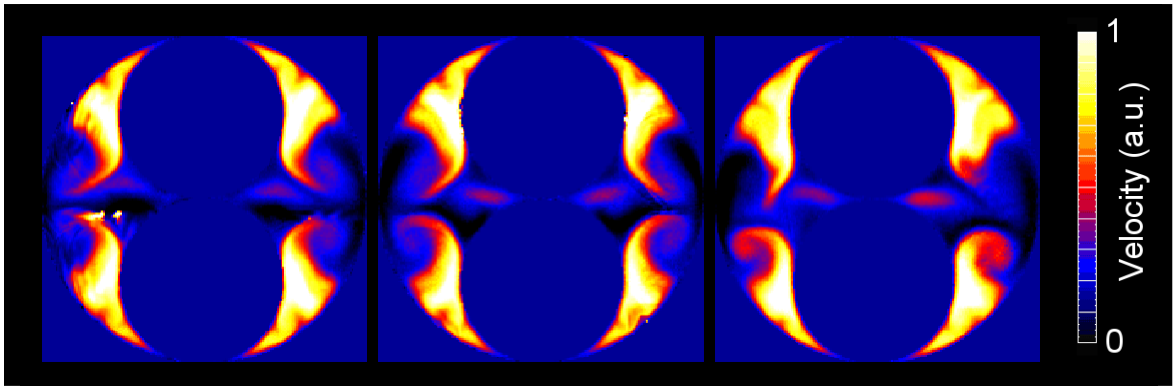


(a) (b) (c)

Figure 7.16 – Sample MRI data from the N = 2 column. Slices shown are from three levels near the centre of the column, where flow should be steady. The Reynolds number (based upon the particle diameter) is 216.

The stability analysis method is illustrated in Figure 7.17. First, data were acquired at each of the 69 layers of balls in the column. It was not possible to compare images directly since there was an issue with incorrect alignment, i.e. the images would not precisely line up. Thus pixels may contain data in one slice and not in another, and in comparing them directly we obtain a large difference which is not truly representative of the system. The autocorrelation function was utilised to compare pairs of pixels within each image, and to then compare together the correlation images of different slices, since this would then be independent of the relative positioning of the original data. Correlation images of neighbouring slices in the bed were subtracted from each other, and the absolute values of these images was summed to obtain a single measure of the overall difference between the two images. Plotting this value as a function of bed height highlighted the fact that slices near the top and bottom of the bed showed larger variations than those in the middle of the bed, shown in Figure 7.18. This is taken as evidence of entrance and exit effects.

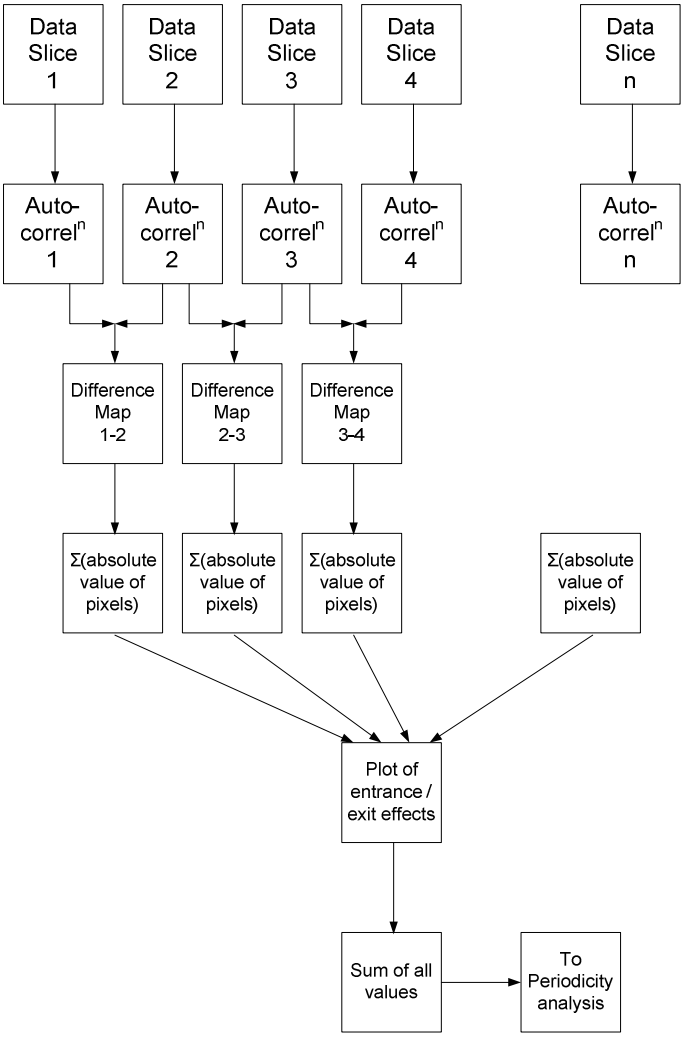


Figure 7.17 – Illustration of the autocorrelation analysis performed by El-Bachir (2005).

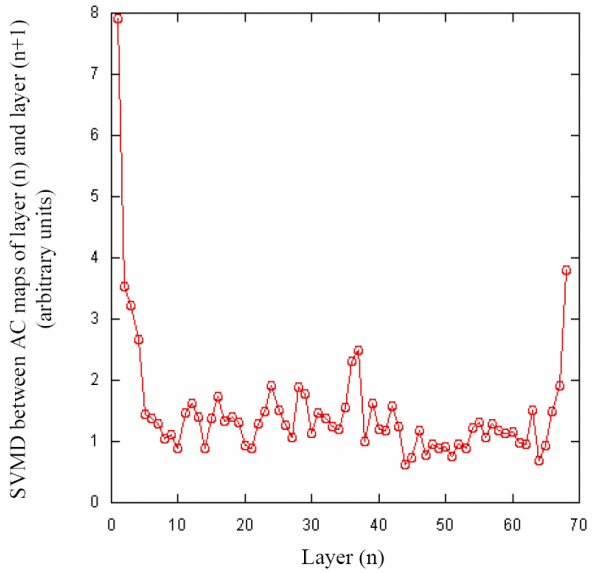


Figure 7.18 – Result of the autocorrelation analysis. Reproduced from El-Bachir (2005).

The drawback with this autocorrelation method is that, as mentioned earlier, the correlation image is geometry dependent and also wraps around the data, which is unrealistic. A new autocorrelation function is now proposed which will overcome the wrap-around effect and the effect of geometry, without the need for post-correction methods. The next sections will outline the creation of this new function, and then compare its performance to that of the original function.

#### 7.5.2.4. Development and validation of a novel autocorrelation function

As a part of this work, a custom correlation function was written to compare only those pixels which contained data, and also to only compare data from pixel pairs within the image, rather than wrapping around the edge of the image. The code was written to allow for 1D or 2D correlations – 1D correlations were of interest to allow temporal correlation analysis for individual pixels (Section 7.5.2.7).

The code uses a discreet form, rather than an integral function, but it will be validated against the original function to show that it produces equivalent results. The procedure is outlined in Figure 7.19. First a jump vector is chosen; this is defined as an (x,y) displacement, whose values range from –X to +X, and –Y to +Y (where X and Y are the dimensions of the data array). Each starting pixel is selected in turn, and we jump from here according to the current jump vector and extract the value of the pixel at this displacement. These two values are multiplied together, and summed for all possible starting pixels. This value is then inserted into the correlation image at the coordinates defined by the parameters in the jump vector (x,y). This is then repeated for all jump vectors, building up the correlation image. For example, if we have a jump vector of displacement 5 pixels in x and 10 in y, we scan through all possible starting pixels, multiply together the starting and end pixels, add all these results together and inset them at coordinate (5, 10). This new formulation neglects jumps that would extend outwith the data area, and also jumps which do not have data at both the starting and ending pixels (as defined by the data geometry). Because not all jumps are possible, the summations we calculate are normalised based upon the total number of jumps which were possible for that particular jump size.

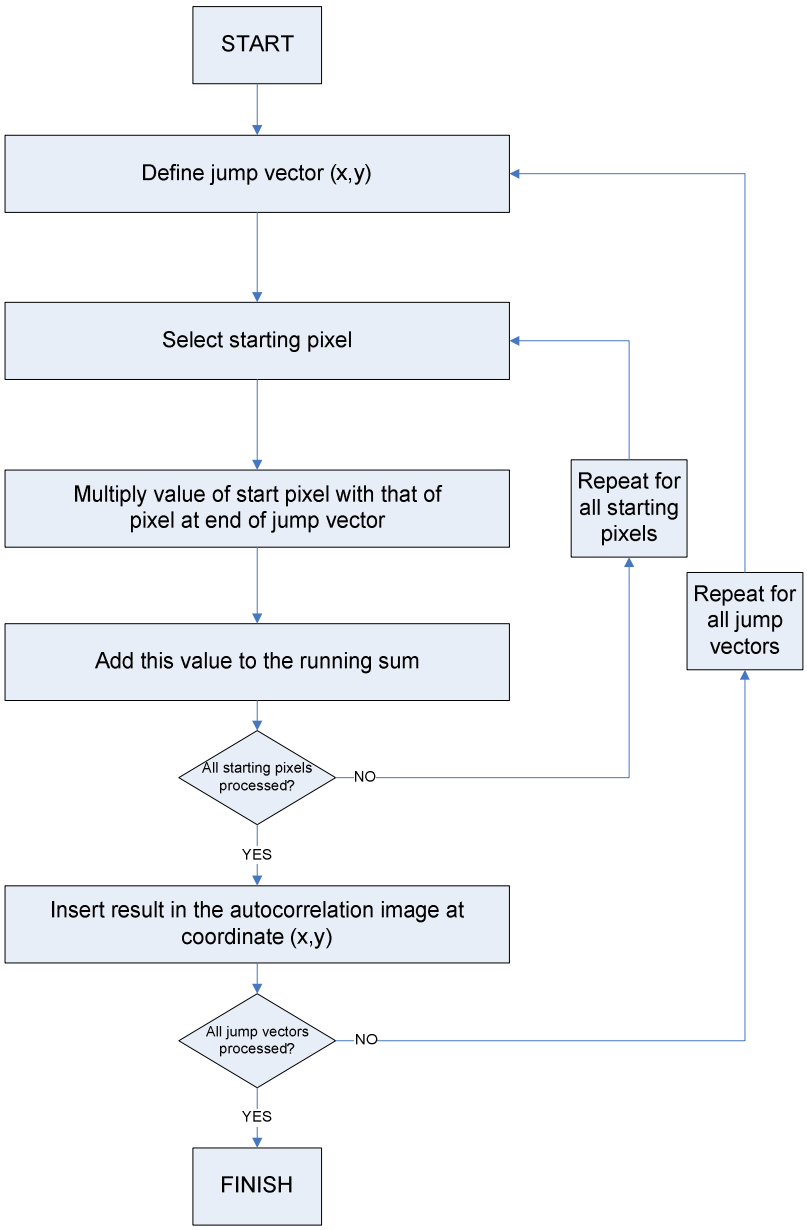


Figure 7.19 – Flowsheet illustration of the novel custom autocorrelation function, which is based on a discreet processing method.

The procedure of El-Bachir (2005) was re-created in order to compare the effect of using four variations of autocorrelation functions. Firstly, using the original autocorrelation function as is; second, removing the geometry factor from the result by subtracting an autocorrelation image of the masked geometry; thirdly using the new autocorrelation function (which is discreet and does not wrap around); and fourthly, using the new autocorrelation function on data pixels only.

The results of this investigation are shown in Figure 7.20. If we compare (a) with (c), and (b) with (d), we can see close similarities between the results for the continuous and discreet autocorrelation functions, and similar features may be identified in each case.

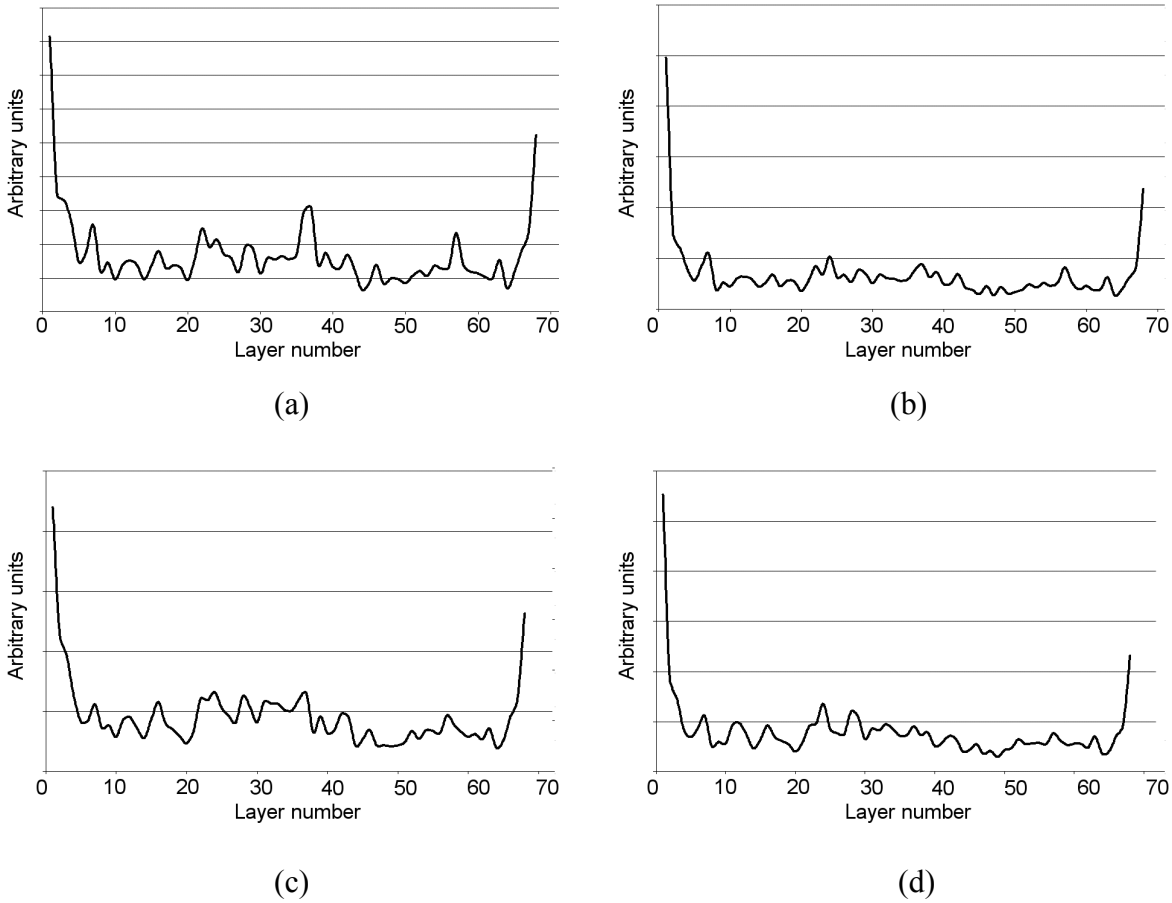


Figure 7.20 – results of the stability analysis of the N = 2 bed, using four different methods. (a) using the original autocorrelation function; (b) using the original autocorrelation function with the original method for removing geometry; (c) using the new autocorrelation function; and (d) using the new function with the new geometry method.

The autocorrelation functions were then applied to a second experimental data set, where 50% of the entrance cross section had been blocked. These data should therefore show more variation at the entrance than the previous results. The results of the autocorrelation analysis are shown in Figure 7.21. Here we see excellent agreement between the two functions [(a) with (c), and (b) with (d)], and this confirms the validation of the new discreet autocorrelation method.

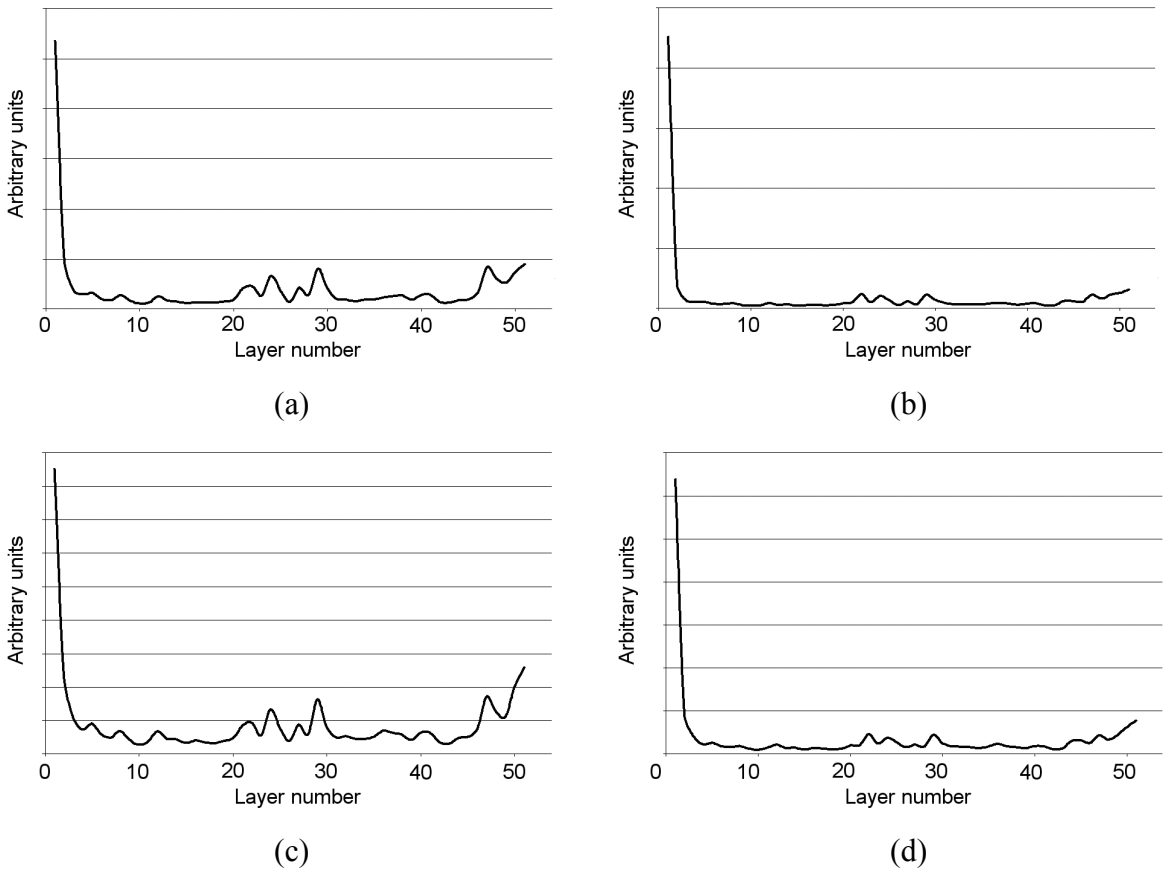


Figure 7.21 – results of the stability analysis of the N = 2 bed under blocked entrance conditions, using four different methods. (a) using the original autocorrelation function; (b) using the original autocorrelation function with the original method for removing geometry; (c) using the new autocorrelation function; and (d) using the new function with the new geometry method.

## 7.5.2.5. Extension of the novel autocorrelation function

We have shown that the custom autocorrelation function is for the most part equivalent to the original autocorrelation function, and therefore validated its applicability. However, the custom function has several advantages – we have far more control over which data are selected for the calculation, and there is no wrap-around effect. Moreover, it can be extended to allow us to perform calculations with different "kernels" to identify different textural properties within our data. For example, we have chosen to *multiply* the pixels at the start and end of the jumps because it best mirrors the analytical convolution method of Eq. 7.9. However, we

could equally have used a different kernel, such as calculating the *difference* between the pixel pairs. This would result in a similarity-weighted image which may be of more value than weighting the image by the square of the pixel values. An example of the difference between these two kernels is shown in Figure 7.22. It is clear from this example that we can see more detail in the new autocorrelation image, and therefore deduce more about the structure of the image, in particular its self-similarity.

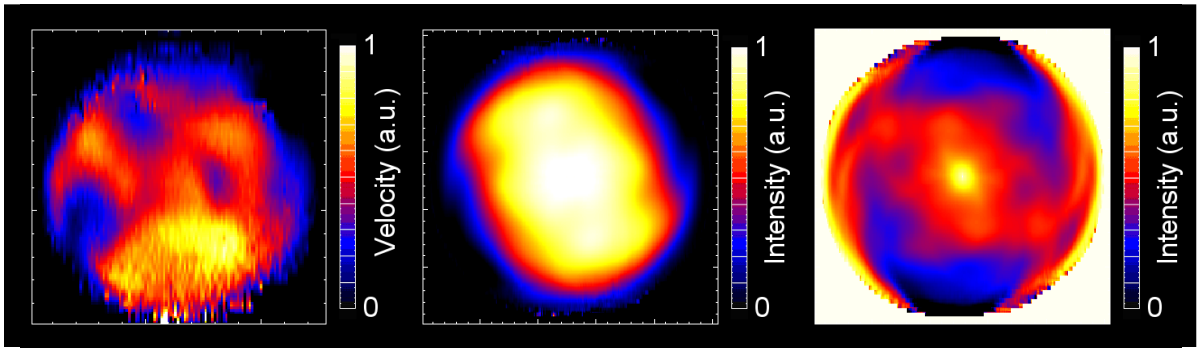


(a) (b) (c)

Figure 7.22 – Example of the autocorrelation function with a subtractive kernel; (a) is some raw data, a MS-COTTAGE velocity image of flow in a pipe at a Reynolds number of 1660, (b) is the autocorrelation image obtained using the original multiplicative kernel, and (c) is the autocorrelation image obtained using the new, subtractive kernel.

This new correlation function was then applied to the data of Figure 7.13 to attempt to quantify the characteristic features, and the resultant autocorrelation plots are shown in Figure 7.23. We can see from these that the bright spot in the centre of the image varies in size, corresponding to a different characteristic feature size in each case. The well-distributed data (Figure 7.13a) give a large self-similar region, whereas the poorly-distributed data (Figure 7.13c) give a small self-similar region corresponding to smaller features. To quantify this characteristic size, radial distribution functions were calculated, shown in Figure 7.24. The radial distribution function shows the average intensity at increasing radial position, weighted by the number of pixels present at that radius.

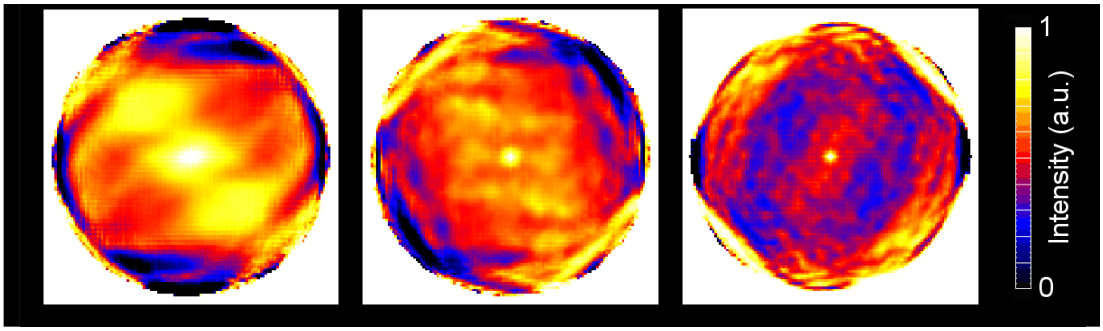


(a) (b) (c)

Figure 7.23 – Auto-correlation images of the data in Figure 7.13, using the subtractive kernel.

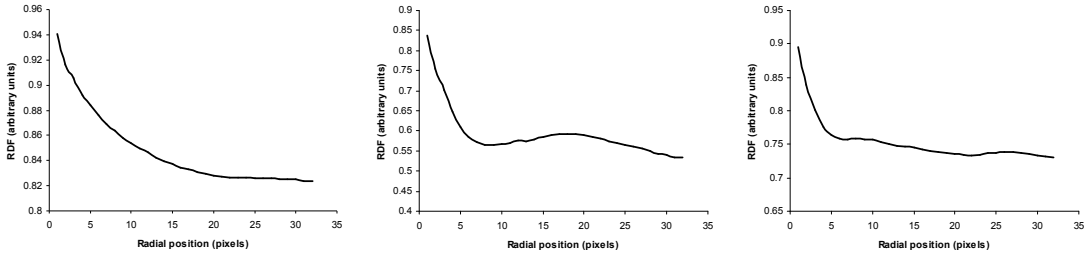


(a) (b) (c)

Figure 7.24 – Corresponding radial distribution functions (RDF) of the data in Figure 7.23.

Three main types of behaviour are observed in the profiles; (a) a single, global minimum; (b) one or more local minima; and (c) overlapping peaks with no local minimum. These are illustrated with synthetic data in Figure 7.25. In the case of a system exhibiting a local minimum (Figure 7.25b), the first peak is due to self-correlation, and the second peak is due to a correlation between neighbouring regions. In this case, the distance corresponding to the first local minima is more relevant than the global minimum since we are interested in the size of individual features (rather than the distance between similar features). For the case of overlapping peaks (Figure 7.25c), we are unable to extract the position of the first local minimum and so the point of inflection is the best approximation of this characteristic distance. Therefore, in order to have a robust technique that could be applied to any profile, it was decided to choose the first point of inflection as our reference point. This will give a meaningful result in all cases which can be compared between data sets, although it may not strictly correspond to the actual size of the features in the data.

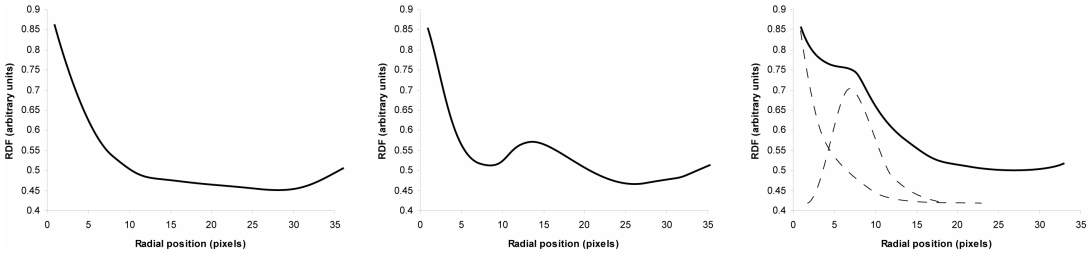


(a) (b) (c)

Figure 7.25 – Synthetic data illustrating the three types of RDF profiles obtained from real autocorrelation data; (a) a single, global minimum; (b) one or more local minima; and (c) overlapping peaks with no local minimum. The dashed lines in (c) indicate possible separate peaks which combine to give the composite curve.

It was decided to define a characteristic length based on the position of the half-height between two points of reference; namely the maximum value and the first point of inflection. Applying this characteristic length measurement to the profiles of Figure 7.24 yielded results of 5.24, 3.29 and 2.69 pixels in each case (note: these are characteristic radii, so the effective diameter of features are 10.48, 6.58 and 5.38 pixels). These numbers do not give a measure of any feature in particular, but rather give a combined average feature size for the details within the data. We may use this statistic measure to rank data in terms of characteristic feature size, and therefore rank the quality of distribution in each case. This method is applied to a range of monolith data in Chapter 10.

## 7.5.2.6. Cross-correlations

There is no reason that we must always perform an *auto*-correlation, i.e. of an image with itself, and indeed the procedure can be adapted to allow cross-correlations where two images are compared pixel-by-pixel for similarities. This can be used to check for subtle differences, or also to check for particular patterns within an image. For example, if we have a feature or shape that we wish to locate or track, then we can use a brute-force method and try correlating all positions of the target pattern with the image, as shown in Figure 7.26 and Figure 7.27. We will obtain a high correlation for the occurrences which correspond to a high probability of match. This is an

alternative method to the Hough transform and Chi$^2$ methods which are presented later in Section 7.6.

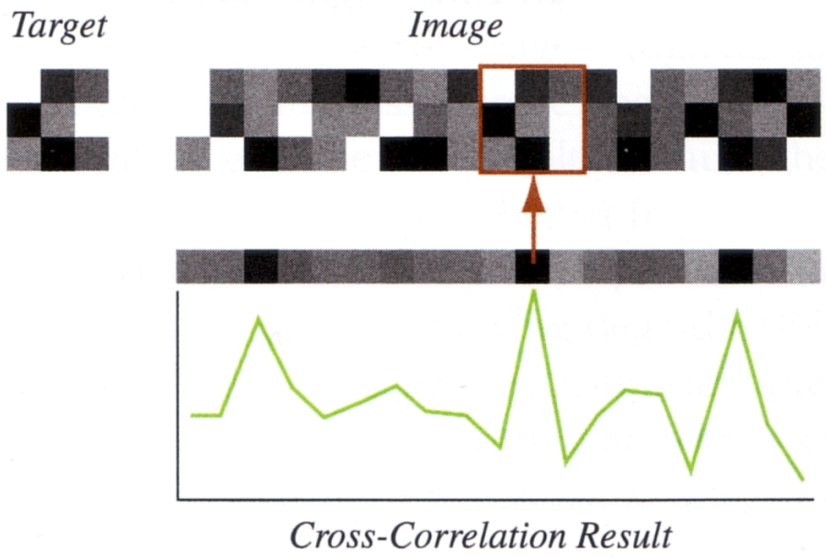


Figure 7.26 – Cross-correlation template matching, showing the target sample and data, with the cross-correlation value shown for each trial position. A high value indicates a good probability of a match.

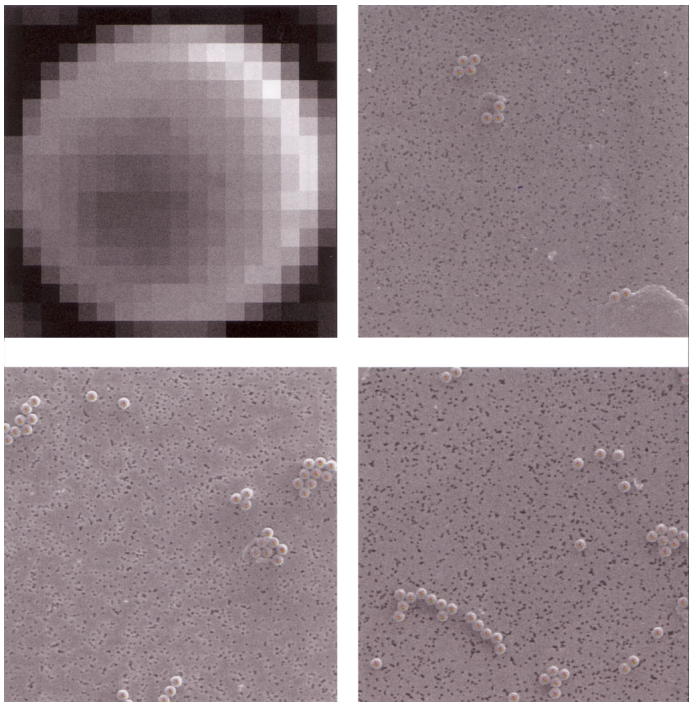

Figure 7.27 – Cross-correlation example, used to detect and count latex spheres on a filter. The first image shows a representative latex sphere, to be used as a search template, and the remaining three images show the identified locations with red spots. Reproduced from Russ (2002).

## 7.5.2.7. Temporal correlations

The autocorrelation or cross-correlation techniques may be adapted to identify temporal correlations in a series of data. For example, if we take a series of channel profiles and examine the intensity fluctuations over time, we obtain data such as Figure 7.28. This shows data acquired with a bubble sparger setup (Section 8.2.1) at a flow rate of 100 cc/min through stagnant water, in a 200 cpsi monolith. Ninety repetitions were acquired, each 5 seconds apart.

The figure shows data from a single row of channels with time in the vertical direction. As can be seen, some channels are continuously filled with liquid, while others fluctuate, indicating the movement of slugs and bubbles in that channel.

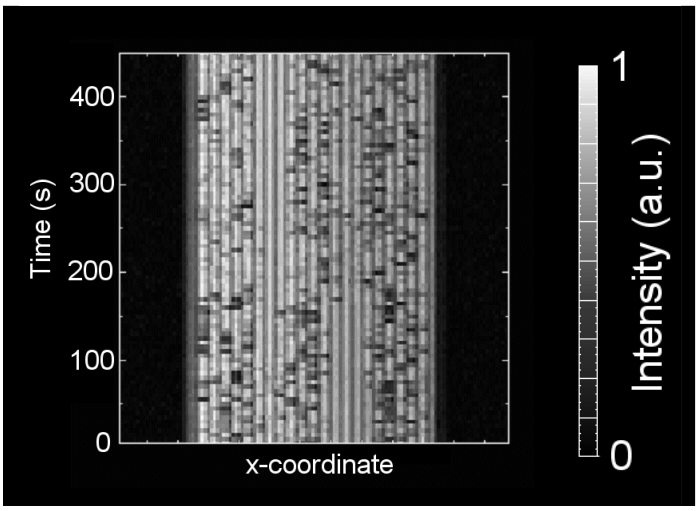


Figure 7.28 – Example of data fluctuations over time. Data are acquired for a 200 cpsi monolith, with the bubble sparger setup (Section 8.2.1) at an air flow rate of 100 cc/min.

By taking a 1D correlation of this data in the time direction, we can deduce two quantities; (i) a measure of the degree of fluctuation present, and (ii) the frequency of such fluctuations, which may be related to the bubble / slug frequency. This frequency analysis depends on the sampling rate – in this case, we acquire data every 5 seconds, so (by the Nyquist rule) we may only observe frequencies of 0.1 Hz or slower. Faster, repetitive imaging techniques (such as COTTAGE, Section 5.4) will allow the observation of higher frequency data.

## 7.5.3. Texture analysis using the GLCM

The Grey-Level Co-occurrence Matrix (GLCM) is an alternative method for quantifying the texture within an image, consisting of a 2D histogram information measure. This section will outline the definition and construction of the GLCM. Then a novel extension of this method to generate a texture-enhanced image is presented, and this is then applied to quantifying texture within turbulent flows.

### 7.5.3.1. Construction of the GLCM

The GLCM and its applicability to texture analysis was first outlined by Haralick *et al.* (1973). A co-occurrence matrix is simply a 2D histogram array in which both the rows and the columns represent a set of possible image values. A histogram is a 1st order statistical representation of the grey-level values in a picture, and the GLCM is a 2nd order representation, describing the position of pixels having similar grey-level values (or more generally, arbitrary local features). It is not simply a statistical probability measure of the image but contains information on the joint probabilities of all pair-wise combinations of similar intensities (Figure 7.29). Haralick *et al.* (1973) report that their automated algorithms are able to correctly identify images of specific sandstones to an accuracy of 89%, aerial photography of land usage to 82% and satellite images of land usage to 84%, by training their algorithms with test data.

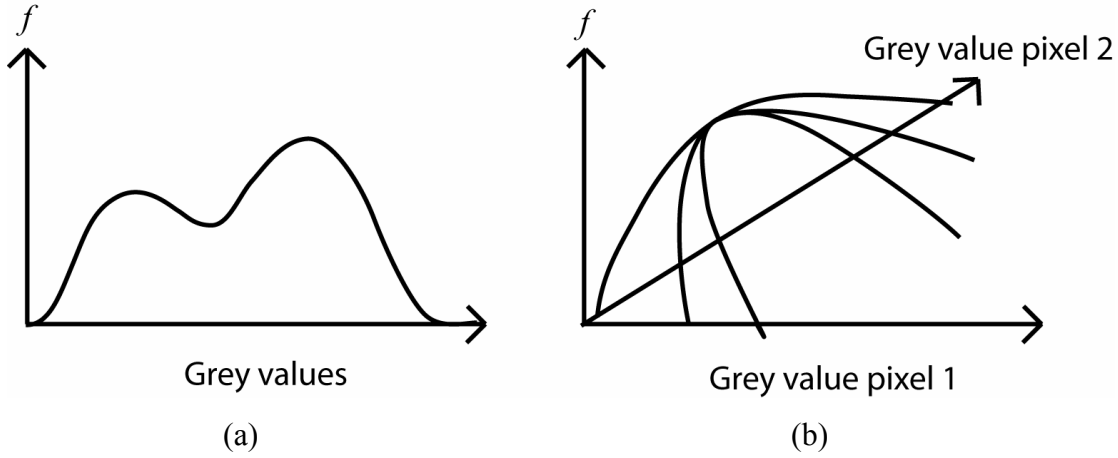


Figure 7.29 – Illustration of (a) first order and (b) second order histogram statistics.

Further manipulation of the information in the GLCM matrix allows the extraction of features such as "contrast" and "entropy", amongst others, which are of statistical use in quantifying the change in distribution quality observed in our data, as will be shown later in this section.

The image being analysed must first be quantised since if each value is a floating point number, there may never be any points with the same value. So, a grey-scale level of typically 64 is used. Since the GLCM has the dimensions of this quantisation level (i.e. 64×64), the choice of quantisation level affects the degree of accuracy obtained. Too few grey levels will result in a highly pixelated GLCM, and the information obtained will be highly averaged, reducing the classification accuracy. However, if we choose too many grey levels then we will have a sparse GLCM array and noise will affect the results obtained.

The GLCM is constructed as follows (illustrated in Figure 7.30)

- First a displacement vector is specified, **d**(dx,dy)
- All pixel pairs this displacement vector apart are computed, and the number of pixel pairs having values $i$ and $j$ are counted, giving $n_{ij}$ and these are plotted in the GLCM at coordinate $P_d[i,j]$ (The representation $P_d$ is given to the GLCM array)
- So, $n_{ij}$ is the number of occurrences of the pixel values $(i,j)$ lying at distance **d** in the image. It should be noted that the GLCM is not symmetric since the number of pairs of pixels having grey levels $[i,j]$ does not necessarily equal the number of pixel pairs having grey levels $[j,i]$.
- The GLCM is typically normalised by dividing each entry by the total number of pixel pairs. The normalised GLCM, $\boldsymbol{N}[i,j]$ is defined by Eq. 7.10.

$$N[i,j] = \frac{P[i,j]}{\sum_i \sum_j P[i,j]} \qquad \textbf{(Eq. 7.10)}$$

  which normalises the co-occurrence values to lie between 0 and 1, and allows them to be thought of as probabilities.

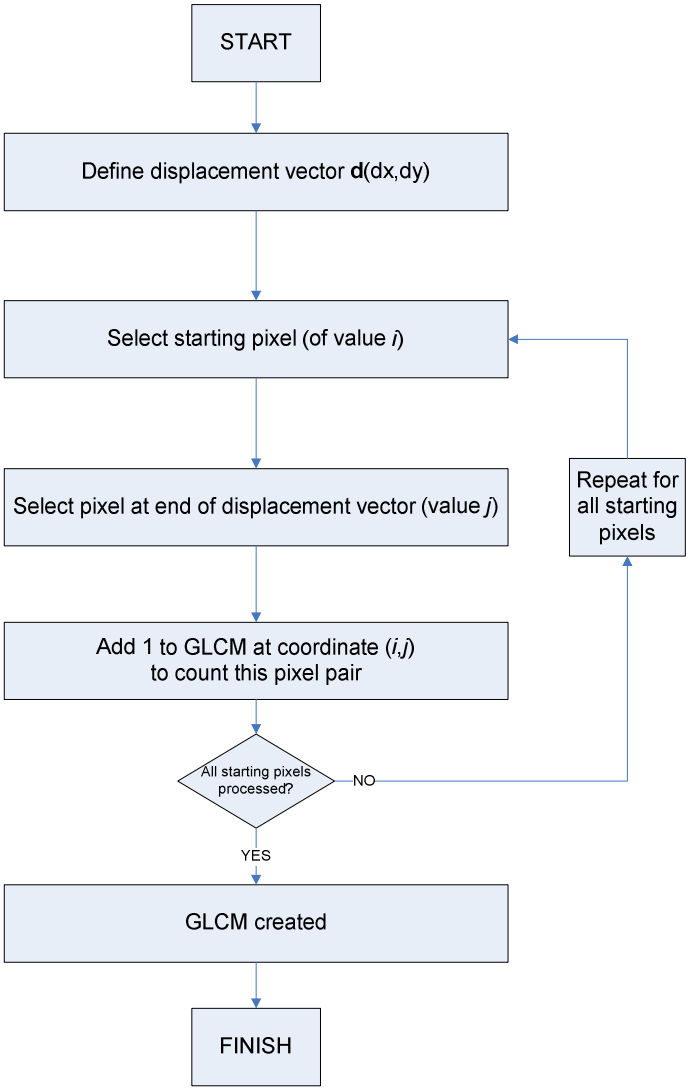


Figure 7.30 – Flowsheet illustration of the GLCM creation process.

An example of the creation of a GLCM is shown in Figure 7.31.

| | | | | |
|---|---|---|---|---|
| 2 | 1 | 2 | 0 | 1 |
| 0 | 2 | 1 | 1 | 2 |
| 0 | 1 | 2 | 2 | 0 |
| 1 | 2 | 2 | 0 | 1 |
| 2 | 0 | 1 | 0 | 1 |

$$P_d = \begin{array}{ccc|c} 0 & 2 & 2 & 0 \\ 2 & 1 & 2 & 1 \\ 2 & 3 & 2 & 2 \\ \hline 0 & 1 & 2 & \end{array} \quad i \quad j$$

Figure 7.31 – Example of the creation of the GLCM, for the case of **d** = (1,1). There are 16 pairs of pixels in the image which satisfy the spatial separation **d** (1,1). Since there are only three grey levels, $P_d[i,j]$ is a 3×3 matrix as shown.

Some observations of the properties of the GLCM are that:

- If the original image contains randomly distributed intensity values, then the GLCM will tend to have similar values in every cell. If the original image contains pixel pairs of the same intensity corresponding to the displacement vector, then the GLCM will contain a large spike.
- An image with coarse texture will tend to have large values along the diagonal of the GLCM, whereas fine texture will cause values to spread out from the diagonal.

One issue with the GLCM is how to choose the displacement vector **d**. Typically one would choose a distance characteristic of the features in the image in question. However, we require a more robust method which will allow us to select those spatial displacements that best capture the structure of the texture. One option is to use a statistic calculation such as Eq. 7.11 to select the vector that gives the maximum amount of textural (structural) information (Zucker and Terzopoulos, 1980).

$$\chi^2(d) = \sum_i \sum_j \frac{P_d[i,j]^2}{r_i c_j} - 1, \text{ where } r_i = \sum_j P_d[i,j] \text{ and } c_j = \sum_i P_d[i,j] \qquad \textbf{(Eq. 7.11)}$$

This is essentially a statistical measure of whether the elements of $P_d$ are independently distributed or correlated. The value of **d** that gives the maximum value of $\chi$ is chosen.

To understand the importance of selecting the spatial relationship properly, consider a texture with a dominant banded structure in the horizontal direction, but with a random structure in the vertical direction (Figure 7.32). If we build a GLCM using a vertical spatial relationship between pairs of pixels, then it would appear as if the underlying texture were completely random. On the other hand, if we use a horizontal relationship we would observe that the texture was highly structured. Computing a GLCM as an average over several directions would yield a description somewhere between these two. Since our data does not typically have a preferred orientation, it was decided to take an average of vectors in 8 directions (0°, 45°, 90°, 135°, 180°, 225°, 270° and 315°), and sum the resultant GLCMs together. The $\chi^2$ statistics were calculated at various step lengths to find the optimum which corresponds to the characteristic length scale for the data.

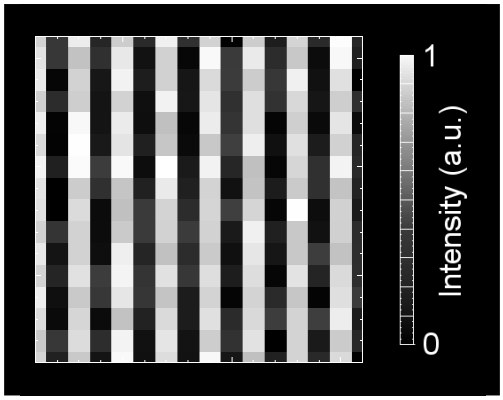


Figure 7.32 – An example of a data set which has textural structure in the horizontal direction, but random variation in the vertical.

The GLCM captures properties of a texture, but they are not directly useful for further analyses, such as the comparison of two textures. Numeric features are computed from the co-occurrence matrix that can be used to represent the texture more compactly. Many different statistical measures are proposed in the literature (BroNielsen, 1997; Clausi, 2002; Wirth, 2002a), and they are summarised in Eq. 7.12 - Eq. 7.19.

- **Moments** – the order $k$ element difference moment is defined as:

$$Mom_k = \sum_i \sum_j |i-j|^k \; P_d[i,j] \qquad \textbf{(Eq. 7.12)}$$

- **Contrast** – a measure of the local variations present in an image. If there is a large amount of variation in an image, the $P[i,j]$s will be concentrated away from the main diagonal and the contrast will be high. Typically, $k = 2$, $n = 1$, making it the same as the **2nd Difference Moment**, and is also known as **Inertia**.

$$C(k,n) = \sum_i \sum_j (i-j)^k P_d[i,j]^n \qquad \textbf{(Eq. 7.13)}$$

- **Homogeneity** – a homogeneous image will result in a co-occurrence matrix with a combination of high and low $P[i,j]$s. Where the range of grey levels is small, the $P[i,j]$s will tend to be clustered around the main diagonal. A heterogeneous image will result in an even spread of $P[i,j]$s. This is the same as the **1st Inverse Difference Moment**.

$$C_h = \sum_i \sum_j \frac{P_d[i,j]}{1+|i-j|} \qquad \textbf{(Eq. 7.14)}$$

- **Entropy** – a measure of information content. It measures the randomness of intensity distribution. Entropy is highest when all entries in $P[i,j]$ are of similar magnitude, and small when the entries in $P[i,j]$ are unequal.

$$C_e = -\sum_i \sum_j P_d[i,j] \ln P_d[i,j] \qquad \textbf{(Eq. 7.15)}$$

- **Correlation** – a measure of image linearity. Correlation will be high if an image contains a considerable amount of linear structure.

$$C_c = \sum_i \sum_j \frac{(i-\mu_i)(j-\mu_j)P_d[i,j]}{\sigma_i \sigma_j} \qquad \textbf{(Eq. 7.16)}$$

$$\mu_i = \sum_i i P_i[i], \quad P_i[i] = \sum_j P[i,j] \quad \text{Hence, } \mu_i = \sum_i \sum_j i P_d[i,j]$$

$$\sigma_i^2 = \sum_i (i-\mu_i)^2 P_i[i], \; \sigma_i^2 = \sum_i \sum_j (i-\mu_i)^2 P_d[i,j]$$

- **Uniformity** – also known as the **Angular Second Moment**, the square root of uniformity is also known as **Energy**.

$$C_u = \sum_i \sum_j P_d[i,j]^2 \qquad \textbf{(Eq. 7.17)}$$

- **Dissimilarity** – this is the same as the **1st Difference Moment**.

$$C_d = \sum_i \sum_j P_d[i,j] |i-j| \qquad \textbf{(Eq. 7.18)}$$

- **Inverse difference moments:**

$$C_{ID} = \sum_i \sum_j \frac{P_d[i,j]}{1+|i-j|^k} \qquad \textbf{(Eq. 7.19)}$$

This series of equations sum all pixels in the GLCM to give a single number, which is of use in calculating statistical measures.

### 7.5.3.2. A novel extension of the GLCM texture information

The GLCM contains a good deal of information on the textural properties of the image, and so rather than just extracting a single number from the GLCM it is proposed that it would be more useful to interpret the textural features contained within the GLCM to highlight areas of interest within the original image. This process is not documented in the literature and is a novel implementation with useful properties, particularly in its ability to amplify features in an image, thereby simplifying further classification.

The proposed procedure is outlined in Figure 7.33. Instead of summing over all values of $i$ and $j$ (as is done in Eq. 7.12 – Eq. 7.19) we calculate the contribution from each $[i,j]$ pair and this is re-inserted back into the GLCM. The GLCM now contains information modified to reflect the property being calculated. The GLCM creation process is repeated, scanning through the original image for pixel pairs, but instead of putting $n_{ij}$ into the GLCM as before, the value from the GLCM is extracted and replaced into the original image. This results in the frequency of that pixel pair (as measured in the GLCM) being inserted into the original image. Hence, the most common features within the image will have a high intensity, with low-frequency features having a low intensity. This results in an image which resembles the original, but for example using the contrast equation, regions of high contrast will be exaggerated relative to regions of low contrast.

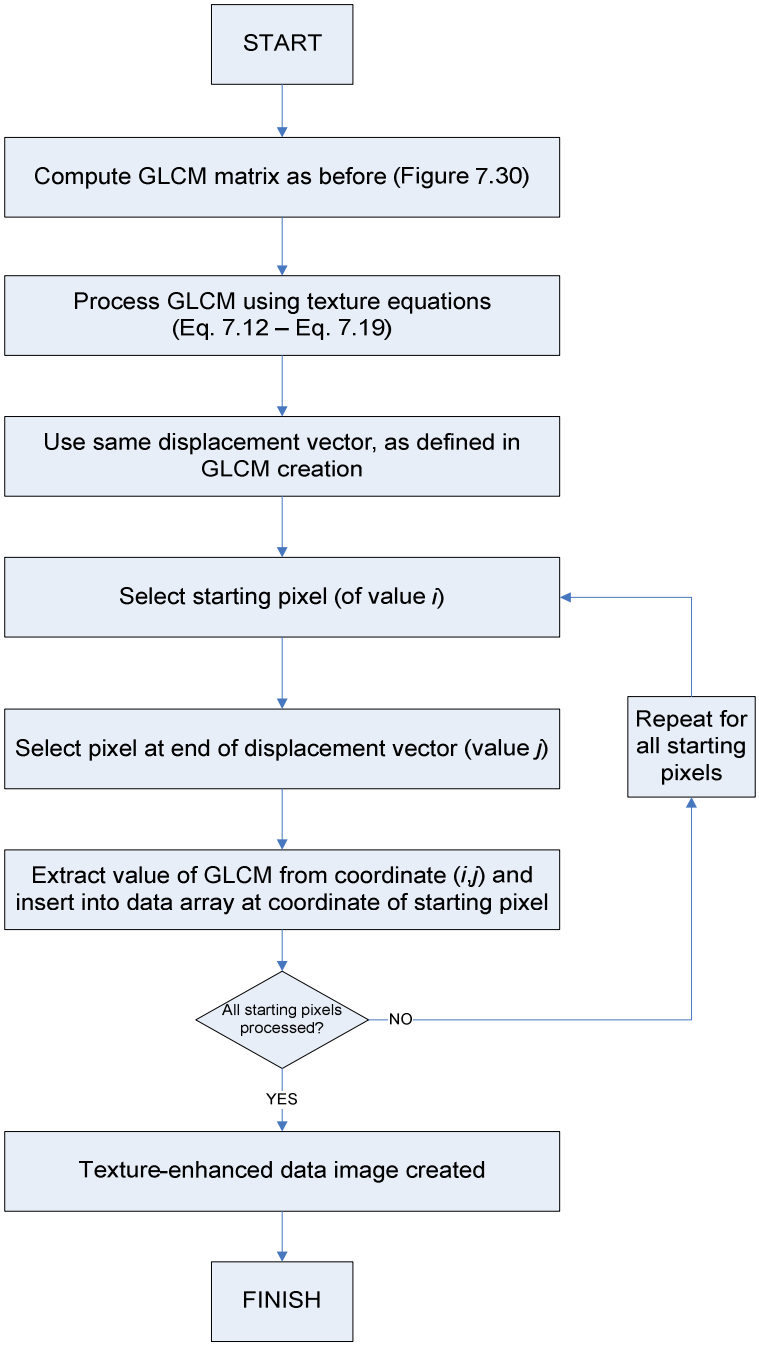


Figure 7.33 – Flowsheet illustration of the process to create a texture-enhanced data array, using the GLCM texture kernels.

A dataset of unsteady laminar flow in a 5 cm diameter tube was used as a test example. Data were acquired with the RARE pulse sequence (Section 2.3.1), and contrast is due to flow contrast. Insufficient length was allowed for true laminar flow to develop so the near-parabolic flow is skewed; in these particular data the maximum flow rate is towards the lower right quadrant rather than in the centre. The results of processing this data at a step length of 8 pixels are shown in Figure 7.34. It can be

seen that the different texture kernels give different results, each amplifying different types of feature present within the original data. By amplifying features of interest in this way, we ease the application of further data processing techniques to our data, allowing differences between images to be more readily apparent.

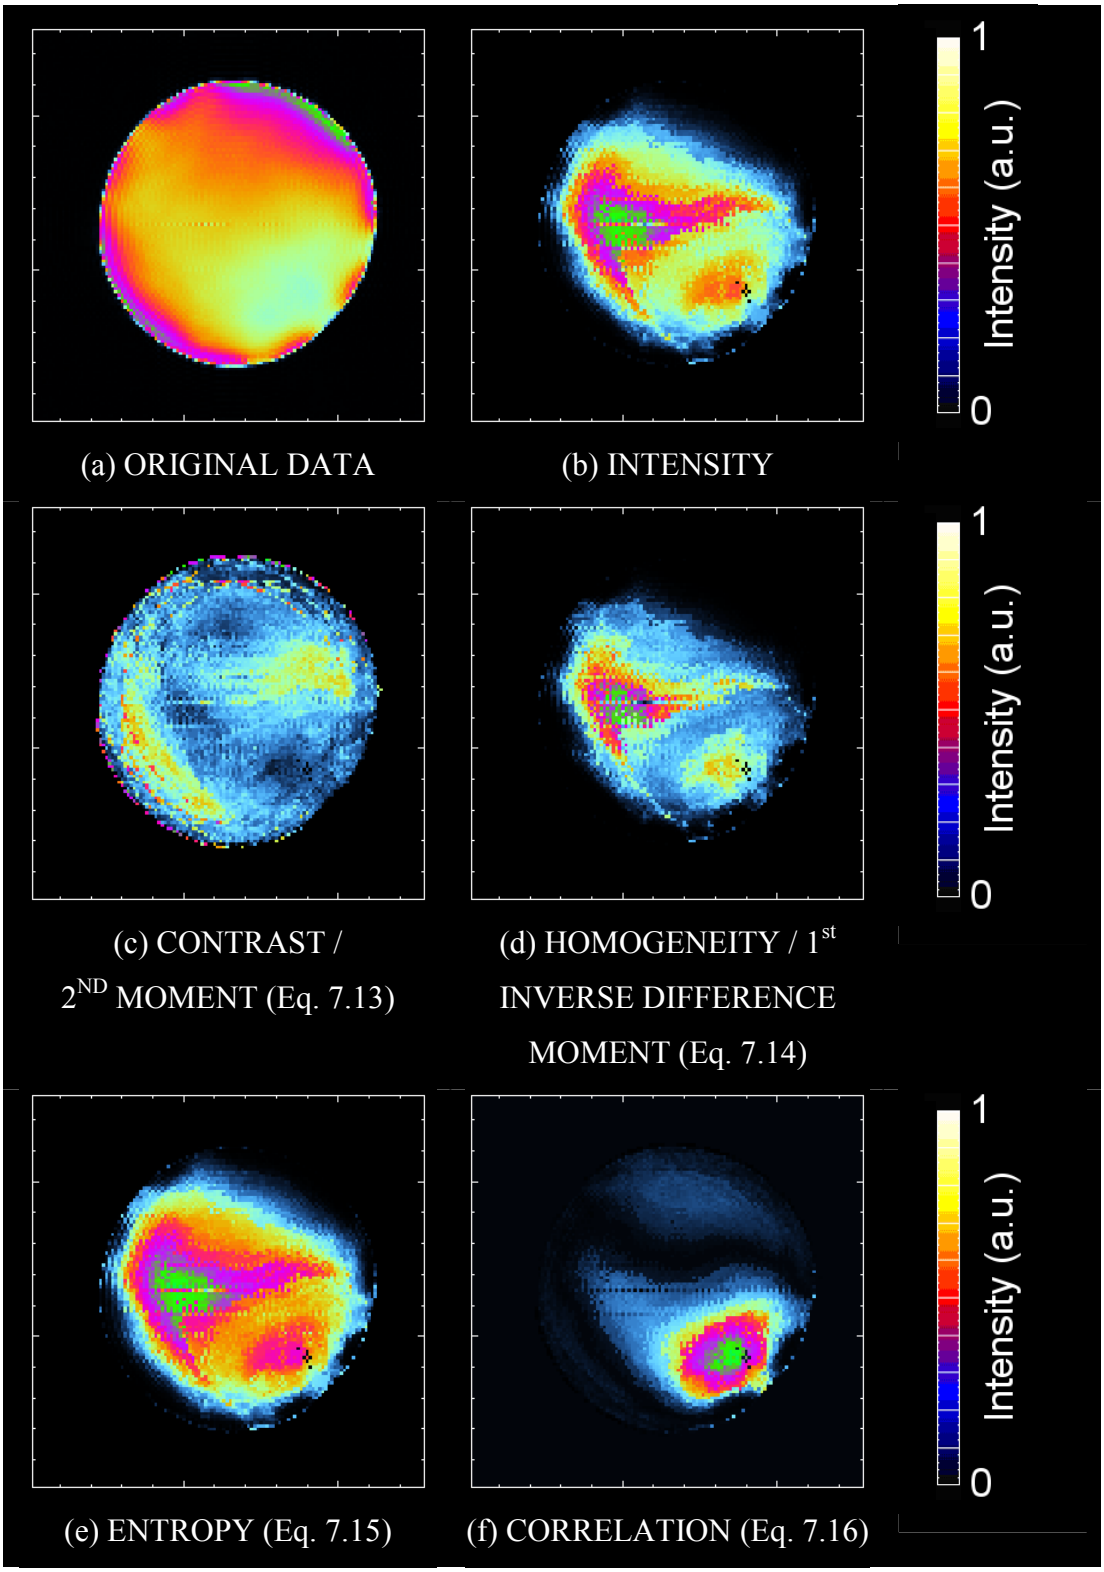

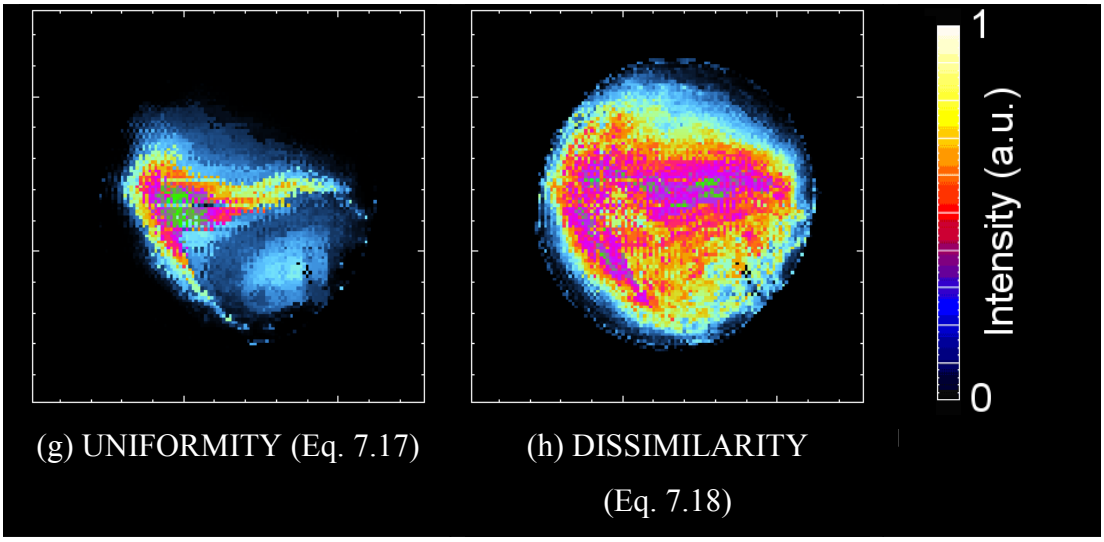


Figure 7.34 – Examples of the GLCM texture enhancement procedure; (a) shows the original input data (RARE spin-density image with flow contrast, of unsteady flow in a pipe); (b) is the texture enhancement with no manipulation of the GLCM; (c)-(h) show the texture enhancement produced by using the various texture kernels on the GLCM.

Many of the images of Figure 7.34 bear a resemblance to each other, and it was reported by Clausi (2002) that they do not all need to be calculated for a comprehensive textural description. The author reports that there are in essence three main types of statistical measure – smoothness statistics, homogeneity statistics, and the correlation statistic.

- Moments of the non-zero probability entries about the GLCM diagonal measure the degree of smoothness of the texture. Contrast, dissimilarity, inverse difference moments are some of these. Of these, Clausi chooses contrast since it produces the clearest results.
- Homogeneity statistics (entropy, uniformity, homogeneity) measure the uniformity of the non-zero entries in the GLCM. Entropy apparently consistently produces the best results, and is chosen.
- The correlation feature produces results that are not correlated with other statistics, and so is added to the group for the new information it brings.

So, the Contrast, Entropy and Correlation processes are suggested as key statistics when analysing textural features. It is therefore proposed to use these to compare

distributions in the monolith reactor, the results of which are summarised in Chapter 10.

There are some additional points regarding the need for care in implementing the GLCM. Not all jumps are possible – near the edge of the image, we cannot jump in all directions since our data does not wrap-around. Hence, when averaging together the values from displacements at each angle, the results were normalised based on the total number of jumps that were possible. There is also a problem in specifying step lengths that are very long, since little information will be obtained – the number of possible jumps is greatly reduced, and this affects the weighting of the final image. Also, the origin of the possible jumps are not homogeneously spread over the image space, i.e. if we imagine a jump of 70 pixels, then for a 128×128 pixel image, the centre of the image cannot be used as the starting point, and so no data is obtained at this point in the image. This is illustrated in Figure 7.35. This can only be avoided by keeping the step length short relative to the dimensions of the data array, and the $\chi^2$ selection method (Eq. 7.11) will naturally avoid these low-information cases.

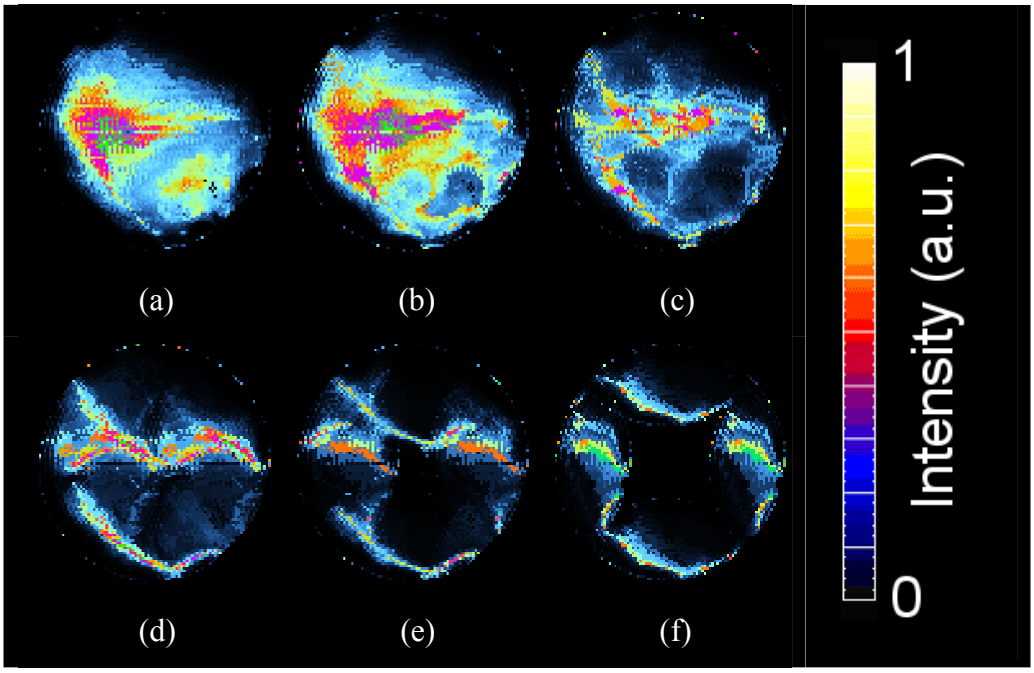


Figure 7.35 – The data of Figure 7.34a processed with the homogeneity function (Eq. 7.14), with step lengths of 10, 20, 30, 40, 50 and 60 pixels.

### 7.5.3.3. Applications of the GLCM texture analysis

This section will summarise a novel observation of the GLCM texture enhancement process, and then show an example of its application to the quantification of turbulence in turbulent MRI data.

An interesting feature of the contrast procedure with a step size of 1, is that it produces a contour map with the number of contours drawn equal to the grey-level quantisation, as shown in Figure 7.36. This figure was generated by applying the contrast texture kernel with a step size of 1 and a grey-level quantisation of 32 to the data of Figure 7.34a, thereby giving 32 contours. This contour-generation feature is explained as follows. On reducing the grey-level quantisation, we reduce the colour depth of the data, which gives a banded effect. In processing with the contrast kernel we therefore identify local regions of high contrast, corresponding to the edges of the banded regions. This then appears as a contour map in the texture-enhanced image.

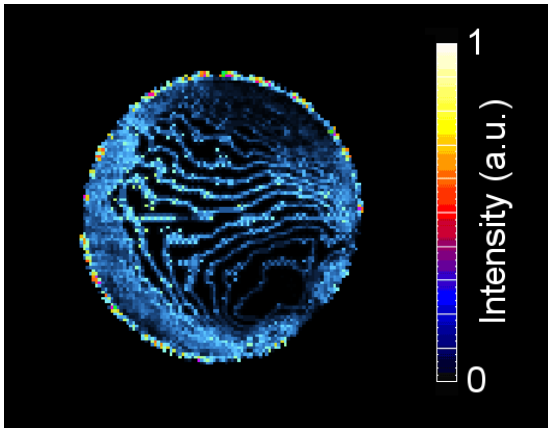


Figure 7.36 – A contour map, obtained by running the contrast procedure with a step size of 1 and a grey-level quantisation of 32 on the data of Figure 7.34a.

The contrast weighting feature was identified as one of the most useful methods for analysing flow heterogeneity. This is because the more turbulent or uneven an image is, the longer the perimeter of the contour lines will be. Hence, by summing the value of the contrast-weighted image, we can obtain a measure of the degree of complexity of the image. This technique has been applied to rapid GERVAIS (Section 5.6) images of turbulent flow in a pipe, and the results will now be summarised. A sample x-velocity data set is shown in Figure 7.37a, along with the contrasted-weighted

image in Figure 7.37b which shows the identification of inhomogeneous regions. MRI velocity data were acquired in the x-, y- and z-directions using GERVAIS in 60 ms (20 ms per velocity component), for turbulent flow in a 29 mm diameter pipe. The field of view is 40 mm, giving a resolution of 313 μm × 313 μm. Superficial flow velocities ranged from 2.5 to 23 cm/s, corresponding to Reynolds numbers from 700 to 6500.

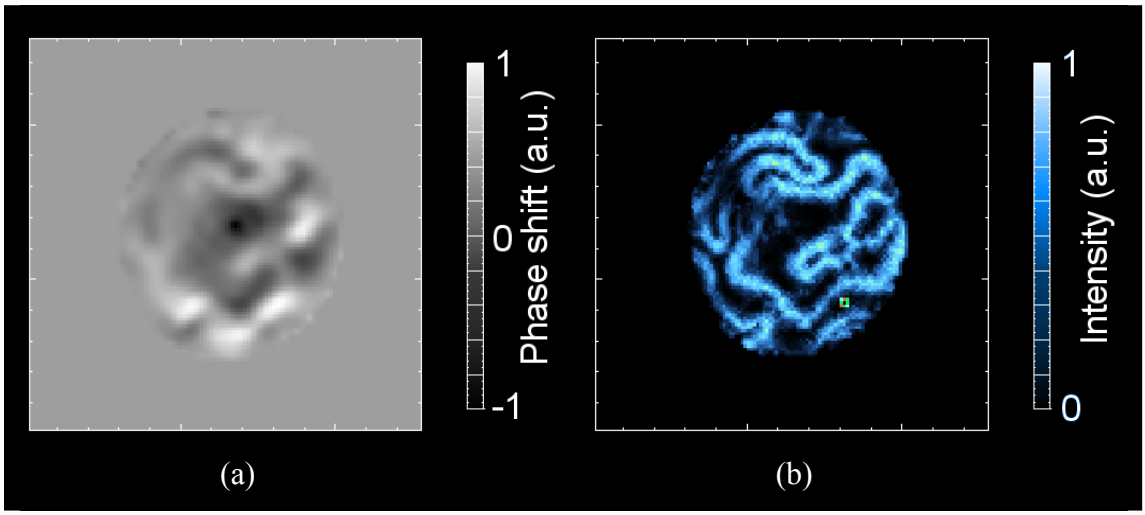


Figure 7.37 – (a) Sample x-component GERVAIS velocity image; (b) the corresponding contrast-weighted image generated for use in this analysis.

On increasing the superficial velocity from 2.5 to 23 cm / s, we see the flow change from steady parabolic flow, through unsteady flow, to fully developed turbulent flow. The effects of this transition in the x-, y- and z-velocity data are shown in Figure 7.38. The calculations were performed by texture-enhancing 100 turbulent images at each flow rate. The values of these images were summed to give a "contrast index", and the standard deviation of the contrast index is shown as an error bar to give an indication of the variation present within these 100 images.

We see a generally linear increase in the complexity of the x- and y-images, however in the z-direction there is a clear increase at around 5 l/min, corresponding to a Reynolds number of approximately 3500 (this is likely to represent a transition point, perhaps from unsteady flow to turbulent flow). It is seen that the standard deviations are much higher, in relative terms, for the x- and y- directions than for the z-

dimension, indicating that the flow is more consistent in the z-dimension, and more unsteady in the x- and y- dimension.

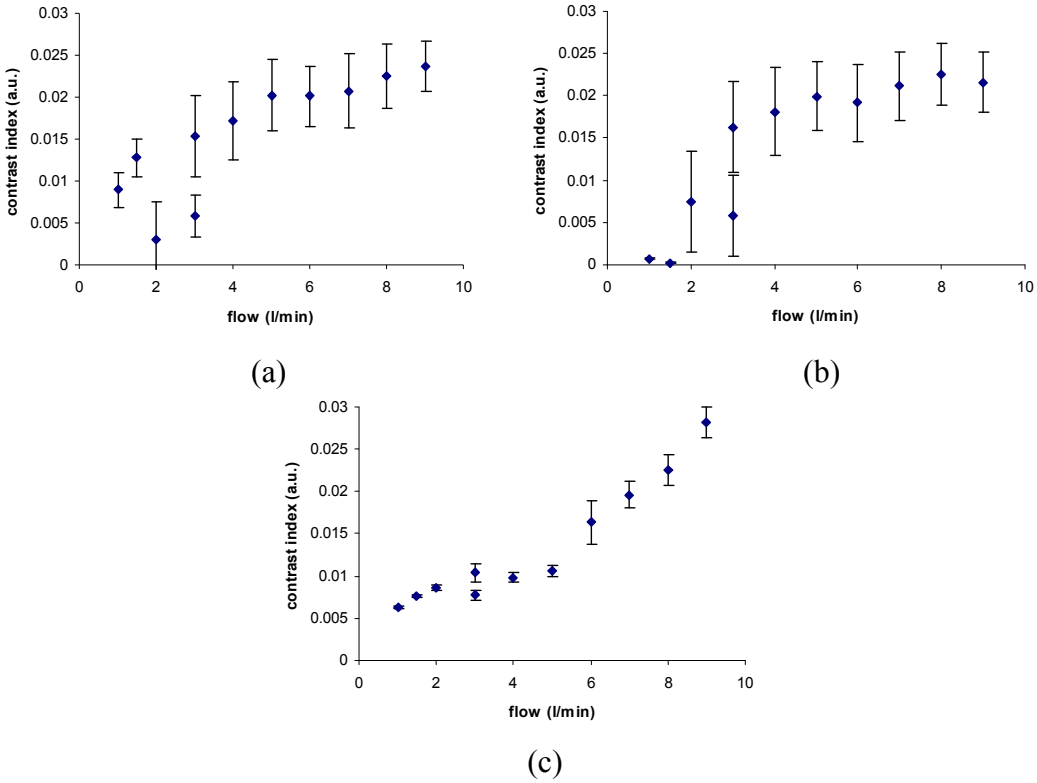


Figure 7.38 – The results of applying the contrast-weighted GLCM texture analysis to turbulent GERVAIS data, for each of the (a) x-, (b) y- and (c) z- directions at increasing flow rate.

So in summary, we have implemented a novel extension of the GLCM to enhance the texture present within an image, and identified a special case of the contrast GLCM which allows the creation of contour plots. This has then been applied to the quantification of turbulence in turbulent flow.

Further reading on other forms of texture analysis can be found in Wirth (2002a), where they cover graylevel difference statistics, runlength statistics, edge enhancement, Hurst coefficients and texture segmentation. Also of interest may be developments in machine vision, i.e. robotic navigation, (Kelly, 1996).

## 7.6. Feature processing

This section is concerned with the location of features within an image, which is of use for image recognition and feature tracking applications. Specifically, these techniques are of interest to track slug profiles in SEMI-RARE data (Section 5.2), in order to measure the velocity within monolith channels. Two methods are investigated, firstly that of the Hough transform, and secondly a custom $Chi^2$ tracking algorithm. These are investigated in application to synthetic data, and then to real MRI data to evaluate their performance under non-ideal conditions.

### 7.6.1. Construction of the Hough transform

The Hough transform (Hough, 1962) refers to a series of algorithms which try to locate parametrically described features within an image, thereby identifying and recognising patterns within an image. These patterns are systematically searched for and any "hits" are recorded in a voting-system procedure. Applications of this technique include feature identification such as finding lines, circles or any parametric shape within an image, for example locating pellet outlines or spherical packings without the need for a complete outline to be present. The original application of the Hough transform was to track paths of charged particles in bubble chambers where the arcs of the paths are often very close together. The technique was designed to separate overlapping tracks and aid in the sifting of the large amounts of data to identify tracks of potential interest for further processing. It was one of the first attempts to automate a visual inspection task, previously requiring hundreds of man-hours to execute.

For example, if we are interested in locating the occurrence of straight lines within an image we may parameterise our search in terms of a slope and an intercept. Hough space will then be 2-dimensional, with each dimension corresponding to our respective parameters. An example of the geometric construction is shown in Figure 7.39. The lines (a)-(c) are sampled at approximately 10 points, and for every point a line is drawn in Hough space using the coordinates of the sampled point as the parameters of the line, as shown in plots (d)-(f). It is an exact theorem that if the series of sampled points lies on a straight line, then in Hough space the lines will

intersect, and the coordinates of the knot correspond to the parameters of the original line in real space. By taking the coordinate of these knots in Hough space and plotting lines back in real space using these parameters, we recover the original lines.

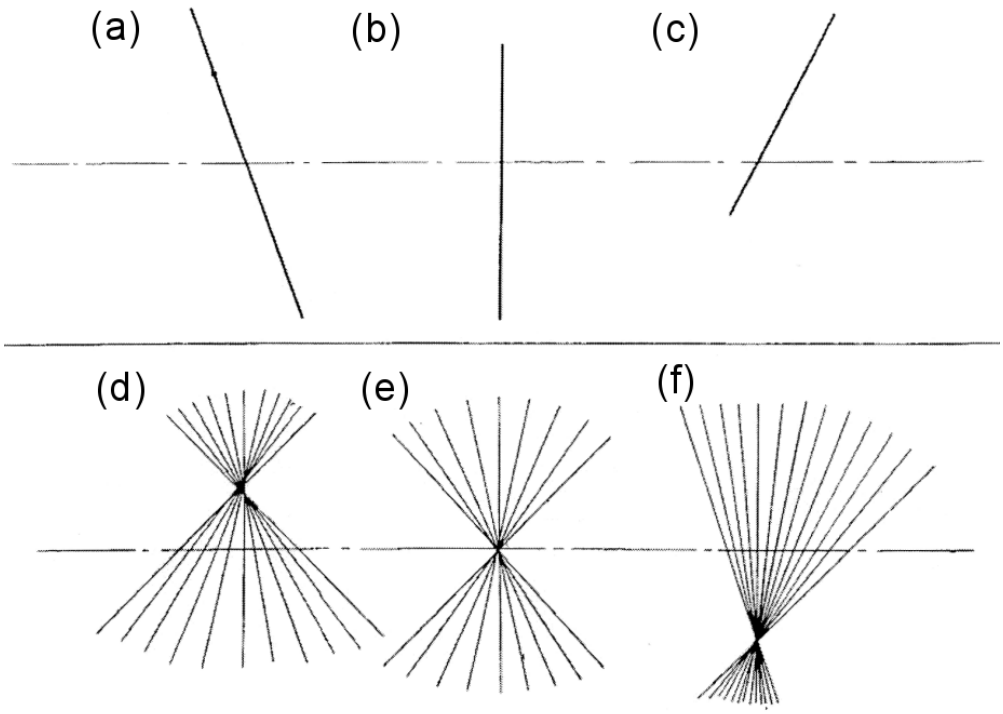


Figure 7.39 – Illustration of the Hough transform for straight lines (a)-(c) three straight lines in real space, (d)-(f) corresponding plots in Hough space, the position of the intersection indicates the most likely parameters to describe the original line. Reproduced from Hough (1962).

If we decide to uniquely define our line with a slope and an intercept (as Hough did in his original paper), then for the case of a vertical line, our Hough space would have to extend to infinity. Duda and Hart (1972) suggested that it was far better to parameterise by drawing a line from the origin to the closest point on the line, and to describe this new line in terms of an angle, $\vartheta$, from the x-axis and a length of the line, P. Hence, it can be shown that the new parametric description of our original line may be written as $P = x \cos\vartheta + y \sin\vartheta$. This formulation is illustrated in Figure 7.40. As a side note, parametric transformations had been around for some time before the Hough transform, and in fact Deans (1981) showed that it was simply a special case of the Radon transform (Radon, 1917).

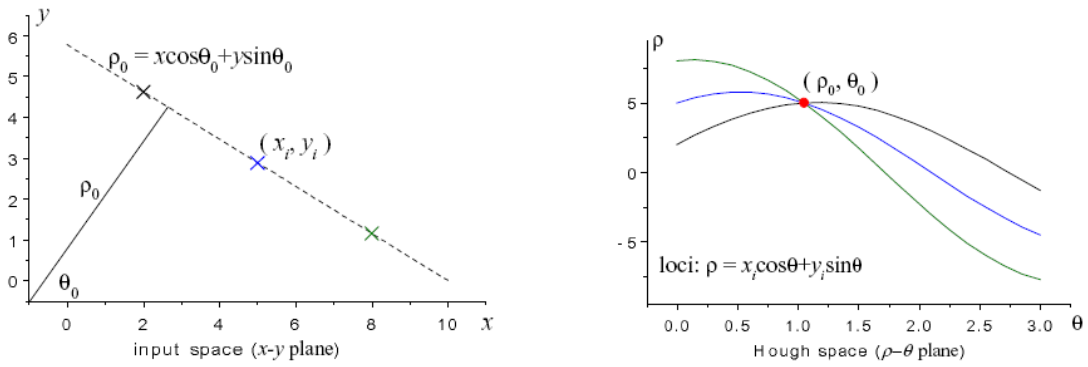


Figure 7.40 – Illustration of the slope-locus line parameterisation proposed by Duda and Hart (1972). Reproduced from Lin and Otobe (2001).

The advantage of the Hough transform is that we do not require a complete shape; it will work on partial shapes just as well since the result is identified by probability of a match. Partial information is sufficient. This is one of the reasons that the Hough transform has become a well-established technique for shape and motion analysis in images containing noisy, missing and extraneous data. However, this is also a potential downside to the technique, in that it cannot, for example, detect the start and end of a line – it simply returns the parameters of the line, and if plotted will extend across the whole of the image space.

There are many features which make the Hough transform desirable:

- Spatially extended patterns in real space are transformed such that they produce spatially compact features in a space of possible parameter values.
- Each image point is treated independently, and therefore the algorithm is suitable for parallel processing. Computational power has always been an issue with image analysis applications, but this is rapidly becoming easier and real-time applications are now a possibility.
- Its independent combination of evidence means that it can recognise partial or slightly deformed shapes. The size and spatial location of the parameter peak provides a measure of the similarity of shape and model.
- The Hough transform should be resistant to the addition of random noise data, since random image points are unlikely to contribute coherently to a single element of parameter space.

- The algorithm can search for multiple instances of a shape with the image in a single processing pass – each instance simply produces a distinct peak in Hough space.

The main disadvantages of the Hough transform are the computational power and storage requirements. The computational power is no longer such an issue, but for each parameter in our model shape, we require another dimension in Hough space. This can rapidly become an intractable problem, with the exponential expansion of memory requirement with additional parameters. Techniques exist to minimise the number of parameters, and to determine groups of parameters in several sequential stages. Attempts to mitigate this problem form the bulk of the work associated with Hough transform techniques. Please see Illingworth and Kittler (1988) and Leavers (1993) for more information on optimisation and parameter selection methods. In this work we will only deal with two-parameter searches.

In fact, the Hough transform may be extended to the detection of any feature which can be described parametrically. For example, for detection of cylindrical pellets within an image, we may require a rotation and scaling parameter. An example of the identification of circular objects is shown in Figure 7.41 with the original data on the left and the corresponding Hough space on the right, where the parameter sought is the position of the circle of a specific radius. The identified circles are then overlaid over the original image. This extension of the algorithm to detect shapes in their entirety (rather than their component edges) is referred to as the generalised Hough transform (GHT), and is covered in more detail by Ballard (1981) and Leavers (1993).

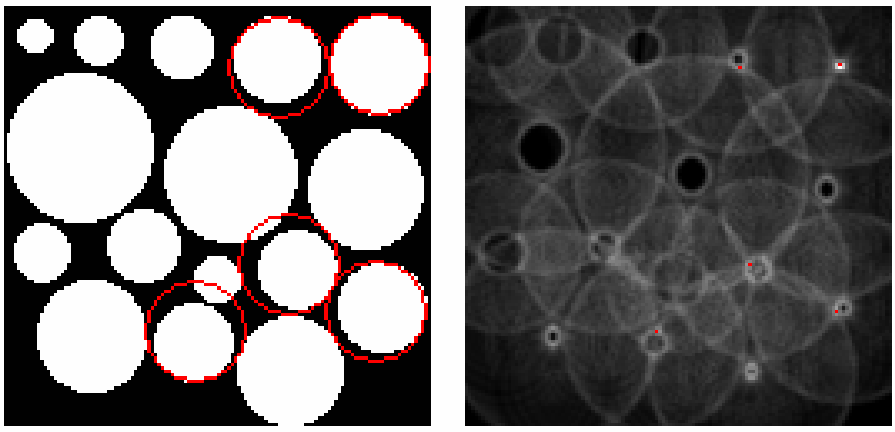

(a) (b)

Figure 7.41 – The application of the parametric Hough transform to locate circles of a certain radius. (a) the original data, and (b) the corresponding Hough space. The identified maxima are overlaid over the original data.

Applications of the Hough transform include machine vision (detection of obstacles or following a path), the recognition of Hebrew characters – which almost entirely consist of straight lines – to an accuracy of 99.6% for machine printed characters and 86.9% for handwritten (Kushnir *et al.*, 1983; Kushnir *et al.*, 1985), fingerprint feature analysis (Lin and Dubes, 1983), satellite image processing and identification of functional or cosmetic defects, among many other image processing applications. In general, any application where large numbers of images must be visually inspected for specific features or patterns, may benefit from the Hough transform.

## 7.6.2. Implementation and viability of the Hough transform

A Hough transform for identifying straight lines was implemented as part of this work. An example was chosen to test the functionality and robustness of the procedure. The plot in Figure 7.42a shows a test input to the Hough transform procedure, showing five randomly drawn lines which we wish to detect. The corresponding Hough space is shown in (b), with $\vartheta$ along the x-axis and P along the y-axis. It can be seen that there are five bright local maxima and these are created by the overlap of many drawn lines, each of which corresponds to a point in real space.

Using the information on the position of these maxima, and re-plotting in real space yields (c), where it can be seen that the original plot is reconstructed with a high degree of accuracy.

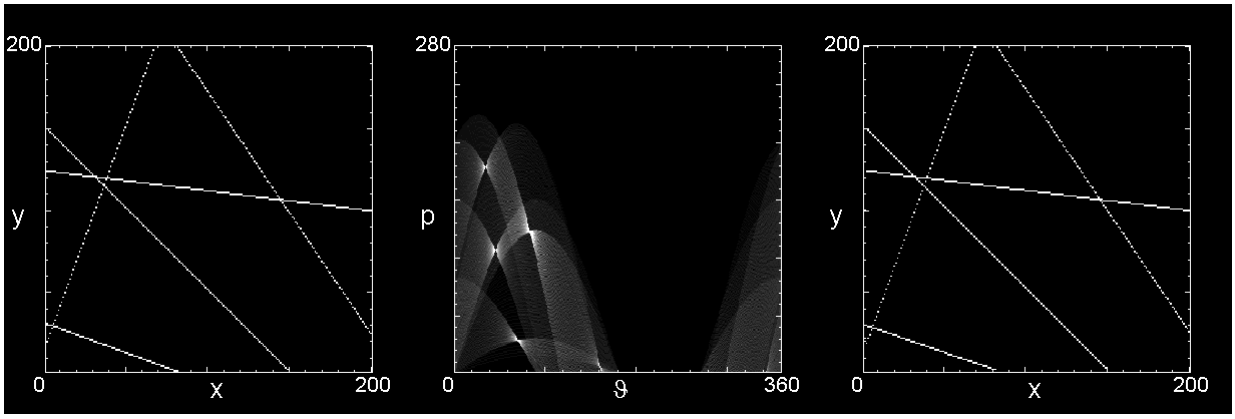


(a) (b) (c)

Figure 7.42 – (a) shows a test input to the Hough transform procedure; (b) shows the corresponding Hough space; (c) using the information from these five maxima, the original plot is reconstructed.

If the automatic detection of local maxima in Hough space is desired, this is probably best performed by some curve fitting / slope analysis to distinguish real maxima from stray points. It was not a problem in this case, however more complex systems would benefit from more rigorous detection methods. Additional information on peak detection techniques is available in Illingworth and Kittler (1988).

As has been previously mentioned, one of the advantages of the Hough transform over other fitting techniques is that it may be applied to partial data, since it uses a "voting system" to statistically identify the most likely shape parameters. The Hough transform was applied to partial information, and Figure 7.43a shows the data of Figure 7.42a with selected regions deleted. The corresponding Hough space in (b) is missing certain lines, however the maxima are still apparent and appear in the same positions. Thus, it is possible to recreate the original image as shown in (c). A small error is observed, whereby the two parts of the same line have been identified as separate lines (with a difference in angle of one degree). This is still an acceptable outcome, and perhaps with a better peak-detection algorithm the maxima could be more clearly identified.

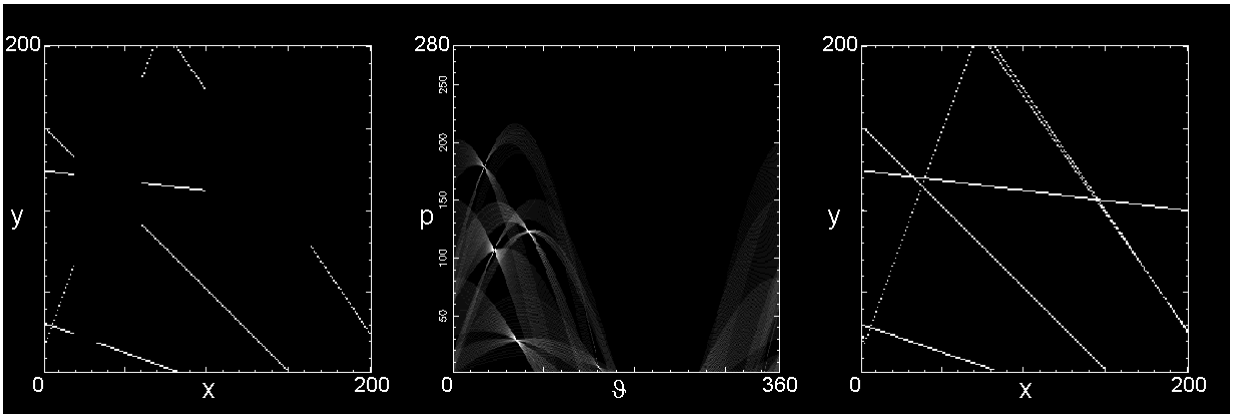


(a) (b) (c)

Figure 7.43 – The results of the Hough transform on partial information. (a) is an image of partial straight lines; (b) shows the corresponding Hough space; and (c) shows the reconstruction of the original data.

A useful application of the Hough transform for our purposes was identified as that of profile tracking – in imaging monoliths we may use the SEMI-RARE sequence (Section 5.2), which allows us to acquire successive snapshots of the same series of slugs. If we can compare the position of these profiles over time, we will then have a measure of the velocity within the channel. A Hough transform which would compare two profiles and estimate the displacement between the two was encoded. This was done by assuming that the profile did not change significantly in shape but simply shifts linearly left or right, and that the amplitude of the profile may decrease (since we will have attenuation or relaxation effects between images). Hence, we have a template shape (the first profile) which we are trying to match up to some data (the second profile). This allows us to construct a two-dimensional Hough space with parameters of displacement and scale. A synthetic test is shown in Figure 7.44a and Figure 7.44b, showing the "displacement" and "attenuation" of two successive Gaussian shapes. The resultant Hough plot in Figure 7.44c clearly shows a maximum, corresponding to a displacement of 80 pixels and a scale factor of 0.5, which is correct.

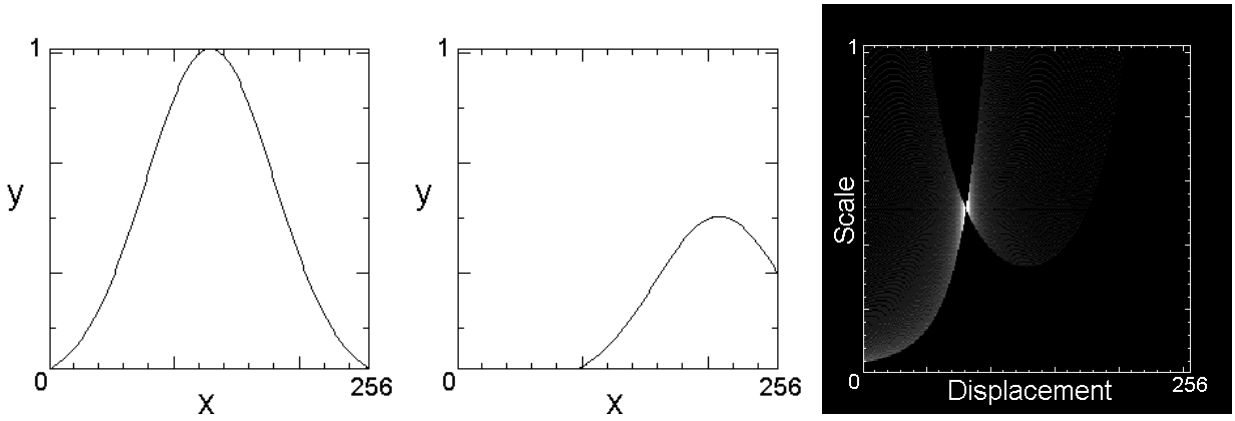


Figure 7.44 – Illustration of profile tracking, implemented using the Hough transform; (a) and (b) are successive Gaussian shapes with different linear displacements and attenuation; (c) shows the Hough plot, which has clearly identified the displacement and attenuation difference between the inputs.

Hence, the Hough tracking method has been seen to work well on synthetic, noise-free data. It was then applied to real data to test its performance in the presence of noise and other non-idealities. The profiles used for this are from a 200 cpsi monolith reactor under two-phase flow provided by a spray distributor (BGQ1740 nozzle, setup described in Section 8.2.3) and a flow of 138 l/hr of water. Data were acquired with the SEMI-RARE pulse sequence, giving four successive images of the slug profiles within the channels. Each data image is taken 80 ms apart. Successive slug profiles from a monolith channel are shown in Figure 7.45. The displacement of the profiles to the right, and the decrease in intensity is apparent but not immediately quantifiable by eye. The data also contain significant noise.

Hough transforms were created and are shown in Figure 7.46, comparing (a) the first and second images, (b) the first and third, and (c) the first and fourth images. The data is not as clear as for the artificial data – the noise in the images jitters the position of the lines we draw sufficiently to prevent them lining up coherently and producing a clear maximum. The lack of data points adds to the problem since we have only 32 points in our data profiles, and thus only 32 lines to be drawn in Hough space.

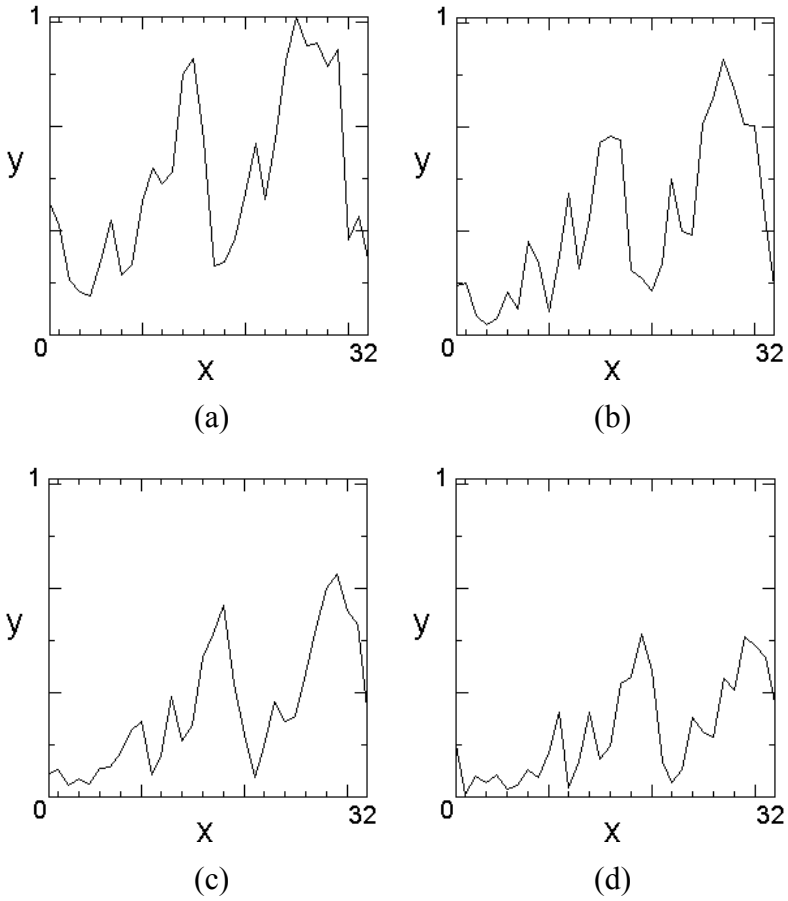


Figure 7.45 – Four successive profiles from a monolith channel are shown, at increasing acquisition times. The linear shift to the right and the decay in amplitude can be seen.

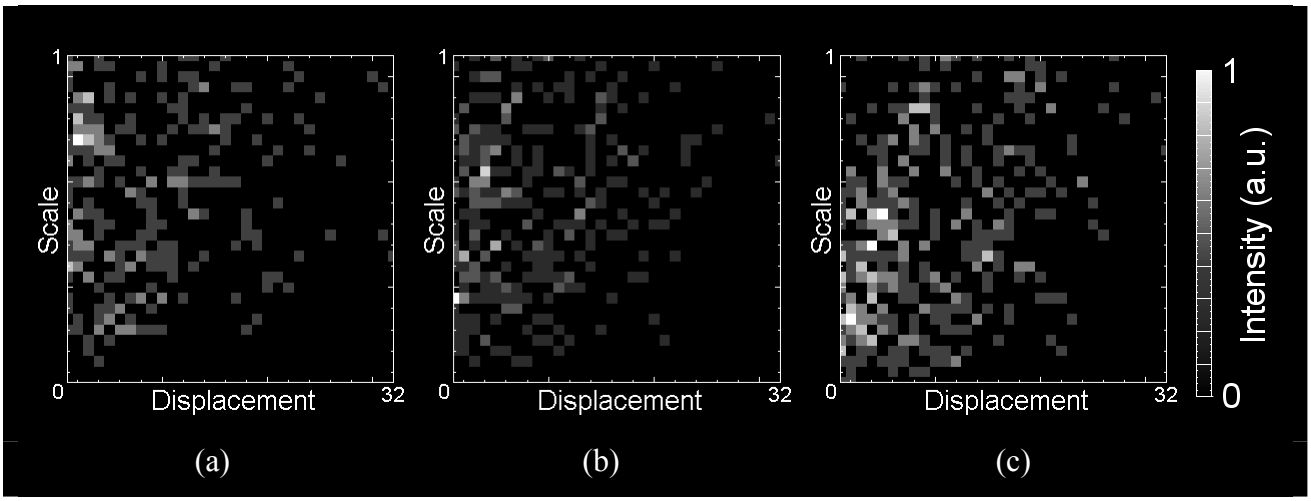


Figure 7.46 – Hough transforms for profile tracking using (a) the first and second profiles; (b) the first and third images; and (c) the first and fourth images.

It was decided to interpolate the profiles to 256 points using linear interpolation, and then to re-apply the Hough analysis to this new data. The results shown in Figure 7.47 are a little clearer, and local regions of higher intensity are observed. The data is still noisy, but the position of a region of high intensity can be observed moving down and to the right (highlighted in red), indicating an increase in position and attenuation between successive images.

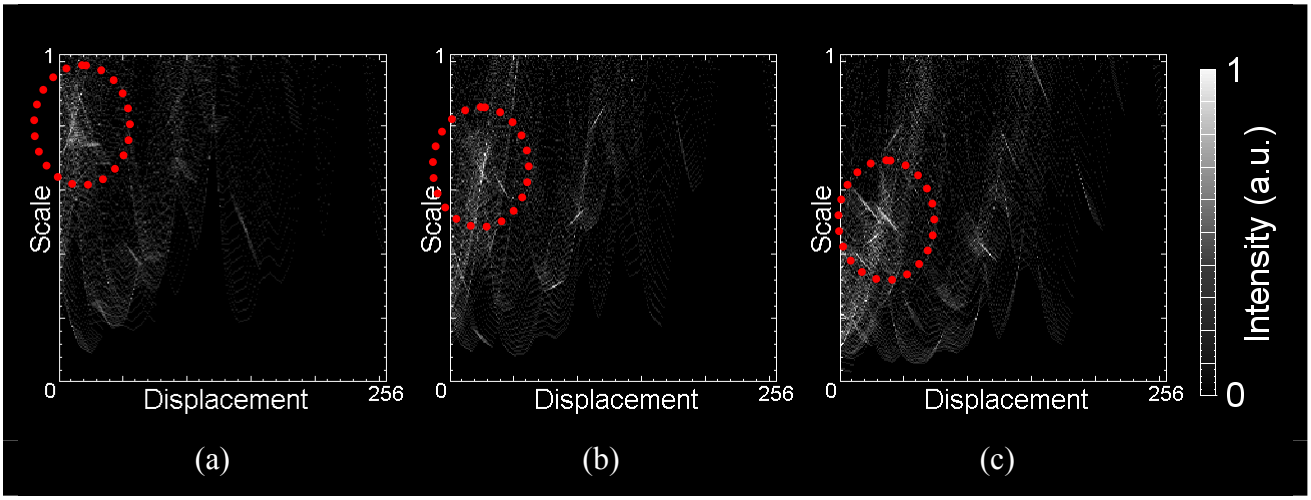


Figure 7.47 – Hough transforms for profile tracking obtained from zero-filling the data of Figure 7.45 (interpolated from 32 to 256 points) using (a) the first and second profiles; (b) the first and third images; and (c) the first and fourth images.

The literature, in general, claims that this technique is robust and resistant to noisy data (Leavers, 1993). Filtering may be of use, either in pre-processing the raw data or for the treatment of Hough space. For example, statistical weighting can be given to features corresponding to short straight lines in order to distinguish from lone image points, which are more likely to be noise data.

However, this robustness to noise is not apparent from the analysis performed as part of this work. The interpretation of Hough space is little simpler than that of comparing the raw data profiles by eye and ruler, so this method is not a marked improvement with regards to measurement or quantification of the data. It was decided that the Hough transform was not sufficiently robust to justify its use in this application. An alternative method was proposed, namely that of a $Chi^2$ profile fitting algorithm, which is detailed in the next section.

## 7.6.3. Chi$^2$ profile fitting

An alternative method of shape detection was implemented, which involved calculating a Chi$^2$ measure of the error at each possible displacement, and choosing the parameters which give the minimum error. This is a brute-force method but should be much more robust to noise than the Hough transform. The operation is performed by taking the first profile, shifting it and multiplying it. This is then compared to the second profile and an error is calculated by summing the square of the difference at every point. A wide range of displacements and scale factors are investigated in turn, and the best fit is given by the lowest value of the errors calculated. This has the advantage over the Hough transform in that it is simple to return a single value for the best fit, rather than a spread of probabilistic values.

The Chi$^2$ fitting algorithm works very well with the simulated data of Figure 7.44a and Figure 7.44b. The error plot is shown in Figure 7.48, and the global minimum is quite distinct, corresponding to a best fit of 80 pixels and the scale factor of 0.5, which is correct.

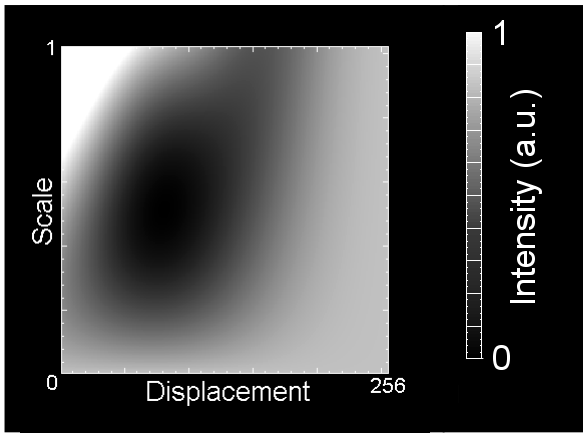


Figure 7.48 – Error plot showing the results of the Chi$^2$ fitting algorithm when applied to the synthetic problem of Figure 7.44.

Applying the method to the real data profiles of Figure 7.45 yields the error plots shown in Figure 7.49. As can be seen, there are distinct minima which are far clearer than the features in the Hough space plot earlier. The position of the minima is seen to move to the right (larger displacement) and downwards (lower intensity) as

expected. Performing the operation on the higher-resolution interpolated plots yields the three plots in Figure 7.50. These are identical in appearance, but allow for more precision in the measurement of the coordinates of the minima.

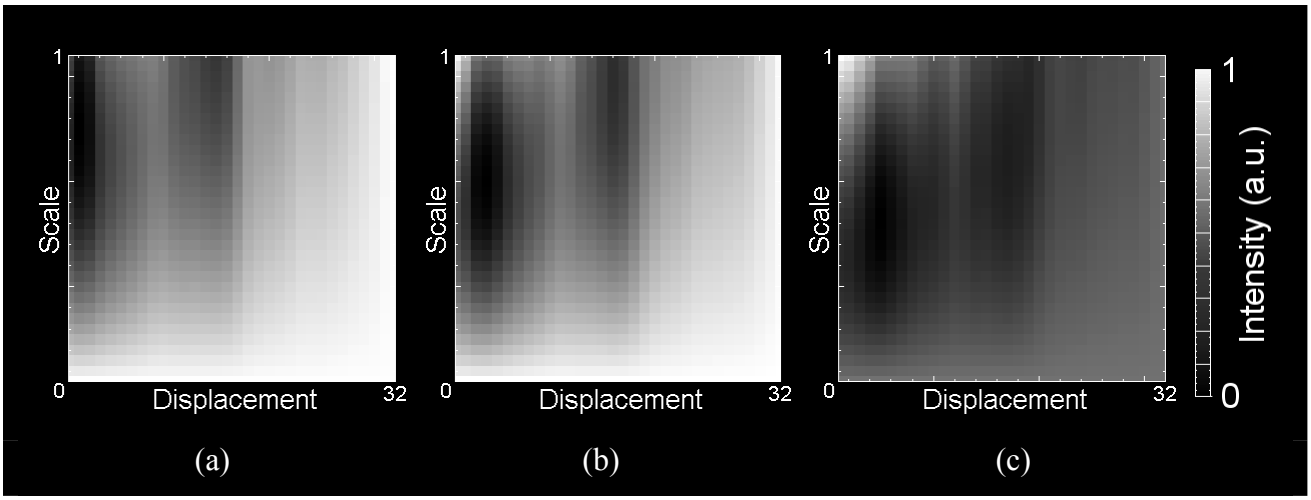


Figure 7.49 – Chi$^2$ error plots using the data of Figure 7.45; (a) the first and second profiles; (b) the first and third images; and (c) the first and fourth images.

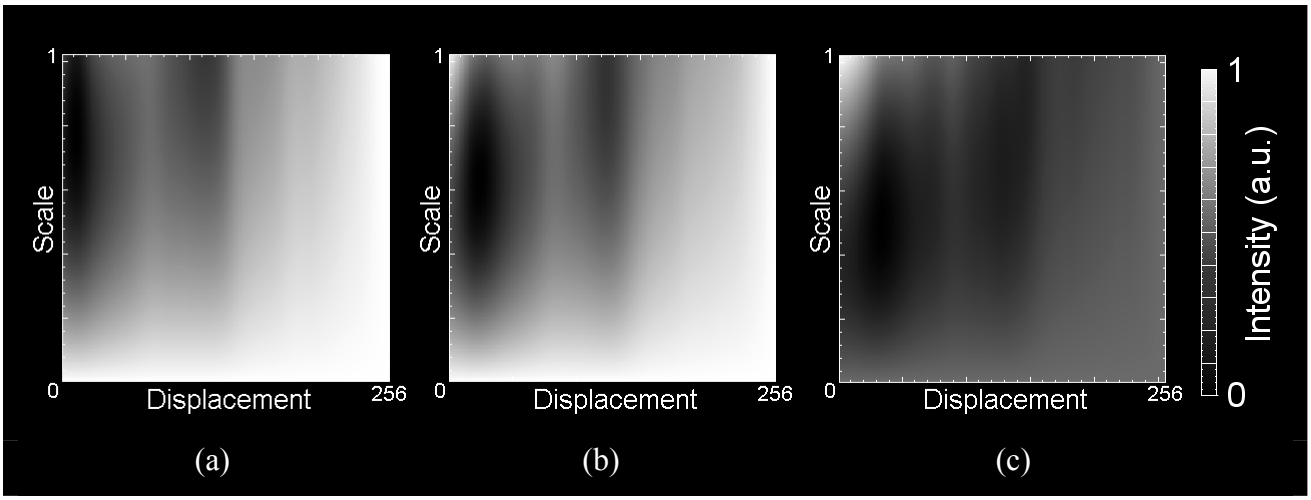


Figure 7.50 – Chi$^2$ error plots obtained from the zero-filled data of Figure 7.45 (interpolated from 32 to 256 points) using (a) the first and second profiles; (b) the first and third images; and (c) the first and fourth images.

For the profiles being investigated here, the results are as shown in Table 7.1 and Table 7.2. Velocities are calculated from knowledge of the time between images and the pixel size.

Table 7.1 – Velocity measurements from the plots of Figure 7.49. The average is 12.8 ± 2.5 mm/s.

| Profiles compared | Displacement (pixels) | Scale factor | Velocity (mm/s) |
|---|---|---|---|
| $1 \rightarrow 2$ | 1 | 0.781 | 10.0 |
| $1 \rightarrow 3$ | 3 | 0.625 | 15.0 |
| $1 \rightarrow 4$ | 4 | 0.500 | 13.3 |

Table 7.2 – Velocity measurements from the plots of Figure 7.50. The average is 13.6 ± 1.0 mm/s.

| Profiles compared | Displacement (pixels) | Equivalent displacement (pixels) | Scale factor | Velocity (mm/s) |
|---|---|---|---|---|
| $1 \rightarrow 2$ | 10 | 1.25 | 0.762 | 12.5 |
| $1 \rightarrow 3$ | 23 | 2.88 | 0.613 | 14.4 |
| $1 \rightarrow 4$ | 33 | 4.13 | 0.480 | 13.8 |

Both methods give similar results, but the interpolated data give more precise values, particularly where the displacements are only on the order of a few pixels. The scale factor seemed less sensitive to the resolution of the data, and is seen to decay as expected. The $Chi^2$ method appears to be relatively insensitive to noise, and so is an ideal method for tracking slug profiles in channels.

The downside to this method is that it is relatively computationally intensive, even when re-written in C. For example, to process 64 channels, in 24 slices with 32 repetitions each, and at 5 different flow rates took just over a week on an AMD 2500+ processor.

## 7.7. Phase unwrapping

This final section will briefly cover the correction of phase wraps within velocity-encoded images. Phase wrapping occurs when the phase measured by the pulse sequence exceeds $2\pi$. Since this corresponds to a complete revolution of the spin magnetisation vector, the measured phase then wraps around and is measured as the

true phase *minus* $2\pi$. This can be seen in Figure 7.51 where the velocity increases up to a value of $2\pi$, and then there is a sudden decrease to zero and further increase beyond that. The data is of laminar flow in a 50 mm diameter pipe at a flow rate of 43 cc/s (a superficial velocity of 2.3 cm/s). Data are acquired with the spin-echo, phase-encoded pulse sequence (Section 2.4.1) with an acquisition time of 4 minutes. The field of view is 55 mm, giving a resolution of 859 μm × 1719 μm.

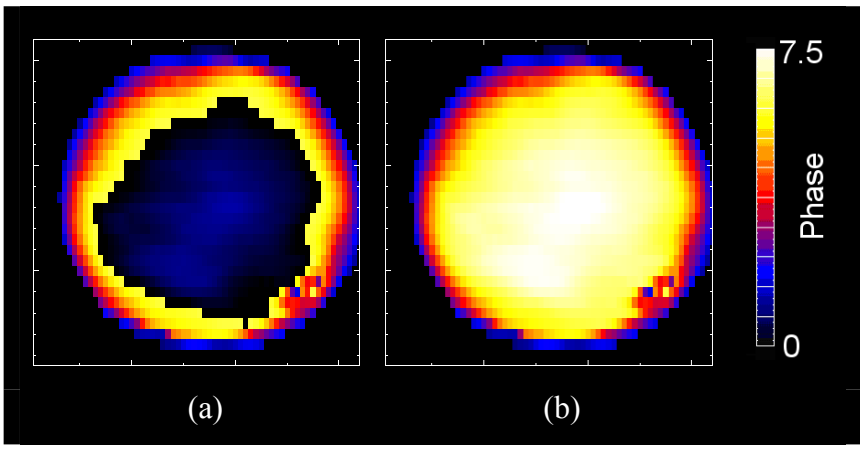


Figure 7.51 – Example of phase-wrapping within a laminar flow data set; (a) shows a velocity map, the phase-wrapping artefact can be clearly seen, (b) the phase-corrected data.

The occurrence of phase-wraps can be avoided by using a lower strength of phase-encoding gradient, in such a way that the maximum velocity in the system does not exceed a phase of more than $2\pi$. The reason why we might choose to acquire phase-wrapped data and then correct it, is that we would be measuring a larger dynamic range and therefore obtain higher precision of the phase measurement.

The unwrapped phase, $\theta_j$, is related to the wrapped phase, $\phi_j$, plus and integer multiple of $2\pi$, as shown in Eq. 7.20.

$$\theta_j = \phi_j + 2\pi \cdot n \qquad \textbf{(Eq. 7.20)}$$

Often, it is clear where phase-wrapping has occurred, particularly in well-defined systems with clear velocity gradients within the data. Problems occur for systems where individual voxels are not necessarily related to the neighbouring ones, and phase unwrapping is not then possible.

A value of $2\pi$ can be manually added to wrapped regions, however this soon becomes time-consuming to manually correct many data images and so automatic methods were sought. Matlab contains a basic one-dimensional phase-unwrap algorithm, which splits a 2D image up into a series of 1D profiles, and the neighbouring pixels in these profiles are compared to highlight pairs which have a significant difference between them. If they differ by roughly $2\pi$, the phase is corrected. This method is fairly robust, but it was observed that the algorithm was defeated by noisy data, as this led to incorrect identification wraps and 1D "zippers" across the image where $2\pi$ had been incorrectly added or subtracted. An example is shown in Figure 7.52 and it can be seen that where there is most noise in the data the performance of the algorithm is very poor.

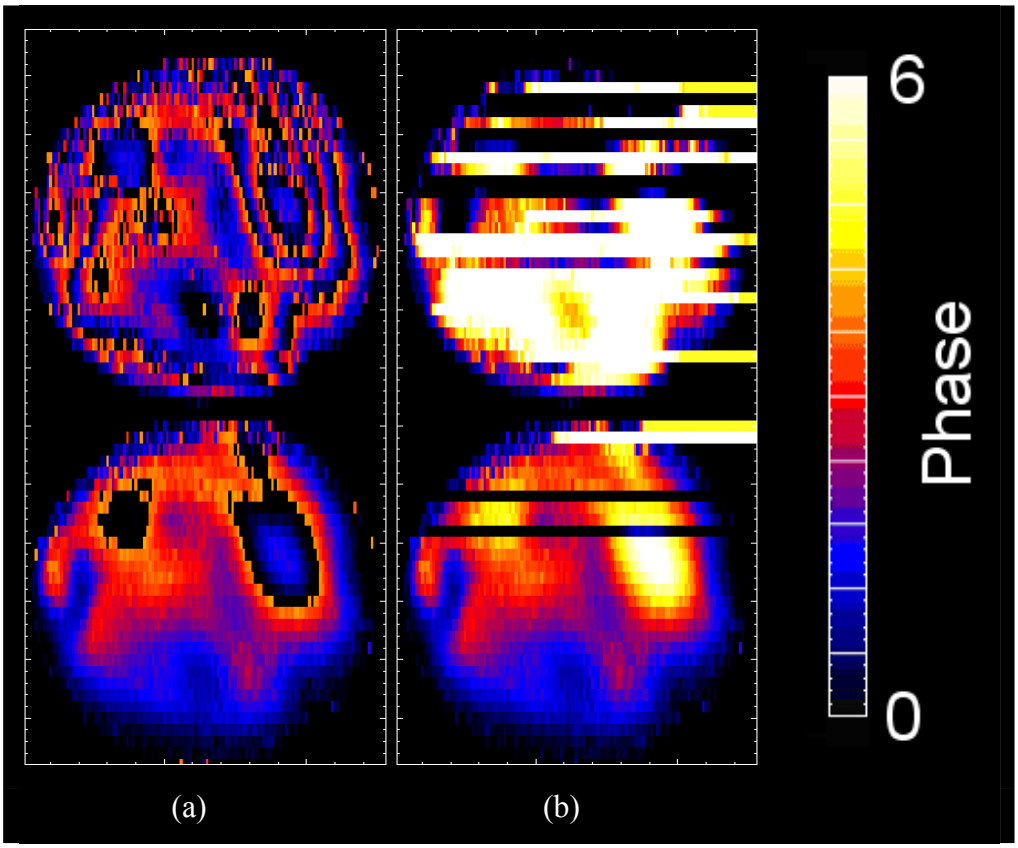


Figure 7.52 – Illustration of the Matlab phase-correction algorithm; (a) shows the uncorrected phase data, and (b) shows the results of applying the Matlab phase-correction algorithm. The data are from the MS-COTTAGE sequence (Section 5.4) applied to unsteady flow in a 50 mm diameter pipe for single phase liquid flow at a Reynolds number of 1664. The first two cross-sectional images are shown, acquired over a time of 162 ms. The field of view is 51 mm, giving a resolution of 398 μm × 1594 μm.

.The Matlab artefact is a nature of the 1D progressive algorithm, and a 2D algorithm would take into account the neighbouring pixels and be far less easily misled. The method of Jenkinson (2003) was implemented to unwrap noisy 2D data, and the method will be briefly outlined here. Figure 7.53 shows the main steps in the PRELUDE procedure (Phase Region Expanding Labeller for Unwrapping Discrete Estimates); (a) is a magnitude (intensity) image of the system, in this case a brain; (b) is the corresponding velocity image, which contains phase-wraps; (c) is a region-segmented image, which is created as part of the phase correction algorithm, and (d) is the result of the phase unwrapping, giving a continuous velocity image. The segmentation and phase-unwrapping procedure will now be outlined.

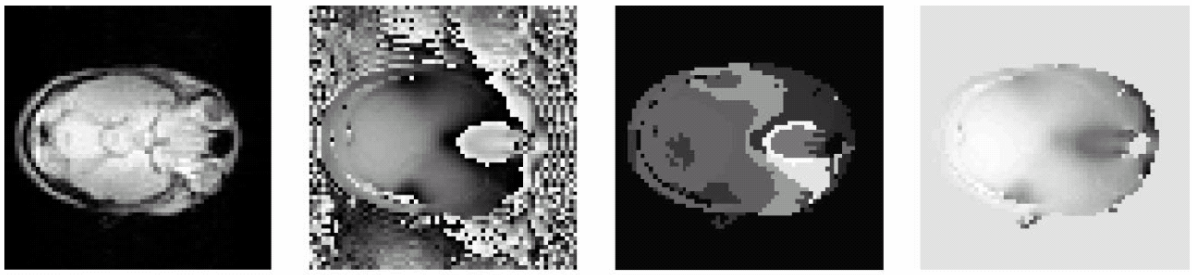

Figure 7.53 – Illustration of the PRELUDE algorithm, applied to a brain velocity image; (a) is the magnitude image of the system; (b) is the phase-wrapped velocity image; (c) is the segmented region analysis created by the PRELUDE algorithm, and (d) is the phase-corrected velocity data. Reproduced from Jenkinson (2003).

The algorithm works by segmenting the image into regions of similar value. For example, we can divide the range of $-\pi$ to $+\pi$ into six regions, as shown in Figure 7.54.

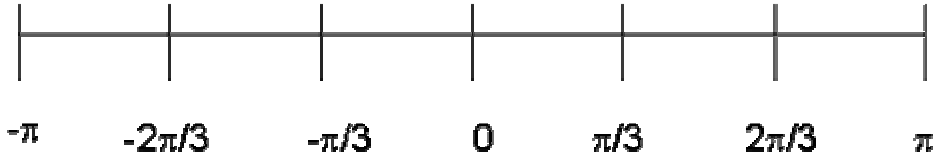


Figure 7.54 – The segmentation of the range from $-\pi$ to $+\pi$

By then classifying the velocity data of Figure 7.53b with these regions, we obtain the image of Figure 7.53c. If we then examine the perimeters between segmented

regions, we can calculate a cost function for the image which penalises the phase difference between neighbouring regions, such as that in Eq. 7.21 which compares two neighbouring regions (*A* and *B*).

$$C_{AB} = \sum_{j,k \in N(j)} (\theta_{Aj} - \theta_{Bk})^2 = \sum_{j,k \in N(j)} (\phi_{Aj} - \phi_{Bk} + 2\pi M_{AB})^2 \qquad \textbf{(Eq. 7.21)}$$

where $j$ is the index of a voxel in region $A$, $k$ is the index of a voxel in region $B$, and $M_{AB} = M_A - M_B$ where $M_A$ and $M_B$ are integer offsets for regions $A$ and $B$.

Thus, if the neighbouring regions are close together in phase value, the relative contribution to the cost function will be low. However, if the regions differ by a large amount (due to phase wrapping) then the contribution will be high, and the cost will be reduced by adding or subtracting $2\pi$ to one of the regions. This is repeated for all pairs of regions until the cost function reaches a minimum.

Therefore this algorithm takes into account the 2D neighbourhood of pixels when calculating phase-wraps, and is also robust to noise unlike the Matlab algorithm. These properties make it ideal for our uses and we now have a robust method for phase correction. An example of the algorithm applied to our MRI data is shown in Figure 7.55, and the results are excellent. The first three images in a MS-COTTAGE sequence are shown, acquired over a time of 242 ms. The experimental setup is the same as for the data in Figure 7.52 but at a lower Reynolds number of 610.

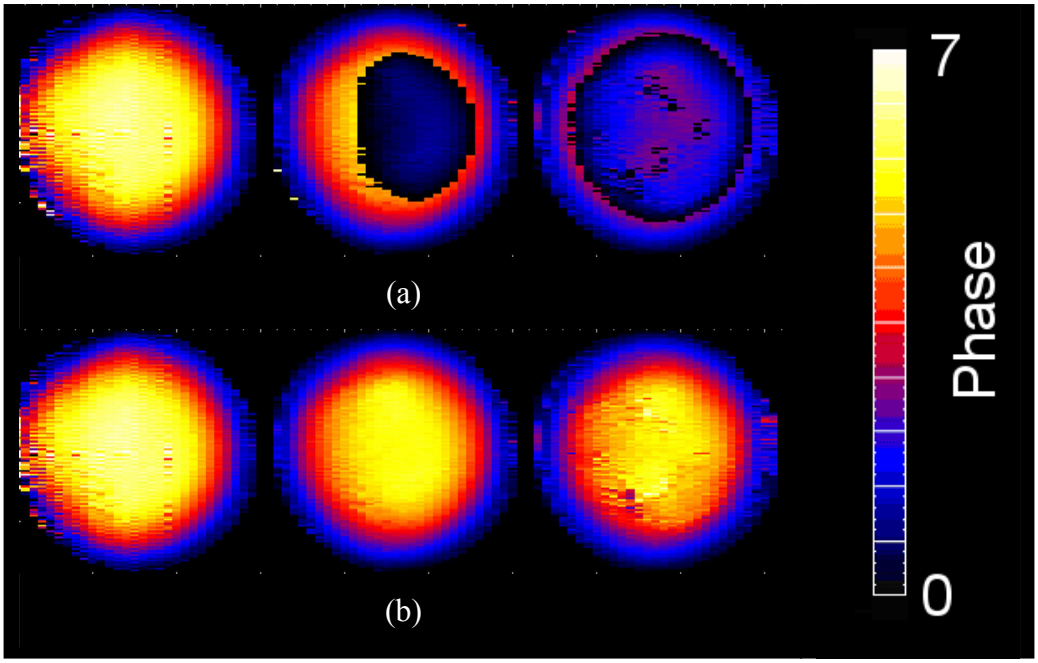


Figure 7.55 – Application of the PRELUDE algorithm to some real MRI data of laminar flow in a 5 cm diameter pipe; (a) phase-wrapped data, (b) unwrapped data using the PRELUDE algorithm.

## 7.8. Conclusions

In this chapter, a range of varying image processing techniques have been presented. Some are traditional techniques, while others are novel developments based on traditional ideas.

The basics of image processing, such as binary gating and smoothing functions were presented, and these are used extensively in the remainder of this work. Background levelling techniques were introduced for the correction of artefacts sometimes seen in MRI images.

Texture processing (and feature enhancement) was then covered, with the aim of quantifying the distribution of features within an image and also their characteristic sizes. This was specifically for quantifying fluid distribution within the monolith reactor, but these techniques would find applications in many different systems. A novel, discreet autocorrelation function was developed and verified against a traditional autocorrelation function. This was then adapted and used to identify characteristic feature sizes within a series of images. The GLCM was then introduced

as a 2D histogram, capable of measuring different textural features within an image. This was implemented, and a novel adaptation was used to enhance the original images by biasing them with various textural quantities. A novel contour-generation algorithm was discovered and used to quantify a series of turbulence-containing MRI data.

Feature tracking and identification was then considered. The traditional Hough transform was implemented, and modified to apply to profile tracking. However, it was found to be unsuitable due to its poor performance in the presence of noise. The brute-force $Chi^2$ tracking method was much more successful, and was able to track noisy MRI data without setback.

Lastly, a 2D phase-unwrapping algorithm for the correction of MRI velocity data was outlined, and this was used extensively in this work. It was found to be far superior to Matlab's 1D method.

Chapters 8-11 will now apply this wealth of image analysis techniques to real MRI data, and attempt to characterise the behaviour of the monolith reactor under various flow conditions.

## 7.9. Nomenclature

| Symbols | Definition |
|---|---|
| $B$ | Information measure, related to entropy |
| $C$ | Contrast (GLCM measure) |
| $C_{AB}$ | Cost function, related to phase unwrapping |
| $C_c$ | Correlation (GLCM measure) |
| $C_d$ | Dissimilarity (GLCM measure) |
| $C_e$ | Entropy (GLCM measure) |
| $C_h$ | Homogeneity (GLCM measure) |
| $C_{ID}$ | Inverse difference moment (GLCM measure) |
| $C_u$ | Uniformity (GLCM measure) |
| $\mathbf{d}$ | Displacement vector in real space |
| $g(x,y)$ | A function in real space (coordinate system $x$,$y$) |
| $g_a$ | Subset of pixel values |
| $G(u,v)$ | Autocorrelation function (coordinate system $u$,$v$) |
| $H$ | Entropy measure |
| $i,j$ | GLCM coordinate space |
| $k$ | Moment, as defined in the creation of the GLCM |
| $\mathbf{k}$ | Reciprocal space vector |
| $M_{AB}$ | Defined as $M_A$-$M_B$ |
| $M_A$,$M_B$ | Integer offsets for regions A, B |
| $Mom_k$ | $k^{th}$ moment |
| $n$ | An integer |
| $n_{ij}$ | Frequency of occurrence of pixel pair $i$,$j$ |
| $N$ | Total number of pixels |
| $N(i,j)$ | Normalised GLCM matrix |
| P | Distance between the origin and the closest point on a line |
| $P_i$ | Pixel at position $i$, |
| $P_d(i,j)$ | GLCM coordinate ($i$,$j$) |
| X,Y | The length of a data array in dimensions $x$ and $y$, respectively. |
| | |
| $\vartheta$ | Angle between the origin and the perpendicular of a line |
| $\phi_j$ | Unwrapped phase angle |
| $\theta_j$ | Wrapped phase angle |
| $\chi^2$ | Textural information measure |
| $\Phi(\omega)$ | PSDF |
| $*$ | Convolution symbol |

| Abbreviations | Definition |
|---|---|
| COTTAGE | Continuous translation tracking with alternating gradient encode |
| GERVAIS | Gradient echo rapid velocity and acceleration imaging sequence |
| GHT | Generalised Hough transform |
| cpsi | Cells per square inch |
| GLCM | Grey level co-occurrence matrix |
| IDA | Image and data analysis software |
| LoG | Laplacian of Gaussian |
| MS-COTTAGE | Multi-slice COTTAGE |
| MRI | Magnetic resonance imaging |
| NMR | Nuclear magnetic resonance |
| PSDF | Power spectrum density function |
| PRELUDE | Phase region expanding labeller for unwrapping discrete estimates |
| RARE | Rapid acquisition with relaxation enhancement |
| RDF | Radial distribution function |
| RF | Radio frequency |
| SEMI-RARE | Single-excitation, multiple image RARE |
| TINA | TINA is no acronym / there is no alternative |

# 7.10.References


Albers, R. E. (2005). "Characterization of the performance of an industrial monolith reactor by accurate mapping of temperature differences." Catalysis Today **105**(3-4): 391-395.

Ballard, D. H. (1981). "Generalizing the Hough Transform to Detect Arbitrary Shapes." Pattern Recognition **13**(2): 111-122.

Bauer, T., S. Roy*, et al.* (2005). "Liquid saturation and gas-liquid distribution in multiphase monolithic reactors." Chemical Engineering Science **60**(11): 3101-3106.

BroNielsen, M. (1997). Rigid registration of CT, MR and cryosection images using a GLCM framework. Cvrmed-Mrcas'97. **1205:** 171-180.

Clausi, D. A. (2002). "An analysis of co-occurrence texture statistics as a function of grey level quantization." Canadian Journal of Remote Sensing **28**(1): 45-62.

Darakchiev, R. and C. Dodev (2002). "Gas flow distribution in packed columns." Chemical Engineering and Processing **41**(5): 385-393.

Deans, S. R. (1981). "Hough Transform from the Radon-Transform." Ieee Transactions on Pattern Analysis and Machine Intelligence **3**(2): 185-188.

Duda, R. O. and P. E. Hart (1972). "Use of Hough Transformation to Detect Lines and Curves in Pictures." Communications of the Acm **15**(1): 11-&.

Edwards, D. P., K. R. Krishnamurthy*, et al.* (1999). "Development of an improved method to quantify maldistribution and its effect on structured packing column performance." Chemical Engineering Research & Design **77**(A7): 656-662.

El-Bachir, M. S. (2005). MRI Study of Flow in Narrow Packed Beds. Dept of Chemical Engineering. Cambridge, University of Cambridge**:** 286.

Haralick, R. M., Shanmuga.K*, et al.* (1973). "Textural Features for Image Classification." Ieee Transactions on Systems Man and Cybernetics **SMC3**(6): 610-621.

Heibel, A. K., J. J. Heiszwolf*, et al.* (2001a). "Influence of channel geometry on hydrodynamics and mass transfer in the monolith film flow reactor." Catalysis Today **69**(1-4): 153-163.

Heibel, A. K., T. W. J. Scheenen*, et al.* (2001b). "Gas and liquid phase distribution and their effect on reactor performance in the monolith film flow reactor." Chemical Engineering Science **56**(21-22): 5935-5944.

Hough, P. V. C. (1962). "Method and Means for Recognizing Complex Patterns." US Patent 3 069 654.

Illingworth, J. and J. Kittler (1988). "A Survey of the Hough Transform." Computer Vision Graphics and Image Processing **44**(1): 87-116.

Jenkinson, M. (2003). "Fast, automated, N-dimensional phase-unwrapping algorithm." Magnetic Resonance in Medicine **49**(1): 193-197.

Jiang, Y., M. R. Khadilkar*, et al.* (1999). "Two-phase flow distribution in 2D trickle-bed reactors." Chemical Engineering Science **54**(13-14): 2409-2419.

Kelly, A. (1996). "Introduction to Mobile Robots - Perception 3: Algorithms for Perception." http://www.frc.ri.cmu.edu/~alonzo/course/per_alg.pdf.

Kushnir, M., K. Abe*, et al.* (1983). "An Application of the Hough Transform to the Recognition of Printed Hebrew Characters." Pattern Recognition **16**(2): 183-191.

Kushnir, M., K. Abe, *et al.* (1985). "Recognition of Handprinted Hebrew Characters Using Features Selected in the Hough Transform Space." Pattern Recognition **18**(2): 103-114.

Leavers, V. F. (1993). "Which Hough Transform." Cvgip-Image Understanding **58**(2): 250-264.

Lin, W. C. and R. C. Dubes (1983). "A Review of Ridge Counting in Dermatoglyphics." Pattern Recognition **16**(1): 1-8.

Lin, X. and K. Otobe (2001). "Hough transform algorithm for real-time pattern recognition using an artificial retina camera." Optics Express **8**(9): 503-508.

Radon, J. (1917). "Über die Bestimmung von Funktionen durch ihre Integralw-erte längs gewisser Mannigfaltigkeiten." Ber. Verh. König. Sächs. Ges.Wiss. Leipzig, Math.-Phys. Kl **69**: 262-217.

Russ, J. C. (2002). The Image Processing Handbook: Fourth Edition, CRC Press.

Savitsky, A. and M. J. E. Golay (1964). "Smoothing and Differentiation of Data by Simplified Least Squares Procedures." Anal. Chem. **36**: 1627-1639.

Sederman, A. J. (1998). MRI Studies of Single and Multi-phase Flow in Porous Media. Department of Chemical Engineering. Cambridge, University of Cambridge.

Shannon, C. E. (1948). "A Mathematical Theory of Communication." The Bell System Technical Journal **27**: 379-423, 623-656.

Siu, N. (2004). "A Model-independent Method for Automatic Correction of Non-uniformity in MRI data." http://www.ece.ubc.ca/~nsiu/doc/eece544projectreport.pdf.

Thacker, N. A., A. J. Lacey, *et al.* (2002). "Validating Field Homogeneity Correction Using Image Information Measures." www.tina-vision.net.

Vokurka, E. A., N. A. Thacker, *et al.* (1999). "A fast model independent method for automatic correction of intensity nonuniformity in MRI data." Jmri-Journal of Magnetic Resonance Imaging **10**(4): 550-562.

Wang, Y. F., Z. S. Mao, *et al.* (1998). "Scale and variance of radial liquid maldistribution in trickle beds." Chemical Engineering Science **53**(6): 1153-1162.

Wirth, M. A. (2002a). "Texture Analysis." www.uoguelph.ca/~mwirth/cis6320/lec6_notes.pdf.

Wirth, M. A. (2002b). "Image Segmentation: Edge-based." http://www.uoguelph.ca/~mwirth/cis6320/lec5_notes.pdf.

Zucker, S. W. and D. Terzopoulos (1980). "Finding Structure in Co-Occurrence Matrices for Texture Analysis." Computer Graphics and Image Processing **12**(3): 286-308.

# Chapter 8 – Monolith Configurations and Holdup Measurement Techniques

This chapter will describe the experimental setups used in this work to investigate the monolith reactor. Chapters 9-11 will probe various characteristics of these setups using the novel MRI techniques and image analysis algorithms developed as a part of this work, previously described in Chapters 5 and 7. This chapter will then focus on the measurement of liquid holdup by MRI techniques, applied to a range of monolith configurations with the aim of extracting operational information on the flow behaviour under different conditions. The ability to measure holdup at high flow rates and in the presence of dynamic two-phase flow will be investigated.

**CHAPTER 8 – MONOLITH CONFIGURATIONS AND HOLDUP MEASUREMENT TECHNIQUES**



## 8.1. Introduction

This chapter will begin by describing the various experimental flow setups used in this work (Section 8.2), and then use MRI to acquire holdup data (Section 8.3) at a range of flow rates and in various process configurations, from the simple bubble sparger to the far more challenging turbulent froth reactor. The applicability of the methods for each of the configurations will be investigated.

## 8.2. Experimental monolith configurations

This section will describe the two-phase flow configurations used as a part of this research into monolith reactors. These cover a range of different flow configurations, fluid distribution types, and superficial velocities which will allow us to probe the limitations of the MRI techniques, and to investigate the extent to which we may characterise these challenging systems. The setups will be presented for two-phase systems of increasing fluid velocity; the bubble sparger (air upflow through stagnant water), co-current upflow (air and water upflow), showerhead distributor (air and water downflow), spray distributor (air and water downflow), and finally the froth reactor (rapid air and water downflow). The flow rates investigated for each configuration are summarised in Table 8.1.

Table 8.1 – Summary of the monolith flow configurations, and the relevant gas and liquid flow rates.

| Configuration | Flow direction | Gas flow rate (cc/min) | Liquid flow rate (cc/min) |
|---|---|---|---|
| Bubble sparger | Up | 100 - 500 | N/A |
| Co-current upflow | Up | 100 - 1000 | 0 - 1900 |
| Showerhead | Down | 100 - 600 | 1000 - 4150 |
| Spray distributor | Down | 100 - 400 | 1000 - 3300 |
| Froth reactor | Down | 100 - 1000 | 2000 - 8000 |

### 8.2.1. Bubble sparger

A flow loop consisting of a Perspex column of 50 mm i.d., containing a cylindrical, square-channel monolith was positioned within the vertical bore of the magnet. In the bubble sparger configuration, gas is bubbled upwards through stagnant liquid, depicted in Figure 8.1. The air is supplied by a compressed air line and was bubbled through a sintered glass frit of diameter 20 mm (Figure 8.2) at a vertical distance of 59 cm below the base of the monolith to produce bubbly flow. The superficial gas flow direction was along the axis of the magnet, vertically upwards. To reduce the surface tension similar to conditions in real reactors, 0.5 wt% of propan-1-ol was added to the water. This allows bubbles to enter channels more easily, and prevents some of the maldistribution problems observed in systems where bubbles are much larger than the channels. This effect has also been observed by Pinto *et al.* (1997).

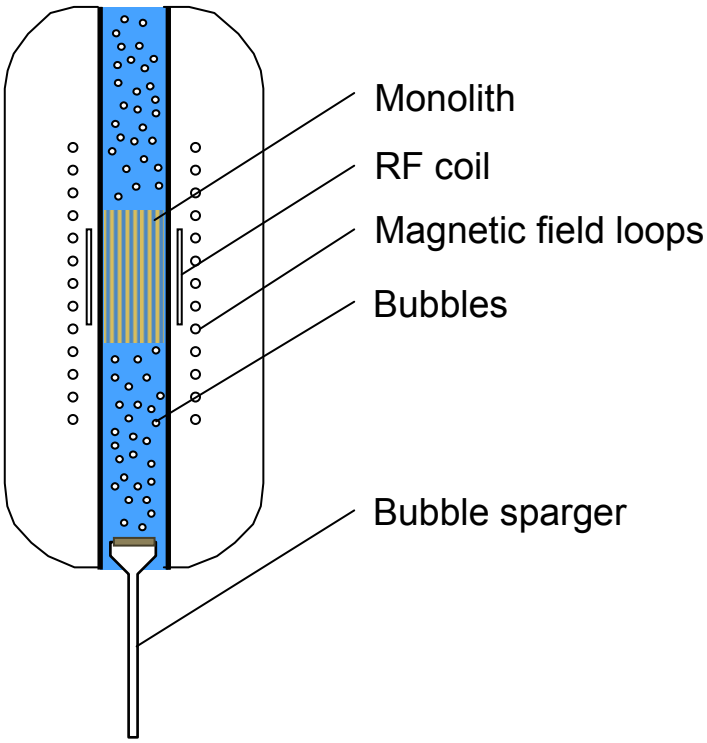


Figure 8.1 – Schematic of the setup for two-phase flow through the monolith using the bubble sparger setup.

Figure 8.2 – A photograph of the bubble sparger used in this work, consisting of a sintered glass frit. The diameter of the frit is 20 mm.

Two monoliths of different cell densities were used (shown in Figure 8.3), namely one of diameter 42 mm and length 0.15 m rated at 200 channels per square inch (cpsi), and the second of diameter 49 mm and length 0.15 m rated at 400 cpsi, with

each channel having square cross-sections of side 1.7 mm and 1.2 mm respectively. The monolith was held in place with a rubber 'O'-ring and a Teflon wrapping which ensured that no fluid bypassed the monolith at the sides.

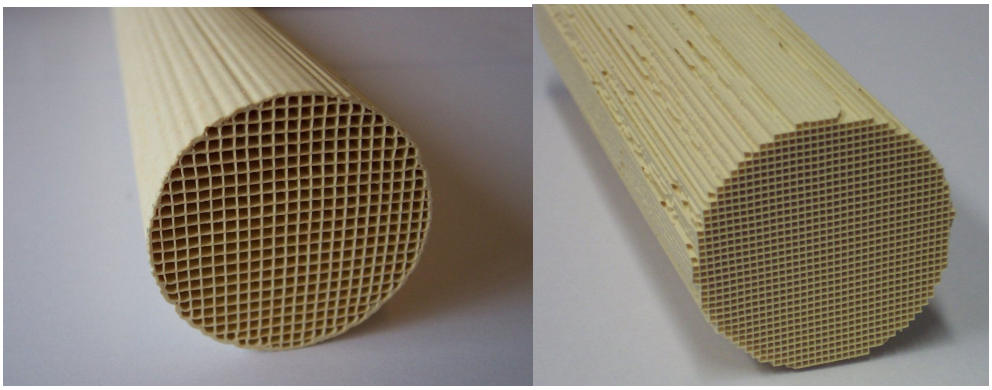

Figure 8.3 – Photographs of monolith sections used in this work; (a) shows a 200 cpsi monolith of 42 mm diameter, and (b) shows a 400 cpsi monolith of 49 mm diameter.

The bubble sparger has the lowest flow rate of the two-phase flow systems, and hence produces the clearest and most useful MRI results. It was also possible to operate the bubble sparger in co-current flow mode – this was done by flowing water around the side of the bubble sparger. It will be shown later in Section 11.2.5 that the flow in this configuration is more stable than air flow alone. However, the low-throughput of the system means that it is not industrially relevant, and so should be regarded primarily a test system for proving MRI techniques. By probing such a system we can discover many interesting characteristics which are of application in interpreting data from more complex systems. The aim is to quantify industrially-relevant two-phase experimental setups, such as the showerhead distributor, the spray distributor, and the froth reactor, which will now be described in turn.

## 8.2.2. Showerhead distributor

This is the most common fluid distribution method found in the literature, and typically consists of a disc with holes drilled in it to provide the fluid distribution. The showerheads used in this work were constructed from Perspex with a diameter of 45 mm, and the holes were drilled on a triangular pitch. Air was flowed around the side of the distributor. Four showerheads were investigated; 19 holes of 1 mm,

19 holes of 2 mm, 31 holes of 1 mm, and 31 holes of 2 mm, as shown in Figure 8.4. The distance between the showerhead and the top of the monolith could also be varied, and experiments were performed in the range 7-21 cm. It will be shown later (Section 10.2) that varying these design parameters gives very different fluid distributions. We can therefore use these results to optimise our design parameters. A schematic of the setup is shown in Figure 8.5.

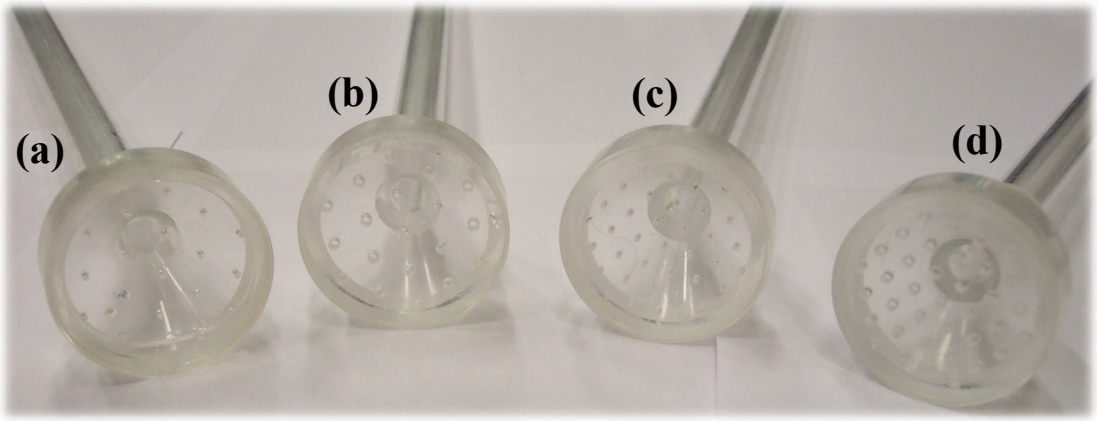


Figure 8.4 – A photo of the four showerheads used as part of this work; (a) 19 holes of 1 mm, (b) 19 holes of 2 mm, (c) 31 holes of 1 mm, and (d) 31 holes of 2 mm.

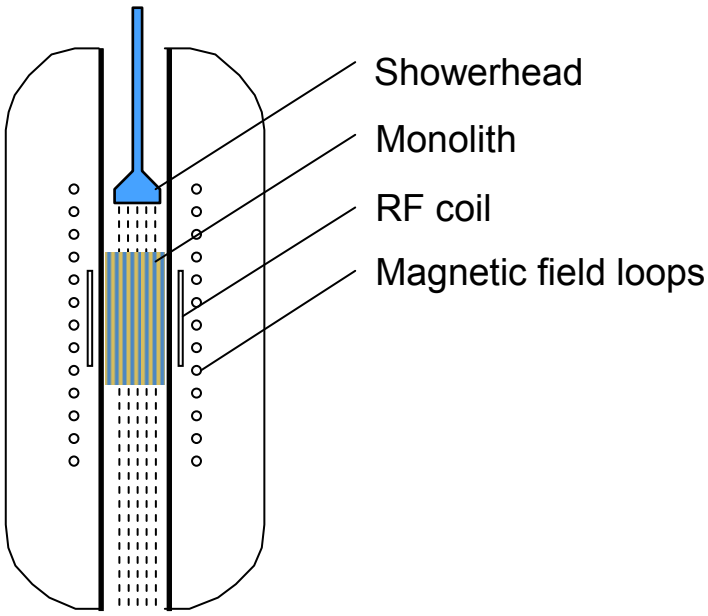


Figure 8.5 – Schematic of the setup for two-phase flow through the monolith using the showerhead distributor for liquid distribution.

### 8.2.3. Spray distributor

This setup is similar to the showerhead distributor, but uses a brass spray nozzle to produce a cone of spray with half-angle of 30°. This is expected to provide a far more even fluid distribution than the showerheads, and also give a more consistent range of droplet sizes and a narrower range of velocities within the monolith channels.

Two different spray nozzles were used, provided by PNR UK Ltd.[1], with model numbers BGQ1740 and DAQ1370 (shown in Figure 8.6). The digits in the model number correspond to the flow rate produced at a pressure of 3 barg; 7.40 l/min and 3.70 l/min respectively. The latter nozzle has restricted flow, and should therefore produce drops of smaller diameter. The effects of this will be investigated later (Section 10.3 and Section 11.2.6). The spray nozzles were positioned 6 cm above the top of the monolith, to ensure that the cone of fluid adequately covered the cross section. At this distance, the edge of the spray cone just touched the sides of the Perspex column at the top of the monolith. Air was flowed around the spray nozzle. A schematic is shown in Figure 8.7.

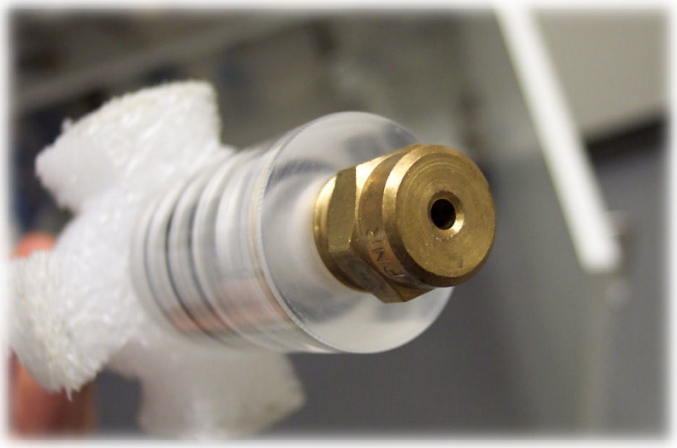

Figure 8.6 – Photo of a spray nozzle used in this work, model number BGQ1740.

[1] PNR UK - www.pnr.co.uk

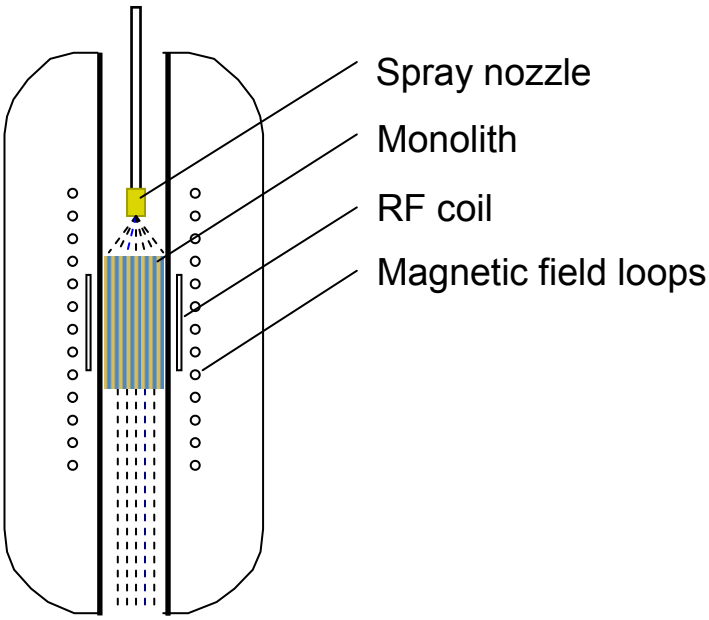


Figure 8.7 – Schematic of the setup for two-phase flow through the monolith using the spray distributor for liquid distribution.

### 8.2.4. Froth reactor

The froth reactor is the most challenging experimental setup used in this work. It consists of a turbulent froth region 52 cm in length, where air and water are mixed under high flow through a narrow orifice. This produces a fine bubble mixture, and there is sufficient energy imparted by the pressure drop across the orifice to create a stable froth region. This two-phase mixture then passes down through the monolith. See Figure 8.8 and Figure 8.9 for illustration. It is expected that most of the MRI pulse sequences will simply fail due to the high velocities (~30 cm/s) and rapidly-changing, dynamic nature of the flow.

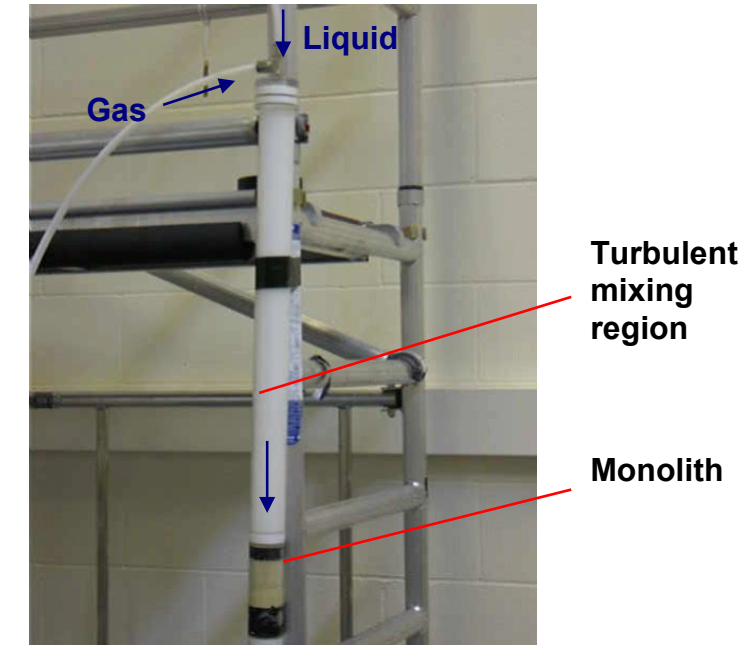


(a)

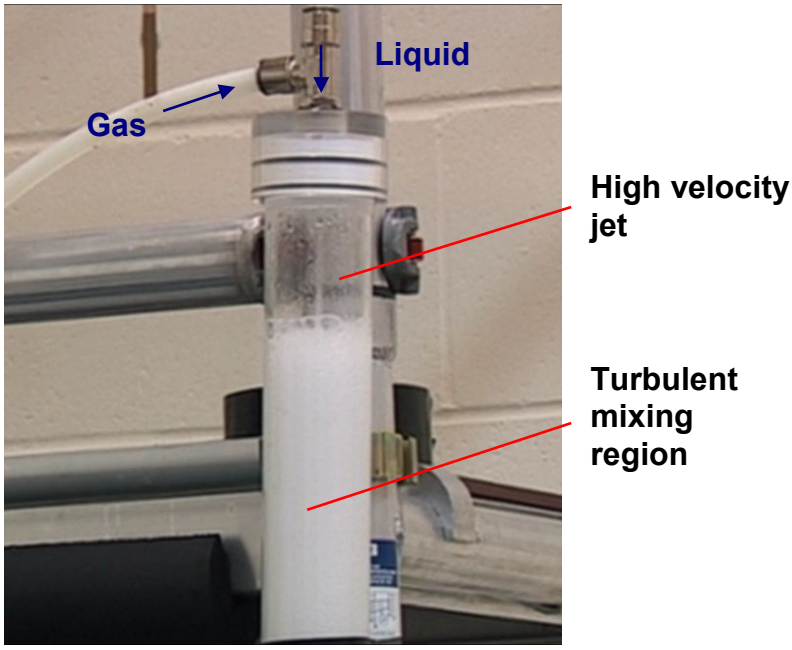


(b)

Figure 8.8 – Photos of the froth reactor in operation; (a) wide-shot, and (b) close-up.

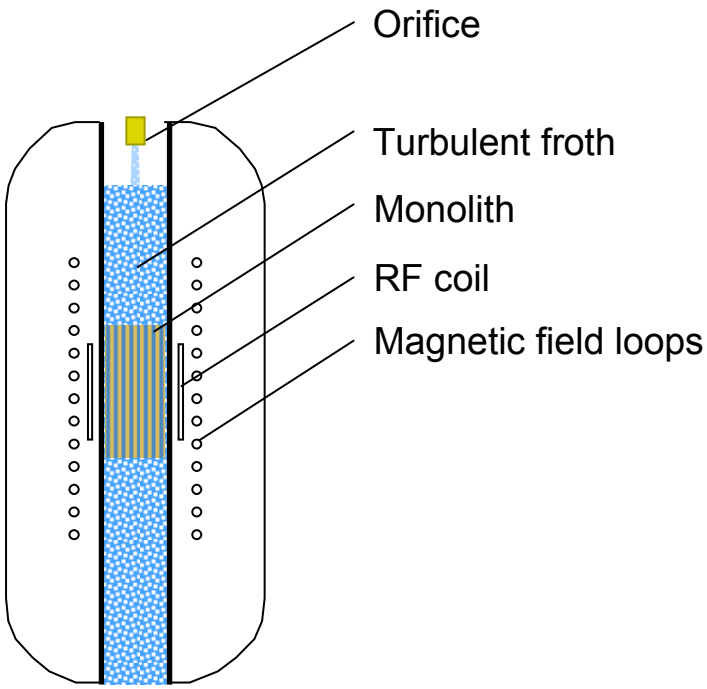


Figure 8.9 – Schematic of the setup for two-phase flow through the monolith using the spray distributor for liquid distribution.

Two different jet nozzles were used in this work; these were simple orifices of 3 mm and 5 mm diameter. These give different jet velocities for the same flow rate, and therefore impart different energies into the turbulent froth for any given flow rate. Using a narrower jet will therefore give the froth more energy and allow the formation of a stable mixture at lower flow rates. This will allow us to investigate the system more easily with MRI.

We have now described the reactor configurations used in this work, and the following section will address the characterisation of the liquid holdup in these systems using MRI techniques.

## 8.3. Liquid holdup measurements

This section will address the measurement of liquid holdup for the various process configurations of the monolith reactor described in the previous section. Experimental MRI results will be presented and discussed, comparing the behaviour of the different reactor configurations. Two methods are proposed for the direct measurement of liquid holdup within the monolith reactor, and the merits of each will now be discussed in turn.

### 8.3.1. Liquid holdup from an FID experiment

Since MRI measures the density of nuclear spins within a sample, it can be used to measure the liquid content directly. The overall holdup in the monolith reactor at any one time can be calculated by comparing the measured intensity of an NMR signal relative to that of a fully water-saturated monolith (highest possible liquid holdup) and that of a drained monolith (noise level). The most reliable MRI experiment to perform is a non-spatially-resolved FID acquisition since this will be very fast (around 2 ms) and will not be significantly affected by relaxation contrast or by velocity effects (spins do not have sufficient time to leave the imaging region during the acquisition time). This gives as near a measurement of the instantaneous spin-density as possible. The FID signal is then Fourier transformed to obtain a spectrum, and then the area under the spectrum is a direct measure of the spin-density of the system, and therefore is directly proportional to the quantity of water present within the coil.

A series of experiments were performed on the various monolith configurations using a standard pulse-acquire FID experiment (Section 2.1.3). A hard 90° pulse was applied, and the subsequent return to equilibrium was acquired over a time of around 2 ms. This was repeated 32 times to obtain an average result, which is particularly important for the dynamic systems where the holdup may vary. Around 5 seconds was allowed between excitations to allow for $T_1$ signal recovery. The signals were then Fourier transformed and the areas under the spectra were measured.

Some examples of the application of this liquid-holdup measurement technique are shown in Figure 8.10 and Figure 8.11. It can be seen from Figure 8.10a that the single-phase flow system measures a liquid holdup of around unity, as expected. The

deviation from unity arises from velocity attenuation, and is seen to increase with liquid flow rate. However, it is seen to be a small effect. Figure 8.10b shows the holdup measured in a bubble sparger setup (Section 8.2.1), with upward bubbly flow through stagnant fluid at a range of air flow rates. Figure 8.10c shows measurements from a spray-distributor setup (Section 8.2.3), and again a clear trend with flow rate may be observed. The liquid holdup for this system is much lower than for the bubble sparger.

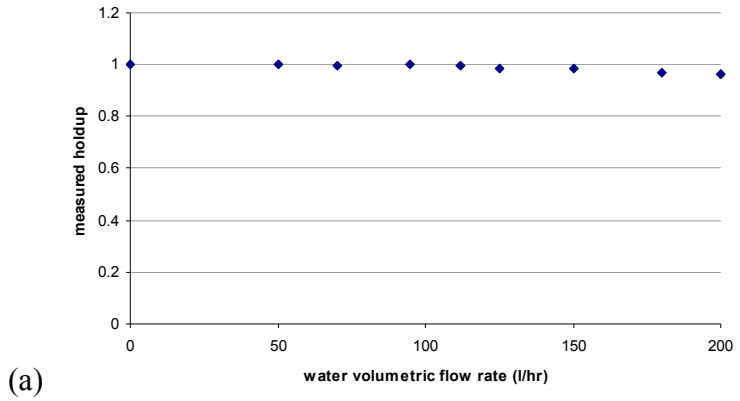


(a)

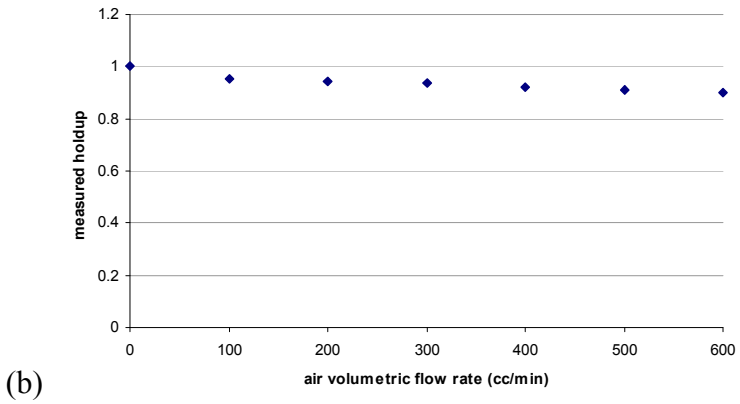


(b)

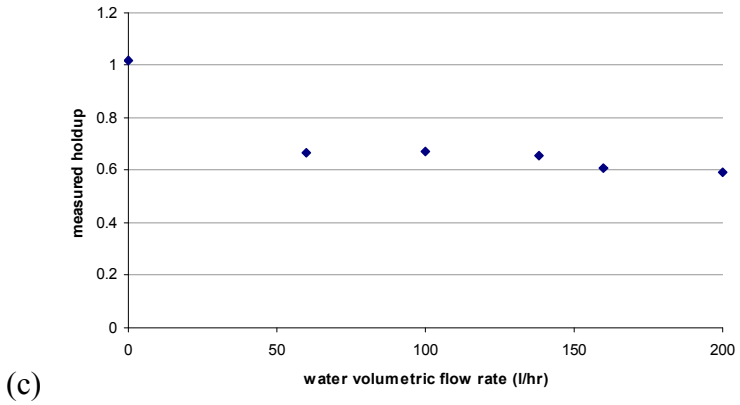


(c)

Figure 8.10 – Results of the liquid-holdup measurement technique in a 200 cpsi monolith for (a) single-phase liquid flow at a range of liquid flow rates, (b) upward bubbly flow through stagnant liquid at a range of air flow rates, and (c) co-current gas/liquid downward flow using a spray distributor (BGQ1740), at a range of liquid flow rates and with the air flow rate held at 400 cc/min.

The spray distributor data show little change in liquid holdup with increasing water volumetric flow rate, but it might have been expected that the liquid holdup would increase. It is proposed that this observation is possibly due to an increased stability of the downflow geometry, whereby if a channel receives more water it will then drain faster (higher hydrostatic head) which reduces the liquid holdup in the channel. Therefore, the system is able to compensate to some extent for variations in the local conditions and retain a constant liquid holdup.

Some examples of two-phase flow, varying both liquid and gas flow rates, are shown in Figure 8.11. Figure 8.11a shows liquid holdup results for the bubble sparger configuration and Figure 8.11b shows results for the froth reactor at a range of gas and liquid flow rates. It can be seen that increasing the gas flow rate typically decreases the liquid holdup in the systems, and a consistent trend is observed.

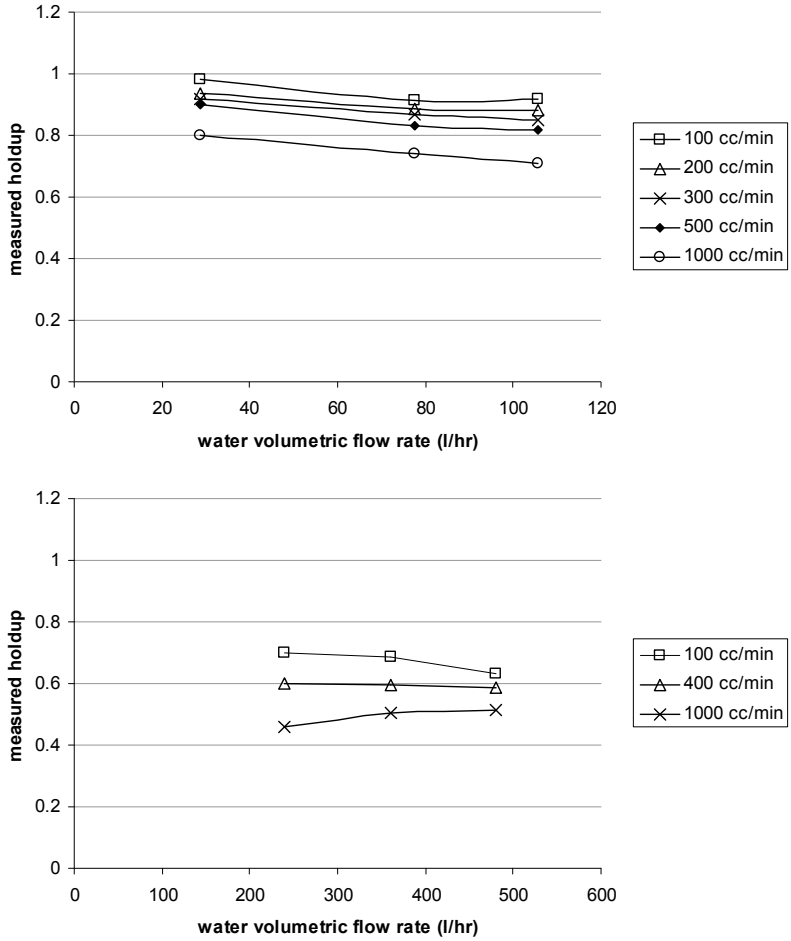


Figure 8.11 - Results of the liquid holdup measurement technique for two-phase flows in a 200 cpsi monolith for (a) co-current upward flow using a bubble sparger, and (b) downflow in a high-flow rate froth reactor.

Other factors such as the effect of monolith cell density may be investigated. Figure 8.12 shows results for the bubble sparger setup using different cell densities, and we can see that the higher cell density monolith has a lower liquid holdup. This trend can be explained by considering the flow behaviour within the channels; bubbles within narrower channels will move more slowly due to surface tension effects, and so for a given air flow rate the air holdup within the reactor must be higher, as seen in the holdup measurement.

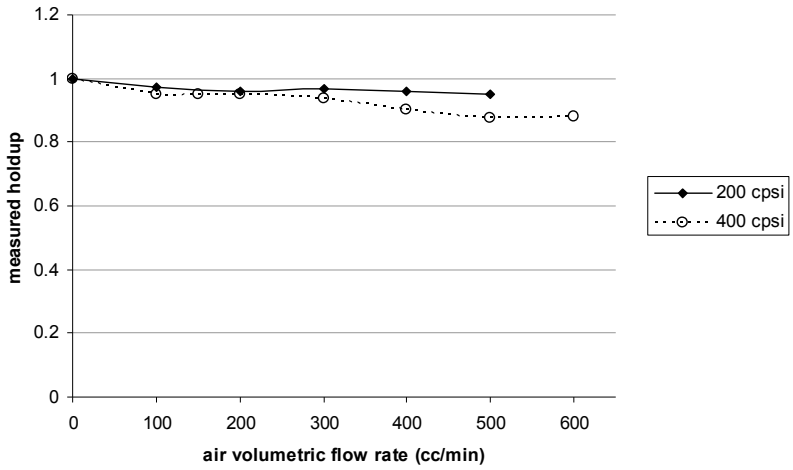


Figure 8.12 – Liquid holdup measurements of the bubble sparger rig, using monoliths of different cell densities.

## 8.3.2. Liquid holdup from 2D images

A second method for the determination of liquid holdup under flowing conditions is to take a 2D slice through a row of channels, such as is shown in Figure 8.13. This is a standard RARE image, with a field-of-view of 5.06 cm × 5.06 cm, and a resolution of 395 μm × 791 μm. The data were acquired over 146 ms. The slice thickness is 1.7 mm, which corresponds to the thickness of a row of channels. By comparing the intensity of this image to that of a fully-liquid-saturated monolith, we can deduce the holdup in each individual channel. The advantage of this technique is that it allows us to resolve the local holdup of each channel, and later it will be shown that we can correlate the local holdup with other variables such as bubble size and local channel velocity (Section 11.2.2).

The disadvantage of this technique, is that we can only study a single row of channels at a time, and so our sample may not be representative of the whole monolith. Also, problems will be encountered at high flow rates where velocity attenuation will affect the intensity, and also spins will leave the field-of-view during the acquisition time. This will result in lower intensities, which will be interpreted as a lower liquid holdup than is actually the case. However, for the relatively low-flow rate systems such as the bubble sparger, this method produces valuable insights into the local hydrodynamics of the system.

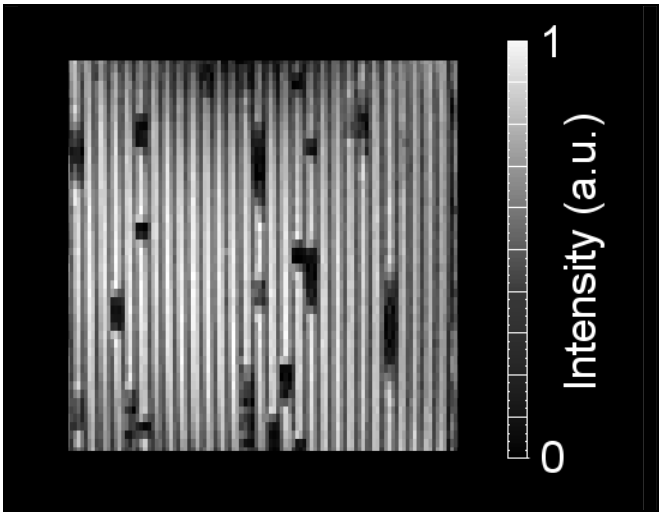


Figure 8.13 – A slice though the bubble sparger system at a flow rate of 100 cc/min. The contrast within the image can be used to determine the local liquid holdup within each channel.

## 8.3.3. Further discussion

It should be noted that the first method for measuring liquid holdup (by FID measurement) measures the total liquid content within the RF coil, rather than specifically the liquid content within the monolith reactor. In practice, there is generally a stagnant space surrounding the monolith which is filled with water but does not flow (thereby avoiding fluid-bypass of the monolith). An example of this effect is shown in Figure 8.14. The holdup we are interested in measuring is related to the spin density within the monolith itself, but the contribution from the liquid-filled region in the surrounding annulus will bias our measurement. Therefore, the true liquid holdup within the monolith channels should be lower than the measured

results of Section 8.3.1. This effect will now be discussed, and a correction method will be proposed and compared to the local holdup measurements of Section 8.3.2.

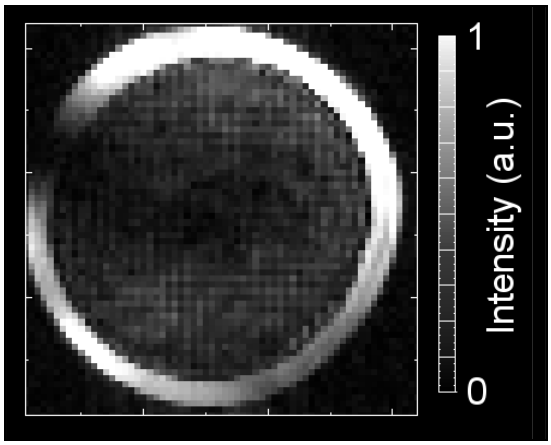


Figure 8.14 – Example of a typical MRI cross-section through the monolith under spray distributor operation. Field-of-view is 5.06 cm × 5.06 cm. The presence of a liquid-filled stagnant zone can be clearly seen.

The annulus area will be of a stable intensity (totally liquid-filled, near-zero motion) and can be assumed to be independent of the flow rate through the monolith reactor. This contribution to the measured intensity can be subtracted from the original FID-based holdup measurement (Section 8.3.1). By following the algebra shown in Eq. 8.1, we may convert our overall holdup measured with the FID-based method into a measure of the liquid holdup of the monolith reactor only. As can be seen from Eq. 8.1, the conversion factor is based upon a single combined parameter which can be calculated from measurable parameters.

$$\text{Measured FID Holdup} = \frac{\text{Flowing Intensity}}{\text{Zero flow Intensity}} = \frac{I_a + I_c}{I_a + I_{c,\max}} = \frac{I_a}{I_a + I_{c,\max}} + \frac{I_c}{I_a + I_{c,\max}}$$

$$\Rightarrow \qquad \text{True Holdup} = \frac{I_c}{I_{c,\max}} = \frac{\text{Measured FID Holdup} \times \left(I_a + I_{c,\max}\right)}{I_{c,\max}} - \frac{I_a}{I_{c,\max}} \qquad \textbf{(Eq. 8.1)}$$

$$= \text{FID Holdup} \times \left(\frac{I_a}{I_{c,\max}} + 1\right) - \frac{I_a}{I_{c,\max}}$$

where $I_c$ is the intensity contribution from flowing channels, $I_{c,max}$ is the intensity contribution from zero-flow channels, and $I_a$ is the intensity contribution from the stagnant annulus region. These parameters can be determined from a 2D slice through the liquid-filled monolith, such as that in Figure 8.15. $I_a$ must be calculated

for a 3D annulus, and the intensity of any voxel in this annulus is taken to be the same as that of an annular pixel in the 2D image. Similarly, $I_{c,max}$ is calculated for a cylindrical volume, with voxel intensity calculated as the average monolith intensity measured from the 2D image.

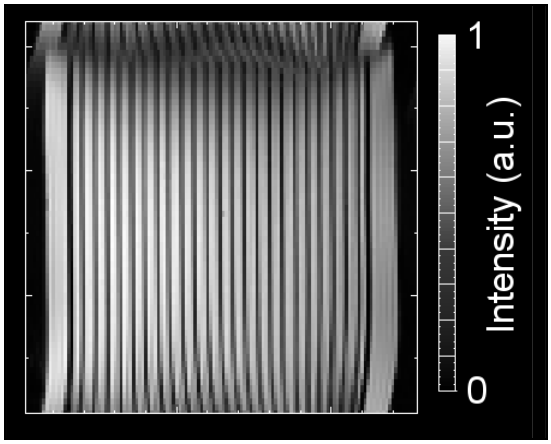


Figure 8.15 – A 2D RARE image of the liquid-filled 200 cpsi monolith. The intensity within the image can be used to determine the necessary parameters ($I_a$, $I_{c,max}$) to correct the FID holdup measurements.

A comparison of the corrected and uncorrected data are shown in Figure 8.16a and Figure 8.16b, respectively. These graphs show the data from the two types of liquid-holdup measurement, and it can be seen that the agreement is far better after the application of this derived correction factor.

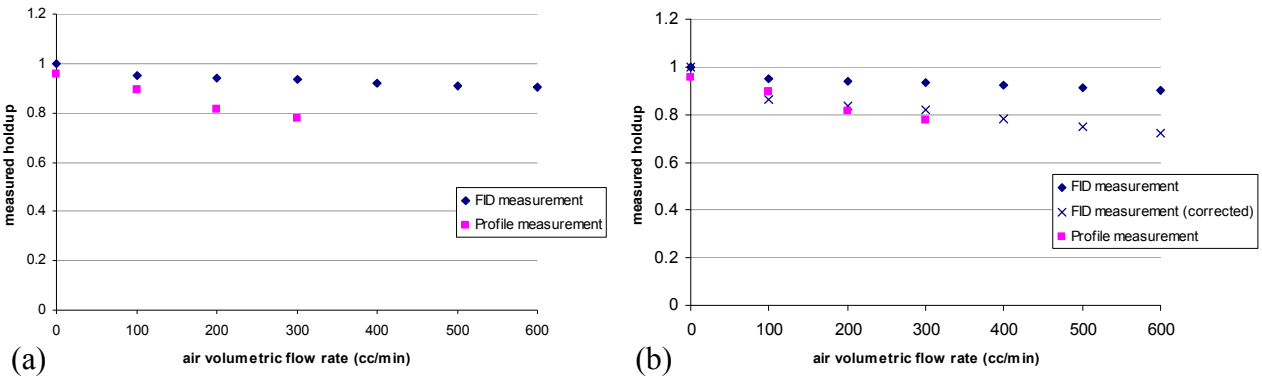


Figure 8.16 – A comparison of the holdup measurements taken by the FID-based measurement of Section 8.3.1 with that of the profile-based measurement of Section 8.3.2; (a) shows the comparison of the data types acquired for the same system, and in (b) the FID-based data has been corrected for the contribution from stagnant regions within the imaging coil.

Therefore, we have shown that the profile-based holdup measurement matches well with the FID holdup measurement, once this has been corrected for errors inherent in the system. The FID-based measurement is much more robust and is a better method to be used for the measurement of holdup in a system, provided the geometry and stagnant regions within the imaging coil can be taken into account.

## 8.4. Conclusions

This chapter has introduced the various monolith process configurations, and illustrated the extent to which we can measure the holdup within these configurations. Two methods for the measurement of holdup were proposed, one of which is very rapid (2 ms) and gives a single measure for the entire reactor volume, and the second has a longer acquisition time (146 ms) but allows the holdup measurement to be localised to individual channels. The first is applicable in very high flow rate systems (up to around 15 cm/s), however the second will break down at relatively low velocities due to loss of spins from the imaging region. It has not been used in systems containing velocities above 5 cm/s. It has been noted that some systems have stagnant annulus regions around the monolith itself, and for this case a correction parameter has been proposed based on the relative intensities and volumes of the flowing and non-flowing regions within the reactor.

## 8.5. Nomenclature

| **Symbols** | **Definition** |
|---|---|
| $I_a$, $I_c$, $I_{c,max}$ | Intensity contribution from flowing channels, from the stagnant annulus area, and from channels in zero-flow |
| $T_1$ | Spin-lattice relaxation constant |

| **Abbreviations** | **Definition** |
|---|---|
| CDC | Co-current downflow contactor |
| cpsi | Cells per square inch |
| FID | Free induction decay |
| MRI | Magnetic resonance imaging |
| RARE | Rapid acquisition with relaxation enhancement |
| RF | Radio frequency |

## 8.6. References

Pinto, M. D. T., M. A. Abraham, *et al.* (1997). "How do bubbles enter a capillary?" Chemical Engineering Science **52**(11): 1685

# Chapter 9 – Bubble and Slug Size Analysis

This chapter will demonstrate the application of the novel MRI techniques described in Chapter 5 to the monolith reactor, with the aim of extracting operational information on the flow behaviour under different conditions. It will be shown that using various data processing techniques we can measure bubble and slug sizes / frequencies in these challenging systems. The extent to which we can study high flow rates is investigated.

**CHAPTER 9 – BUBBLE AND SLUG SIZE ANALYSIS**



## 9.1. Introduction

This chapter will use the novel MRI techniques described in Chapter 5 to acquire data from the monolith reactor under different configurations. The image processing techniques described in Chapter 7 will be implemented to extract the maximum amount of useful information from the data, including bubble and slug sizes and frequencies. The applicability of the methods to each of the configurations will be investigated, discussing the limitations of the methods.

As mentioned in Section 3.7.3, controlling the lengths of bubbles and slugs within the monolith channels is critical to the optimal operation of the monolith reactor. This chapter will begin by describing the MRI experimental parameters used in this work (Section 9.2), and then the data processing methods are presented (Section 9.3). Some discussion and verification of assumptions in the method will follow (Sections 9.4). Finally, the results are presented (Section 9.5), with a discussion of the extent to which we can study these challenging two-phase systems. The observations will then be discussed and compared to observations in the literature which use traditional measurement methods (Section 9.6).

## 9.2. MRI experiments

By utilising the RARE pulse sequence (Section 2.3.1) we can obtain snapshots of the flow within a row of channels and process this information to gain insight into the flow conditions present. An example is shown in Figure 9.1a showing a slice through a single row of channels. This image was acquired in 146 ms and has a slice thickness of 0.6 mm. The MRI technique is only applicable to low flow rate systems such as the bubble sparger (with typical velocities up to 5 cm/s), since the slug interfaces need to be well defined and unaffected by blurring over the acquisition time. It is found that the velocities encountered in the showerhead and froth reactor configurations are too fast to obtain useful data for slug length analysis. The spray distributor can have low enough velocities but the slug profiles are less well-defined, perhaps due to the droplet-nature of the spray which may not have sufficient time to coalesce to form Taylor flow. An example for the spray distributor system is shown in Figure 9.1b.

Therefore, the bubble sparger system was used to acquire a series of data at a range of gas flow rates. The data sets are 128×64 pixels, and 90 repetitions were taken for each flow rate (100-300 cc/min). The field of view was 5.06 cm × 5.06 cm, with a slice thickness of 0.6 mm. A time of 5 seconds was allowed between repetitions to allow for $T_1$ recovery.

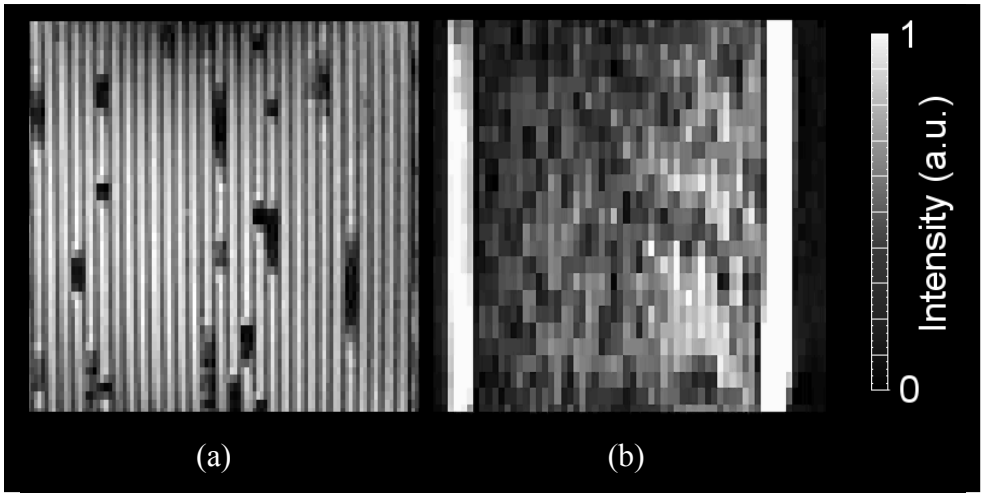


Figure 9.1 – (a) MRI data showing a slice through a single row of channels for the bubble sparger system. Data are 128×64 pixels with a resolution of 395 μm × 791 μm. The data were acquired over 146 ms; (b) MRI data showing a slice through a single row of channels for the spray distributor system. Data are 64×64 pixels with a resolution of 859 μm × 859 μm. The data were acquired over 156 ms.

## 9.3. Data analysis procedure

The data analysis procedure for processing raw MRI data and extracting bubble and slug lengths is outlined in the flowsheet of Figure 9.2, and each step will now be discussed in detail.

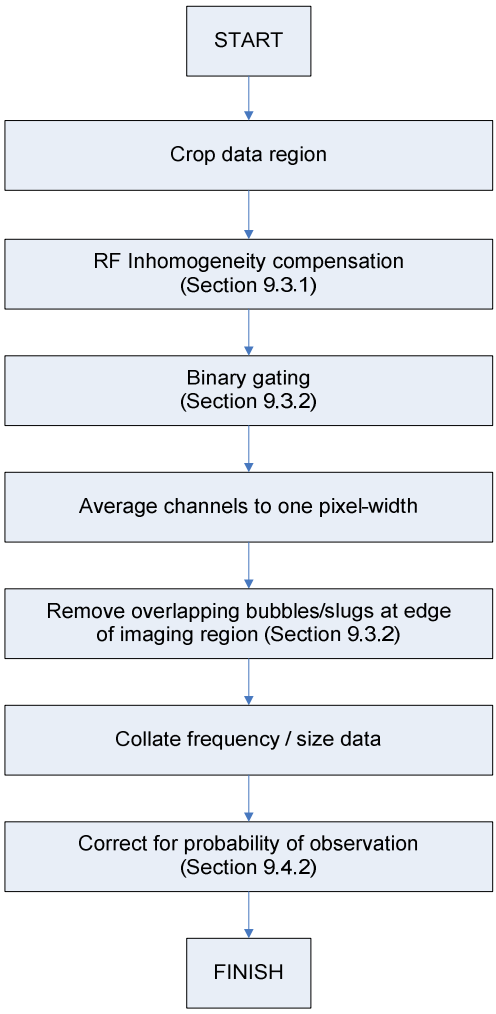


Figure 9.2 – Flowsheet of the data analysis procedure for processing raw MRI data and extracting bubble and slug size distributions.

## 9.3.1. Data processing

An example data set is presented in Figure 9.3a, showing the presence of Taylor flow within the channels of the monolith at a gas flow rate of 100 cc/min. Figure 9.3b is a corresponding data set for the liquid-saturated monolith acquired under the same parameters – these data may be regarded as a measure of the inhomogeneity in the system. The effects of RF coil inhomogeneity are apparent, and we can correct for this by dividing through the data by the homogeneity measure. The results of this correction procedure are shown for a single profile in Figure 9.4, and for the whole dataset in Figure 9.5. This correction factor is crucial if we are to binary gate the data to segregate gas and liquid regions. From Figure 9.4 we see that binary gating is now possible since we can define a single gating value to be used for the whole array.

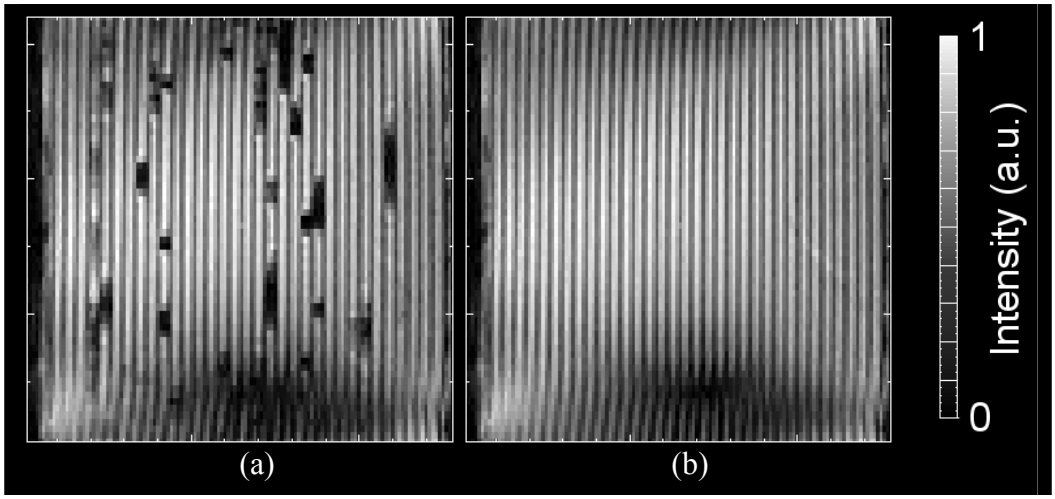


Figure 9.3 – (a) An example data set, showing the presence of Taylor flow within the channels of the monolith. The gas flow rate is 100 cc/min; (b) a liquid-saturated monolith acquired with the same acquisition parameters, showing the effects of RF coil inhomogeneity.

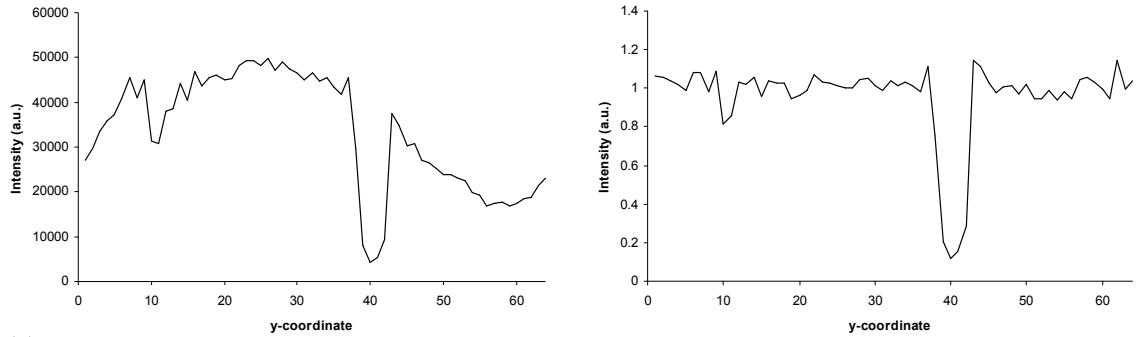


(a) (b)

Figure 9.4 – (a) shows a single channel profile from Figure 9.3a; and (b) shows the data after correction for local inhomogeneities.

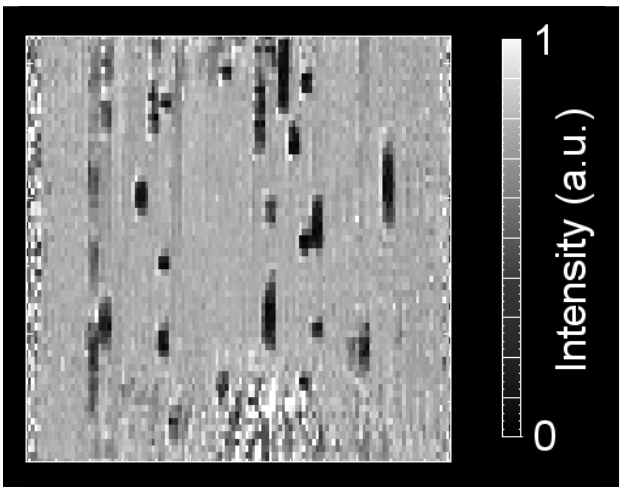


Figure 9.5 – The effects of correcting for the RF coil inhomogeneity, showing the corrected data of Figure 9.3a.

By processing the data in this manner we have lost the distinction of the channel walls. Figure 9.5 is a measure of the differences between the flowing data and the liquid-filled data, and therefore wall pixels have values around unity (as do liquid-filled pixels); the only pixels which should be non-unity are gas-filled pixels which should contain a value of near zero. However, since we are interested in the fluid rather than the walls it is desirable to mask out the wall regions. These are identified by projecting the original data of Figure 9.3b in the y-direction, giving the plot of Figure 9.6. The position of the monolith walls are clearly identified as local minima in the intensity profile. If we assume that each wall is one pixel in width, then we can mask the data as shown in Figure 9.7. The actual width of the walls is on the order of 210 μm, and since the MRI resolution is 395 μm this seems a reasonable assumption. Minimal data will be lost (we lose some from spins very close to the wall) and if we are only interested in the position of the bubble/slug interface, this will not be affected.

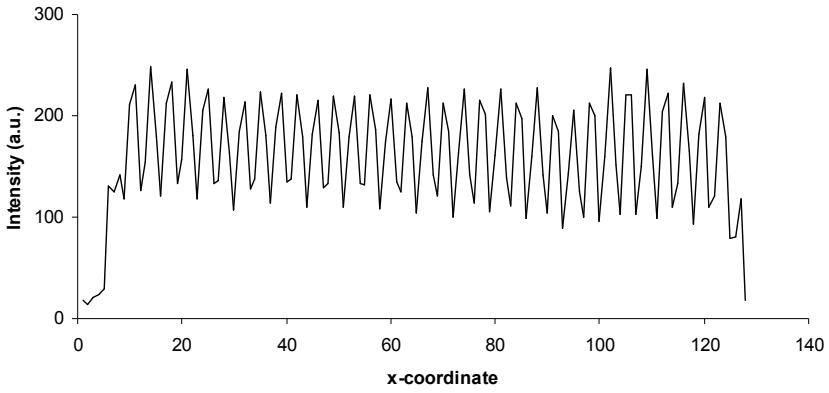


Figure 9.6 – A projection of the data of Figure 9.3a in the y-direction. The position of the monolith walls can be clearly identified.

Identification of the wall pixels allows us to produce the data of Figure 9.7a, and cropping the top and bottom regions (where the signal-to-noise ratio is poor), results in Figure 9.7b.

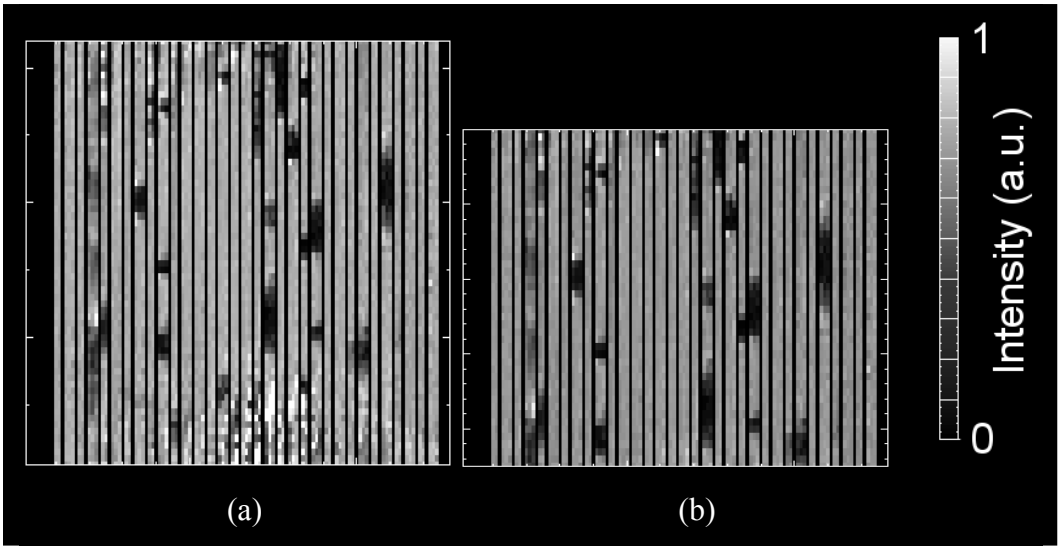


Figure 9.7 – (a) The results of masking out the monolith wall regions, and (b) the result of cropping the regions of poor signal-to-noise from top and bottom.

Some channels of Figure 9.7b are seen to be two pixels wide and some are three pixels wide, depending on the position of the wall. To remove any possible meniscus effect, it was decided to average together data within individual channels to a single pixel-width for processing (Figure 9.8).

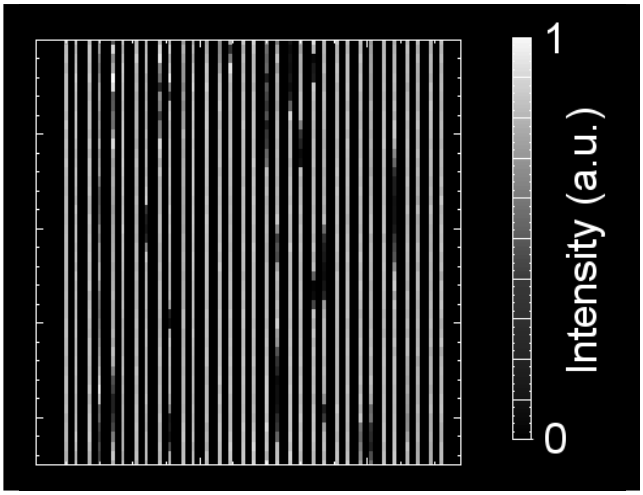


Figure 9.8 – The results of averaging together data within individual channels to one pixel-width for further processing.

## 9.3.2. Binary gating and bubble / slug identification

The data must now be binary gated to define the position of the bubble / slug interfaces. Some typical channel profiles are shown in Figure 9.9. From Figure 9.9(a-c), it can be seen that the noise level is around 0.1, and the maximum intensity fluctuates around 1. Therefore, the binary gating level should be set at a value of 0.55 to locate the half-height of the bubble / slug interface. The results of binary gating are shown in Figure 9.10. Further discussion and justification of this choice of gating level is presented in (Section 9.4.1).

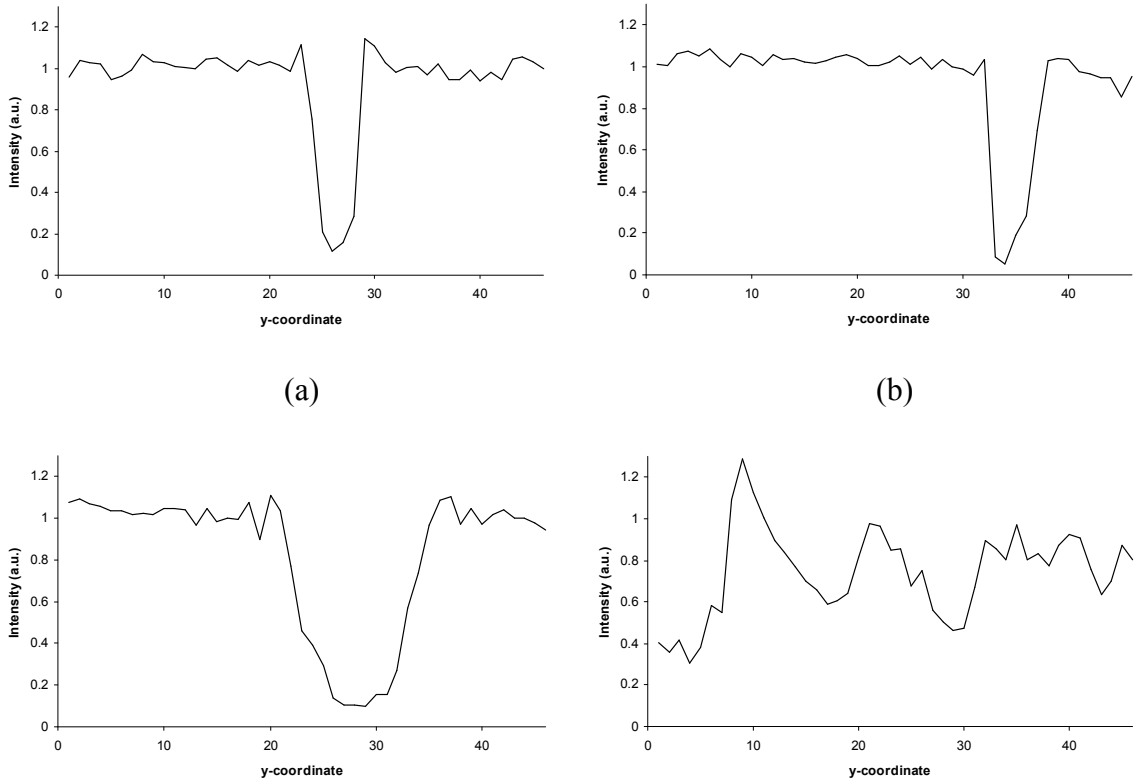


Figure 9.9 – A selection of channel profiles, taken from the data of Figure 9.8b; (a-c) shows the presence of well-defined bubble profiles; (d) shows a less well-defined channel profile. The meaning and implications of this type of data shown in (d) is discussed later in Section 9.4.1.

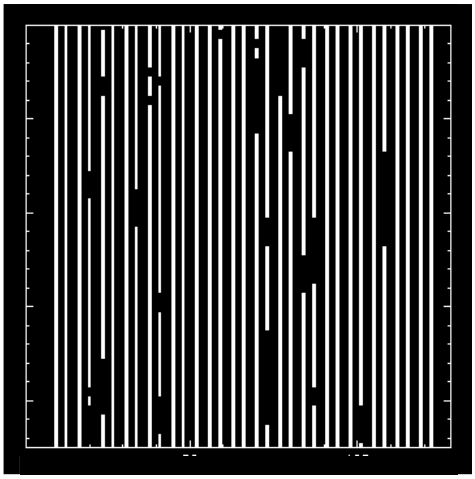

Figure 9.10 – The result of binary gating the sample data. Liquid is shown as white; walls and air bubbles are shown in black.

The data were then processed to identify bubbles and slugs, and to extract information on their lengths. First of all, we will examine the slugs. Only those slugs which are entirely contained within the field-of-view are measured, otherwise the true length of the slug would be longer than is apparent from our measurement. The results of this identification analysis of slugs are shown in Figure 9.11. For the identification of bubbles we need to invert the binary-gated data such that bubbles are now in white (as shown in Figure 9.12a), and then we perform the identification in the same manner (Figure 9.12b). Again, any bubbles which overlap the edge of the field-of-view are removed. The bubbles and slugs may now be measured and counted in a robust manner.

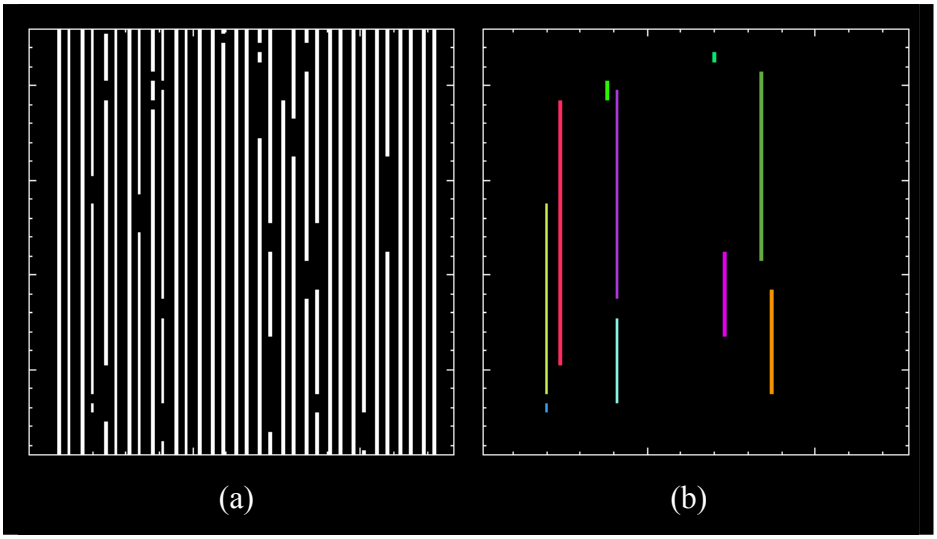


Figure 9.11 – (a) shows the binary-gated data, and (b) the identification of slugs, where each separate slug is shown by a different colour.

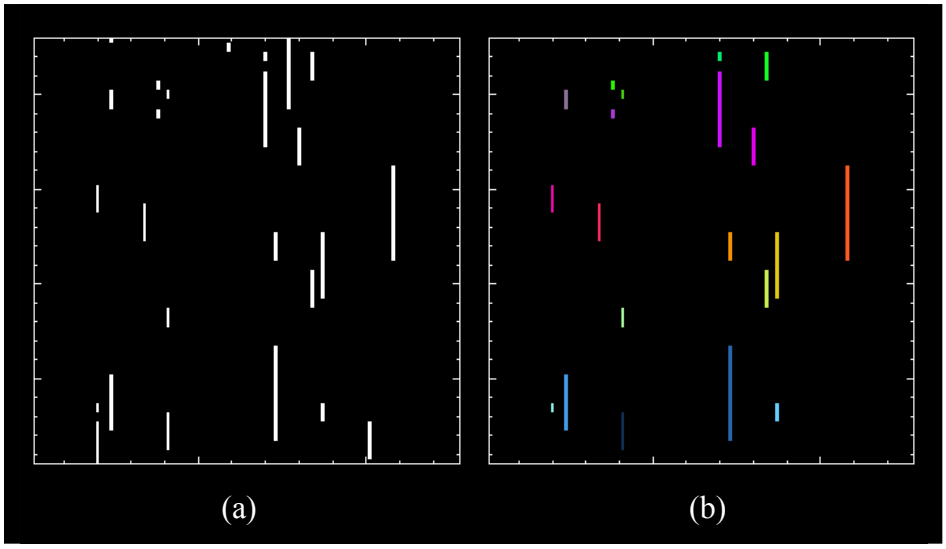


Figure 9.12 – (a) shows the inverted binary-gated data (with bubbles in white), and (b) the identification of bubbles, where each separate bubble is shown by a different colour.

## 9.4. Discussion of analysis procedure

However, before we present the results there are several aspects of the analysis to consider in order to better understand the sensitivity of the method we have chosen. First, we must investigate the choice of gating level – this directly affects the lengths of the bubbles and slugs and the choice of this value must be justified. Then we must address the question of whether the quantities we measure are in fact representative of the system itself – we encounter difficulties in observing very long slugs since they overlap the edge of the field-of-view, and will not be measured. Lastly, we will investigate the occurrence of less well-defined slug profiles and discuss their implication and hydrodynamic nature. Once we are satisfied that the procedure is relatively robust, final bubble and slug length / frequency results will be presented in Section 9.5.

### 9.4.1. Sensitivity of the gating level

This section will investigate the sensitivity of the choice of gating level on the results. This gating level ultimately defines the boundaries of the bubbles and slugs, and therefore their lengths, and so is crucial to the data analysis. We previously defined the gating level at 0.55, corresponding to the half-height between the noise level of

0.1 and the signal level of 1. The slope observed at a bubble / slug interface is a motion-blurring artefact, which is related to the velocity within the channel. By measuring the length of the bubbles and slugs at the half-height, we should obtain the most accurate measure of their *true* lengths.

If we vary the gating level between 0.4 and 0.7, we can get a feel for how sensitive the results we obtain are to the choice of gating level. For example, if we investigate the 100 cc/min flow rate case we can investigate how altering the gating level affects the total number of bubbles and slugs detected, as shown in Figure 9.13a and Figure 9.13b.

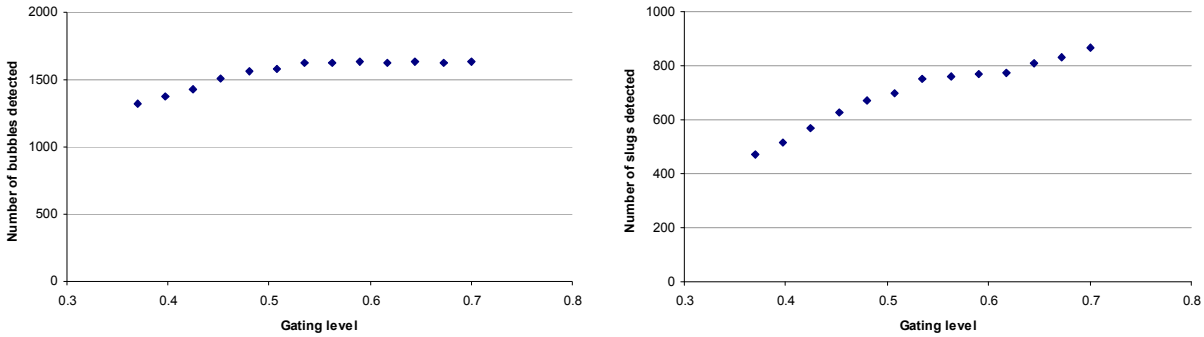


Figure 9.13 - Variation of (a) the number of bubbles detected and (b) the number of slugs detected as a function of gating level, for the bubble sparger system at a flow rate of 100 cc/min of air.[(1)]

The trends are interesting because it might be expected that changing the gating level should have no effect on the *number* of bubbles detected, only their length. If the bubble profiles consist of drops in intensity to near zero for the entire length of the bubble, then it will always intersect our gating level whatever it may be set at. But the trend observed from the profiles is that the number of detected bubbles and slugs changes with gating level, and this shows that there is some other effect. It is likely that this is due to the type of profile described in Section 9.3.2 where the slug interface is not well defined. In this case, changing the gating level will result in

[1] It is observed that the number of bubbles observed does not tally with the number of slugs observed. It might have been expected that these should match, however there is an issue with our limited field-of-view, which results in more bubbles being identified than are slugs. This issue is further addressed in Section 9.4.2.

different numbers of intersections and hence this can explain the trends observed in the gating sensitivity analysis. As can be seen from the plots of Figure 9.13a and Figure 9.13b, a gating value of 0.55 (as was previously suggested as a sensible value) gives a stable value of the number of bubbles detected, and a short plateau in the number of slugs detected. It appears to be a good choice.

Changing the gating level has the following effect on our results, as shown in Figure 9.14. The distribution changes from an exponential decay with lots of very small bubbles to one with more, large bubbles and exhibiting a maximum.

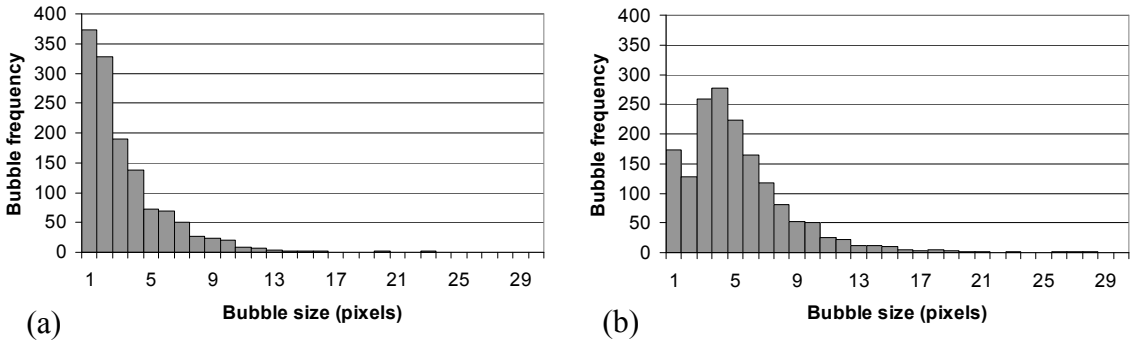


Figure 9.14 - Bubble size frequency charts for the bubble sparger at an air flow rate of 100 cc/min at gating levels of (a) 0.4 and (b) 0.7.

## 9.4.2. Measurement Effects

Since we neglect slugs which overlap the edge of the field-of-view, we can only measure slugs that are sandwiched between two bubbles. This means that as we look for longer and longer slugs, we are unlikely to see them since they cannot be detected by our method. The probability of observing them becomes vanishingly small as the slug length approaches the dimension of the field-of-view. Thus, we would expect to observe an apparent decrease in frequency with increasing slug length, even if the distribution were uniform in actuality. It is proposed that the error in our measurement can be estimated as a probability of successful measurement as a function of slug length. This argument holds true for the observation of bubbles too, but the effect is expected to be smaller due to the nature of the slugs being several times longer than the bubbles.

It is proposed to model the bubble (or slug) position in the manner shown in Figure 9.15. In processing our data we have selected a field of view consisting of 44 pixels in the flow-direction, however any bubbles which touch the edge of the field-of-view are removed from the analysis, so the useful field-of-view is 42 pixels. If we assume that the bubble is equally likely to be at any position, then for a one-pixel bubble we have a range of 42 detectable positions out of 44 possible positions, and hence the likelihood of detection for a one-pixel bubble is $(42/44)$ = 95%. Similarly, for a two-pixel bubble there are 41 useful positions out of 44, giving a probability of detection of 93%. Calculating these probabilities as a function of actual length yields the plot of Figure 9.16. It can be seen that as the bubble / slug length increases, the probability of detection decreases. Above 42 pixels in length, the bubbles and slugs cannot be detected.

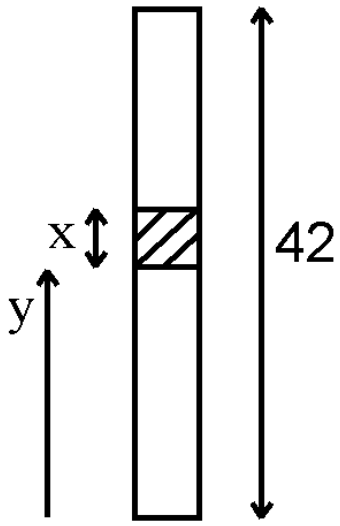


Figure 9.15 – Schematic of the position of a bubble (or slug) within the viewing area. The length of the bubble is *x* and the position of its leading edge is *y*. The total length of the field-of-view is 42 pixels.

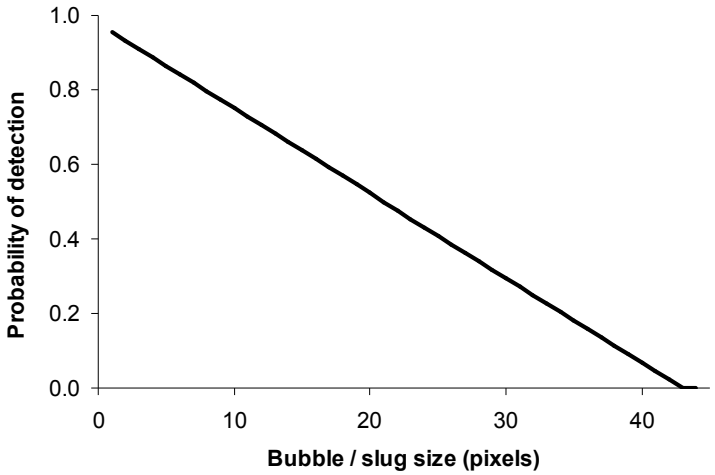


Figure 9.16 – Probability of detection as a function of actual bubble / slug length.

Hence, using this probability factor we can now adjust the observed bubble and slug distribution plots. i.e. if we detect 10 one-pixel bubbles, then we can say that we will likely have 10.5 bubbles in reality, and if we have 5 ten-pixel bubbles then we will have 6.6 in reality. For a small sample size this kind of approximation may be dubious, but when applied to a large enough dataset the results should tend towards the correct value(2). Of course, we run the risk of greatly skewing the tail end of the graph, since if several very long bubbles are detected their effect will be over-amplified, so care should still be taken in interpreting the results. The results of applying this correction factor are illustrated in Figure 9.17 which shows bubble and slug distributions before correction, and Figure 9.18 which shows the effect of correction for this probability of observation.

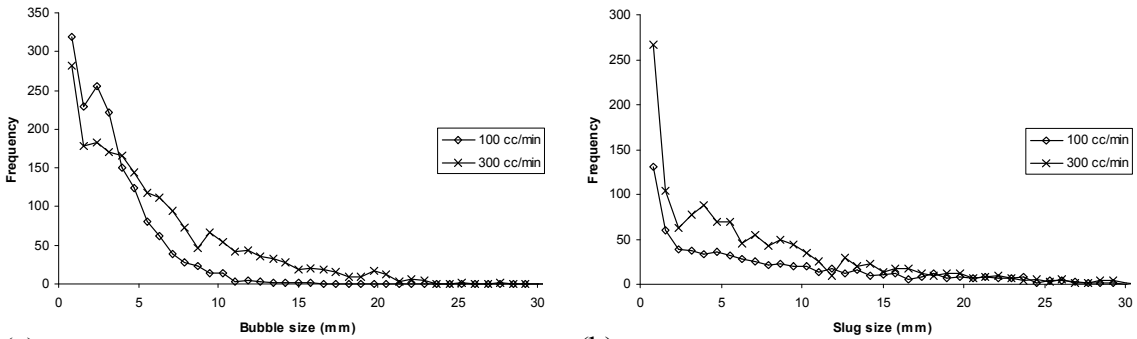


(a) (b)

Figure 9.17 – Results of the length/frequency analysis for the bubble sparger at 100 cc/min and 300 cc/min for (a) bubbles, and (b) slugs.

[2] For reference, our data sets consist of 33 channels, and 90 repetitions at each flow rate, giving a sample size of 2970 observations. These 2970 channels contain on the order of 1600 bubbles and 800 slugs.

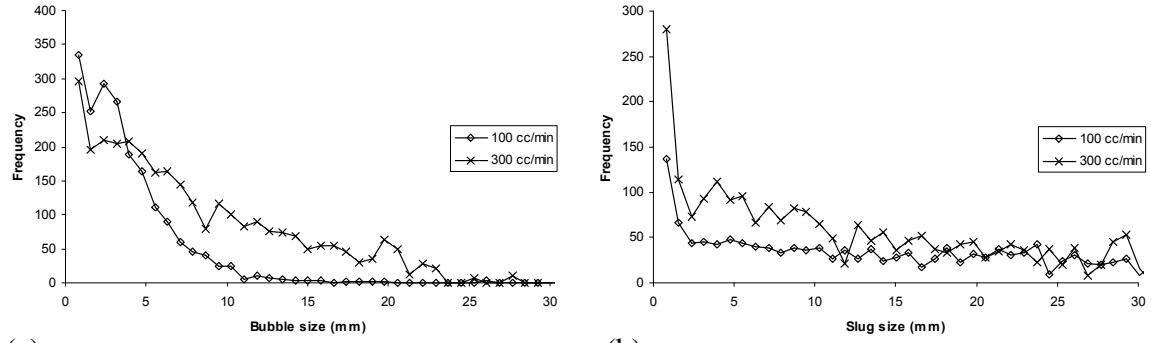


(a) (b)

Figure 9.18 – Results of the length/frequency analysis corrected for the probability of observation for the bubble sparger at 100 cc/min and 300 cc/min for (a) bubbles, and (b) slugs.

As can be seen, there is relatively little effect on the bubble distribution on applying the probability factor associated with measurement of the system. This is due to the fact that the bubbles are generally short, and so most of the bubbles will be detected without difficulty. The probability correction makes the slug distributions much less exponential and nearer uniform, indicating that the length of slug obtained is more or less independent, with perhaps a slight bias towards shorter ones.

### 9.4.3. Discussion – poorly-defined profiles

This section will discuss the types of profile illustrated in Figure 9.9d, which is rather more awkward to analyse than most of the profiles observed due to the poor-definition of the bubbles and slugs. Possible explanations for the occurrence of these type of profiles could include very fast flow, where the bubbles are actually overlapping within the acquisition time or perhaps slugs have not had time to coalesce and form Taylor flow (i.e. we have small droplets). It might be estimated that there are three significant bubbles in this particular profile (since there are three significant dips in intensity). However, the actual length of the bubbles is much more difficult to estimate. Gating at our chosen gating level of 0.55 will clearly not give an appropriate measure for the length of these bubbles. We will now further investigate the poorly-defined bubble / slug interface of the type shown in Figure 9.9d.

In order to probe whether these profiles are due to fast-moving channels or some other effect, it is necessary to first identify their occurrence and then to correlate this to the local channel velocity and the local holdup. In order to quantify the occurrence of these unusual profiles the following procedure is proposed to identify them; first we draw lines across the intensity plots shown in Figure 9.19 at values of 1/3 and 2/3 of the intensity profile height (as measured from the noise level to the nominal height of the profile), which corresponds to values of 0.4 and 0.7 (see Eq. 10.2 and Eq. 10.3).

$$noise + \frac{1}{3}(1 - noise) = 0.1 + \frac{1 - 0.1}{3} = 0.4 \qquad \textbf{(Eq. 10.2)}$$

$$noise + \frac{2}{3}(1 - noise) = 0.1 + \frac{2 \times (1 - 0.1)}{3} = 0.7 \qquad \textbf{(Eq. 10.3)}$$

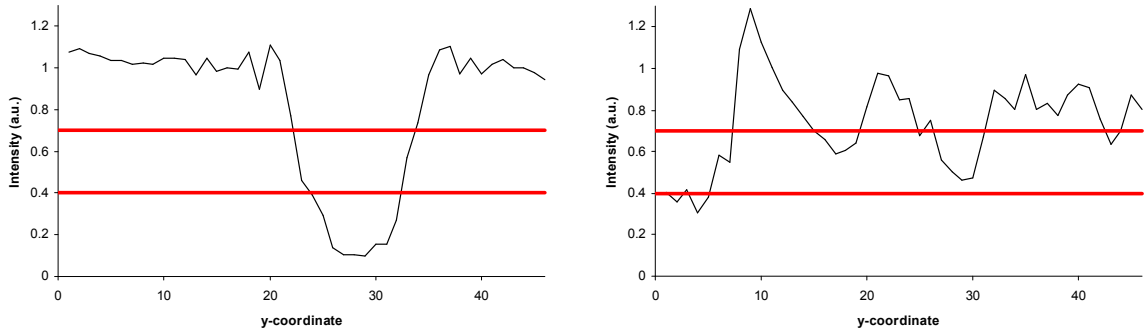


Figure 9.19 – Sample profiles overlaid with lines of value 0.4 and 0.7, which correspond to 1/3 and 2/3 of the height, as measured from the noise level to the nominal height of the profiles.

We can see that the profile of Figure 9.19a has the same number of intersections for the upper gating level as for the lower one, but that the profile of Figure 9.19b has many more intersections for the upper gating level than for the lower one. Hence we now have a robust criterion for identifying such profiles. We can now tally the number of such profiles as a function of total flow rate, and this is summarised in Table 9.1. It can be seen that the occurrence of these profiles increases with flow rate. The number of these profiles is relatively small, and so they are not expected to have a great contribution to the overall flow description. However, they are significant enough to warrant further attention.

Table 9.1 – Number of identified channels as a function of air flow rate

| Flow rate (cc/min) | Number of Identified Channels | Percentage of Total Channels |
|---|---|---|
| 100 | 199 | 7 |
| 150 | 291 | 10 |
| 200 | 344 | 12 |
| 300 | 373 | 13 |

A more useful analysis is to correlate the occurrence of these types of profile with the local channel velocity (as measured by the $Chi^2$ tracking technique of Section 7.6.2) and with the local holdup. If we examine Figure 9.20, we can see that the profiles which have poorly-defined slug interfaces typically contain lower liquid hold-ups, and Figure 9.21 shows that they have higher velocities than the other channels. Therefore, it seems logical to say that the poorly-defined profiles are a result of high gas holdup and high velocity, which suggests we are in a different flow regime than Taylor flow. The lower liquid holdup means that the slugs may be broken up into smaller droplets, and there will be increased motion-blurring due to the higher velocities, which is consistent with the data acquired.

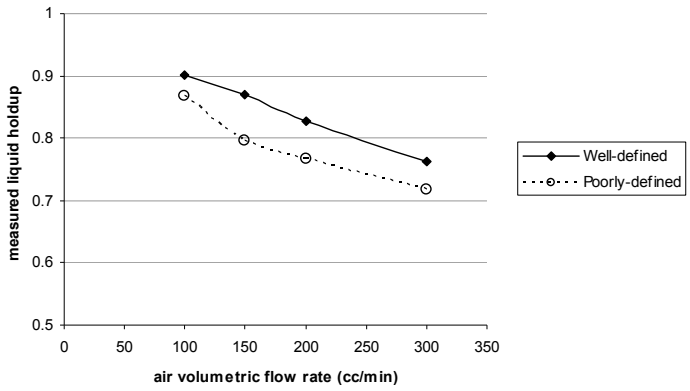


Figure 9.20 – A plot of average holdup with gas flow rate for the bubble sparger system, showing separate results for well-defined slug profiles and for poorly-defined slug profiles.

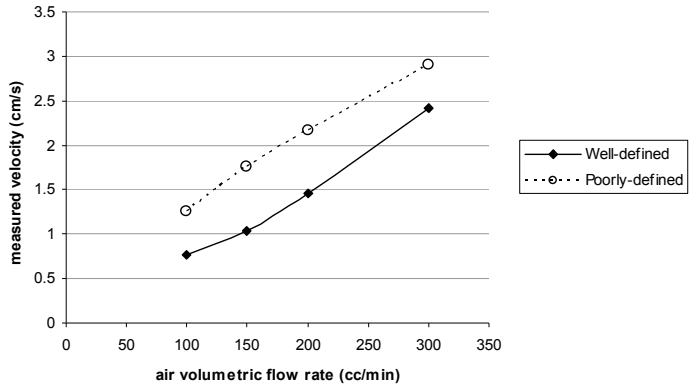


Figure 9.21 – A plot of average velocity with gas flow rate for the bubble sparger system, showing separate results for well-defined slug profiles and for poorly-defined slug profiles

Removing the unusual bubble profiles from our analysis has the effect shown in Figure 9.22 (for bubbles) and Figure 9.23 (for slugs). It is seen that the number of one-pixel bubbles detected is dramatically reduced. This implies that the one-pixel bubble is mainly a manifestation of the fact that the gating level intersects these profile several times, dissecting any bubble that may or may not really exist into smaller sections. Applying the same analysis to the slugs shows a similar story. The number of 1-3 pixel slugs is reduced, giving a more uniform distribution.

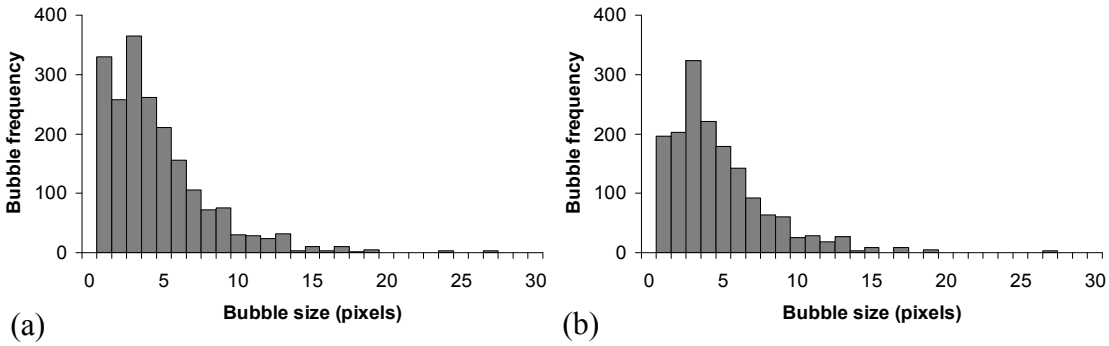


Figure 9.22 – The effect of removing poorly-defined profiles on the bubble size distribution at a flow rate of 100 cc/min; (a) shows the results for all channels, (b) shows the results neglecting the poorly-defined profiles.

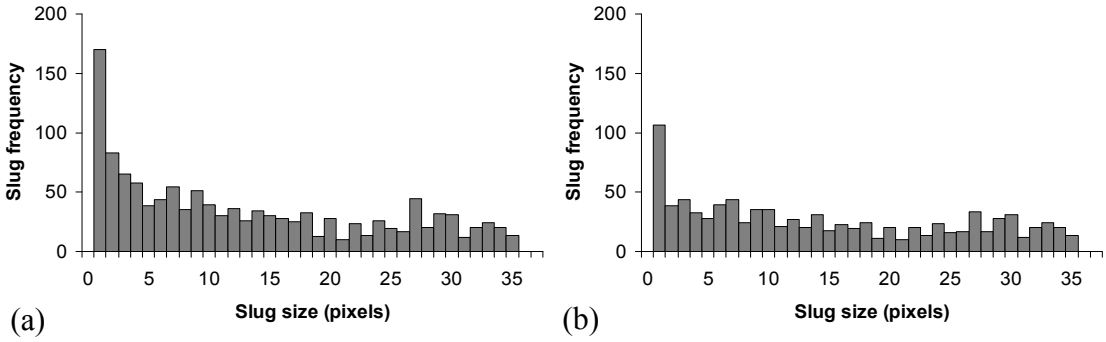


Figure 9.23 – The effect of removing poorly-defined profiles on the slug size distribution at a flow rate of 100 cc/min; (a) shows the results for all channels, (b) shows the results neglecting the poorly-defined profiles.

## 9.5. Results

In summary, we have presented an analysis method for measuring the bubble and slug lengths within the channels of a monolith using in-situ MRI data. We have investigated the sensitivity of the method, including our choice of gating level and derived a correction for the likelihood of detection as a function of increasing slug length. We have identified a proportion of the channels which are not operating in Taylor flow, which conduct on the order of 10% of the total flow rate. These profiles have been removed from the bubble / slug length analysis since we cannot obtain meaningful length measurements from them. The results are shown in Figure 9.24 and Figure 9.25 for 100 cc/min and 300 cc/min.

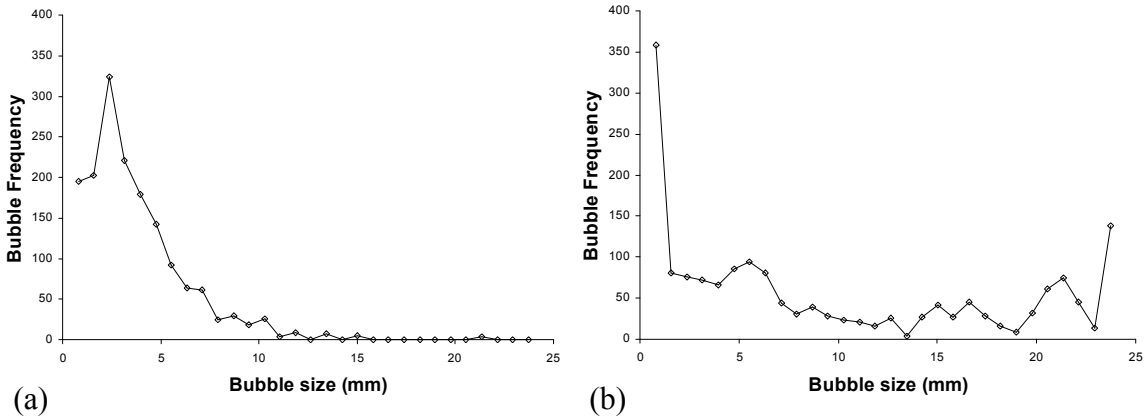


Figure 9.24 – Final bubble size distributions for air flow rates of (a) 100 cc/min, and (b) 300 cc/min.

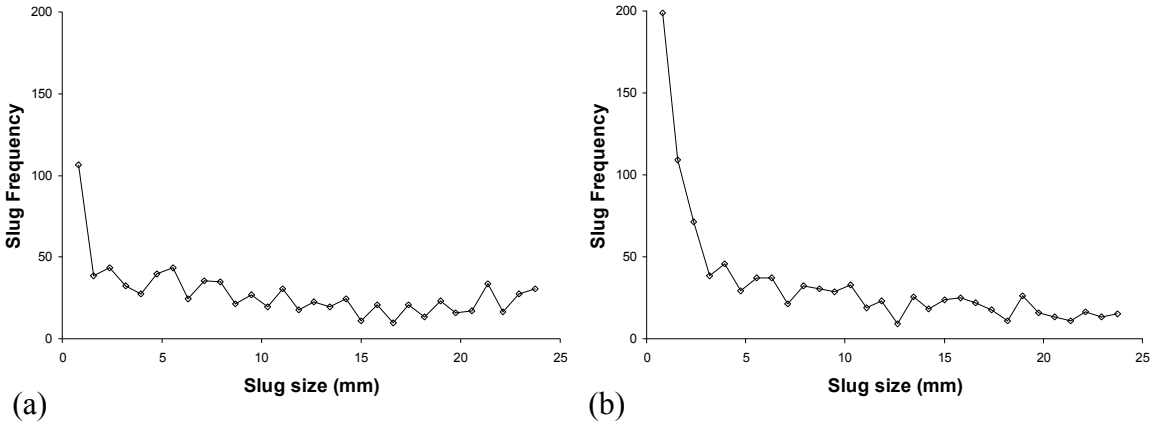


Figure 9.25 – Final slug size distributions for air flow rates of (a) 100 cc/min, and (b) 300 cc/min.

The general exponential decay seen does not seen to vary with flow rate, however at larger flow rates there are many more slugs present. It is seen from a comparison of the bubble distributions at different flow rates in Figure 9.24 that increasing the flow rate produces fewer small bubbles and more, larger bubbles. The total number of bubbles generally increases with increasing flow rate, though there is no straightforward correlation as shown by the data in Figure 9.26.

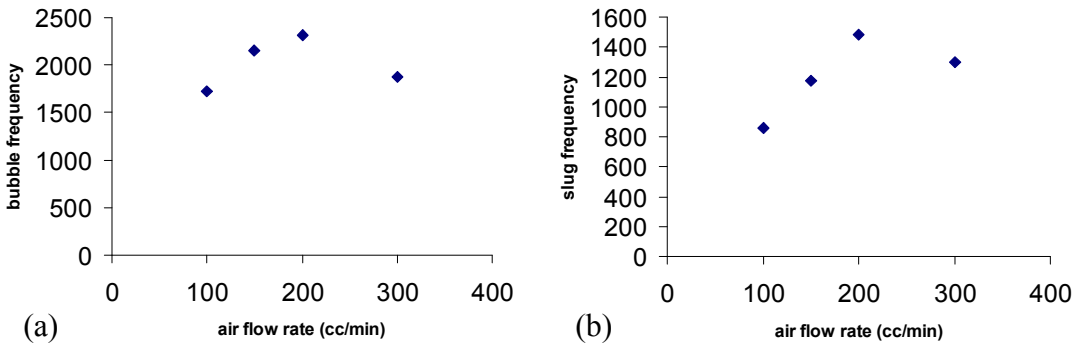


Figure 9.26 – (a) Plot of the number of bubbles detected as a function of air flow rate, and (b) a plot of the number of slugs detected as a function of flow rate.

## 9.6. Comparison with the literature

As mentioned in Section 3.7.1, there are a number of observations of slug and bubble lengths in the traditional literature, most often as a result of observing single capillaries optically under well-controlled conditions. This section will discuss the results of this work in comparison to these observations.

Previous work by Sederman *et al.* (2003) reported bubble size and velocity distributions for MRI experiments on the same system as used in this work. However, their results are based upon higher density monoliths, so the absolute values would not necessarily be expected to match. Their results of bubble lengths as a function of flow rate in a 400 cpsi monolith are shown in Figure 9.27. This compares favourably to the trend observed in this work for the 200 cpsi monolith.

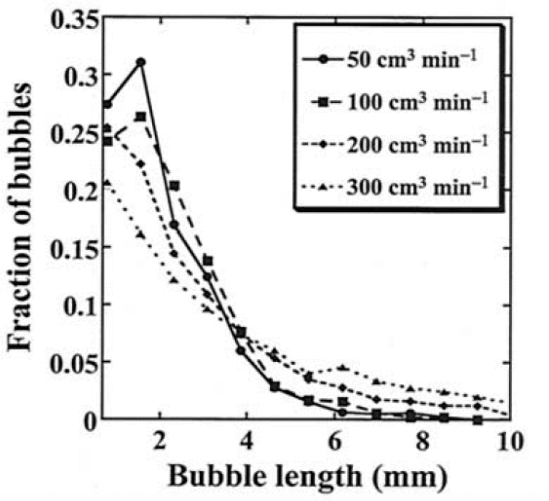


Figure 9.27 – Distributions of bubble length for a 400 cpsi monolith obtained from processing SEMI-RARE data. Reproduced from Sederman *et al.* (2003).

Heiszwolf *et al.* (2001) and Kreutzer *et al.* (2001) observed that, for various diameters of capillaries, the slug lengths were around 2-5 times the diameter of the channel, depending on the gas and liquid flow rates. For the 200 cpsi monolith used in these experiments this corresponds to slug lengths of 3.5-8.5 mm. Referring to the results of Section 9.5, we can see that the majority of the slugs we have observed are within this length range, however there are still many shorter slugs and many longer ones. There is not a clear cut-off point in our distribution and the frequency is mostly

uniform with increasing slug length. Therefore, the reported correlation for capillaries does not appear to hold for monoliths.

Some apparently robust correlations were reported by Thulasidas *et al.* (1995) who found a correlation between the ratio of bubble length to unit cell length as a function of the ratio of liquid flow to total liquid flow. These are shown in Figure 9.28. The system that the authors used was a very well-controlled system, and each experiment produced a line of equal-sized bubbles and equal-sized slugs. In the monolith system we may have bubbles of various lengths within the same channel, and so the question of how to define a unit cell becomes important. From the diagram in their paper, the authors seem to suggest that (in the direction of flow) the bubble leads the corresponding slug.

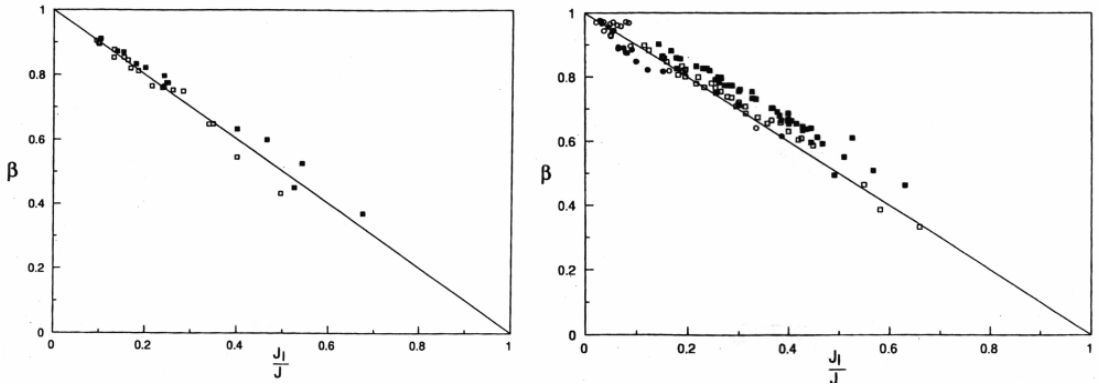


(a) (b)

Figure 9.28 – Ratio of bubble length to unit cell length as a function of ratio of liquid flow to total flow rate for (a) circular and (b) square capillaries. Data for (■) 100 centistoke silicone oil and (□) 50 centistoke silicone oil, both in bubble-train flow. Reproduced from Thulasidas *et al.* (1995).

Our MRI data was reprocessed by identifying pairs of bubbles and slugs (the direction of the channel flow was calculated first and the bubble / slug pairs were chosen such that the bubble was at the leading edge of the unit cell). The ratio of the liquid flow to the total flow was calculated by measuring the liquid holdup within the channel, since this is the equivalent measure of the authors' results. The results of this analysis are presented in Figure 9.29.

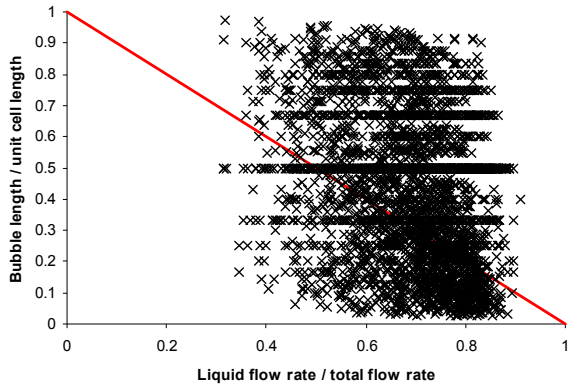


Figure 9.29 – Plot of the ratio of bubble length to the unit cell length, as a function of the ratio of the liquid flow rate to the total flow rate. The correlation reported by Thulasidas *et al.* (1995), is shown in red.

It can be seen that the correlation is very poor. It suggests that the bubble length is not in fact a function of the flow rate through the monolith channels. It appears that while the correlation may hold for a single channel capillary under controlled, driven flow it is not applicable to such a parallel channel system as the monolith.

The authors also reported a correlation between the volume fraction and the ratio of the liquid flow rate to the total flow rate (Figure 9.30).

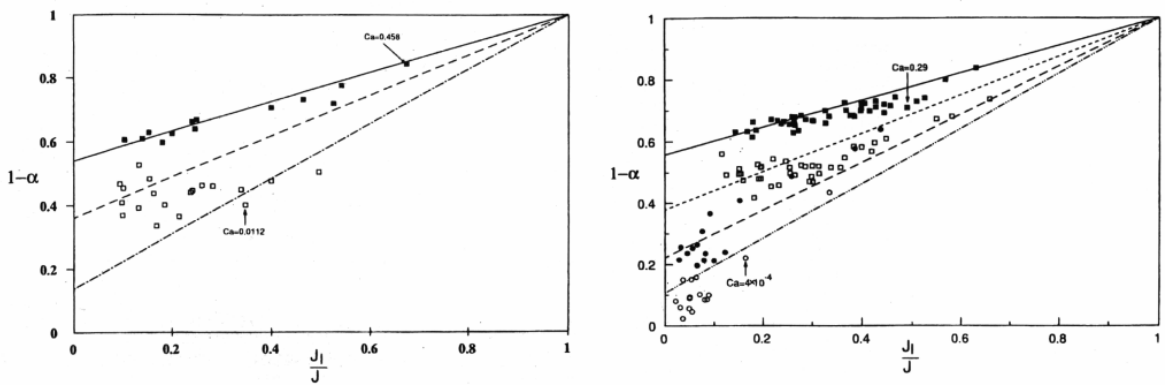


(a) (b)

Figure 9.30 – Volume fraction of liquid as a function of ratio of liquid flow rate to total flow rate for (a) circular and (b) square capillaries. Data for (■) 100 centistoke silicone oil and (□) 50 centistoke silicone oil, both in bubble-train flow. The lines are drawn at constant Ca (top Ca = 1.0; middle Ca = 0.10; bottom Ca = 0.01). Reproduced from Thulasidas *et al.* (1995).

However, in regard to a single channel these two quantities are very similar; for plug flow, one would expect to produce a plot of $y = x$. The deviation from this must come from the presence of downflow in the corners of the channel, and the film flow around bubbles. Since our MRI experiment must measure the flow through the channel by assuming plug flow, unfortunately our experimental method cannot observe any such effect.

Part of the problem in comparing single-capillary observations to parallel-channel systems such as the monolith is that the capillary systems are constrained by the driven flow rates, i.e. the overall gas and liquid flow rate to the capillary is controlled, and so the longer a slug is, the more pressure there is on the air to enter the channel. Similarly, the longer a bubble we have the more pressure there is for water to enter. The flow of air and water are not independent, as they are in the monolith system. We have no reason to predict a link between velocity and slug length for monolith channels since we do not have control of the superficial gas and liquid velocities in individual channels.

## 9.7. Liquid-filled channels

Another parameter of interest in the monolith system is the number of totally liquid-filled channels. These are not desirable since we require both gas and liquid to be present within a channel for reaction to occur. By totalling the number of such channels and so we can get a feel for the number of channels that are actually in Taylor flow, and the number that are liquid-filled at any one time.

A graph of the number of liquid-filled channels measured as a function of air flow rate is shown in Figure 9.31. This graph shows clearly that as the flow rate is increased the number of totally liquid-filled channels decreases, apparently levelling-off at high gas flow rates. This indicates that at higher flow rates more of the channels contain gas, and this is therefore a better system for performing three-phase reactions.

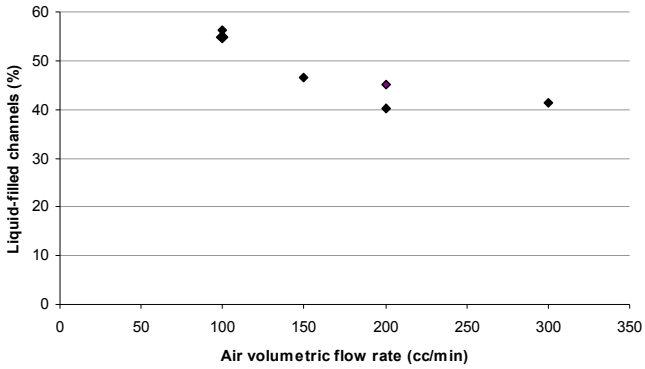


Figure 9.31 – A plot of the number of liquid-filled channels with air flow rate, for the bubble sparger system.

## 9.8. Conclusions

This chapter has demonstrated a data processing procedure to analyse bubble and slug distributions for the case of the bubble sparger configuration of the monolith reactor. MRI data are acquired with the RARE pulse sequence, which is rapid enough to allow measurements in dynamic two-phase systems. However, above velocities of around 5 cm/s, measurements become difficult and hence this method has not been applied to other monolith configurations.

A discussion of the choice of a suitable gating level was presented, as well as some consideration of the types of channel profiles observed, and measurement effects are discussed. Bubble and slug size / frequency results for a range of air flow rates between 100 and 300 cc/min were then presented. These results suggest that slug frequencies appear to be independent of slug length, whereas bubble lengths tend to increase with air flow rate. A comparison with correlations available in the literature was not possible due to the limitations of the detail captured in the MRI data. Also, the literature tends to focus on single-capillary observations, and some reasons for the possible invalidity of such a comparison are discussed.

This concludes our analysis of the bubble and slug hydrodynamics of the bubble sparger system. The following chapter will discuss the question of fluid distribution to the monolith reactor, and illustrate methods for quantifying the quality of the fluid distribution.

## 9.9. Nomenclature

| Symbols | Definition |
|---|---|
| $T_1$ | Spin-lattice relaxation constant |

| Abbreviations | Definition |
|---|---|
| Cpsi | Cells per square inch |
| MRI | Magnetic resonance imaging |
| RARE | Rapid acquisition with relaxation enhancement |
| RF | Radio frequency |
| SEMI-RARE | Single-excitation, multiple image RARE |

## 9.10. References


Heiszwolf, J. J., M. T. Kreutzer, *et al.* (2001). "Gas-liquid mass transfer of aqueous Taylor flow in monoliths." Catalysis Today **69**(1-4): 51-55.

Kreutzer, M. T., P. Du, *et al.* (2001). "Mass transfer characteristics of three-phase monolith reactors." Chemical Engineering Science **56**(21-22): 6015-6023.

Sederman, A. J., M. D. Mantle, *et al.* (2003). "Quantitative 'real-time' imaging of multi-phase flow in ceramic monoliths." Magnetic Resonance Imaging **21**(3-4): 359-361.

Thulasidas, T. C., M. A. Abraham, *et al.* (1995). "Bubble-Train Flow in Capillaries of Circular and Square Cross-Section." Chemical Engineering Science **50**(2): 183-199.

# Chapter 10 – Quantification of Flow Maldistribution

This chapter will demonstrate the application of a novel image processing technique described in Chapter 7 to the monolith reactor, with the aim of quantifying the quality of the fluid distribution within the monolith reactor under different conditions. MRI data acquired with the well-established RARE MRI pulse sequence will be processed, though the technique is applicable to many types of data which contain features or textural information. The quality of the distribution obtained with showerhead distributors and spray distributors is discussed.

**CHAPTER 10 – QUANTIFICATION OF FLOW MALDISTRIBUTION**



## 10.1. Introduction

The homogenous nature of the fluid distribution to the monolith is a critical factor in its operation. This section will cover the analysis of the quality of the distribution under different operating conditions. The method of Section 7.5.2.5 will be used, involving a custom autocorrelation function combined with a characteristic length analysis. This has been previously shown to have advantages over traditional standard deviation measurements and other techniques for distribution quantification.

We will begin by varying some of the geometry design parameters available for the showerhead distributor (Section 8.2.2) and discuss the effects on the distribution observed. The effect of gas and liquid flow rates will also be investigated. Data from

the spray distributor (Section 8.2.3) will then be analysed and compared to the performance of the showerhead distributor.

## 10.2. Showerhead geometry design parameters

Implementing the autocorrelation-based method of Section 7.5.2.5, we can estimate the characteristic feature length within a data image, and use this to classify the distribution observed. For example, a series of experiments was performed using different designs of showerheads, with varying gas and liquid flow rates, and with the showerhead at varying heights above a monolith of 200 cpsi. The results are shown in Figure 10.1 in order of the characteristic length, and the associated process configurations are summarised in Table 10.1. It can be seen that the homogeneity of the fluid distribution varies significantly from one experiment to another, and the analysis was used to order the results from most homogeneous (largest characteristic length) to least homogeneous (shortest characteristic length).

Images were acquired with the RARE pulse sequence (Section 2.3.1) with a field-of-view of 51 mm × 51 mm giving a resolution of 797 μm × 797 μm. The slice thickness was set at 40 mm. Fifty repetitions were taken, each around 5 seconds apart to allow for $T_1$ recovery. The acquisition time for each image was 150 ms. The data were acquired with a thick slice to ensure that we captured signal from fast-moving fluid, and the fifty repetitions were taken to average-out any dynamic fluctuations in the system and give a fair measure of the quality of the fluid distribution and velocity distribution present. It is important to note that the contrast in the images is not only a measure of the homogeneity of the fluid distribution, but is also velocity weighted and so will have some measure of the velocity homogeneity in the system. The data contain combined information on the fluid and velocity distribution, both of which should be evenly distributed in an ideal reactor system. Therefore, the data are still useful qualitatively, since we desire the minimisation of maldistribution whatever its source may be.

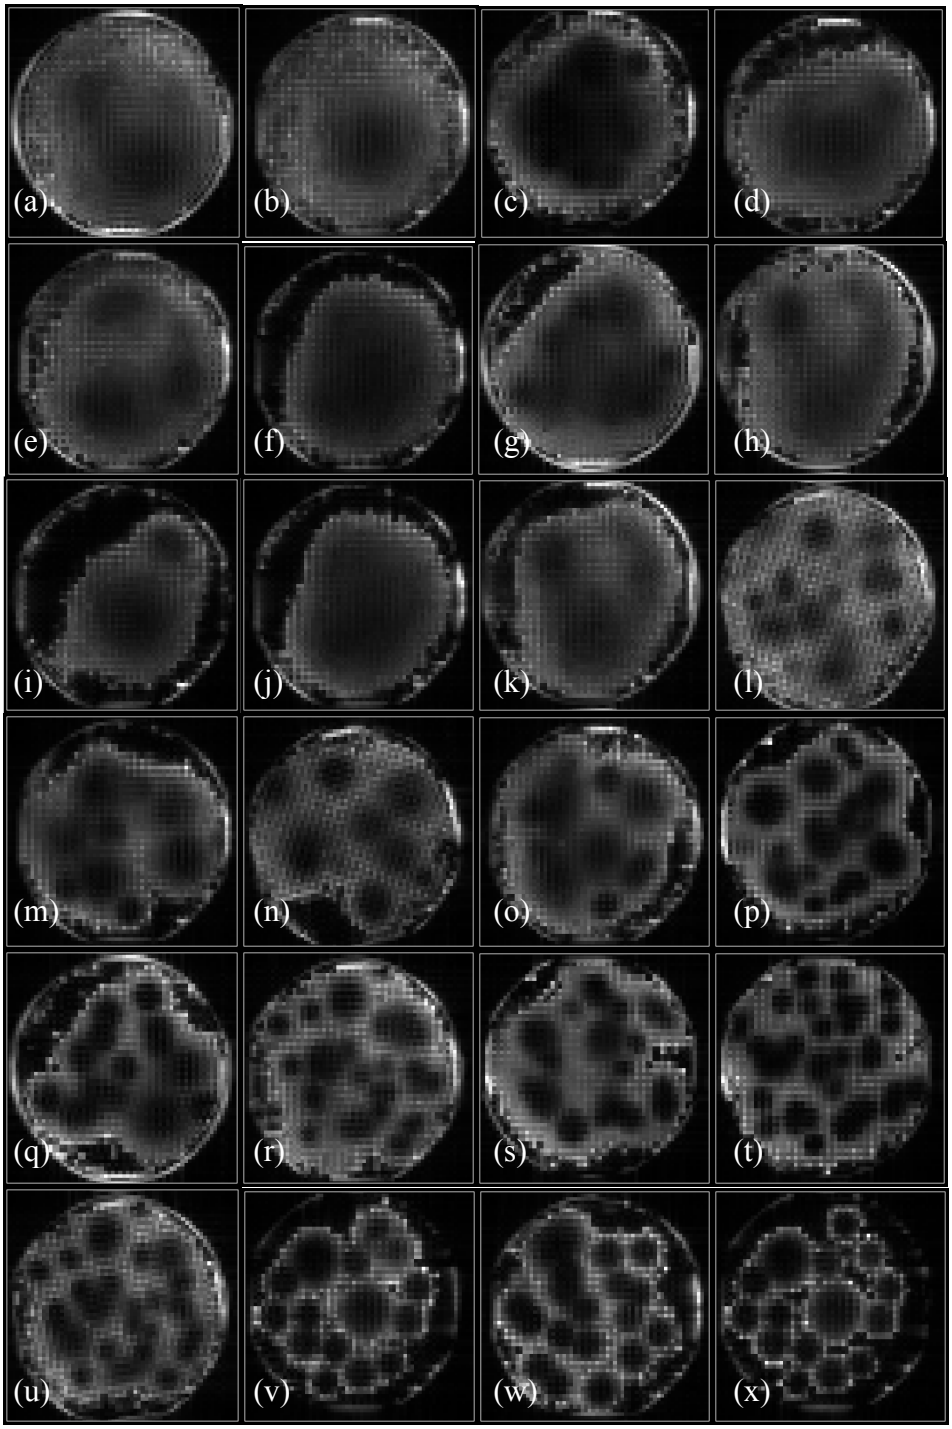


Figure 10.1 – The results of applying the classification analysis to data from a range of experimental conditions using showerhead distributors. The data are shown in order of most homogeneous to least homogeneous (a-x).

Table 10.1 – Experimental parameters for the data of Figure 10.1 (note: 1 pixel = 797μm).

| Figure number | Characteristic Radius (pixels) | Number of holes | Diameter of holes (mm) | Distance above monolith (cm) | Liquid flow rate (l/hr) | Gas flow rate (cc/min) |
|---|---|---|---|---|---|---|
| (a) | 10.97 | 19 | 1 | 21 | 80 | 200 |
| (b) | 9.38 | 19 | 1 | 14 | 80 | 200 |
| (c) | 7.85 | 31 | 2 | 14 | 250 | 600 |
| (d) | 7.09 | 19 | 1 | 14 | 80 | 200 |
| (e) | 6.91 | 19 | 1 | 14 | 80 | 410 |
| (f) | 6.48 | 19 | 1 | 14 | 80 | 200 |
| (g) | 6.35 | 19 | 1 | 21 | 100 | 420 |
| (h) | 6.17 | 19 | 1 | 7 | 80 | 600 |
| (i) | 5.43 | 19 | 1 | 14 | 80 | 420 |
| (j) | 5.19 | 19 | 1 | 14 | 80 | 600 |
| (k) | 4.98 | 19 | 1 | 7 | 80 | 420 |
| (l) | 4.70 | 31 | 1 | 14 | 70 | 420 |
| (m) | 4.69 | 19 | 1 | 14 | 95 | 410 |
| (n) | 4.66 | 31 | 1 | 14 | 70 | 200 |
| (o) | 4.32 | 19 | 1 | 14 | 100 | 200 |
| (p) | 4.06 | 31 | 1 | 14 | 100 | 420 |
| (q) | 3.99 | 19 | 1 | 21 | 120 | 200 |
| (r) | 3.94 | 31 | 1 | 14 | 100 | 200 |
| (s) | 3.88 | 31 | 1 | 14 | 100 | 420 |
| (t) | 3.81 | 31 | 1 | 14 | 150 | 420 |
| (u) | 3.80 | 31 | 1 | 14 | 100 | 420 |
| (v) | 3.63 | 19 | 1 | 7 | 120 | 200 |
| (w) | 3.60 | 19 | 1 | 14 | 115 | 410 |
| (x) | 3.08 | 19 | 1 | 7 | 120 | 240 |

It is clear that we now have a robust technique that is capable of classifying the homogeneity of intensity within an image. The order of the quality of the distributions is much as if they had been manually arranged, but the factor of human error has been removed. This analysis technique allows us to investigate the effect of changing design parameters on the fluid distribution, and thereby optimise the performance of the monolith reactor by selecting appropriate geometry and flow rates.

The individual effects of the various design parameters will now be investigated. Figure 10.2 shows the variation in characteristic radii as a function of water and air flow rates for the distributor with 19 holes of 1 mm at a distance of 14 cm above the monolith. We can see from Figure 10.2a that the characteristic radius decreases (and therefore the homogeneity decreases) with increasing liquid flow rate. The reason for this trend is believed to be due to the nature of the individual jets; at low flow rates the jets are observed to oscillate and there is droplet breakup well above the monolith

surface; in contrast at higher flow rates, the jets are faster and more sharply defined and hence there is less break-up and spreading of the drops. This observation suggests that high velocities will concentrate the flow to narrow points at the surface of the monolith, which is what is observed in the MRI data. From Figure 10.2b we can see the effect of increasing the gas flow rate is to also decrease the characteristic radii, and therefore decrease the homogeneity. The effect is less pronounced than the effect of changing the water flow rate. The mechanism relating gas flow rate and changes in the fluid distribution is uncertain.

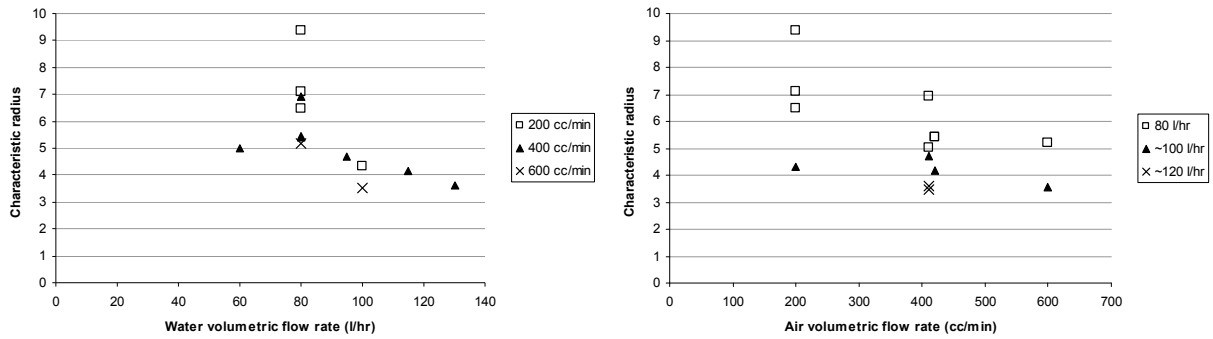


(a) (b)

Figure 10.2 – Plots of characteristic radii resolved as a function of (a) water flow rate and (b) gas flow rate, for the distributor with 19 holes of 1 mm at a distance of 14 cm above the monolith. The same data are plotted in (a) and (b), but resolved in different ways.

Figure 10.3 illustrates the effect of changing the distributor position above the monolith surface. It can be seen that increasing the distributor height increases the homogeneity of the image. This is explained by the increased droplet break-up, which helps distribute the fluid more evenly over the monolith cross-section. From examining the data as a function of liquid flow rate we can clearly see that the optimum configuration is to have a low flow rate and a distributor positioned high above the monolith.

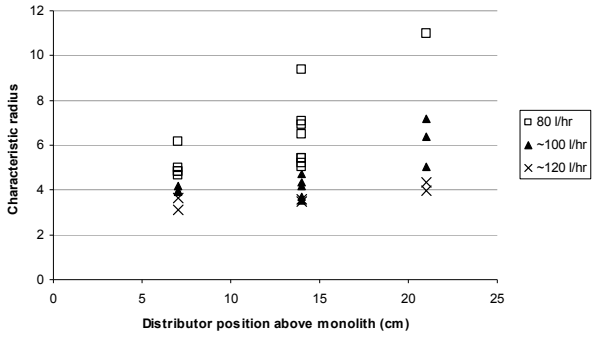


Figure 10.3 – Plots of characteristic radii as a function of distributor position, for the distributor with 19 holes of 1 mm.

## 10.3. Comparison of showerhead and spray distributors

This distribution quantification technique was then used to characterise the spray distributor setup (Section 8.2.3), which was expected to have a far more even fluid distribution than the showerhead setup, due to its higher control over the fluid distribution and a narrow range of velocities over the target area. A typical x-y cross-section for spray operation is shown in Figure 10.4, highlighting the marked improvement over the showerhead data of Figure 10.1. Data are acquired with the same parameters as for the showerhead MRI data in the previous section, and on the same monolith (200 cpsi).

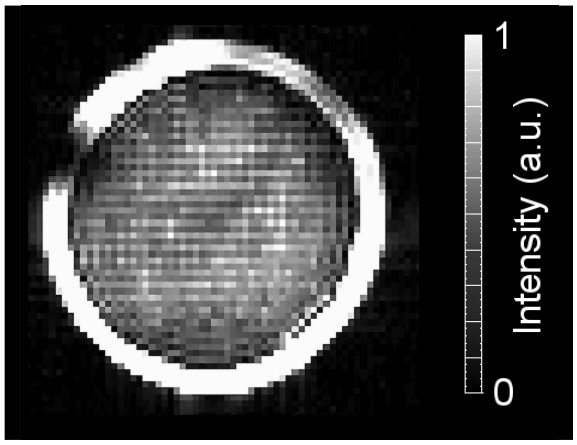


Figure 10.4 – A sample image of flow in a spray distributor system, illustrating the more homogeneous nature of the fluid distribution compared to those of Figure 10.1.

Figure 10.5 compares the results of applying the technique to the showerhead (Figure 10.5a) and spray distributor system (Figure 10.5b). Characteristic radii are shown as a function of liquid flow rate, with the air flow rate held constant at 400 cc/min. It can be seen that the characteristic radii are in general quite large for the spray distributors, indicating good homogeneity, with few measurements dropping below a radius of 6 pixels. In comparison most of the showerhead measurements fall below this value.

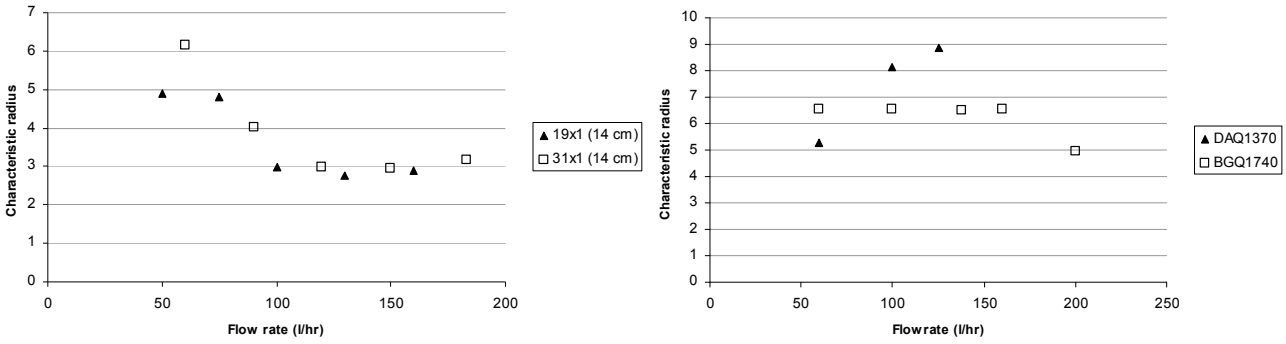


(a) (b)

Figure 10.5 – A comparison of the homogeneity analysis as a function of liquid flow rate for (a) the showerhead distributors, and (b) the spray distributors. The air flow rate was held constant at 400 cc/min in both cases.

Figure 10.5b contains data for the two types of spray nozzle[1]. The DAQ1370 nozzle is seen to produce a more homogeneous distribution, and this is perhaps due to the smaller droplets it produces. However, there are not many data points on which to base this conclusion. From the BGQ1740 data, it appears that increasing the liquid flow rate does not adversely affect the homogeneity of the distribution until quite high flow rates, which is in contrast with the showerhead data where increasing the flow rate slightly caused significant maldistribution.

[1] See notes in Section 7.2.3. The DAQ1370 nozzle has a smaller orifice, and so for the same flow rate should produce smaller droplets, with a higher initial velocity than the BGQ1740 nozzle.

## 10.4. Conclusions

In summary, we have applied the novel method for homogeneity quantification to real reactor setups successfully, and have used it to quantify the performance of different distributor setups. The spray nozzles have been shown to be superior to the traditional choice of showerheads at nearly all flow rates and perform significantly better at high flow rates, which are industrially relevant. The only disadvantage with the spray nozzles is that they produce a higher pressure drop and therefore incur an additional running cost. Depending on the application and the reaction kinetics (i.e. recycle ratio of the reactor), it may or may not be economically justifiable to use showerheads instead.

## 10.5. Nomenclature

| **Symbols** | **Definition** |
|---|---|
| $T_1$ | Spin-lattice relaxation constant |

| **Abbreviations** | **Definition** |
|---|---|
| cpsi | Cells per square inch |
| MRI | Magnetic resonance imaging |
| RARE | Rapid acquisition with relaxation enhancement |

# Chapter 11 – Two-phase Velocity Measurement in the Monolith Reactor

This chapter will demonstrate the application of the novel velocity-imaging MRI techniques described in Chapter 5 to the monolith reactor, with the aim of extracting operational information on the flow behaviour under different process conditions. The applicability of each technique will be considered in turn, and the limits of the techniques will be investigated. Velocity distributions will be presented and correlations with holdup and flow direction will be discussed.

**CHAPTER 11 – TWO-PHASE VELOCITY MEASUREMENT IN THE MONOLITH REACTOR**

## 11.1.Introduction

This chapter will implement the novel MRI techniques described in Chapter 5 to acquire data from the monolith reactor under different process configurations, ranging from the simple bubble sparger to the far more challenging turbulent froth reactor. As well as a homogeneous fluid distribution, it is desirable to operate with a narrow range of velocities in the monolith channels in order to control the residence time distribution of the reactor. The following sections will cover two time-of-flight techniques (SEMI-RARE tracking and line-excitation tagging), followed by two rapid, velocity-encoded MRI sequences (COTTAGE and MS-COTTAGE). Each method will be outlined in turn, and then applied to a range of monolith operating conditions, indicating the applicability of the methods for each of the configurations and describing their accuracy and limitations under extreme flow conditions. This work will measure individual channel velocities and obtain information from these challenging systems even at high flow rates and in the presence of dynamic two-phase flow.

## 11.2.SEMI-RARE

By using the SEMI-RARE technique (Section 5.2) we obtain a series of successive images from a single excitation, and by comparing the position of features, such as bubble interfaces, from one image to the next we are able to directly measure the velocity within that channel. The robust $Chi^2$ tracking method (Section 7.6.3) was implemented to measure the apparent velocity, and these results are compiled into velocity / frequency plots. The following experimental setups were investigated – the bubble sparger (described in Section 8.2.1), co-current upflow (Section 8.2.1), the spray distributor (Section 8.2.3) and the froth reactor (Section 8.2.4). These setups have very different flow rates and fluid distribution methods, and are expected to yield contrasting results. The following sections will discuss the application of the SEMI-RARE technique to each of the setups in turn.

## 11.2.1. Bubble sparger

The bubble sparger configuration (Section 8.2.1) was set up with a 200 cpsi monolith, and a series of SEMI-RARE data at varying flow rates were acquired. The data were acquired at 128×64 pixels, corresponding to a resolution of 398 μm × 797 μm. The field of view was 51 mm × 51 mm, and the slice thickness was 1.7 mm, which corresponds to the depth of a single channel. Sets of four images were acquired from each excitation, with each being acquired over 156 ms. Hence the echo times associated with each image are taken as 78, 234, 390 and 546 ms, which are the mid-points of each acquisition. Ninety repetitions were taken at each flow rate to ensure we have a statistically significant number of measurements to process.

The Chi$^2$ tracking method of Section 7.6.3 was applied to each channel in turn, tracking the movement of features (such as bubbles). The results at a range of air flow rates are shown in Figure 11.1. A tri-modal distribution is observed; we have upflow due to the upward movement of air through the monolith channels; we have downflow due to the downward recirculation of fluid in water-containing channels; and we have stagnant channels which do not move. It is seen that at higher air flow rates, the upward and downward nodes stretch towards larger velocities, as might be expected.

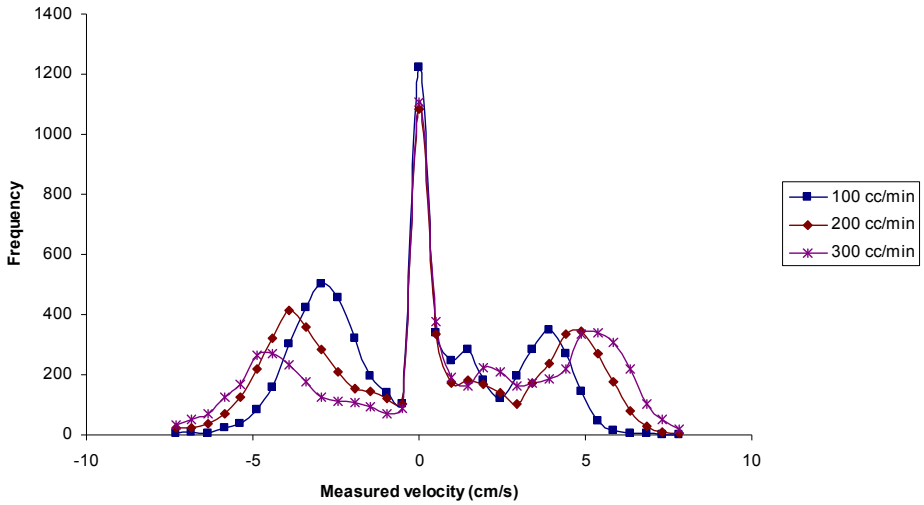


Figure 11.1 – Comparison of velocity/frequency measurements of the bubble sparger setup at a range of gas flow rates using the 200 cpsi monolith. Positive velocity corresponds to upflow.

At first inspection, it might be expected that the Chi$^2$ tracking method would fail for the case of totally liquid-filled channels, since there are no features to track. However, the method still works because of the inflow of fresh spins into the acquisition slice. These are observed as a region of low intensity entering the channels, and hence we can use the rate of inflow of these fresh spins to measure the velocity of that channel.

The data of Figure 11.2 shows some similar data from the work of Sederman *et al.* (2003) which used a 300 cpsi monolith. The data were re-processed with the Chi$^2$ tracking method. The results show the same tri-modal results, particularly at higher flow rates. A larger spike is observed at zero, which is probably due to the higher Ca number (we have the same surface tension, but a smaller diameter of channel here), and so the flow is more liable to become stagnant within individual channels.

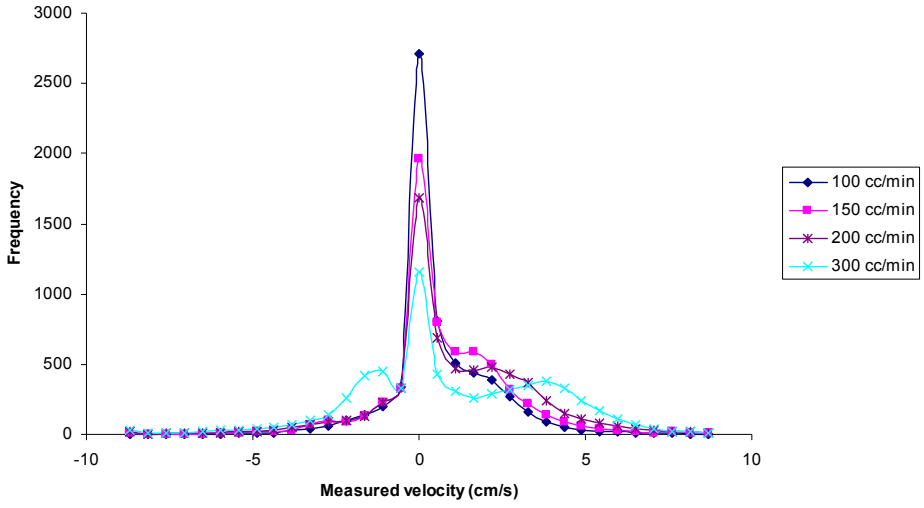


Figure 11.2 – The data acquired by Sederman *et al.* (2003) on a 300 cpsi monolith are processed using the same tracking method as that of Figure 11.1.

The results reported by Sederman *et al.* (2003) are shown in Figure 11.3. This shows the same data but processed in a different manner - the authors used a time-consuming manual method for measuring the bubble lengths, and hence fewer bubbles were counted. However, their results for the 300 cpsi monolith compare favourably with the results obtained using the new processing method (Figure 11.2).

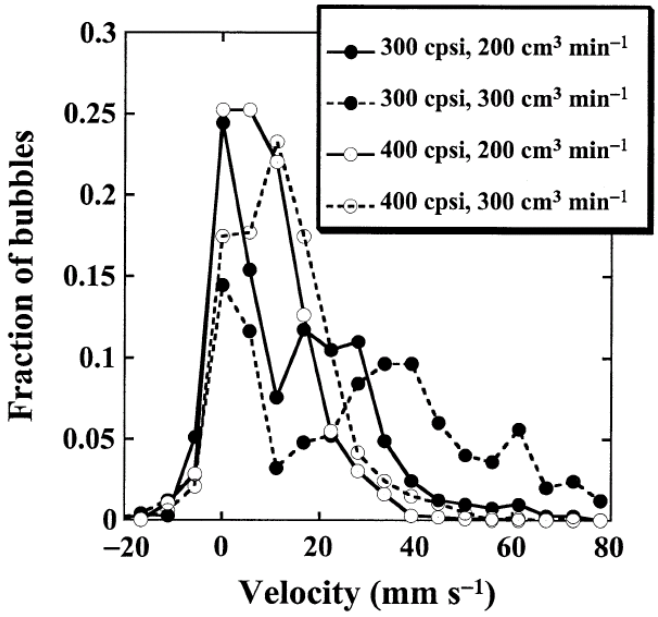


Figure 11.3 – Distributions of bubble velocity for 300 and 400 cpsi monoliths at a range of air flow rates. Reproduced from Sederman *et al.* (2003).

## 11.2.2. Velocity / holdup correlation

We can further process the 200 cpsi bubble sparger data by correlating the local channel holdup (as measured by the proportion of high intensity to low intensity within a channel) against the measured local velocity. It is perhaps to be expected that channels containing more air (lower liquid holdup) will be more inclined to move upwards, and also at a higher velocity. The plot of Figure 11.4 shows the data (from all flow rates of 100-300 cc/min) for each channel. These data were combined into one plot since the flow behaviour of the individual channels should be relatively independent of the overall flow rate. Since we are inspecting the channels one-by-one, we can treat them as being independent of the system which they are in. The important factor is the flow rate of the *individual* channel and how this is related to its *local* holdup.

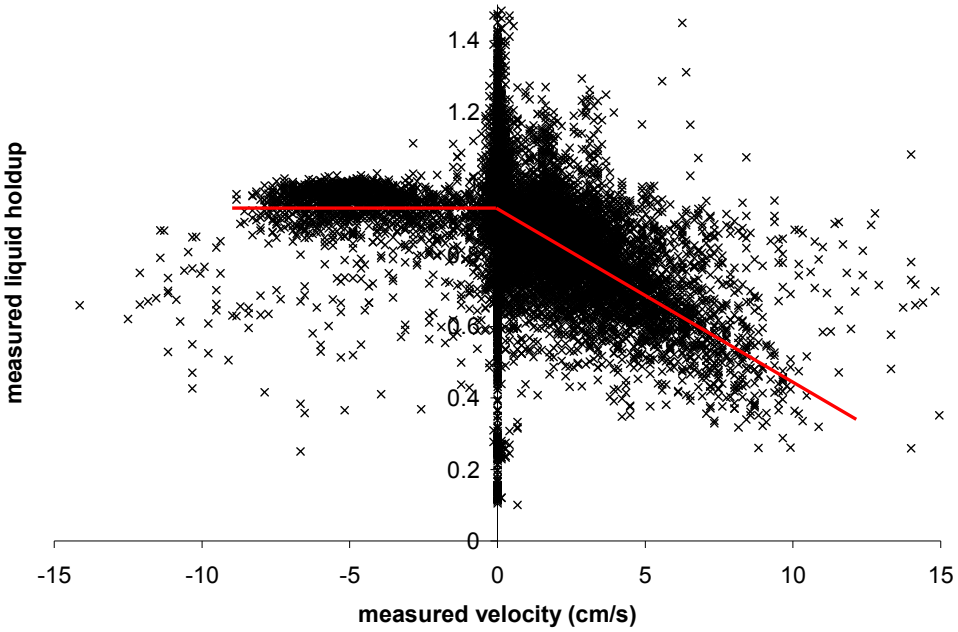


Figure 11.4 – A scatter plot of the data from the 200 cpsi bubble sparger system. Each channel in treated in turn, and plotted according to its local liquid holdup and measured velocity.

Although there is a good degree of scatter in the plot of Figure 11.4, we may still observe trends within the data. There appear to be two separate flow regimes, as indicated by the red lines overlaid on the data. To further investigate the results, we can separate out the data corresponding to totally liquid-filled channels from that of channels which contain some air. These two plots are shown in Figure 11.5.

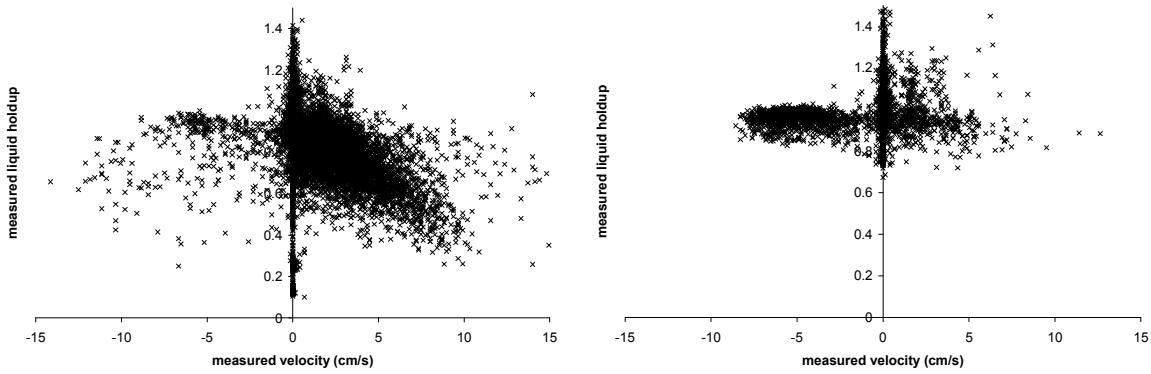


(a) (b)

Figure 11.5 – The data of Figure 11.4 resolved by hydrodynamic regime; (a) shows data for channels containing one or more bubbles, and (b) shows data for channels containing no bubbles (i.e. liquid filled).

It can be seen from Figure 11.5 that very few of the channels which contain air travel downwards (negative velocity) and it is the liquid-filled channels which contribute most to the negative velocities observed. In channels which are moving upwards, there is a clear correlation between having a lower liquid holdup and having a higher upwards velocity, which is as expected.

We can extract the slope of the line in Figure 11.4, and build an empirical relation between local holdup and velocity for this specific monolith system, as shown in Eq. 11.4.

$$\text{Holdup} = 1 - 0.0567 \times \text{Velocity}$$
$$\Rightarrow \text{Velocity} = \frac{1 - \text{Holdup}}{0.0567} \qquad \textbf{(Eq. 11.4)}$$

## 11.2.3. Channel flow rate / direction correlation

We may also investigate the ability of the channels to change hydrodynamic operation – for example, once a channel is in upflow, is it capable of changing direction? This was investigated by resolving the velocity data such that we can observe the velocity in each channel as a function of time. Figure 11.6 shows the results of this analysis, with time in the vertical direction. It can be clearly seen that the velocity in each channel is relatively constant, and that there are no channels which change their flow characteristics. The same flow orientation is even observed between experiments, which indicates a long-term correlation in the operation of the channels. These experiments had been performed in succession, without stopping and re-starting the system. If the system had been purged and re-started between experiments then the correlations would have been strong within individual experiments only, not necessarily between experiments as well.

This stability observation reinforces the comments of Section 3.8, which stated that once a channel is in upflow, there is no force acting to counteract this. The presence of liquid-filled or stagnant channels will lead to under-utilisation of the catalyst since gas and liquid are not contacting in these places. Hence, we desire more fluctuations

in the behaviour of the channels than is actually seen experimentally. It will be shown later that co-current flow, and other experimental set-ups, are superior in this aspect.

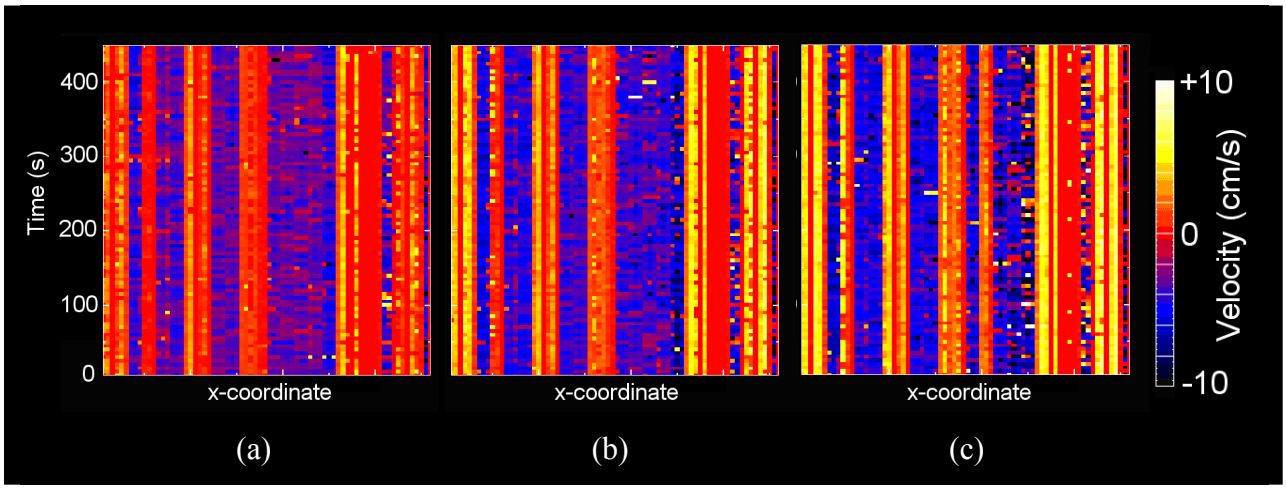


Figure 11.6 – The velocity data of the bubble sparger system (Figure 11.1), resolved by individual channels, and shown as a function of time (on the y-axis). Air flow rates are (a) 100 cc/min, (b) 200 cc/min, and (c) 300 cc/min.

## 11.2.4. Experimental repeatability

The repeatability of the bubble sparger system was investigated, and this will now be addressed. We have seen that the bubble sparger is sensitive to the initial hydrodynamics, and will remain in this mode of operation throughout an experiment. For example, the data of Figure 11.7 shows the velocity / frequency measurements for a flow rate of 500 cc/min of air. Line (A) was performed directly after several previous experiments, whereas line (B) was flushed through with water before acquiring data. As can be seen, the velocity distribution varies significantly between the two experiments. From observation of the profiles from the two data sets it is seen data (A) contain few moving channels all of which have high gas holdup, whereas data (B) contain more moving channels with slugs clearly defined. This observation is perhaps due to the result of only observing one row of channels in the monolith system, and is not necessarily representative of the whole system.

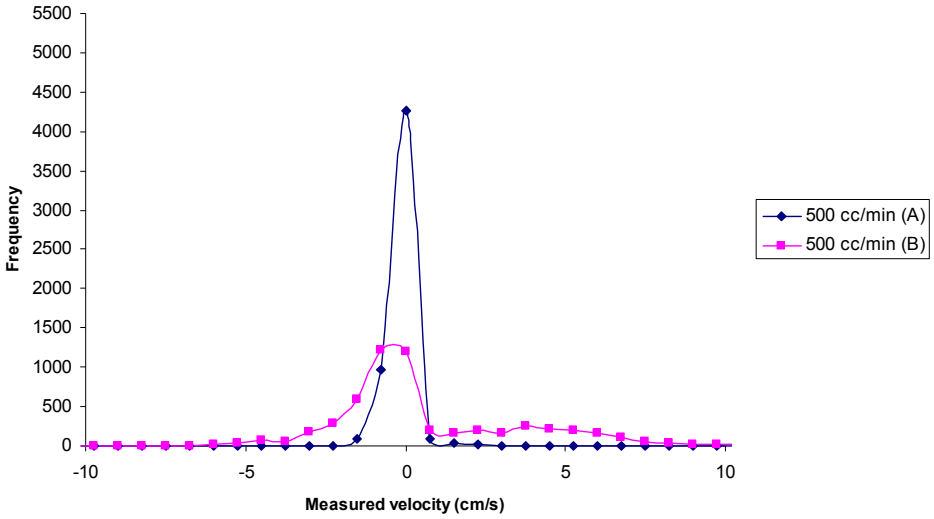


Figure 11.7 – Repeatability study of bubble sparger data, showing results for the same experimental set up and flow rates, but under different startup conditions.

It is proposed that co-current flow would be a less static system, i.e. that the additional flow pressure would enable channels to switch from stationary, to flowing, and back again, which is not apparently a common occurrence with the current bubble sparger configuration. Co-current flow will be addressed in the next section.

## 11.2.5. Co-current flow

The experimental setup is identical to that of the bubble sparger, but we now operate co-currently with water flowed upwards with the air from the bubble sparger (Section 8.2.1). Experiments were performed at water flow rates of 29, 78 and 108 l/hr, and air flow rates of 100, 200, 300, 500 and 1000 cc/min to investigate the performance of the monolith reactor in each case. The results of the SEMI-RARE tracking analysis are plotted, resolved by air flow rate (Figure 11.8), and by water flow rate (Figure 11.9) respectively.

From these data, it is apparent that increasing either the gas or liquid flow rates results in a move from a plot which has a spike at zero and a decay in frequency towards higher velocities, to a distribution which has two clear nodes; one stagnant and one at higher velocity. It is proposed that the best distribution is that given by the highest gas and liquid velocities (108 l/hr and 1000 cc/min respectively), and this is in

keeping with the observations of Section 3.8 which state that higher flow rates are in general more stable. The uniformity is better at these high flow rates, aside from some non-moving channels.

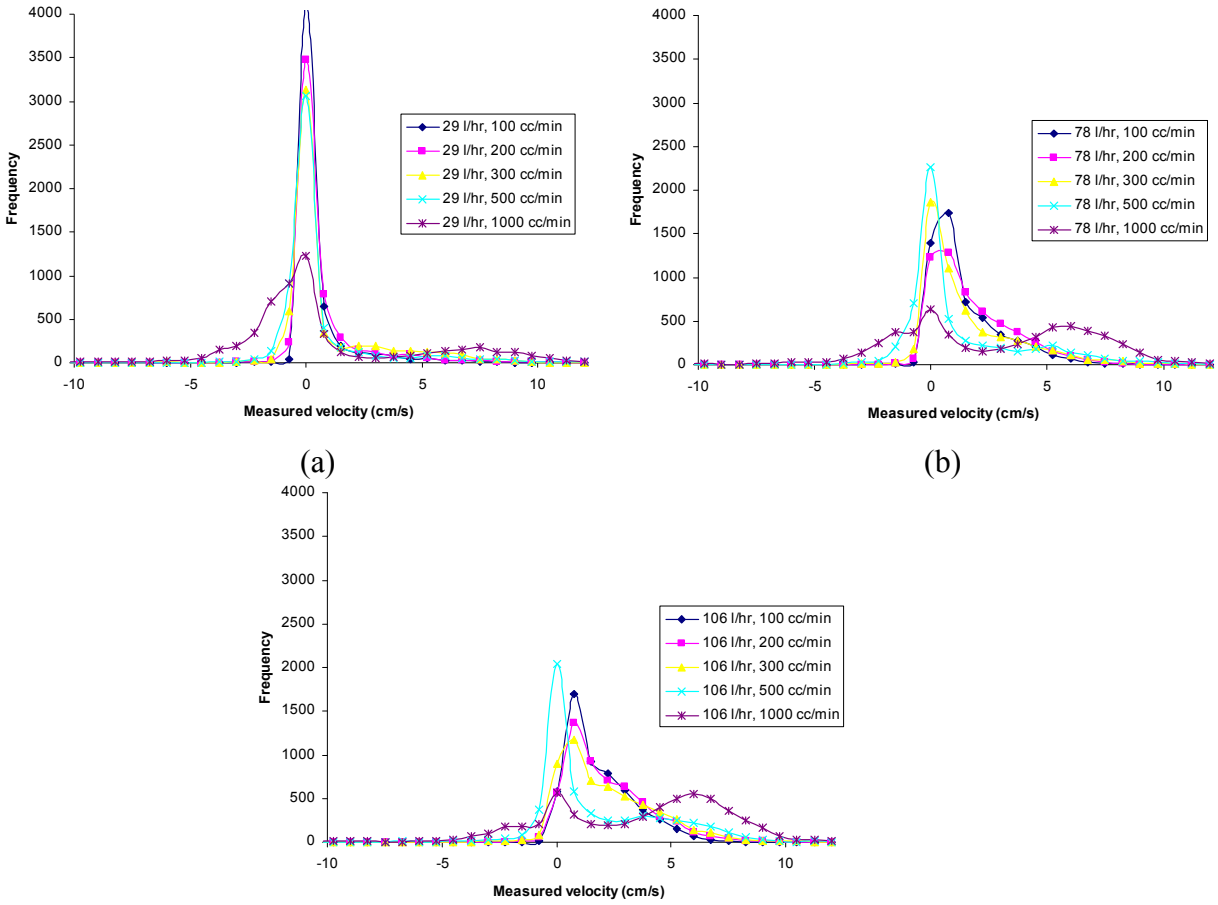


Figure 11.8 – Comparison of the results of the velocity measurement of slugs under co-current upflow conditions. The data are resolved by air flow rate for liquid flow rates of (a) 29 l/hr, (b) 78 l/hr, and (c) 108 l/hr. Positive velocity corresponds to upflow.

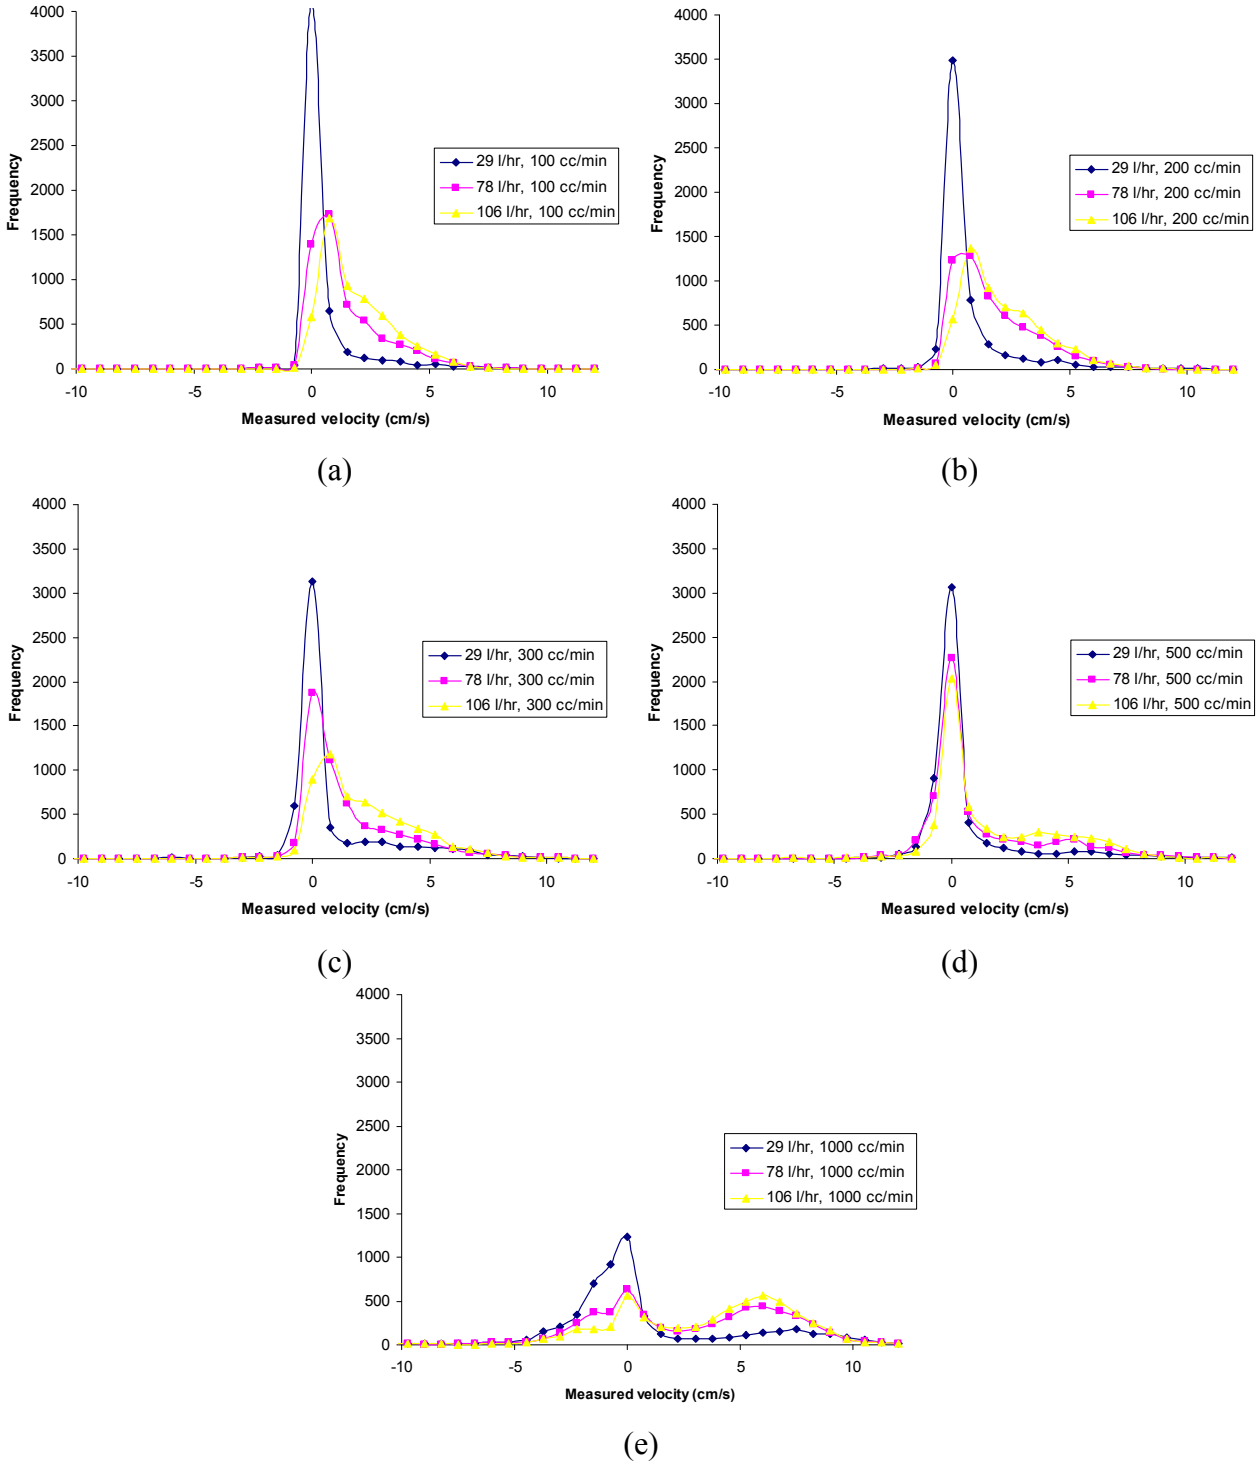


Figure 11.9 – Comparison of the results of the velocity measurement of slugs under co-current upflow conditions. The data are resolved by water flow rate for air flow rates of (a) 100 cc/min, (b) 200 cc/min, (c) 300 cc/min, (d) 500 cc/min, and (e) 1000 cc/min. Positive velocity corresponds to upflow.

The flow rate / direction correlation method of Section 11.2.3 was calculated for the co-current flow configuration, and the results are shown in Figure 11.10. It can be seen that the co-current flow configuration is better able to change the hydrodynamics within individual channels, and fewer channels are stuck in the same mode of operation over the length of the experiment.

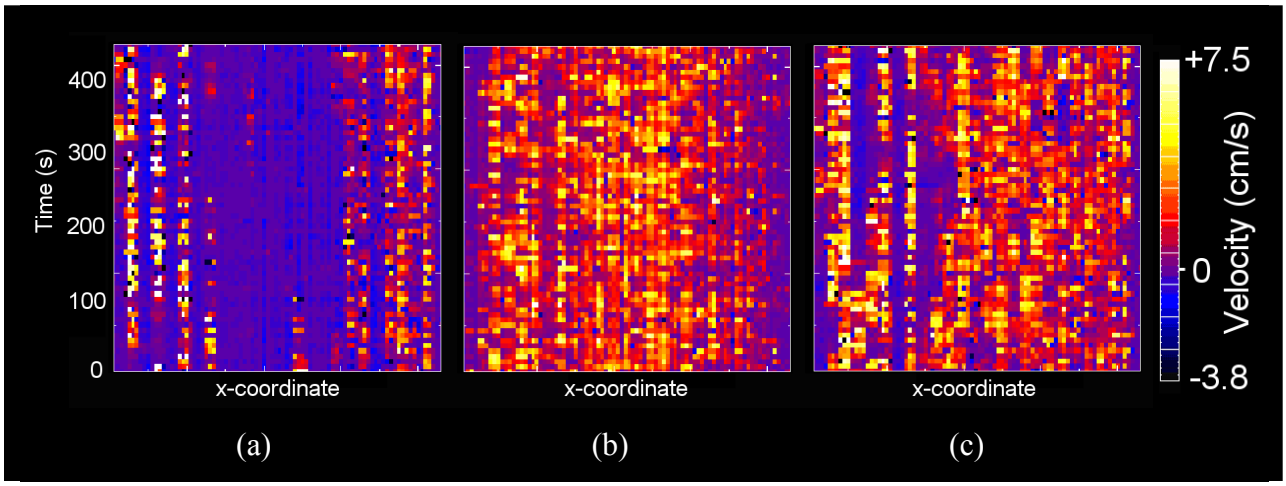


Figure 11.10 – The velocity data of the co-current flow configuration (Figure 11.9b), resolved by individual channels, and shown as a function of time. The air flow rate is constant at 200 cc/min and the liquid flow rates are (a) 29 l/hr, (b) 78 l/hr, and (c) 108 l/hr. Positive velocity corresponds to upflow.

## 11.2.6. Spray distributor

The Chi$^2$ tracking technique was then applied to the spray distributor, which has a lower liquid holdup and higher velocities. It is expected to have a more homogeneous liquid distribution, and therefore a narrower velocity distribution.

MRI data were acquired, as for the bubble sparger system of Section 11.2.1. A sample image is shown in Figure 11.11. The channels are not obviously in Taylor flow, as bubble and slug interfaces are difficult to distinguish with our given resolution[1]. However, there is definite "texture" within each channel, and this texture can be easily tracked from one SEMI-RARE image to the next using our Chi$^2$ tracking algorithm. The inability to define bubble and slug boundaries only affects

[1] The resolution can be increased, however such data would take longer to acquire and this would increase the amount of motion blurring, actually making it harder to define slug interfaces.

the measurement of bubble and slug lengths. In fact, we may still estimate the local holdup in the system by investigating the overall intensity within the channels.

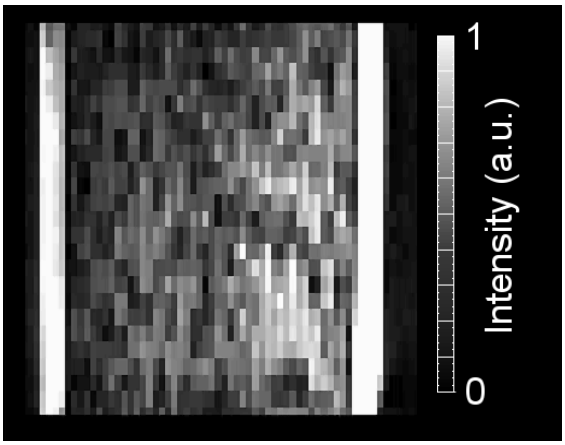


Figure 11.11 – A sample SEMI-RARE image of two-phase flow through a row of channels in the 200 cpsi monolith, using a spray distributor.

The results of tracking velocity profiles for two different spray nozzles at a range of flow rates are shown in Figure 11.12. It can be seen from Figure 11.12 that the velocities observed are relatively unimodal, with a spread around a central value. The spread is far narrower than most other fluid distribution techniques, and so gives the best control over the residence time within the monolith. There are relatively few channels which do not move, particularly at higher flow rates. The distribution shifts to higher velocities at higher flow rates, and the distribution broadens slightly. The two graphs are for two different nozzle types – the nozzle of Figure 11.12a has a higher flow than Figure 11.12b for a given pressure drop, and so it is expected that (a) will give larger droplets than (b). However, this appears to make relatively little difference to the nature of the velocity distribution; droplet sizing is expected to have a far greater effect on slug sizes, however.

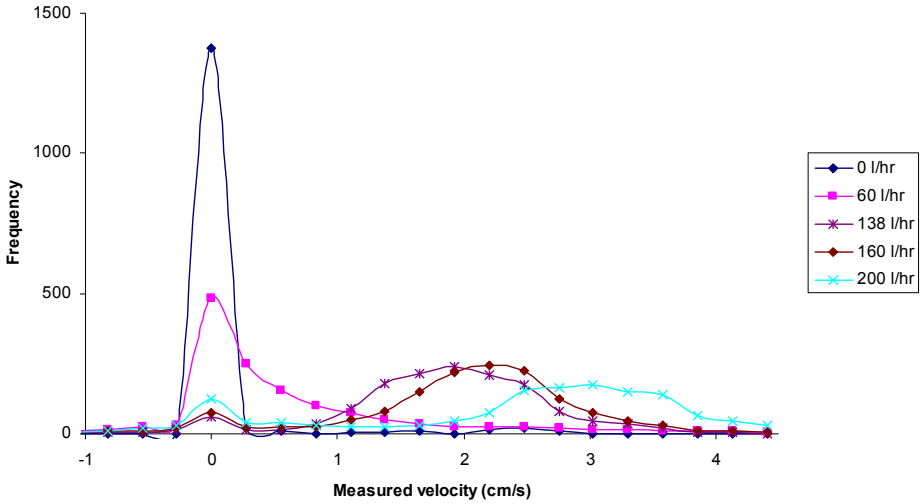


(a)

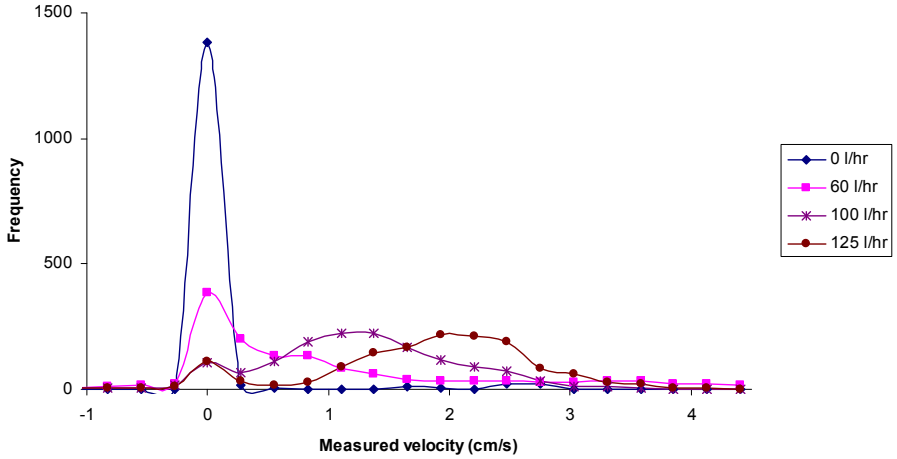


Figure 11.12 – Measurement of the velocity distribution in the monolith reactor for the spray distributor setup. The gas flow rate was held at 400 cc/min and the liquid flow rate was varied between zero and 200 l/hr. (a) shows the velocity / frequency plot for the BGQ1740 spray nozzle, and (b) shows the results for the DAQ1370 nozzle.

The flow rate / direction correlation of Section 11.2.3 was calculated for the spray distributor configuration, and the results are shown in Figure 11.13. It can be seen that the spray distributor has even flow in every channel, and there are velocity fluctuations present equally in each channel. Only the lowest flow rate experiment showed signs of maldistribution.

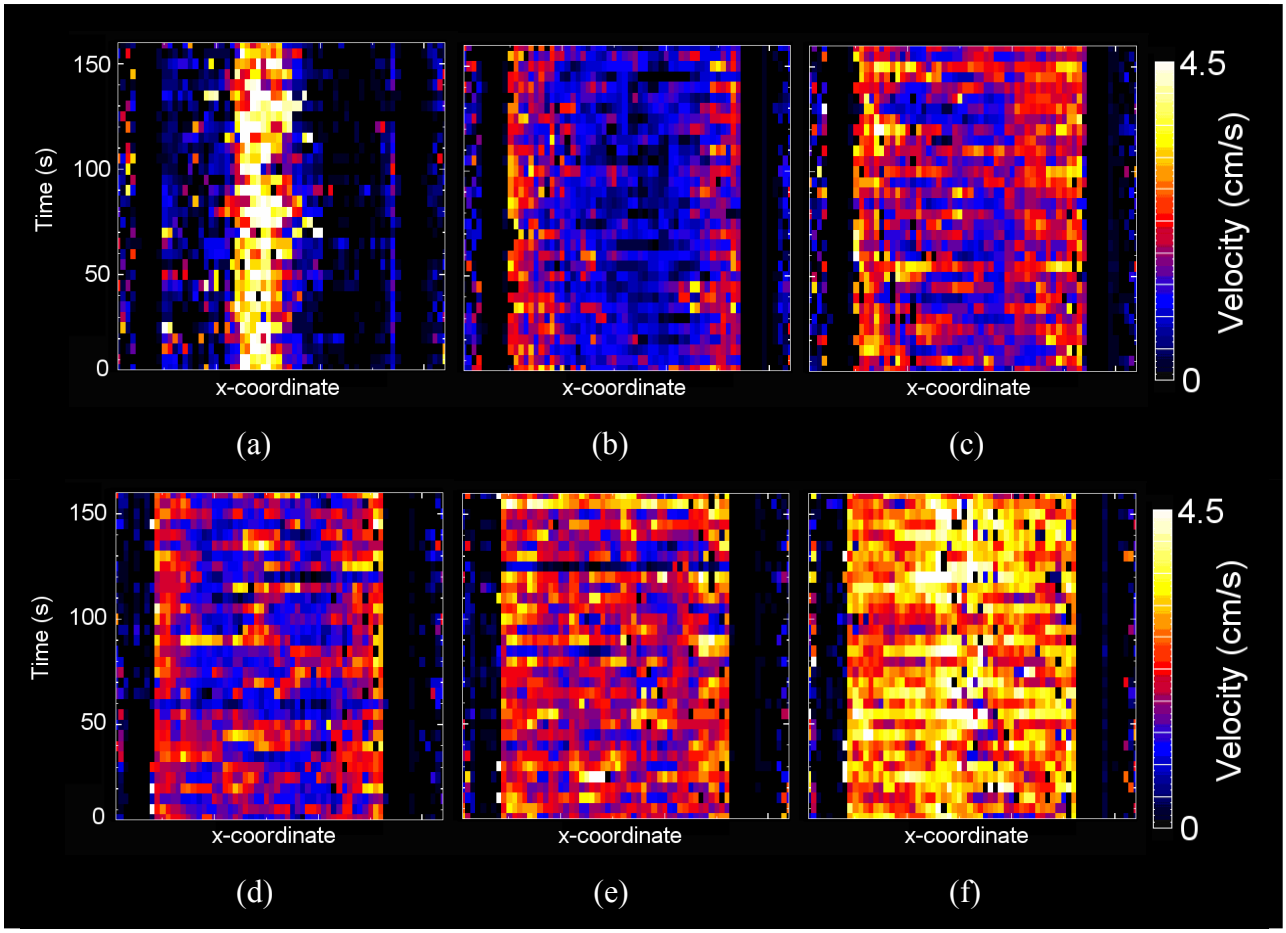


Figure 11.13 – The velocity data of the spray distributor configuration (Figure 11.12), resolved by individual channels, and shown as a function of time (on the vertical axis). The air flow rate is constant at 400 cc/min and the liquid flow rates are (a) 60 l/hr, (b) 100 l/hr, and (c) 125 l/hr with the DAQ1370 nozzle; and (d) 138 l/hr, (b) 160 l/hr, and (c) 200 l/hr with the higher capacity BGQ1740 nozzle.

## 11.2.7. Froth reactor

The velocity-tracking technique was finally applied to the froth reactor (Section 8.2.4). This configuration has very high velocities (up to 30 cm/s), with a high air holdup, and is the most challenging to investigate with MRI. As a consequence, the data obtained is relatively poor; the signal is attenuated strongly and blurred over the 150 ms acquisition time. However, the first two images from the SEMI-RARE sequence are sufficient to obtain a velocity measure. Subsequent images are in general too badly degraded to give meaningful results.

MRI data were acquired with a field of view of 51 mm × 51 mm giving a resolution of 398 μm × 797 μm. The slice thickness was 1.7 mm, which corresponds to the depth of a single channel. Sets of four images were acquired from each excitation, with each being acquired over 156 ms. Hence the echo times associated with each image are taken as 78, 234, 390 and 546 ms, which are the mid-points of each acquisition. Sixty-four repetitions were taken at each flow rate, to ensure we have a statistically significant number of measurements to process. Sample data are shown in Figure 11.14. Signal-to-noise is very poor in this system, however we may still track the velocity of individual channels since we can observe the rate of inflow of fresh spins (zero signal). The validity of processing this noisy data will be discussed with the presentation of results.

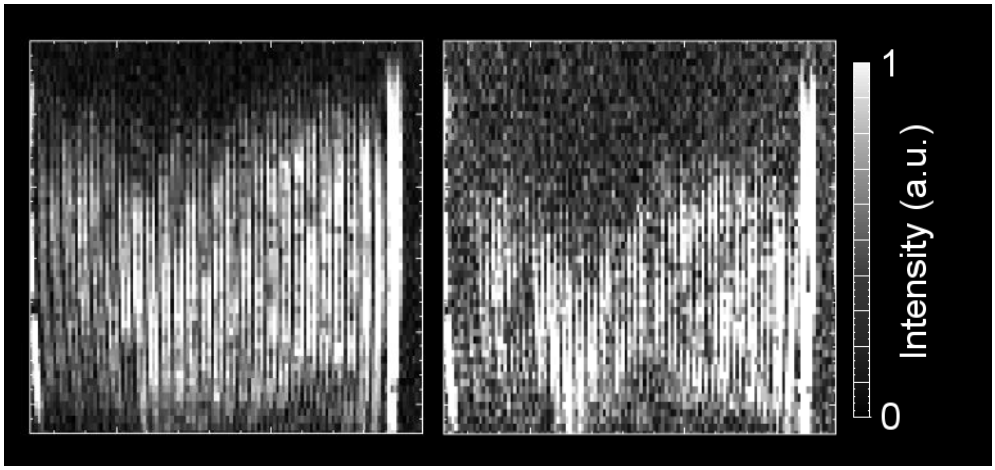


Figure 11.14 – Sample MRI experimental data for the froth reactor with a 5 mm jet operating at a flow rate of 4 l/min of water and 100 cc/min of air. The first two images in the SEMI-RARE acquisition are shown.

The results of analysing the froth reactor data using the $Chi^2$ tracking method (Section 7.6.3) are given in Figure 11.15. Figure 11.15a is for the froth reactor with a 3 mm jet nozzle, and Figure 11.15b is for a 5 mm jet nozzle. The narrower nozzle will give a higher pressure drop for a given flow rate, and this increased loss in energy manifests itself as an increase in the turbulent mixing within the froth itself – the effect of this is to create a stable froth at lower flow rates than is possible for the larger diameter nozzle. The advantage of being able to obtain a stable regime at low flow rates is that it makes the MRI measurement significantly easier since the velocities are lower.

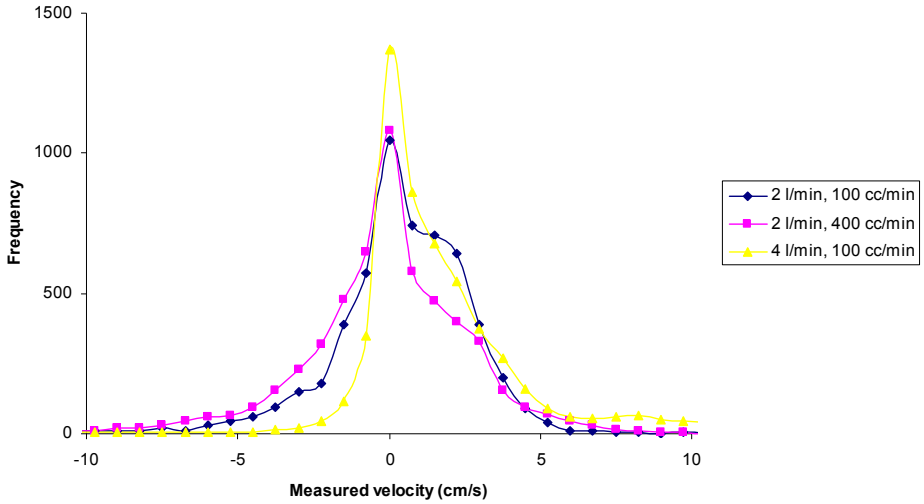


(a)

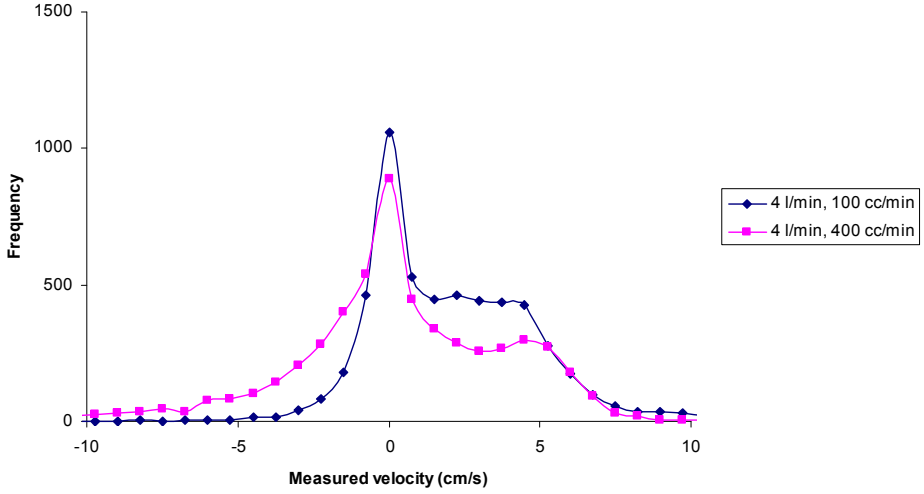


(b)

Figure 11.15 – Comparison of the results of the velocity measurement in the froth reactor with two different jet nozzles; (a) is for a 3 mm jet, and (b) is for a 5 mm jet. The water and air flow rates are indicated, respectively. Downflow corresponds to positive velocities.

By examining these plots, we can see that a higher gas flow rate leads to increased recirculation (shown by more negative flow within channels). Increasing the liquid flow rate appears to narrow the range of velocities observed.

However, by comparing these results with more reliable methods (presented later in Section 11.4.3), we see that the results presented here do not match up. It is clear that the data is in fact too noisy to extract meaningful velocity results, and hence the froth reactor system has illustrated the limitation of this SEMI-RARE tracking method.

The following section will cover the line-excitation MRI technique in application to measuring velocities within our experimental process configurations.

## 11.3.Line-excitation

This section will now examine the performance of the line-excitation method of Section 5.3 in application to the monolith configurations. This is a novel MRI time-of-flight method for tagging spins and measuring their displacement over the acquisition time. Since the spins within the monolith channels are constrained to move in one axis only, this displacement can be directly related to a velocity measurement.

### 11.3.1. Slice thickness

The choice of the excitation slice thickness is important for the line-excitation setup. Figure 11.16 shows the orientation of the excitation and refocusing RF pulses, relative to the monolith structure. This image is reproduced from Chapter 5.

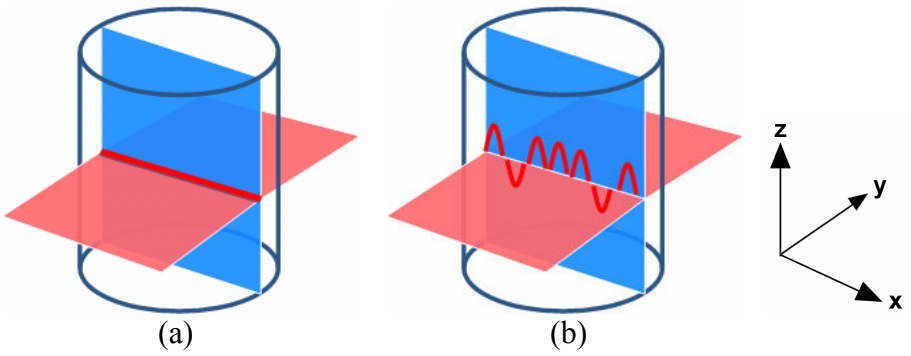


Figure 11.16 – Schematic of the orientation of the excitation plane (red) and refocusing plane (blue). The image is acquired in the refocusing plane. The schematic shows that, (a) with zero flow, a single line is refocused; (b) with flow, the line is displaced according to the axial flow velocity.

Figure 11.17 illustrates the effect of using different excitation slice thicknesses of 5 mm and 1.7 mm in the 200 cpsi monolith. The data shown are for the bubble sparger configuration at a range of air flow rates. Theoretically, the ideal slice thickness is as narrow as possible because then we can directly measure the evolution

of the spins without having to consider the shape of the initial distribution. However, we can see that although thinner slices give clearer parabolic profiles, the signal to noise ratio is significantly lower and hence it is better to use a thick slice even though it complicates the data analysis later.

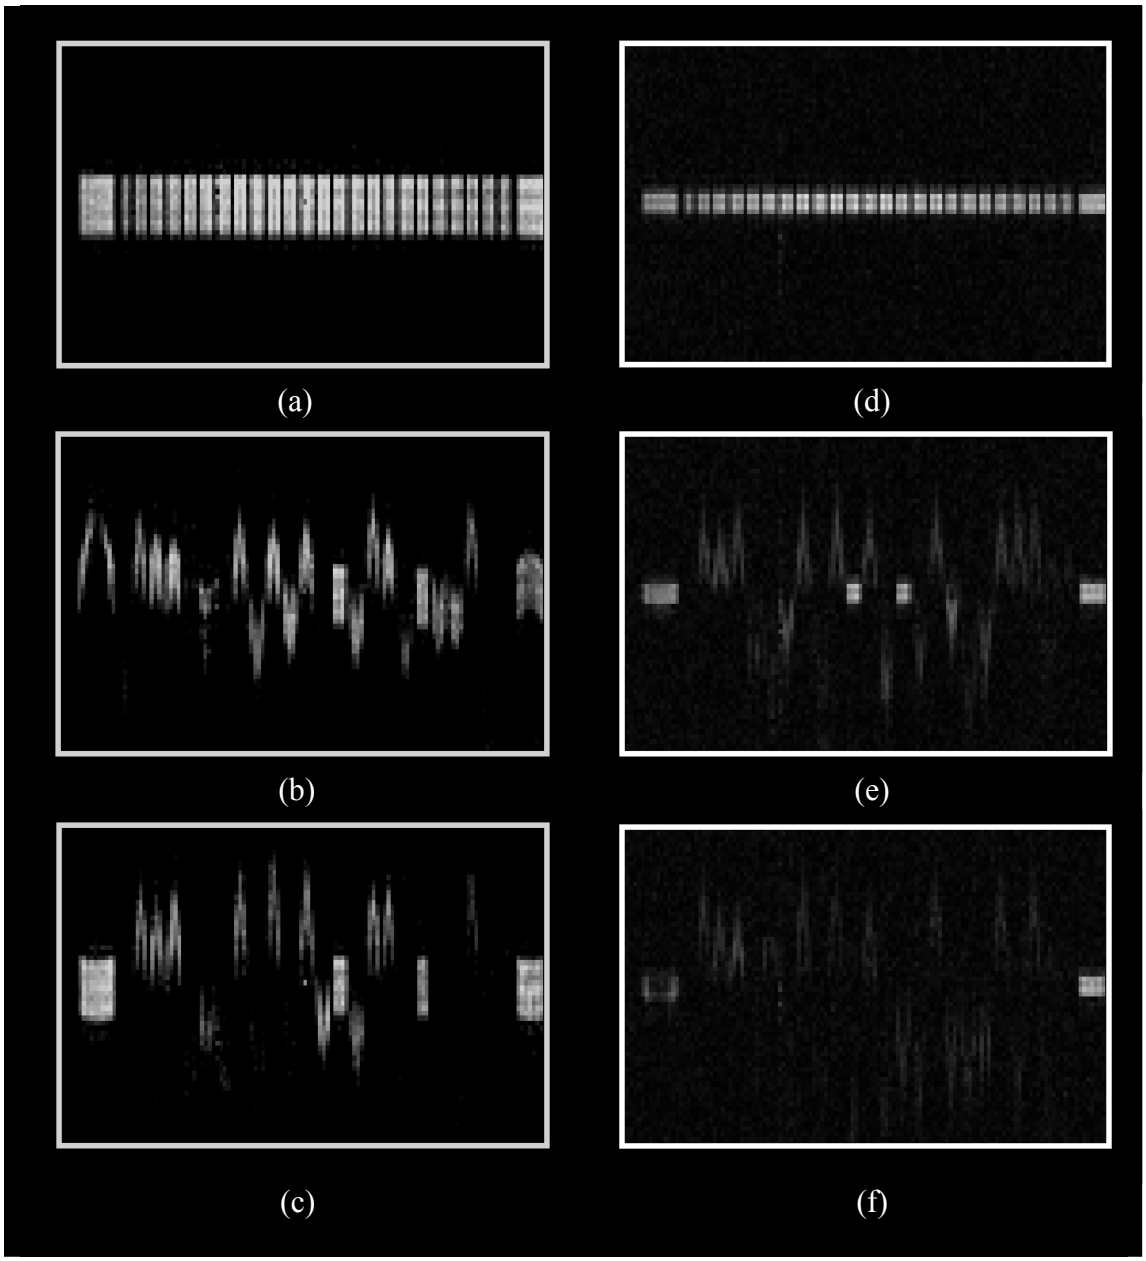


Figure 11.17 – A comparison of different slice thicknesses; (a-c) show a 5 mm thick excitation slice at flow rates of 0, 100 and 200 cc/min; (d-f) are corresponding images for a 1.7 mm excitation slice. In all cases, images are taken with a FOV of 50 mm × 50 mm, and a resolution of 391 μm × 391 μm, and a refocusing slice thickness (into the page) of 1.7 mm. Both the excitation pulse and the refocusing pulse are sinc3 pulses (Section 5.3).

In very high flow rate systems, three main problems are encountered – fluid may leave the imaging coil over the imaging time, the signal-to-noise ratio decreases due to spins being spread over more voxels, and there is also a drop in signal due to diffusive attenuation (shearing). These problems can be overcome by taking fewer phase increments, and hence giving a more rapid imaging time. Fluid no longer leaves the coil and signal-to-noise is improved, but at the sacrifice of resolution.

## 11.3.2. Bubble sparger

Having shown in Section 6.3 that the line-excitation method is still quantitative in application to the monolith reactor, it can now be applied to two-phase flow configurations. This section will illustrate its application to the bubble sparger system (Section 8.2.1), highlight the quality of the data we may obtain, and show the limits of the technique.

An example of the line-excitation sequence in the monolith is shown in Figure 11.18. MRI data were acquired using the SEMI-RARE sequence for the 200 cpsi monolith at gas flow rates between zero and 500 cc/min. Images are taken with a field-of-view of 50 mm × 25 mm, a resolution of 391 μm × 391 μm, an excitation slice thickness of 5 mm, and a refocusing slice thickness of 1.7 mm.

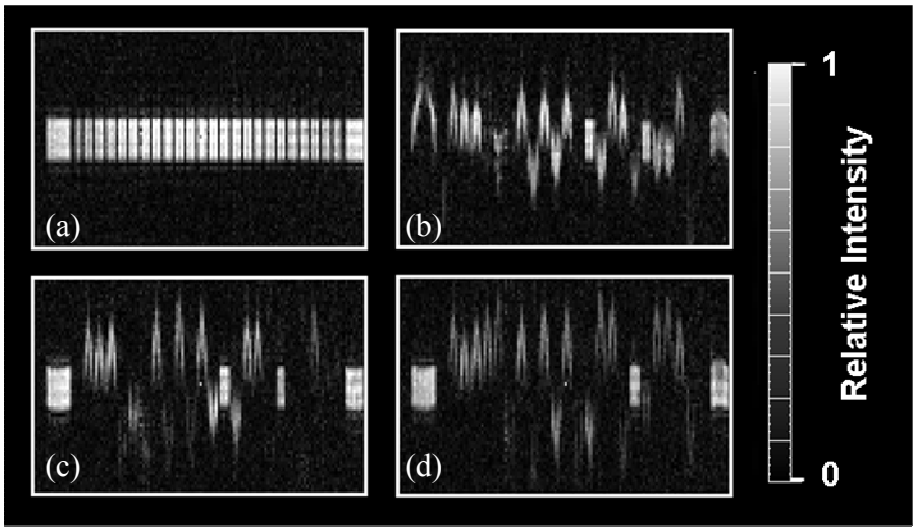


Figure 11.18 – Line-excitation data at a range of flow rates. The setup used the 200 cpsi monolith, and a gas sparger was used to bubble gas upwards through stagnant water. Recirculation is seen in some of the channels. Flow rates were (a) zero, (b) 150 cc/min, (c) 200 cc/min and (d) 300 cc/min of air.

The data were processed by considering each channel in turn and tracking the position of the maximum intensity within these channels. This displacement was then taken as the position corresponding to the mid-point of the acquisition time. Using this method, we can therefore extract velocity / frequency plots from the above data sets, and these are shown in Figure 11.19.

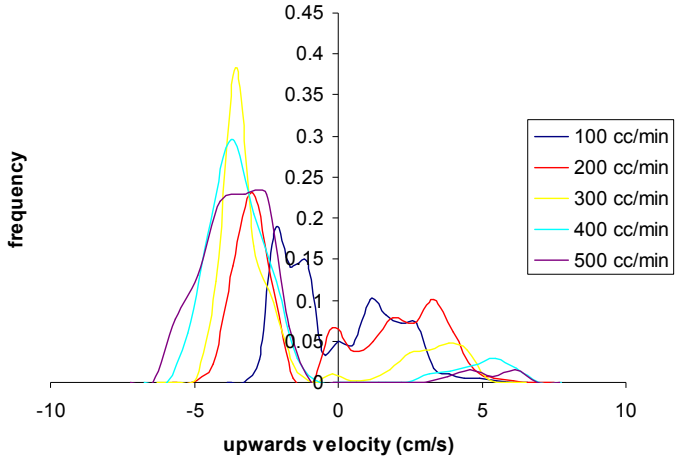


Figure 11.19 – Velocity histogram produced from the line-excitation tagging data, for the bubble sparger system at a range of air flow rates. Data are normalised.

The plot shows that increasing the air flow rate pushes the distribution to higher velocities in both the upward and downward directions. This plot should be compared to that of Figure 11.1. It can be seen that the overall trend is similar, however the line-excitation method lacks information from the upward-moving channels; there should be more channels in upflow than downflow for the bubble sparger system, but this is not what the plot shows. This is attributed to one of the limitations of the line-excitation technique – if we excite a channel which has a high gas holdup, there is a good chance that there may be a gas bubble in the excitation slice. This will result in an error in the tracking method, since we will have little or no signal from that channel. Hence, we may not always be able to track the upwards-moving channels, and this will worsen at higher gas flow rates and higher gas holdups, which is consistent with the observations from the data.

## 11.3.3. Slice thickness investigation

Using the established technique for extracting velocity / frequency plots, we can now investigate the effect of the chosen MRI slice thickness on the results obtained. This is important since we are using the displacement of the spins to track the movement of the fluid – if we do not measure all the spins within a given channel, we may measure an incorrect spreading of the spins.

The data obtained for the 200 cpsi monolith are shown in Figure 11.20 for data acquired with slice thicknesses of 1.7 mm and 1.2 mm. The correct channel thickness is 1.7 mm, and the purpose of this test was to investigate the sensitivity to using an incorrect slice thickness. The signal-to-noise ratio is higher for the thicker slice as would be expected. The profiles themselves are almost indistinguishable by eye, but by processing the data we can investigate subtle differences.

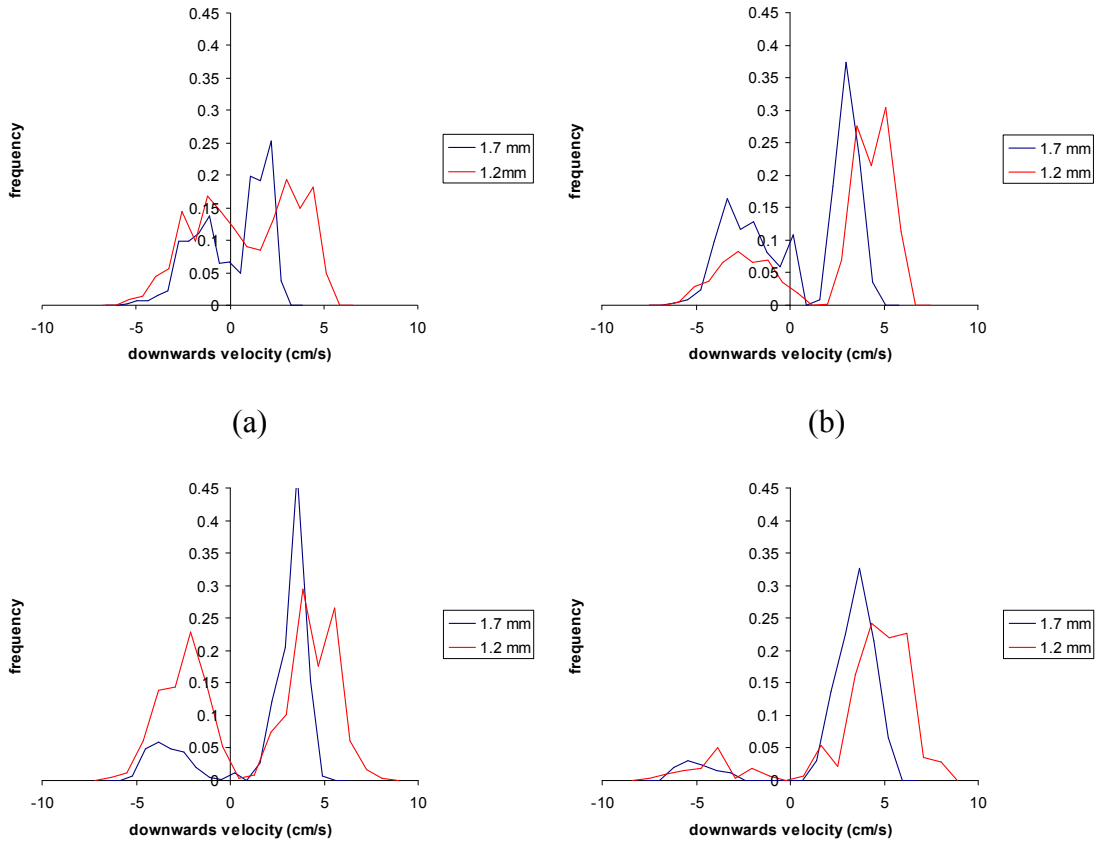

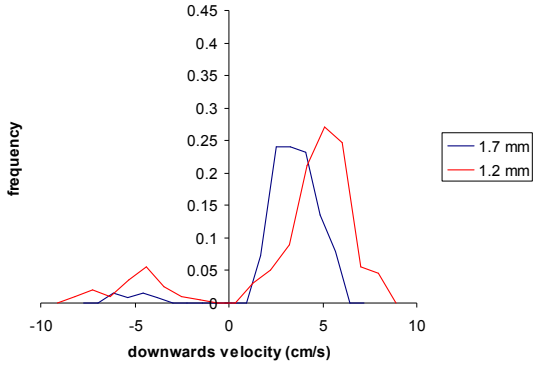


(e)

Figure 11.20 – Histograms of velocity / frequency obtained as a result of processing the line-excitation data. The graphs shown are for a 200 cpsi monolith in the bubble sparger configuration at increasing gas flow rate, showing the results for analyses on 1.7 mm and 1.2 mm refocusing slice thicknesses. Gas flow rates are (a) 100 cc/min, (b) 200 cc/min, (c) 300 cc/min, (d) 400 cc/min, and (e) 500 cc/min.

From these histograms, we can see that the plots for the thinner excitation slice are shifted to higher velocities relative to those from thicker slices which cover the entire channel. This is as expected, since the narrower slice will lose signal from slow-moving spins near the channel walls, thereby giving an unrepresentatively high measurement of the displacement in the channels. This highlights the need for the tailoring of the slice thickness (and shape) to closely match the system of interest.

## 11.3.4. Co-current flow

The line-excitation technique was then applied to co-current flow in the bubble sparger system. This system has a similarly high liquid holdup and only slightly higher velocities than the bubble sparger system, and so was deemed to be a suitable application of the line-excitation sequence. MRI data were acquired with the same parameters as for the bubble-sparger setup, and an example is shown in Figure 11.21.

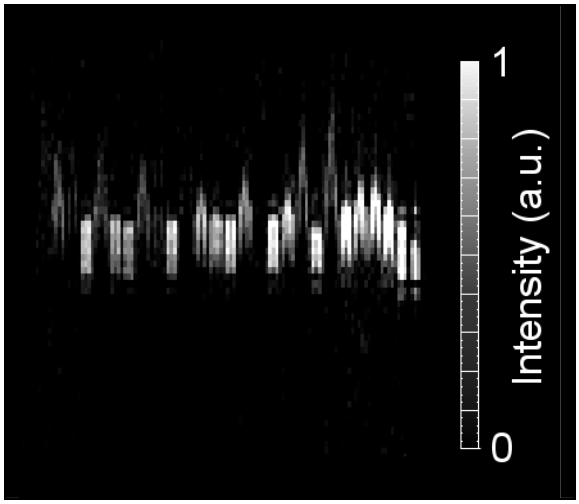


Figure 11.21 – A sample line-excitation acquisition of co-current flow in a 200 cpsi monolith. Gas was supplied at 100 cc/min and water at 78 l/hr.

The results of this analysis are shown in Figure 11.22. It can be seen that the histogram obtained at increasing liquid flow rate shifts as expected.

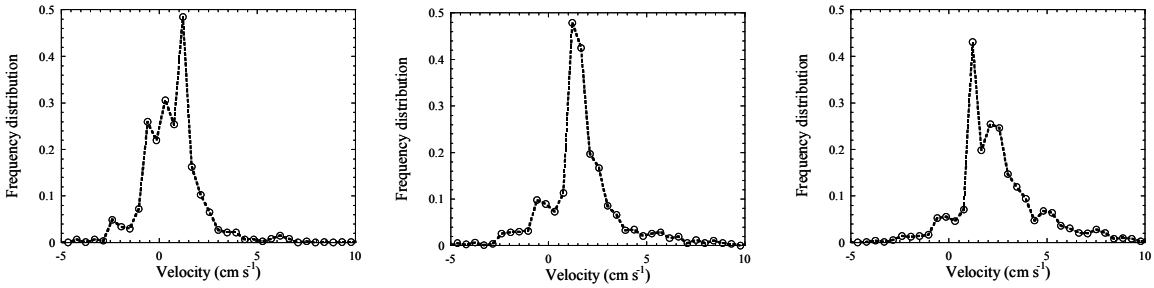


Figure 11.22 – Results of applying the line-excitation method to the co-current bubble sparger configuration at a range of liquid flow rates

It has been shown previously that the co-current flow setup helps to improve the dynamics of the monolith system; rather than having a few channels that are solely in upflow, the distribution is far better. Therefore, there should be fewer channels with a very high gas holdup, and this should allow the technique to better represent the true state of the system. However, at higher flow rates the histogram becomes less well-formed and this is probably due to the higher gas holdup. These results should be compared with the later COTTAGE results of Section 11.4.1.

### 11.3.5. Other monolith configurations

The line excitation technique has not been applied to the showerhead, spray or froth reactors. The reason for this is that the quality of the data in these systems is not sufficient for robust data analysis. The line-excitation method breaks down in the presence of very fast flow because high shear and also the spreading of spins over a larger region, decreasing the voxel signal intensity. It also suffers in systems of high gas holdup since we will have fewer spins to detect, and also there is the possibility of an air bubble being present within the excitation slice, giving zero signal. Therefore the application to the showerhead distributor fails because of the very high local velocities, and the high gas holdup within some channels; the spray distributor fails because of the variable fluid holdup within the channels, which makes it hard to track spins; and the froth reactor fails due to the very high velocities present in the system.

However, for slow-moving systems of high liquid holdup, this technique provides a valuable insight into the instantaneous flow characteristics of the system; for example we are able to take a rapid "fingerprint" of the flow direction and magnitude within the channels of the monolith, which other traditional techniques would not be able to achieve in such dynamic systems. Also, we are purely using the displacement of the spins to measure the displacement, and the absence of velocity-encoding MRI gradients makes this a true one-shot technique.

### 11.3.6. Summary

The line excitation method can suffer from poor signal-to-noise since we are only exciting and reading a narrow line of spins. Consequently, when we acquire an image over, say, 150 ms, we find that the signal is quite weak, and we tend to lose signal from the slowest moving spins closest to the walls due to their shorter $T_2$ (caused in part by interactions with the walls). For this reason, it was decided to develop a faster pulse sequence, COTTAGE, which sacrifices some of the spatial resolution for far higher time resolution. This means that the spins do not have long to relax, and we can probe faster-moving systems without cause for concern. The application of COTTAGE to the monolith system will be addressed in the following section.

## 11.4.COTTAGE

This section will cover the application of the novel MRI pulse sequence, COTTAGE (Section 5.4), to the monolith reactor configurations. The velocity distributions measured with this method are presented, and the potential of this technique is discussed.

### 11.4.1. Co-current flow

The COTTAGE technique was first applied to co-current flow (Section 8.2.1), since we have already shown it to be a superior modification to the bubble sparger setup. Sample datasets are shown in Figure 11.23, for an air flow rate of 100 cc/min and increasing liquid flow rates. Time is on the vertical axis and the colour scheme corresponds to displacement. These MRI data were acquired with a field of view in the x-axis of 51 mm and a resolution of 398 μm, with an echo time of 2.4 ms between successive profiles, plotted up the y-axis. Sinc3 pulses (Appendix 6.9) were used for the excitation slice (5 mm) and the refocusing slice (1.7 mm).

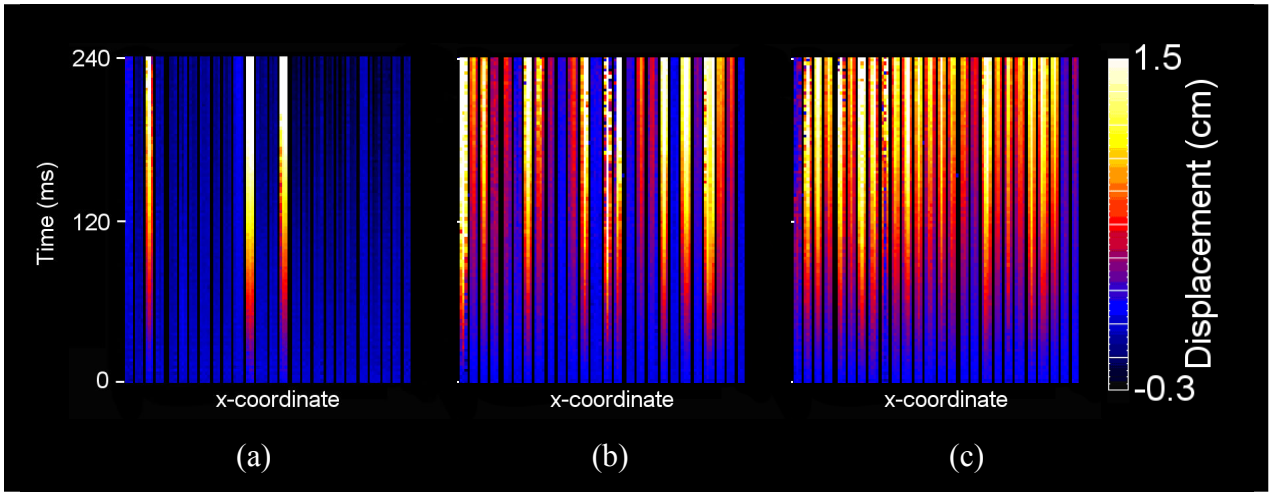


Figure 11.23 – COTTAGE velocity data for the 200 cpsi monolith operating under conditions of co-current gas-liquid upflow, with a constant gas flow rate of 100 cc/min and liquid flow rates of (a) 29, (b)78 and (c) 105 l/hr.

Some selected profiles extracted from the data taken at 78 l/hr are shown in Figure 11.24b, where the channel numbers correspond to those marked in Figure 11.24a. It can be seen that the increase in displacement is approximately linear over the

observation time, and hence it is straightforward to extract velocities from the data by fitting to the gradient of the lines. Some selected profiles at increasing time are shown in Figure 11.24c.

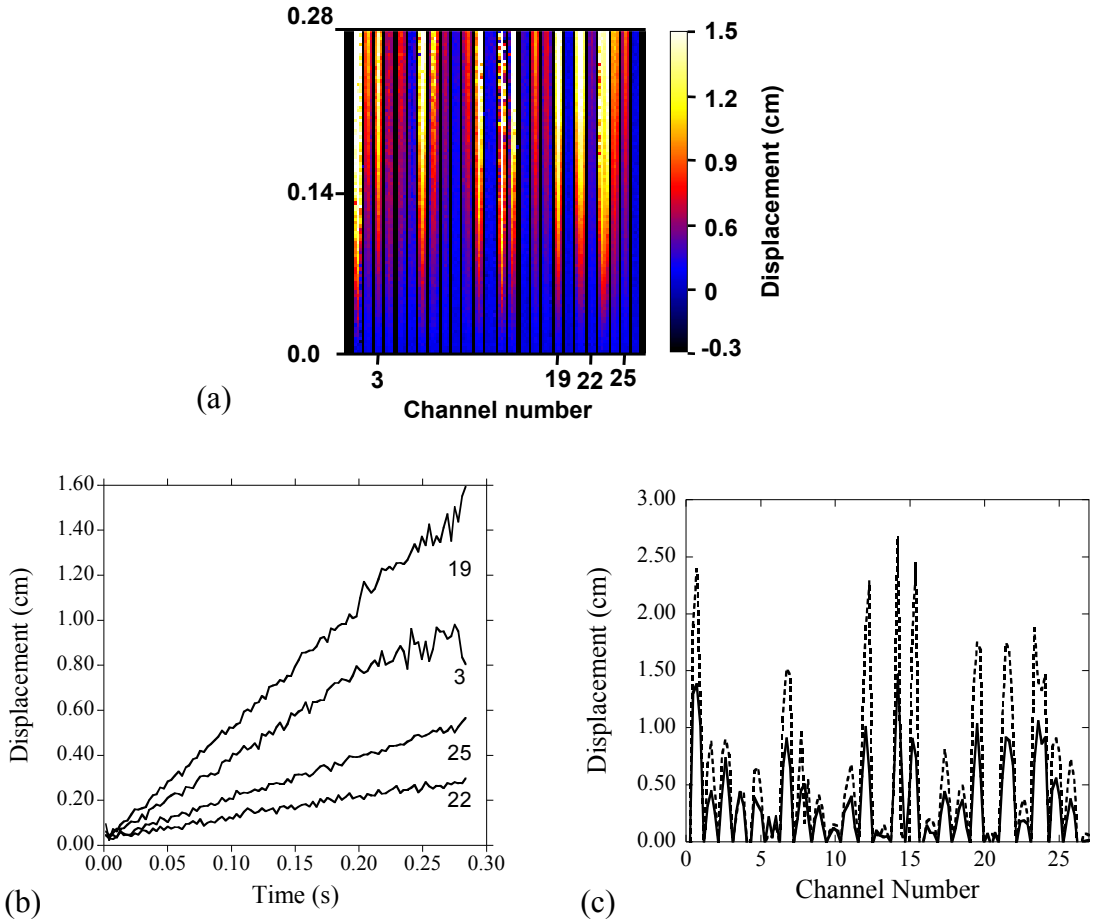


Figure 11.24– (a) A sample COTTAGE dataset acquired for co-current flow in the monolith reactor, at an air flow rate of 100 cc/min and a liquid flow rate of 78 l/hr; (b) sample channel profiles extracted from the COTTAGE dataset, where the corresponding channel numbers are marked; (c) selected profiles taken from the COTTAGE dataset at 140 ms (–) and 280 ms (--) after excitation.

Processing this data leads to the velocity histograms of Figure 11.25, which have been overlaid over previous line-excitation data for the same system. These data show us that the histograms obtained with the line-excitation and COTTAGE methods compare favourably, particularly at lower flow rates. At higher flow rates the COTTAGE data appears to be more robust, whereas unusual artefacts appear in the line-excitation histograms. This suggests that COTTAGE is a better method for measuring the displacement in such monoliths systems.

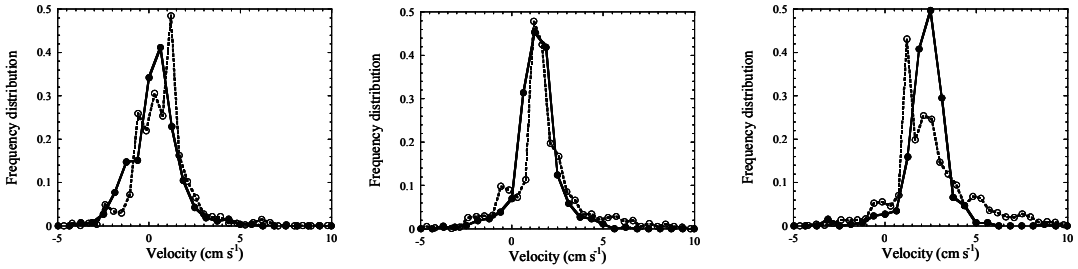


Figure 11.25 – COTTAGE-derived velocity histograms calculated for the monolith operating under conditions of co-current gas-liquid upflow, with a constant gas flow rate of 100 cc/min and liquid flow rates of (a) 29, (b)78 and (c) 105 l/hr. The distributions are normalised. The corresponding line-excitation histograms are overlaid for comparison; (--) line-excitation; (–) COTTAGE.

An explanation for the differences in the techniques is that the line-excitation method suffers from relaxation effects over its relatively long acquisition time, and so the NMR signal is weaker and the interpretation of the contrast becomes more difficult. In contrast, the COTTAGE technique uses a phase-based measurement to give the average displacement, and has far higher temporal resolution. COTTAGE therefore allows us to utilise purely data from early in the acquisition, and so relaxation effects are minimised.

## 11.4.2. Dynamic observations

While much of the flow phenomena observed are at steady state on the order of the acquisitions (100-200 ms), some channels in the co-current flow system are observed to change behaviour over shorter timescales. For example, the plot of Figure 11.26 shows evidence of a channel changing direction during the acquisition time. It begins by moving upwards, and then changes direction and flows downwards. The data is noisy, but the trend is quite clear. This type of information on fast dynamic processes could not have been obtained by any other MRI method, because only COTTAGE has sufficient temporal resolution. For example, a line-excitation acquisition would have averaged together the displacement over the first 156 ms, and hence we would have lost this dynamic information. This highlights the potential of this technique to probe highly dynamic systems.

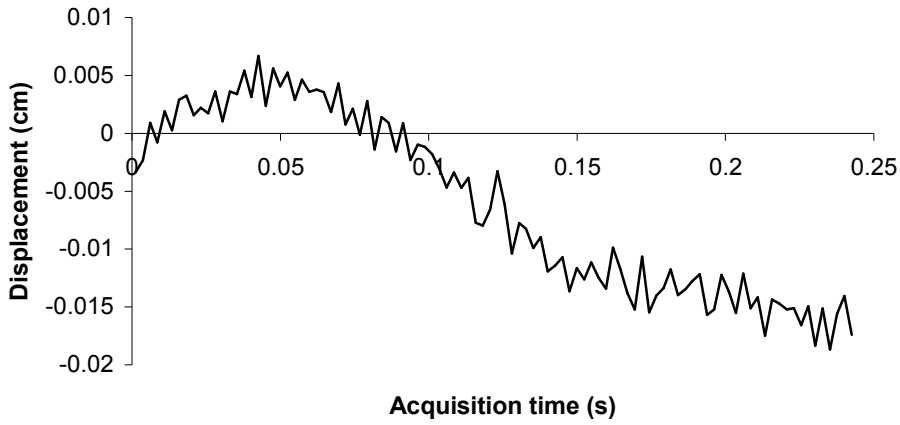


Figure 11.26 – An observation of the flow in a monolith channel changing direction over very short timescales (~100 ms). Profile was extracted from a COTTAGE data set of co-current gas-liquid upflow, with a constant gas flow rate of 100 cc/min and liquid flow rate 29 l/hr.

As mentioned in Section 3.7.4, Pinto *et al.* (1997) reported that the velocities in bubble train flow through a capillary are constant, except for the moment when a new bubble enters the capillary – then there is around a 200 ms period of fluctuation due to a pressure build-up as the bubble is pulled in. It has been shown that the COTTAGE technique is able to study fluctuations on this timescale

## 11.4.3. Froth reactor

The COTTAGE technique has been shown to provide robust data with a high temporal resolution. It will now be applied to the froth reactor configuration, which is the most challenging of the setups. We experience high flow rates and most MRI techniques fail here. Examples of COTTAGE experiments on the froth reactor are shown in Figure 11.27. The experimental MRI parameters are the same as for co-current flow, but with a weaker phase-encoding gradient in order that the increased flow rate will not cause the measured phase to exceed $2\pi$. Sixteen repetitions were taken at each flow rate. The froth reactor was set up with the 5 mm orifice, and the experiments shown are for 240 l/hr and 480 l/hr of water with 100 cc/min of air. It can be seen that the COTTAGE data is excellent at a large range of flow rates, the noise that we see at later times is due to fluid physically leaving the RF coil, and this

is unavoidable. However, even at very high flow rates there are still sufficient points to extract velocities from the COTTAGE data.

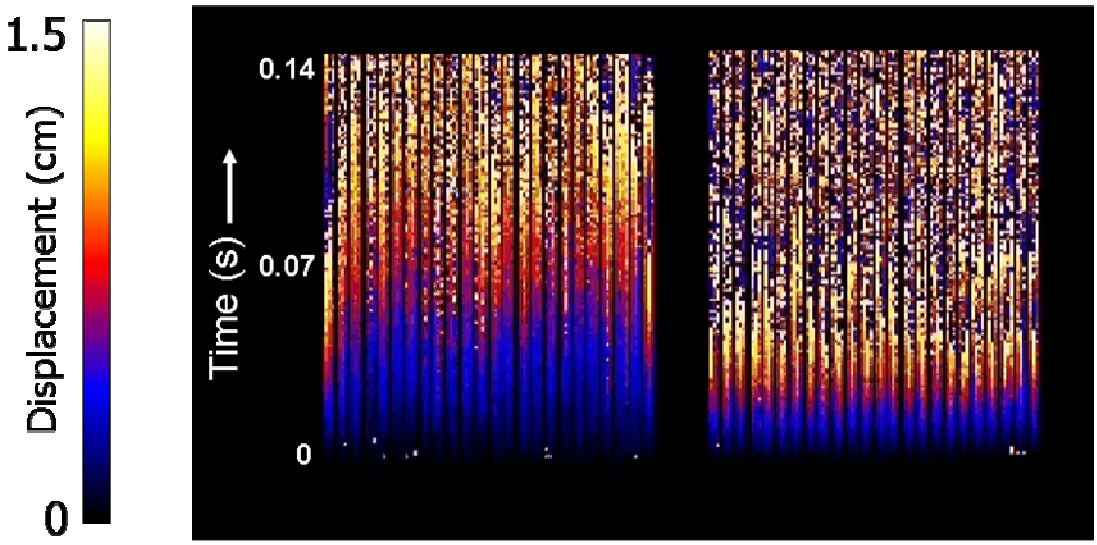


Figure 11.27 – COTTAGE data from froth reactor experiments at 100 cc/min air and (a) 4000 cc/min water, (b) 8000 cc/min water.

An initial concern of using the froth reactor with its high-velocity turbulent jet was that fluid would be pushed through the centre of the monolith much faster than at the edges. However, by examining the velocity distribution using the COTTAGE technique, we can see that the velocity distribution is relatively homogeneous across the monolith cross-section. Therefore, we can say that the turbulent mixing zone must serve to dampen the force of the turbulent jet, and hence the residence time of the fluid within the monolith is relatively unimodal.

Experiments were performed on the froth configuration a 5 mm orifice at air flow rates of 100, 400 and 1000 cc/min, and water flow rates of 4, 6 and 8 l/min. The results are shown in Figure 11.28 and Figure 11.29. The histogram of Figure 11.28 shows the effect of increasing the liquid flow rate; the velocity distribution is seen to shift and broaden slightly. Figure 11.29(a-c) shows the effect of changing the gas flow rate. In each case, increasing the gas flow rate broadens the velocity distribution without shifting it significantly. This would imply that to obtain a unimodal velocity distribution, the best configuration is to have the minimum flow rate of gas while keeping the turbulent mixture stable.

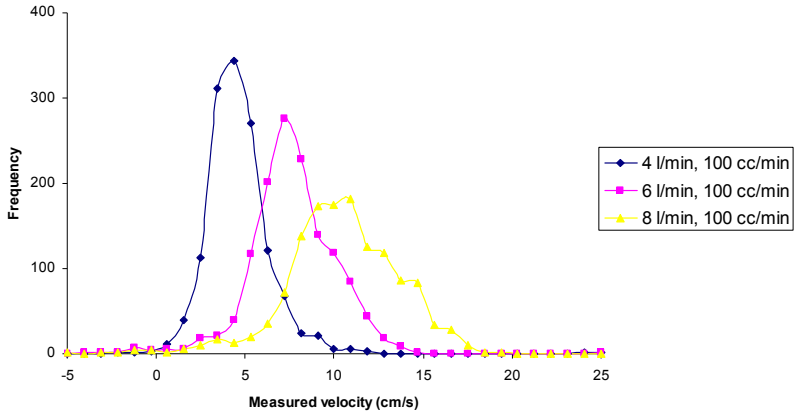


Figure 11.28 – Velocity histograms obtained using the COTTAGE technique on the froth reactor (5 mm orifice) at a range of liquid flow rates.

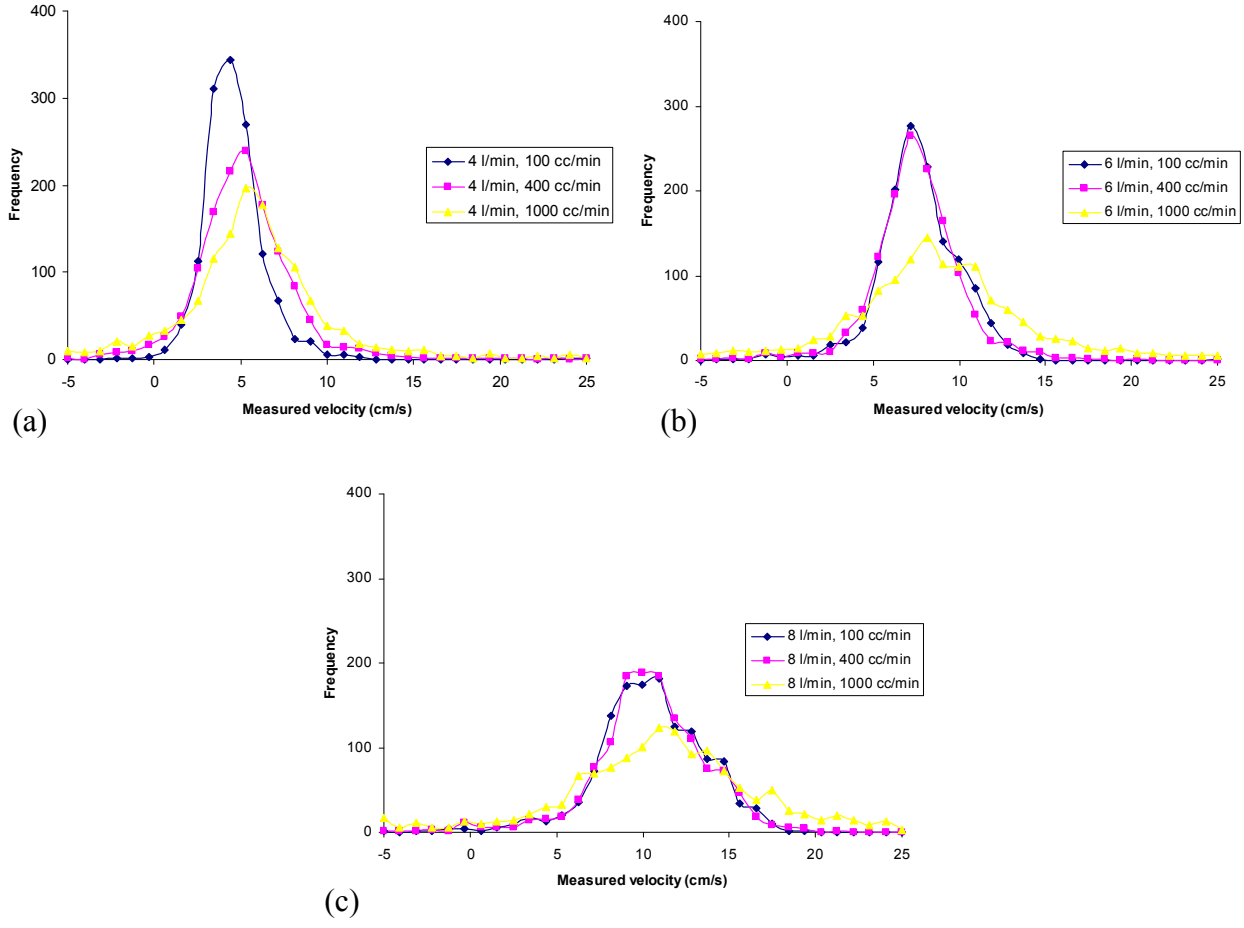


Figure 11.29 – Velocity histograms obtained using the COTTAGE technique on the froth reactor (5 mm orifice) at a range of gas flow rates. Results are shown for water flow rates of (a) 4 l/min, (b) 6 l/min, and (c) 8 l/min.

These velocity histograms are significantly different to those obtained from the froth reactor using the SEMI-RARE tracking technique in Section 11.2.7. This is explained by the limitations of the SEMI-RARE technique making it inappropriate for this flow configuration, and the COTTAGE technique should give far more reliable results due to the more robust nature of the pulse sequence.

Data were also obtained for the froth reactor with the lower-flow rate 3 mm orifice. Experiments were performed at air flow rates of 100, 400 and 1000 cc/min, and water flow rates of 2 and 4 l/min. The results of this analysis are presented in Figure 11.30. A similar trend is observed here, whereby increasing the liquid flow rate shifts the histogram to higher velocities (although the degree of broadening is not so clear for this case). The increase in gas flow rate also broadens the distribution without shifting it significantly.

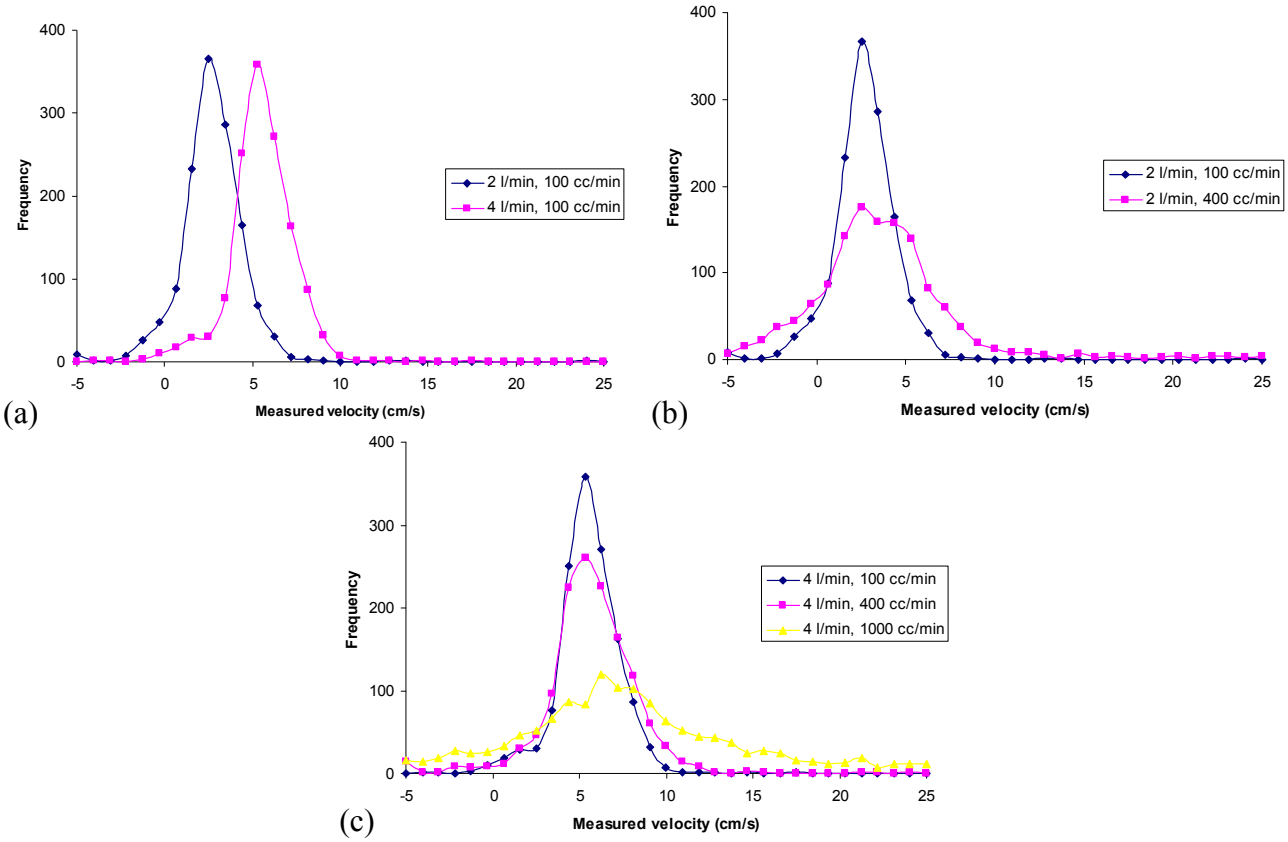


Figure 11.30 – Velocity histograms obtained using the COTTAGE technique on the froth reactor (3 mm orifice) at a range of gas and liquid flow rates. Results are shown for (a) a range of water flow rates while keeping the gas flow constant; (b) at a range of air flow rates with water flow rates held at (b) 2 l/min, and (c) 4 l/min.

## 11.5.Multi-slice COTTAGE

The multi-slice adaptation of the COTTAGE technique (Section 5.5) was then applied to the monolith reactor. As mentioned in Chapter 6, the signal-to-noise related to using this technique in the monolith system is quite poor, and was found to be worse in the presence of two-phase flow. It has only been applied to the bubble sparger system in this work.

### 11.5.1. Bubble sparger

MRI data were acquired using the novel MS-COTTAGE technique, with a field of view in the x-direction of 51 mm and a resolution of 398 μm. 31 slices were acquired in the y-direction, each 1.7 mm wide and 1.7 mm apart. This gives a field-of-view of 53 mm in the y-direction. Each profile was acquired in 3.3 ms, giving a total acquisition time of 102 ms. Sample data are shown in Figure 11.31 for the bubble sparger system at a range of air flow rates.

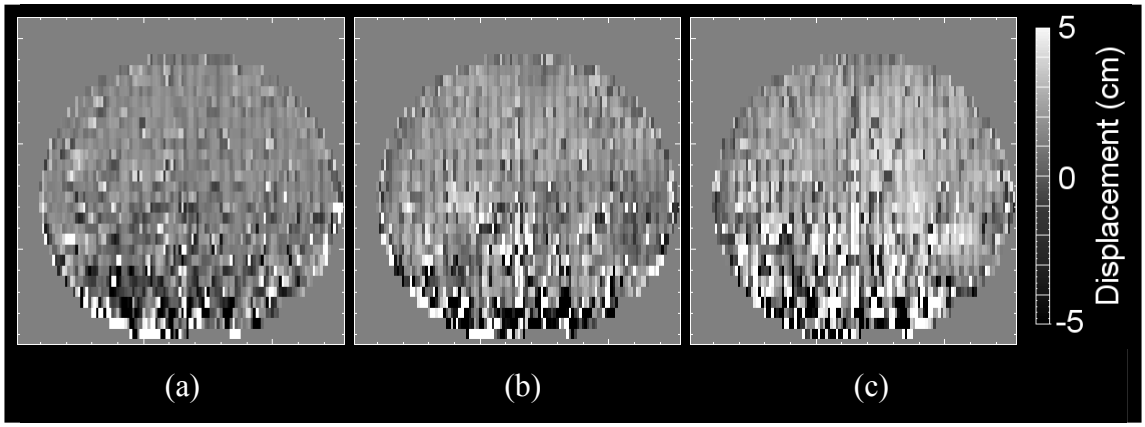


Figure 11.31 – MS-COTTAGE data for the bubble sparger system at air flow rates of (a) 100 cc/min, (b) 200 cc/min and (c) 500 cc/min.

Some velocity histograms were compiled using the data acquired, and are presented in Figure 11.32. The histograms appear quite broad or "smoothed", and it is likely that there is a large degree of noise in the measurements. However, the trend is still consistent with previous measurements, and even if this is not a robust, quantitative technique in application to two-phase flow, it is still a very rapid 2D velocity visualisation technique.

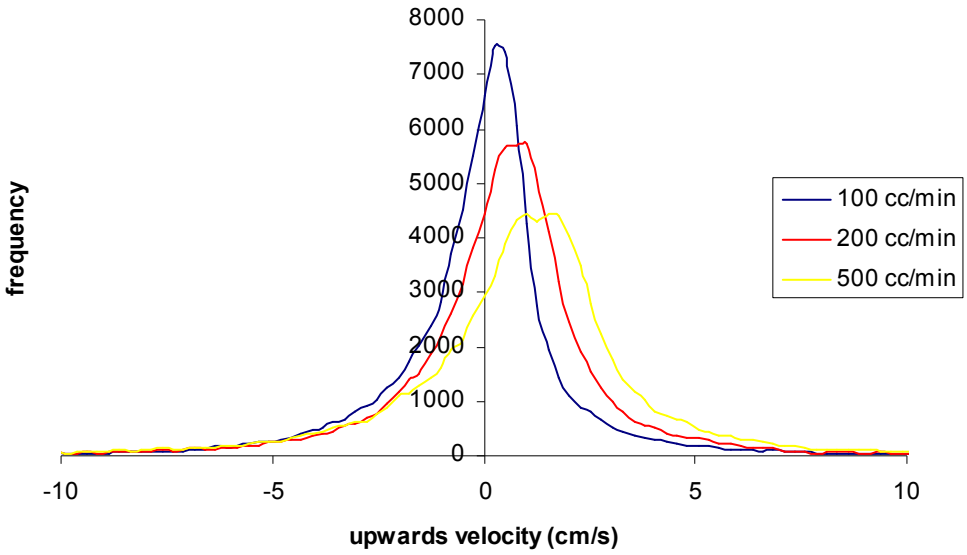


Figure 11.32 – MS-COTTAGE-derived velocity histograms calculated for the monolith operating under bubble sparger configuration, at a range of gas flow rates.

## 11.5.2. Discussion

Signal to noise is generally very poor in these MS-COTTAGE experiments. This could be improved by implementing an APGSTE sequence. With a simple stimulated-echo experiment such as MS-COTTAGE, we assume that the spins experience the same magnetic susceptibility. If the spins instead move to a region of different magnetic susceptibility, then they will be dephased and signal is lost. An APGSTE sequence would compensate for this, and should increase the performance of the sequence for two-phase flow in monoliths.

# 11.6. GERVAIS

It was attempted to perform a GERVAIS experiment in a monolith under two-phase flow conditions, even though it has been shown that GERVAIS is not quantitative in the monolith (Chapter 6) and that its performance will be far worse during two-phase flow due to increased susceptibility effects. However, some data were acquired for the showerhead system, and some example data will now be presented and discussed.

Figure 11.33 illustrates the effect of magnetic susceptibility on the GERVAIS pulse sequence. These data were acquired with a field-of-view of 65 mm and a resolution of 1.01 mm × 1.01 mm. The slice thickness was 40 mm, and fifty images are averaged to remove any fluctuations in the system. Figure 11.33a shows the results obtained from the GERVAIS sequence (the intensity is shown, rather than the phase-component) and the results of the RARE pulse sequences are shown in Figure 11.33b for comparison. Acquisition times are 20 ms and 150 ms respectively.

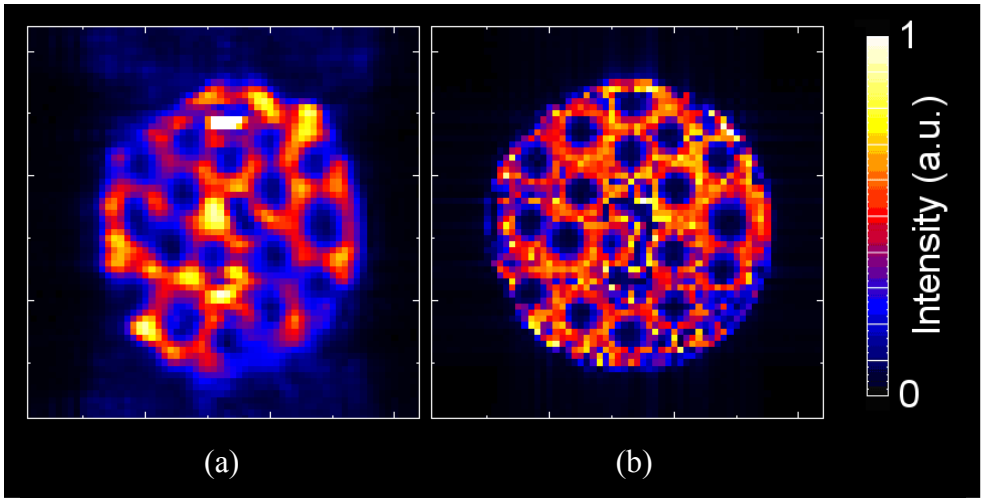


Figure 11.33 – Sample cross-sections through the monolith operating in showerhead configuration; (a) is a GERVAIS intensity image, and (b) is a RARE image.

It can be seen that the GERVAIS intensity image contains physical distortions and fluctuations in the obtained intensity, in comparison to the RARE image. These are predominantly due to magnetic susceptibility effects, which the RARE sequence is far more resistant to. It can be seen that there is a significant loss of intensity where the showerhead jets hit the monolith, and even when the acquisition time is as short as 20 ms, little signal is obtained from the centre of the jets. This causes problems for our analysis, since we will obtain random phase angles from voxels which are swamped by noise. It is possible to gate the GERVAIS data, and only extract velocity data from regions of high intensity. The results of processing the data in this manner are shown in Figure 11.34. It can be seen that the velocity data from regions of high intensity (low flow rate) are flowing steadily and are relatively evenly distributed. In contrast, the data from the low-intensity regions (high flow rate) appear far more noisy, and the data is not reliable.

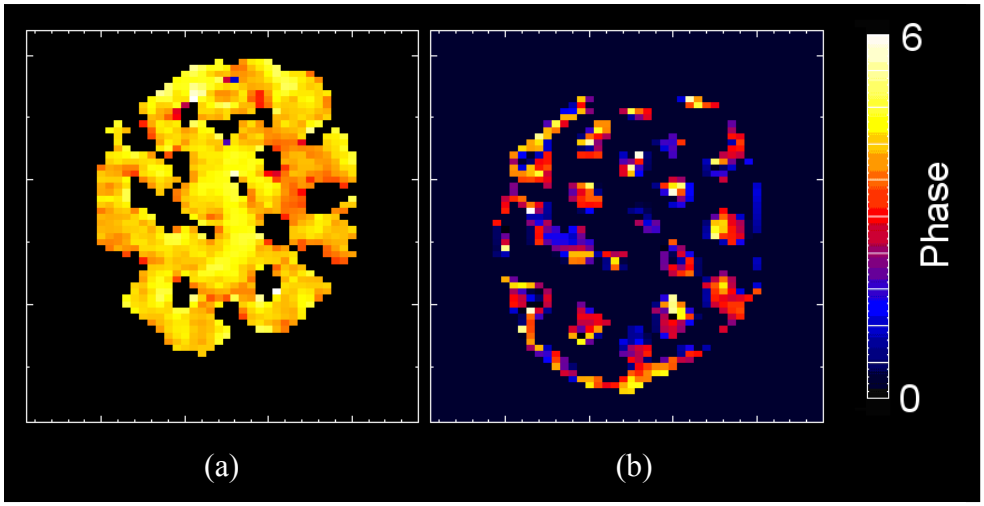


Figure 11.34 – GERVAIS data for the showerhead configuration, showing (a) data from regions of high signal strength, and (b) data from regions of low signal strength.

We have shown that the GERVAIS technique is capable of obtaining data from two-phase flow in the monolith system, however it struggles with regions of very high velocity and cannot provide quantitative velocity measurements. The showerhead system is difficult to analyse due to the simultaneous presence of regions of very high flow rate (>30 cm/s) and regions of low flow rate (< 5cm/s).

## 11.7.Conclusions

This chapter has applied the various MRI velocity techniques to two-phase flow in the various monolith reactor configurations. It has been found that the SEMI-RARE time-of-flight tracking technique could only be used on low-velocity (< 5 cm/s) and preferably high-liquid holdup systems. By examining the flow within individual channels of the monolith, a correlation was found between the local holdup in that channel and its upwards velocity. Co-current flow was found to have a better velocity distribution (in that channels were not limited to one mode of operation, as for the bubble sparger).

The line-excitation method was found to work well at low velocities, and has the advantage of being able to provide a rapid, visual fingerprint of the flow direction and magnitude for a row of channels. It was found to break down at higher velocities, and

high gas hold-ups. This was due to relaxation effects, and the increased possibility of air bubbles entering the excitation slice.

The COTTAGE technique was found to work extremely well in all monolith systems, regardless of flow rate. Velocity distributions were obtained for each process configuration, and either agreed with or were deemed to be superior to results from other techniques. The ability to measure dynamic changes in these challenging systems is something that other techniques simply cannot do.

MS-COTTAGE was found to perform poorly in the presence of two-phase flow, which degraded the signal-to-noise ratio. Possible improvements to the pulse sequence were suggested for further work.

GERVAIS was found to be only a qualitative method in the monolith, and struggled to cope with the very high velocities present in the showerhead device. The technique was not applied to other process configurations, because the other MRI techniques were thought to be more useful due to their potential to be quantitative.

For clarity and for easy reference, the findings of Chapters 8-11 will now be summarised in Chapter 12.

## 11.8. Nomenclature

| **Symbols** | **Definition** |
|---|---|
| $T_1$ | Spin-lattice relaxation constant |
| $T_2$ | Spin-spin relaxation constant |

| **Abbreviations** | **Definition** |
|---|---|
| APGSTE | Alternating pulsed-gradient, stimulated echo |
| Ca | Capillary number |
| COTTAGE | Continuous translation tracking with alternating gradient encode |
| cpsi | Cells per square inch |
| FOV | Field of view |
| GERVAIS | Gradient echo rapid velocity and acceleration imaging sequence |
| MRI | Magnetic resonance imaging |
| MS-COTTAGE | Multi-slice COTTAGE |
| NMR | Nuclear magnetic resonance |
| RARE | Rapid acquisition with relaxation enhancement |
| SEMI-RARE | Single-excitation, multiple image RARE |

# 11.9.References

Pinto, M. D. T., M. A. Abraham*, et al.* (1997). "How do bubbles enter a capillary?" Chemical Engineering Science **52**(11): 1685-&.

Sederman, A. J., M. D. Mantle*, et al.* (2003). "Quantitative 'real-time' imaging of multi-phase flow in ceramic monoliths." Magnetic Resonance Imaging **21**(3-4): 359-361.

# Chapter 12 – Summary of Experimental Findings

This chapter will summarise the experimental findings described over the previous four chapters (8-11). The limitations of the novel pulse sequences are stated, followed by a summary of the flow characteristics of the monolith reactor process configurations.

**CHAPTER 12 – SUMMARY OF EXPERIMENTAL FINDINGS**



## 12.1. Introduction

This chapter will summarise the performance and limitations of the novel MRI techniques described in Chapter 5, and applied to a range of monolith process configurations in Chapters 8-11. The second half of this chapter will then summarise the flow characteristics observed in these systems, and the implications for the design and operation of monolith reactors.

## 12.2. Novel MRI summary

This section will summarise the performance and quantitative nature of the various novel MRI techniques used in this investigation. The type of data which may be measured will be described, along with the limitations of each sequence.

### 12.2.1. Holdup techniques

Section 8.3 described two methods of measuring the instantaneous liquid holdup in the monolith reactor. The first was very rapid (2 ms) and used a simple, non-spatially-resolved FID acquisition to acquire a direct measurement of the spin density present within the RF coil. This was found to be very resistant to velocity attenuation and diffusive attenuation effects, and was used up to flow rates of 15 cm/s with little error (around a 5% underestimate in holdup at the highest flow rates).

The second technique used a standard RARE acquisition to acquire a 2D image in 150 ms. This method was more susceptible to velocity attenuation effects, and is only suitable for relatively low flow rate systems (up to 5 cm/s), however it has the advantage of spatially-resolving the holdup within individual channels.

### 12.2.2. RARE

RARE was used to obtain snapshots of the bubbles and slugs within the channels of the monolith reactor (Chapter 9). However, since the data took 150 ms to acquire, blurring became significant at all but the lowest velocities (up to 5 cm/s). In practice, the technique could only be applied to the low flow rate bubble sparger and co-current flow systems.

### 12.2.3. SEMI-RARE

The SEMI-RARE pulse sequence was used to acquire a series of images from a single excitation, and by tracking features from one image to the next we can extract local velocities. This was applied successfully to the bubble sparger, co-current flow and spray distributors (Section 11.2). It was not necessary to be able to identify slug

interfaces, but any "texture" could be tracked, as was shown for the spray system (where slugs are not resolved). The technique was even applicable to monolith channels containing no air bubbles – we may track the inflow of fresh spins into the channel (this is seen as the movement of a zero-intensity front). However, SEMI-RARE failed to produce sensible results in the high flow rate froth reactor system due to a very poor signal-to-noise ratio.

### 12.2.4. Line-excitation

The line-excitation technique was applied to the bubble sparger and co-current flow systems only (Section 11.3). This was because higher flow rate systems caused attenuation effects, and a loss of signal intensity. Also, systems of high gas holdup were unsuitable for study with this technique since this increased the likelihood of exciting a gas-filled channel, losing information on the spread of the spins.

The results of studying the bubble sparger compared excellently with the results of the COTTAGE technique, except at high gas holdups for the reasons explained above. For suitable systems, it allows us to obtain a rapid fingerprint of the flow and direction within a row of channels in only 150 ms. Also, by further processing of the line-excitation data we can say something about the propagator associated with each channel, and this method is presented later in Chapter 14.

### 12.2.5. COTTAGE

COTTAGE was applied to the co-current and froth reactor setups (Section 11.4). It was found to give sensible results in both systems, up to velocities of 20 cm/s. COTTAGE is more robust to the presence of higher gas holdups than the-line excitation technique, because it uses phase-encoding to measure the velocity rather than the spreading of spins. Therefore, a drop in signal strength is not so destructive.

Dynamic observations were also possible due to its high temporal resolution, allowing the visualisation of the change in direction of individual channels over short timescales (~100 ms). This meant it was possible to study observations such as that

of Pinto *et al.* (1997), who claim that there is a 200 ms velocity fluctuation as a bubble enters a capillary. As such, COTTAGE is one of the most powerful techniques for uniaxial velocity measurement, particularly in porous systems such as the monolith reactor.

### 12.2.6. MS-COTTAGE

The MS-COTTAGE technique was found to have signal-to-noise problems in the monolith reactor, and did not perform as well as had been hoped (Section 11.5). 2D velocity maps were acquired for the bubble sparger system, and the data were sensible but not as clear as using other methods. Further work would be necessary to improve the performance of this technique in porous media before meaningful results could be obtained.

## 12.3. Monolith flow characteristics

This section will now summarise the relative performance and flow characteristics of the various process configurations used in this investigation, addressing observed flow phenomena and operational behaviour, with an aim to optimisation of design for industrial applications.

### 12.3.1. Bubble sparger

The bubble sparger system showed liquid hold-ups of around 80-95%, decreasing very slightly with increasing air flow rate. The effect of channel density was investigated, and it was found that smaller channels gave a higher gas holdup (Section 8.3.1), presumably due to the larger surface tension effects which cause the gas to move more slowly, and hence give a higher gas holdup.

The bubble and slug size / frequency distribution was measured in this system (Section 9.5). It was found that the slug length was relatively uniform, with slugs of all sizes observed equally. At low flow rates, we see many small bubbles (most below 5 mm), and at higher air flow rates we see many longer bubbles.

Not all channels produced clean bubble-slug interfaces, even at these low velocities. Some channels were observed to have poorly-defined interfaces, and through close investigation this was deduced to be caused by increased velocity (30% higher) and increased gas holdup (5% higher) in these channels (Section 9.4.3). This prevented the formation of Taylor flow in these cases. Such channels constituted around 7-13% of the channels under observation, with their number increasing with gas flow rate.

Comparisons with the literature (Section 9.6) contradicted the capillary-based observations of Heiszwolf *et al.* (2001) and Kreutzer *et al.* (2001), who claimed that their slugs were around 2-5 times the diameter of the channel. Our observations show that there is no such peak in distribution for monolith channels. Other correlations of bubble length to slug length, and volume fractions as a function of flow rate, were either not applicable to the monolith system or could not be measured with our techniques.

The number of totally liquid-filled channels was seen to decrease with increasing air flow rate, indicating that more of the channels were becoming active (Section 9.7). This ranged from 55% of the channels to 40% at higher flow rates. Liquid-filled channels will not allow reaction to occur (due to the lack of a gas phase) and hence are undesirable for industrial operation.

The velocity distribution was observed with the SEMI-RARE and line-excitation techniques. SEMI-RARE measurements observed a tri-modal distribution, with some channels in upflow, some in downflow, and some stationary (Section 11.2.1). In higher density monoliths, a larger distribution was observed at zero flow, presumably due to the higher surface tension associated with smaller channels, which would therefore resist movement. Bubble velocities compared favourably with the results of Sederman *et al.* (2003). The line-excitation method observed a similar tri-modal distribution, although the number of channels observed in upflow was fewer; this is believed to be due to higher gas holdups in upflow, which mean that we have fewer measurements which give valid measurements, i.e. this is a factor of the measurement technique rather than representative of the system itself (Section 11.3.2).

Correlating local channel holdup with velocity (Section 11.2.2) showed two distinct flow regimes; liquid-filled channels were predominantly in downflow, whereas channels containing some air were predominantly in upflow, and the velocity associated with these channels was directly proportional to their gas holdup.

The dynamic nature of the bubble sparger was investigated by observing the operation of each channel as a function of time (Section 11.2.3). It was observed that almost all channels were stuck in their mode of operation, i.e. if a channel is currently in upflow, it will continue to operate in this manner. This is clearly an inefficient use of the monolith structure, and is in agreement with the literature that upflow is unstable compared to downflow (Section 3.8). Also, we can deduce that the bubble sparger system is highly dependent on startup conditions, and this was verified by startup under different conditions which gave different distributions (Section 11.2.4).

### 12.3.2. Co-current flow

The co-current flow configuration gave a liquid holdup between 70 and 95%, depending on the gas and liquid flow rates used (Section 8.3.1). Holdup decreased with increasing gas flow rate, and decreased slightly with increasing liquid flow rate.

The velocity distributions were investigated using the SEMI-RARE, line-excitation and COTTAGE techniques. SEMI-RARE measurements observed that at low liquid flow rates the distribution resembles that of the bubble sparger, but on increasing the liquid flow rate the distribution changed to a bimodal distribution with a peak at zero and another at higher velocity (Section 11.2.5). High gas and liquid flow rates were observed to give the best distributions, with fewer stationary channels and with a narrow velocity distribution (and therefore a narrow residence time).

The line-excitation method gave results which compared favourably with the SEMI-RARE technique, again showing the bimodal trend at higher liquid flow rates (Section 11.3.4). The results obtained with the COTTAGE technique again show the same trends, but are considerably more reliable than the other techniques, particularly at higher gas holdups.

Observing the operation of individual channels over time showed a big improvement over the bubble sparger system – channels were now free to change their mode of behaviour, and were no longer dependent on startup conditions (Section 11.2.5). This was believed to be due to the additional flow pressure, overcoming the resistance to flow within channels.

### 12.3.3. Showerhead distributors

The geometry parameters regarding the design of showerhead distributors were investigated, and their performance was rated by using a novel image processing technique for quantifying maldistribution (Section 10.2). It was found that homogeneity of the distribution worsened with increasing liquid flow rate. This was due to the sharpening of the liquid jets, and hence fluid was not distributed evenly. In contrast, at lower flow rates the jets were observed to oscillate and droplet-breakup resulted in better fluid distribution over the cross-section. Increasing the gas flow rate also worsened the fluid distribution, though the mechanism for this is less clear.

Increasing the height of the showerhead above the top of the monolith helped to improve the fluid distribution, and this was again attributed to increased droplet-breakup. Therefore, the optimum showerhead configuration is to have relatively low liquid flow rates, and a high distributor position, ensuring that drops are evenly dispersed.

### 12.3.4. Spray distributors

The liquid holdup for the spray distributor configuration was around 60-70%, decreasing slightly with increasing liquid flow rate (Section 8.3.1). This decrease (rather than the expected increase with liquid flow rate) was believed to be due to a higher hydrostatic head, which caused the channels to drain more quickly, reducing the liquid holdup. This is consistent with predictions of Section 3.8, and is taken as evidence of the increased stability of downflow operation, where the increasing liquid flow is counteracted by an increase in draining rate.

It was not possible to measure the slug sizes in spray operation, because the resolution of the RARE technique was not sufficient. The droplet nature of the spray caused the slugs to be far shorter than for the bubble sparger system, and hence this may indicate increased gas-liquid-solid contacting, improving mass-transfer effects.

The fluid distribution was observed to be far more homogeneous than for the showerhead distributors (Section 10.3). There was relatively little change in the quality of the distribution with increasing liquid flow rate, in stark contrast to the showerheads where increasing the flow rate slightly caused significant maldistribution. Spray distributors are therefore superior in performance to showerhead distributors, and are also more appropriate for industrially-relevant flow rates. However, they come with the disadvantage of increased pumping costs.

Tracking the velocity of the channels with the SEMI-RARE technique (Section 11.2.6) shows that the velocity distribution is narrow, and almost all channels are flowing. This is desirable since we will have a sharp residence time distribution, and good control over the reactor performance. Increasing the liquid flow rate shifted the distribution to higher velocities, and broadened it slightly. Investigating the operation of individual channels with time showed that the velocity distribution was very homogeneous (Section 11.2.6). These results confirm the observations from the literature (Section 3.9), which state that downflow configurations are more stable and easier to scale up than upflow.

### 12.3.5. Froth reactor

The measured liquid holdup for the froth reactor was in the range 45-70% (Section 8.3.1). The holdup decreased with increasing gas flow rate, and increasing the liquid flow rate had little effect on holdup.

The velocity within the froth reactor was investigated using the SEMI-RARE and COTTAGE techniques. The SEMI-RARE technique failed, due to the high velocities and poor signal-to-noise ratio associated with the system (Section 11.2.7). COTTAGE fared much better and was able to extract meaningful velocity data at all

flow rates, up to 20 cm/s (Section 11.4.3). COTTAGE was able to show that the velocity distribution across the monolith cross section was very homogeneous, and that the range of velocities present in the system were unimodal, and fairly narrow. Increasing the liquid flow rate shifted the velocity distribution to higher velocities, and broadened the peak. Increasing the gas flow rate broadened the distribution, without shifting it significantly. Similar results were observed for different orifice diameters. The optimum configuration was taken to be operation at the lowest possible gas flow rate (while keeping the turbulent mixture stable).

## 12.4. Nomenclature

| **Abbreviations** | **Definition** |
|---|---|
| COTTAGE | Continuous translation tracking with alternating gradient encode |
| cpsi | Cells per square inch |
| GERVAIS | Gradient echo rapid velocity and acceleration imaging sequence |
| MS-COTTAGE | Multi-slice COTTAGE |
| RARE | Rapid acquisition with relaxation enhancement |
| SEMI-RARE | Single excitation, multiple image RARE |

## 12.5. References


Heiszwolf, J. J., M. T. Kreutzer, *et al.* (2001). "Gas-liquid mass transfer of aqueous Taylor flow in monoliths." Catalysis Today **69**(1-4): 51-55.

Kreutzer, M. T., P. Du, *et al.* (2001). "Mass transfer characteristics of three-phase monolith reactors." Chemical Engineering Science **56**(21-22): 6015-6023.

Pinto, M. D. T., M. A. Abraham, *et al.* (1997). "How do bubbles enter a capillary?" Chemical Engineering Science **52**(11): 1685-&.

Sederman, A. J., M. D. Mantle, *et al.* (2003). "Quantitative 'real-time' imaging of multi-phase flow in ceramic monoliths." Magnetic Resonance Imaging **21**(3-4): 359-361.

# Chapter 13 – Deconvolution Techniques

This chapter will cover some of the basics behind deconvolution techniques, explain why they are useful and highlight some of the approaches that have been taken to overcome the inherent problems encountered in removing blurring and motion artefacts from data. The conclusions of this review will be applied to real MRI data in order to remove motion-blurring artefacts observed in MRI data, and the results are presented later in Chapter 14.

**CHAPTER 13 – DECONVOLUTION TECHNIQUES**



## 13.1. Introduction

From home photography to world-class astronomy, many images will inevitably contain motion-blurring or out-of-focus artefacts. Fortunately, it is usually possible to

remove or reduce motion-blurs, out-of-focus blurs and optical defects using deconvolution techniques. The goal of the deconvolution process is to reconstruct the original scene from a degraded observation.

For the case of a flowing system being studied with MRI, we find that as a result of motion during the acquisition time, images may appear blurred. This is because the intensity originating from a group of spins is spread over the observed pixels, and is analogous to motion-blurring experienced in photography. This type of artefact is usually considered to be one-dimensional and linear, and can be deconvolved using the methods outlined in this chapter. In application to the line-excitation profiles acquired in the monolith system (described in Section 5.3), we require the deconvolution of the initial excitation shape and also the deconvolution of motion-blurring to obtain a true measurement of the displacement of the spins over time, which is analogous to a propagator measurement. This procedure is outlined in Chapter 14, using some of the techniques described in this chapter.

This chapter will review literature on deconvolution techniques and highlight their advantages and disadvantages. First we will cover direct deconvolution techniques where the blurring function is known (or may be estimated accurately), and then move onto so-called blind deconvolution techniques, which can deconvolve blurring without *a priori* knowledge of the blurring function.

The final section will test the most promising of the deconvolution techniques, to investigate their performance on simulated data in the presence of noise, in preparation to their application to real data in Chapter 14.

## 13.2. Deconvolution

In many imaging applications the observed image is composed of the true image convolved with a blurring function, also known as the transfer function, optical transfer function or point-spread function (PSF). A PSF-based blurring function may arise from factors such as out-of-focus optics, motion, or insufficient bandwidth in electronics. Some example PSFs for different types of blur are illustrated in Figure

13.1. If we consider the convolution of an image of pixel dimension $N \times M$ with a linear motion-blur in image space, then the observed image intensity $y(n,m)$ is given by Eq. 13.1.

$$y(n,m) = \sum_{l=0}^{L} c_l x(n, m-l) \qquad \textbf{(Eq. 13.1)}$$

where $x(n,m)$ is the intensity of the original image and $c_l$ ($l = 0,1 \ldots L$) defines the PSF with L being the extent of motion-blur.

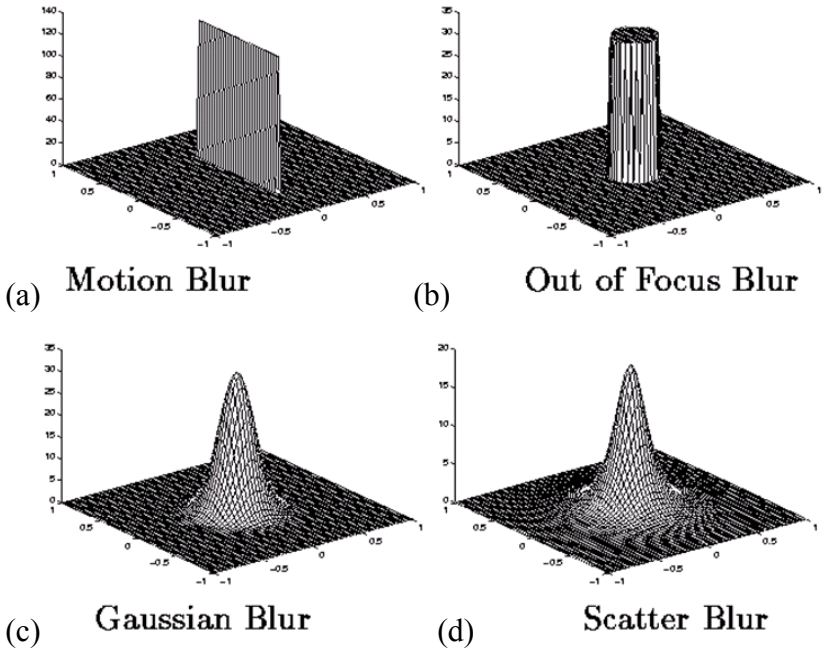


Figure 13.1 – A selection of different blurring artefacts encountered in image processing; (a) linear motion-blur; (b) out-of-focus blur; (c) Gaussian blur and (d) scatter blur. Reproduced from http://adassxiv.ipac.caltech.edu/Presentations/Monday/Chan_O1_1.ppt

Deconvolution is the process used to reverse the effects of convolution on recorded data. The concept of deconvolution is widely used in the techniques of signal processing and image processing, and an outline of the main approaches is presented in this chapter. These techniques were originally developed with application to signal processing, photography, cleaning astronomical images or remote sensing, but they are equally applicable to MRI acquisitions. The foundations for deconvolution were laid by Norbert Weiner of the Massachusetts Institute of Technology in his book (Weiner, 1949).

In general, rather than working in real space, it is more convenient to work with convolutions in terms of operations in Fourier space, as shown in Eq. 13.2.

$$g(x,y) = f(x,y) * h(x,y), \text{ and } G(u,v) = F(u,v) \times H(u,v) \qquad \textbf{(Eq. 13.2)}$$

where $g$ is the recorded signal, and $f$ is the signal that we wish to recover, but has been convolved with some other signal $h$ during the recording process. The function $h$ typically represents the PSF of an instrument, but can be due to other influences such as focussing or motion-blurring, depending on the observed system. $G$, $F$ and $H$ are the respective Fourier transforms of each signal. As can be seen they are simply multiplied in Fourier space, rather than attempting a complicated convolution in real space. The problem of recovering the true image $f(x,y)$ requires the deconvolution of $h(x,y)$ from the degraded image. If we know or can estimate $h$, then we may recover the signal by direct deconvolution, which translates to a simple division in Fourier space.

For example, we may measure the PSF directly by selecting a region of an image containing intensity originating from a point source – the actual recorded image of this point will give a direct measure of the PSF. However, the point-like objects used to measure the PSF in this way have to be small and thus tend to be dim. This approach is only really practicable in astronomy situations where stars appear as small, bright points as seen from the earth. If however, we cannot know or measure $h$ in advance, this case is termed "blind deconvolution" due to the lack of information on the system. To obtain a solution, we typically need to use statistical methods to estimate the form of $h$. These blind methods are discussed later in Section 13.4.

In real, physical systems, the situation is never as precise as Eq. 13.2, and we always encounter measurement noise of some description, which contributes as shown in Eq. 13.3.

$$g(x,y) = f(x,y) * h(x,y) + \varepsilon \qquad \textbf{(Eq. 13.3)}$$

where $\varepsilon$ is additive noise that has contaminated our recorded signal. If we were to assume that a noisy signal is noiseless when we try to make a statistical estimate of $h$, our estimate will be incorrect. In turn, our estimate of $f$ will also be incorrect. The worse the signal-to-noise ratio, the poorer our estimate of the deconvolved signal will

be. However, if we have at least some knowledge of the type of noise in the data (for example that it is white or Gaussian noise), or if we have some statistical measure of the noise, we may be able to improve the estimate of *f*.

Due to the presence of noise in particular, a mathematically optimum procedure for deconvolution does not exist, nor is there a proof that one exists in all cases; it is an imperfect science. The practical techniques that are used have developed under a variety of assumptions and apply to many real situations, producing impressive improvements in image quality under many applications. Of course, it is always better to try to prevent the imperfections in the imaging system as far as possible, since removal of the error by deconvolution will never be as satisfactory.

## 13.3.Direct deconvolution methods

Direct deconvolution methods are only possible where the PSF is known or can be estimated accurately *a priori*. These techniques include the Fourier or Laplace transform, and more sophisticated methods such as Weiner filtering or the Lucy-Richardson algorithm, as well as many others. These methods will be outlined in this section.

### 13.3.1. Fourier and Laplace transforms

The Fourier transform takes advantage of the fact that any one-dimensional function, $f(x)$, can be expressed as a summation of a series of sine and cosine terms of increasing frequency. The Fourier transform, $F(u)$, describes the amount of each frequency term that must be added together to reconstruct the original function. It can be written as Eq. 13.4.

$$F(u) = \int_{-\infty}^{\infty} f(x) e^{-2\pi i u x} dx \qquad \textbf{(Eq. 13.4)}$$

and importantly, it is possible to recover the spatial domain function in the same way, as shown in Eq. 13.5.

$$f(x) = \int_{-\infty}^{\infty} F(u) e^{2\pi u i x} du \qquad \textbf{(Eq. 13.5)}$$

The discrete form of these equations is given by Eq. 13.6.

$$F(u) = \frac{1}{N} \sum_{x=0}^{N-1} f(x) e^{-\frac{i 2\pi u x}{N}} \text{, and } f(x) = \sum_{u=0}^{N-1} F(u) e^{\frac{i 2\pi u x}{N}} \qquad \textbf{(Eq. 13.6)}$$

The dimensions of the discrete Fourier transform data are the same as the original image, but each pixel now contains a real and an imaginary part, hence we require twice the storage capacity. The dimensions of the discrete Fourier transform are no longer infinite, since we are limited by the Nyquist theorem – the highest frequency we may describe is that given by two neighbouring pixel pairs, hence the number of Fourier frequency terms, and hence the dimension (and resolution) of Fourier space, is the same as for the original image.

The Fourier deconvolution methods make use of the equality in Eq. 13.7.

$$F(g) = F(f * h) = F(f) \times F(h) \qquad \textbf{(Eq. 13.7)}$$

where $F(\cdot)$ denotes a Fourier transform. As can be seen, in Fourier space (and similarly in Laplace space) the transforms of the convolved terms are simply multiplied together, and therefore by dividing the transform of the obtained image by the known PSF we may directly obtain the original signal by Eq. 13.8.

$$f = F^{-1}\left[\frac{F(g)}{F(h)}\right] \qquad \textbf{(Eq. 13.8)}$$

The Fourier transform method is simple and straightforward to implement with NMR data, since operations in Fourier space are commonplace. Where used, the Fast Fourier Transform (FFT) is implemented, based on the algorithm of Cooley and Turkey (1965).

However, in dividing one function by another in Fourier space, we have the problem of dividing by very small numbers (or zero) at the edge of Fourier space (corresponding to high frequencies in Fourier space, and hence small features in the real image). Therefore we inevitably have the problem of noise magnification, and this will dominate the result if the noise level is too great (see Section 13.6.1 for an

illustration). Random pixel noise (also referred to as speckle) must be minimised – noise in either the acquired image or particularly in the PSF is significantly amplified. The quantitative measure of the noise level is defined by the signal-to-noise ratio (SNR) and is defined in terms of the standard deviation of the signal and that of the noise (which is usually estimated from the baseline of the data) as shown in Eq. 13.9.

$$\mathrm{SNR[dB]} = 10\log_{10}\left(\frac{\sigma_{image}}{\sigma_{noise}}\right) \qquad \textbf{(Eq. 13.9)}$$

In general, if the SNR is in the range 40-50 dB then the noise is practically invisible. Problems with deconvolution will be experienced where the SNR falls to around 10 or 20 dB (Russ, 2002).

Therefore, as has been stated in the literature (Russ, 2002), deconvolution methods are not generally used without some form of smoothing or filtering function to suppress the noise magnification. This process is often referred to as an apodisation or tapering operation, whereby we bring the edges of Fourier space smoothly to zero. Smoothing functions have the effect of suppressing high-frequency, short scale features from the image, thereby blurring it. We effectively suppress the noise variations in the image at the expense of sharpness of details. We must therefore seek to find a reasonable trade off between signal restoration and noise magnification.

Smoothing functions have been covered previously in Section 7.2.3, where it was stated that the Gaussian is probably the most popular smoothing method, and is used in many image processing applications. A single smoothing pass is generally sufficient to remove most of the noise-related errors. We can control the degree of smoothing by adjusting our choice of smoothing function and its parameters.

We may need to iterate our deconvolution since the presence of noise makes the deconvolution imprecise and non-unique. There are a number of linear and non-linear algorithms to obtain deconvolutions in the presence of noise – such as Weiner filtering and the Lucy-Richardson algorithm. These will now be examined in more detail.

## 13.3.2. Weiner filtering

Weiner filtering was first proposed by Weiner (1949) and further reading on its implementation is available in Press *et al.* (1986) and Russ (2002). Weiner filtering attempts to overcome the effects of noise magnification during the deconvolution process. As we have already seen, the basic inverse filter for deconvolution is given by Eq. 13.10.

$$F(u,v) \approx \left[\frac{1}{H(u,v)}\right] G(u,v) \qquad \textbf{(Eq. 13.10)}$$

With Weiner filtering, an adjustable scalar parameter is added to increase the denominator, as shown in Eq. 13.11.

$$F(u,v) \approx \left[\frac{1}{H(u,v)}\right]\left[\frac{|H(u,v)|^2}{|H(u,v)|^2 + K}\right] G(u,v) \qquad \textbf{(Eq. 13.11)}$$

Therefore we obtain Eq. 13.12.

$$F(u,v) \approx \left[\frac{H^*(u,v)}{|H(u,v)|^2 + K}\right] G(u,v) \text{, since } |H(u,v)|^2 = H^*(u,v)H(u,v) \qquad \textbf{(Eq. 13.12)}$$

The constant $K$ can overcome the singularity that may exist in the above equation when $H(u,v)$ has zero-value points. Theoretically, the factor $K$ is dependent on the statistical properties of the image and their relative noise, defined as Eq. 13.13.

$$K \approx K(u,v) = \frac{W_{nn}(u,v)}{W_{xx}(u,v)} \qquad \textbf{(Eq. 13.13)}$$

where $W_{nn}$ and $W_{xx}$ are the power spectra of the noise and the object respectively. The power spectrum of an image is the square of the magnitude of the real part of the Fourier transform of the original image. We may obtain an estimate for $K$ if we can identify a region which is totally filled with noise pixels, and another which is largely composed of image pixels.

In practice however, $W_{nn}$ and $W_{xx}$ cannot usually be measured precisely (due to the inability to separate noise pixels from image pixels completely) and so $K$ is typically treated as an adjustable parameter which controls the trade-off between sharpening and noise. As a rule of thumb, Gonzales and Woods (1992) suggest taking $K = 2\sigma^2$

for the case of Gaussian noise, where $\sigma^2$ is the variance of the noise. They claim that this generally yields good results. Varying the value of $K$ will adjust the amount of filtering (i.e. blurring) that is encountered. Since $K$ is not always easily determined *a priori* it may be necessary to repeat the calculation with different values of $K$ and select the optimum restored image. Ranges of optimum $K$ in the literature typically vary between 0.001 and 0.1 depending on the system under investigation. Note that when $K = 0$, this corresponds to no smoothing effect, and the algorithm simplifies to the familiar case of $F(u,v) = G(u,v) / H(u,v)$.

### 13.3.3. Incremental Weiner filter

In the presence of noise, the deconvolution cannot be performed exactly, and so an iterative method was proposed. The incremental Weiner filter has an iterative form, and in the presence of noise it can further improve the solution in comparison to the conventional Weiner filter. In the incremental Weiner filter described in Mouyan and Unbehauen (1995), we use an error array as in Eq. 13.14.

$$S(u,v) = G(u,v) - F(u,v)H(u,v) \qquad \textbf{(Eq. 13.14)}$$

and the estimation given by the filter is shown in Eq. 13.15.

$$F_{new}(u,v) = F_{old}(u,v) + \frac{H^*(u,v)S(u,v)}{|H(u,v)|^2 + K} \qquad \textbf{(Eq. 13.15)}$$

This iterative scheme has the properties of reducing the error with each step. This can be seen by denoting $S_1 = G - HF_{old}$ and $S_2 = G - HF_{new}$. We then obtain Eq. 13.16.

$$S_2 = \frac{K}{|H|^2 + K} S_1 \qquad \textbf{(Eq. 13.16)}$$

Because $K > 0$ and $|H|^2 > 0$, then for all $u,v$ we have $S_2 \leq S_1$, and hence the error always reduces. In the noiseless case, $S_2$ will approach zero.

The larger the value of $K$, the more stable the algorithm will be – however this is at the expense of the quality of the deconvolution, which will be more greatly smoothed with higher values of $K$.

## 13.3.4. Lucy-Richardson

This was originally derived from Bayes' theorem in the early 1970s by Richardson and Lucy (Richardson, 1972; Lucy, 1974). The reason for its popularity is its implementation of a maximum likelihood technique and its apparent ability to produce reconstructed images of good quality in the presence of high noise levels. Bayes' theorem describes conditional probability, and is given by Eq 13.17.

$$P(x|y) = \frac{P(y|x)P(x)}{\int P(y|x)P(x)dx} \qquad \textbf{(Eq. 13.17)}$$

The probability $P(x)$ can be identified as the object distribution $f(x)$; the conditional probability $P(y|x)$ can be identified as the PSF centred at $x$, i.e. $h(y,x)$; and the probability $P(y)$ can be identified as the degraded image or convolution, $g(y)$.

This inverse relation permits the derivation of the iterative algorithm of Eq. 13.18.

$$f_{i+1}(x) = \frac{\int h(y,x)g(y)dy}{\int h(y,z)f_i(z)dz} f_i(x) \qquad \textbf{(Eq. 13.18)}$$

where $z$ is a dummy variable and $i$ is the iteration number.

We can rewrite this in terms of convolutions, as shown in Eq. 13.19.

$$f_{i+1}(x,y) = \left\{ \left[ \frac{g(x,y)}{f_i(x,y) * h(x,y)} \right] * h^*(-x,-y) \right\} f_i(x,y) \qquad \textbf{(Eq. 13.19)}$$

If the PSF $h(x)$ is known or may be estimated, then we may find $f(x)$ by iterating the above equation from an initial guess $f_0(x)$. Large deviations of the initial guess from the true solution are lost rapidly in the early iterations, while small details are added more slowly in subsequent iterations. We may also include a non-negativity constraint.

The Lucy-Richardson algorithm claims to be robust in the presence of noise, and so by incorporating it into a deconvolution algorithm one would expect the results to improve. It is easily integrated into the IBD algorithm (Section 13.4.1), using it in place of the standard deconvolution steps. Since Bayes' theorem relates conditional

probabilities, the algorithm takes into account statistical fluctuations in the image, thereby giving the Lucy-Richardson algorithm its ability to reconstruct noisy images. Fish *et al.* (1995) report that it performs well up to levels of around 10% noise (where the noise is defined by the standard deviation divided by the highest intensity in the image). In comparison to similar algorithms using the Weiner filter, Fish *et al.* (1995) report that the Lucy-Richardson algorithm is far more resistant to noise, with the Weiner filter performing poorly at around 0.1% noise.

Of course, the literature is not in total agreement – van Kempen *et al.* (1996) state that the Lucy-Richardson algorithm is very sensitive to noise. It is possible that, depending on what is regarded as a reasonable noise level, the algorithm may not be robust enough.

## 13.4. Blind deconvolution

In most deconvolution situations the PSF is assumed to be known explicitly – for example the PSF may be a function of the imaging equipment, in which case it can be measured or estimated and then deconvolved from every image acquired with that equipment. However, there are many cases when the PSF cannot be measured. It may be impossible or perhaps dangerous to measure, or the PSF may constantly change with every acquisition, in which case it is impossible to measure the corresponding PSF. For example, in the case of imaging astronomical objects through the atmosphere, the distortion due to turbulence cannot practically be determined or predicted. Moreover, if we have equipment where the PSF may change significantly with time, it may be better to estimate the PSF every time rather than use an incorrect, outdated measure, since if it is inexact it will degrade the obtainable resolution of the equipment.

When we face the problem of not knowing the PSF of the motion-blurring *a priori*, this leaves the conundrum of separating the true image from the PSF, both of which are unknowns. At first glance, it appears that we do not have enough information to do this unambiguously. However, there exist a range of "blind deconvolution" techniques, which may assist us by incorporating partial information about the

system. The general concept of blind deconvolution dates back to pioneering works by Oppenheim *et al.* (1968) and Stockham *et al.* (1975).

Blind deconvolution is critical in many fields such as astronomy (Bertero and Boccacci, 2000), remote sensing (Muller, 1988), biological and medical imaging (Adam and Michailovich, 2002), microscopy (Wilson and Hewlett, 1991), *etc*. Without good quality images, quantitative analysis becomes very difficult and so images are often treated prior to analysis. The most widely-known example of deconvolution is the defect in the Hubble Space Telescope's optical system. It had an incorrectly manufactured main mirror which produced poorly-focussed images. Eventually a compensating optical device was installed, but even with the original defect it was possible to obtain high-resolution results by deconvolution with the known PSF of the mirror.

Some important characteristics of the problem of blind deconvolution are:

- The true image and the PSF must be irreducible for unambiguous deconvolution – an irreducible signal is one which cannot be exactly expressed as the convolution of two or more component signals. This is important, since if either the true image or the PSF are reducible, we effectively have three or more signals convolved together – ambiguity arises in deciding which components belong to the true image and to the PSF, and hence the solution becomes non-unique.
- Ideally we would obtain the result of Eq. 13.20.

$$\tilde{f}(x,y) = f(x,y) \qquad \textbf{(Eq. 13.20)}$$

where $\tilde{f}(x,y)$ is the estimated image and $f(x,y)$ is the true image. However, we typically may only obtain a scaled, shifted estimate, such as in Eq. 13.21.

$$\tilde{f}(x,y) = k \times f(x - d_x, y - d_y) \qquad \textbf{(Eq. 13.21)}$$

where $k$, $d_x$ and $d_y$ cannot be determined by blind deconvolution (Lane and Bates, 1987). We may estimate the shift and scale values after the deconvolution by using additional constraints, such as specifying that the mean value of the image is preserved, and we may also use information about

the location of the "support" or extent of the image. The support of the image is defined as the smallest rectangle that can be drawn around the object of interest, and in general the support must be fully contained within the image.

- The case of additive noise is generally assumed to follow the following linear model of Eq. 13.22.

$$g(x,y) = f(x,y) * h(x,y) + n(x,y) \qquad \textbf{(Eq. 13.22)}$$

where $n(x,y)$ is the term due to additive noise, which may be of uniform, Gaussian or Poission distribution. The noise term is important because in practical imaging situations additive noise is not negligible. Common types of noise include electronic noise (in imaging equipment), photoelectric noise due to the statistical nature of light and the capture process (i.e. signal to noise ratio) and quantisation noise resulting from signal digitisation.

Some practical constraints are summarised by Kundur and Hatzinakos (1996):

- The image restoration problem by blind deconvolution is in general ill-conditioned – i.e. a small perturbation of the data may produce large deviations in the resulting solution. The problem is often ill-conditioned owing to the presence of additive noise; during the process of inversion, excessive amplification of noise results at the high frequencies.
- Exact deconvolution is impossible in the presence of noise, since we have only a statistical measure of the noise level. In general this makes $g(x,y)$ irreducible, thus only an approximate deconvolution can be performed.
- The solution may not be unique – since only partial information is used in the optimality criterion, there may be many solutions which have similar likelihood measures but different estimates of the true image and PSF.
- Depending on the application of interest, the blind deconvolution algorithm may be tailored to suit our needs – i.e. in real time imaging we may sacrifice some of the accuracy for reduced computational complexity and convergence speed, whereas in medical imaging the image quality is the priority, and computational time is less important. In general, the different methods and algorithms proposed in the literature have a range of trade-offs of these features.

Blind deconvolution techniques may be separated into two categories:

- identifying the PSF separately from the true image, in order to obtain the true image by classic direct deconvolution techniques. The estimation of the PSF and true image are therefore disjointed procedures.
- Incorporating the identification procedure with the restoration algorithm. Thus we are simultaneously identifying the PSF and the true image at each iteration, and this leads to the development of more complex algorithms.

The ill-posedness of the algorithm, and its sensitivity to noise can be controlled to some extent by regularisation. Regularisation is the name given to any process which improves the ill-posedness of an algorithm. This may be done by reformulating the algorithm, approximating it, or adding additional constraints to ensure convergence. Any approximations must be carefully chosen to ensure that a physically acceptable solution can be obtained while retaining stable convergence from a computational point of view. Regularisation methods were first introduced by Tikhonov and Arsenin (1977) and enthusiastic readers are directed here for more information on the formulation and application of these methods. Mouyan and Unbehauen (1995) state that the discretisation of a continuous deconvolution problem can also alleviate its ill-posedness. Fortunately, our data is very much discretised – typically we have 1D profiles of 64 or 128 pixels.

The following sections will cover some of the main approaches to blind deconvolution, and highlight their applicability with respect to their ability to restore signals satisfactorily, particularly in the presence of noise.

### 13.4.1. Iterative Blind Deconvolution

One of the earliest blind deconvolutions was first suggested by Ayers and Dainty (1988). This is generally considered to be the simplest blind deconvolution algorithm, and is the benchmark to which all later efforts are compared. The flow of the algorithm is illustrated in Figure 13.2. First, we make an initial guess ($\tilde{f}_0$) of the true function, which obeys our constraints and approximates the expected shape. Then we Fourier transform this function, and use it to calculate what the corresponding PSF

(*H*) would be. This is then Fourier transformed and altered such that it obeys the constraints (typically non-negativity). Then we use this to calculate an estimate of the Fourier-transformed true function (*F*), and the true function (*f*). This takes us full-circle, and the iteration then continues until convergence. In general, the obtained solution is not affected by the choice of the initial estimate, although an appropriate initial guess will of course converge faster.

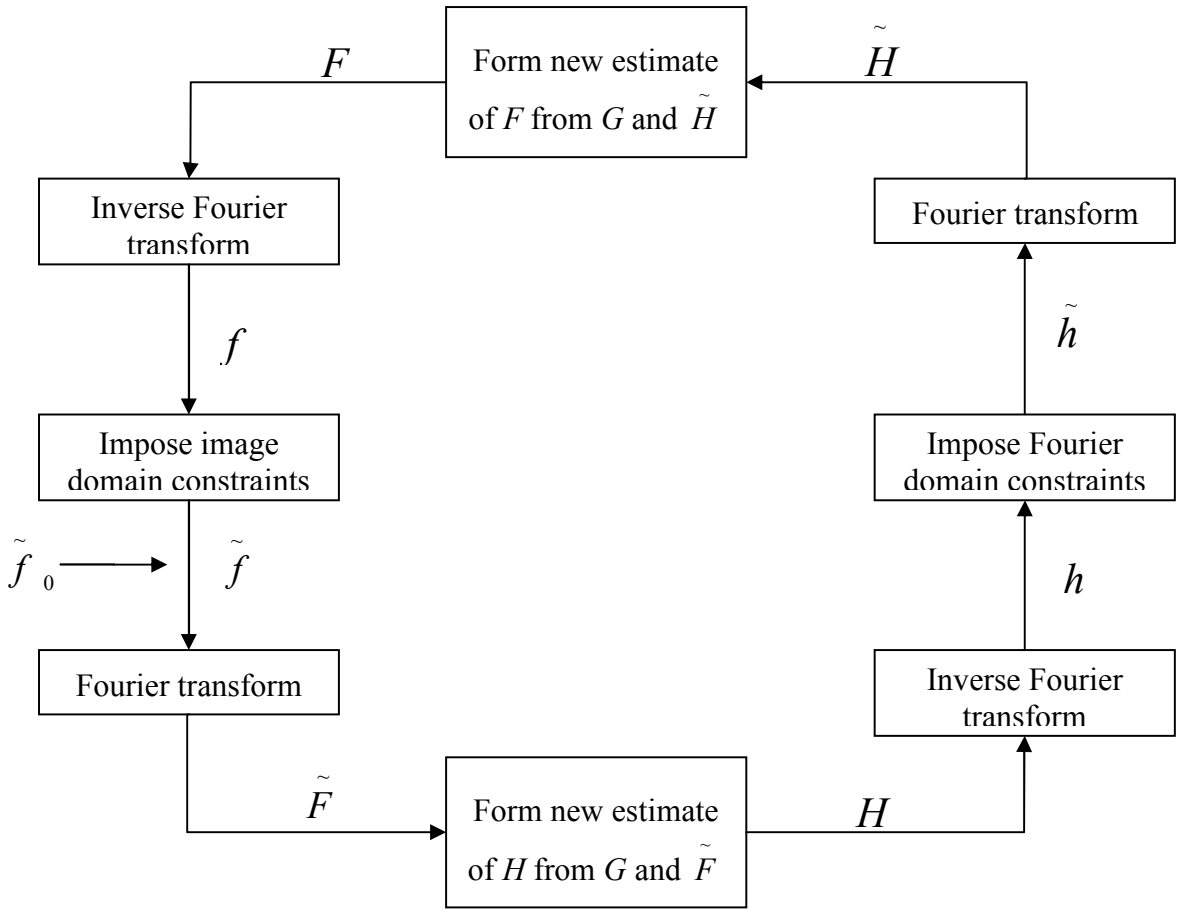


Figure 13.2 – Illustration of the Iterative Blind Deconvolution algorithm.

In their proposed method, not only are the images constrained to be non-negative, but also energy is conserved at each iteration – this was achieved by uniformly redistributing the sum of any negative values in the function over the entire function. This method of energy conservation resulted in improved convergence behaviour.

However, the authors identify two main problems with the technique:

- The inverse transform has associated problems because of the need to invert a function that possesses regions of relatively low value.
- Zeros at particular spatial frequencies in either of the functions result in no information at that spatial frequency being present in the convolution.

The points above must be addressed in order to impose the Fourier domain constraint. At its most basic level, this constraint simply states that the Fourier product of the two function estimates must be equal to the convolution function.

Ayers and Dainty (1988) propose a method to over come these problems; they point out that at each step of the iteration we have two estimates for each Fourier spectrum, i.e. the function $\tilde{F}(u)$ and the estimate $G(u)/\tilde{H}(u)$, which is obtained on imposing the Fourier domain constraint. Both of these estimates have properties in common with the desired deconvolved solution – the former has a non-negative inverse transform, and the latter obviously satisfies the Fourier domain constraint. Therefore at each iteration, the two estimates may be averaged together to obtain a new, composite estimate via a weighting function (Eq. 13.23). This function is not necessary for convergence, and indeed the rate of convergence is dependent on the weighting parameter, $\beta$, whose optimum value is not predictable *a priori*.

$$F_{i+1}(u) = (1-\beta)\tilde{F}_i(u) + \beta\frac{G(u)}{\tilde{H}_i(u)} \qquad \textbf{(Eq. 13.23)}$$

where $0 \leq \beta \leq 1$. The value of the constant $\beta$ is set before the algorithm is run, and denotes a weighting towards one or other of the estimate functions.

Small, confined regions of low or zero value present in the convolution are dealt with by using only the estimate $\tilde{F}(u)$, thus the estimate obtained by imposing the Fourier constraint contributes no new information to the new estimate. This therefore reduces the regions which would otherwise be massively amplified by noise, e.g. Eq. 13.24.

$$\text{if } |H(u)| < \text{noise level, } F_{i+1}(u) = \tilde{F}_i(u) \qquad \textbf{(Eq. 13.24)}$$

Next, the problem of $\tilde{H}(u)$ [or equivalently $\tilde{F}(u)$] having a small modulus – if this is less than the modulus of the convolution spectrum, then instead of performing the linear averaging previously described in Eq. 13.23, the inverses of the two function estimates are averaged as in Eq. 13.25.

$$\frac{1}{F_{i+1}(u)} = \frac{(1-\beta)}{\tilde{F}_i(u)} + \beta \frac{\tilde{H}_i(u)}{G(u)} \qquad \textbf{(Eq. 13.25)}$$

The reasoning behind this type of inverted averaging is that large values are prevented from dominating the image, which is desirable since they correspond to noise which has been amplified during the deconvolution / Fourier transform process. An alternative to this method is to employ one of the noise-resistant direct deconvolution techniques such as the iterative Weiner or Lucy-Richardson filters to alleviate the problem of noise amplification, and division by zero if the PSF has points of zero value.

Shen and Zhang (2000) comment that while the IBD algorithm is computationally simple, it lacks reliability in its performance – it is sensitive to the initial guess, and the algorithm is sometimes unstable. In addition, the uniqueness and convergence properties of the IBD are still unknown in terms of a rigorous mathematical proof, despite extensive work in the area. This lack of theoretical understanding has hindered its further maturation into a widely applicable algorithm in practice.

## 13.4.2. Error metrics

As a side note, whenever we have an iterative process, it is useful to be able to measure the progress of the iteration, and to quantify the quality of the resultant solution. For this purpose, there are various "error metrics" which may be formulated (Mouyan and Unbehauen, 1995).

Error metrics are, at their most basic level, a measure of the error in the iteration process and can therefore be used as a measure of convergent progress. For example, at the most basic level we may have an error metric such as Eq. 13.26.

$$E_1 = E[(m * y - x_0)^2] \qquad \textbf{(Eq. 13.26)}$$

where $E[\cdot\,]$ denotes the expectation.

Error metrics can be used along with maximum likelihood techniques to guide the iteration in a direction so as to efficiently minimise the error, and therefore speed up the convergence process. Kundur and Hatzinakos (1996) quote the percentage mean squared error in a deconvolution as Eq. 13.27.

$$\mathrm{MSE}(\hat{f}) \triangleq 100 \frac{\sum_{\forall(x,y)} \left[ a\hat{f}(x,y) - f(x,y) \right]^2}{\sum_{\forall(x,y)} f^2(x,y)} \qquad \textbf{(Eq. 13.27)}$$

And because the scale of the obtained image is not important, we choose $a$ such that the MSE is minimised, as shown in Eq. 13.28.

$$a = \frac{\sum_{\forall(x,y)} \hat{f}(x,y) f(x,y)}{\sum_{\forall(x,y)} f^2(x,y)} \qquad \textbf{(Eq. 13.28)}$$

## 13.4.3. Maximum entropy / likelihood method.

This method is another popular form of deconvolution, and it stems from research in information theory. Maximum entropy methods (MEMs) function by minimising a smoothness function, calculated from the entropy in an image. The solution selected by this method is that which fits the data (to within the noise level) and also has maximum entropy. The "entropy" of the solution is a measure which, when maximised, produces a positive image with a compressed range in pixel values. This "image entropy" measure should not be confused with "physical entropy". A typical entropy measure is given by Eq. 13.29.

$$\varepsilon = -\sum_k I_k \ln \frac{I_k}{M_k e} \qquad \textbf{(Eq. 13.29)}$$

where $M_k$ is an estimate for the image which incorporates *a priori* knowledge about the object. If the distribution is unknown, a constant baseline may be used. An

iterative procedure must then be performed to generate a solution which has the maximum entropy possible.

The resolution of this method is limited by the signal-to-noise ratio. These methods have not been implemented as a part of this work, but more information on MEM methods can be found in Gull and Daniell (1978), Katsaggelos and Lay (1991), Cornwell and Bridle (1996), and Lam and Goodman (2000).

### 13.4.4. Parametric blind deconvolution

Other techniques where certain properties of the PSF are known in advance may be solved by using parametric deconvolution techniques. For example, if we know that the PSF is a blurring or a linear motion-blur function, then we know the basic shape of the PSF (see Section 13.2). Note that in general, non-linear (i.e. accelerating) blurs cannot be removed – they will be approximated as linear and the error encompassed into the noise term (Jefferies and Christou, 1993). In general, the monolith system can be assumed to be in pseudo steady state over the acquisition time, so the approximation of a linear blur is a reasonable one. The parameters of the PSF can be fitted by iteration, using a maximum entropy technique to find the optimum result. In image processing, we also know that the value of the real function must be non-negative everywhere, and so this is another constraint to the deconvolution.

In this approach, we attempt to describe the PSF using various parameters, and to then find their optimum values by using a maximum likelihood estimator built into our algorithm. Providing that our PSF can be represented in this way, we improve the convergence properties by further constraining the problem. This is sometimes referred to as semi-blind deconvolution, because we are using prior knowledge to aid the deconvolution iterations. For example, we may know that we have a Gaussian blurring function, or a linear motion-blur. In the case of motion-blur, we may estimate the PSF as Eq. 13.30.

$$\mathrm{PSF}(u,v) = \frac{\sin\left(\frac{u}{2}\cdot V_x \cdot T + \frac{v}{2}\cdot V_y \cdot T\right)}{\frac{u}{2}\cdot V_x \cdot T + \frac{v}{2}\cdot V_y \cdot T} \qquad \textbf{(Eq. 13.30)}$$

where $V_x$ and $V_y$ are the components of linear velocity in the $x$ and $y$ directions and $T$ is the total exposure time. For the case of defocused blur, we may use the form of Eq. 13.31.

$$\mathrm{PSF}(u,v) = \frac{J_1\left(a\cdot\sqrt{u^2+y^2}\right)}{a\cdot\sqrt{u^2+y^2}} \qquad \textbf{(Eq. 13.31)}$$

where $J_1$ is a first order Bessel function, and $a$ is a scaling factor proportional to the diameter of the optical system aperture.

The way an iterative parametric algorithm would be implemented is to perform a certain number of iterations, then fit parameters to this result. Then the new parametric model is input back into the iteration and repeated until converged. (Fish *et al.*, 1995). In comparison with the simpler non-parametric methods, Fish *et al.* (1995) reported that a Lucy-Richardson based parametric IBD algorithm performed reasonably at up to 20% SNR, compared to 10% for the non-parametric implementation.

Markham and Conchello (1999) presented a parametric method for the separation of the real image and a blurring PSF. They derive a maximum-likelihood method for the PSF, in which they assume a mathematical model for the PSF that depends on a relatively small number of parameters (twenty) rather than the thousands of pixel values we would otherwise have to deal with. They claim that the method is robust to aberrations in the PSF and to noise in the image.

## 13.4.5. Simulated annealing algorithm

Like the iterative blind deconvolution of Ayers and Dainty (1988), this is a very general algorithm, requiring little *a priori* information on the system. It also requires the true image and PSF to be non-negative and with known finite support.

Simulated Annealing (SA) is a Monte-Carlo global minimisation technique, first proposed by McCallum (1990), and is summarised in Kundur and Hatzinakos (1996). It is named after the analogous processing of metals – when liquid metal is cooled, the random motions of the metal atoms are reduced. If the cooling is done sufficiently slowly, then the metal will reach the absolute minimum energy state related to the complete atomic ordering of the metal. If the cooling is done too fast, then a sub-optimal state may result. In the simulated annealing algorithm, a cost function, $Q$ (analogous to the energy of the metal), is minimised through the random perturbation of parameters. As the iteration progresses, the "temperature" parameter is gradually reduced and the amount of random changes are also decreased.

Just as in the case of annealing metals, the temperature must be reduced sufficiently slowly every iteration in order for the system to reach the global minimum of the cost function. In the limit of infinite precision and infinitely many iterations, SA is guaranteed to reach the global minimum of the cost function in Eq. 13.32.

$$Q[\tilde{f}(x,y),\tilde{h}(x,y)] = \sum_x \sum_y [\tilde{f}(x,y) * \tilde{h}(x,y) - g(x,y)]^2 \qquad \textbf{(Eq. 13.32)}$$

where $Q$ is the cost function and is equivalent to a measure of the standard deviation of the two estimates from the true convolution.

The algorithm begins with initial guesses of $f_0$ and $h_0$, the true image and PSF respectively.

(i) set cycle counter $n_c$ = 1, scan counter $n_s$ = 1, $N_c$ is the maximum number of cycles of the iterative algorithm, $N_s$ is the number of scans performed per cycle.

(ii) For the $k^{\text{th}}$ cycle, set values for $T_k$ (temperature) and $\alpha_k$ (scale of perturbations). These scale with every loop of the cycle, for example we might define $T_k = 0.8T_{k-1}$ and $\alpha_k = 100\sqrt{T_k}$ . See notes below.

(iii) Scale the estimates of the true image and the PSF by $\beta$ and $1/\beta$ respectively, where $\beta$ is chosen such that the scaled functions have the same root mean square value.

(iv) For each pixel in the estimates, perturb the current pixel with a randomly distributed value in the range $[-\alpha_k, \alpha_k]$. If the new pixel value is negative, set it

to zero (non-negativity constraint). Calculate $dQ$, the change in the cost function; if $dQ \leq 0$, then the perturbation is accepted. If $dQ > 0$, then the perturbation is accepted with probability $\exp(-dQ/T_k)$. See notes below.

(v) If $n_s < N_s$, then increment $n_s$ and return to step (iv) to perform another perturbation pass.

(vi) If $n_c < N_c$, then increment $n_c$ and return to step 2, otherwise stop algorithm.

The choice of values for $T$ and $\alpha$ will define the convergence properties of the algorithm and are therefore important. If we reduce $T$ too quickly then we may become trapped in a local minimum; if $T$ is reduced too slowly, the convergence will be excessively slow. The definition of $\alpha$ will also affect the probability of avoiding local minima. In step (iv), some perturbations that increase in the cost function are accepted to prevent the algorithm from becoming trapped in a local minimum – this only works when the temperature is reduced slowly.

In the presence of noise, the cost function cannot be reduced to zero, but instead it plateaus to a level indicative of the standard deviation of the noise. The SA algorithm is very computationally intensive, though it has guaranteed reliability. It is unlikely that it is a practical solution to large images, but it may suffice for our purposes in deconvolution of 1D profiles.

## 13.5. Summary of deconvolution techniques

In summary, there are many ways of performing both direct and blind deconvolutions. Many are concerned with improving computational efficiency, and these can be largely ignored in our application since we are not striving for real-time analysis; image quality is our main concern.

Direct deconvolution is best implemented by either using a Weiner filter (conventional or incremental) or the Lucy-Richardson algorithm which may have better performance in the presence of noise. Both are clearly superior to Fourier deconvolution.

Blind deconvolution techniques are generally more complex, and some are incomprehensible to all but the most hardened mathematician or statistician. However, the simpler methods such as the IBD and Simulated Annealing are straightforward to implement and can produce satisfactory results, which fortunately means that the more sophisticated techniques are only necessary for purists.

## 13.6. Application of deconvolution techniques

This section will briefly cover some of the results of the deconvolution techniques previously described in application to simulated data, paying particular attention to their performance in the presence of (simulated) additive noise. The methods investigated are:

- Fourier deconvolution
- Weiner filtering
- Lucy-Richardson
- IBD
- Simulated Annealing.

We will use simulated data initially to test the features of the algorithms in a controlled manner. If the algorithms perform poorly here, then they will be of no use to real data which is far less ideal. Our simulated signals are a truncated sinc pulse (original function), and a propagator for laminar flow in a square pipe (PSF function) shown in Figure 13.3. (The calculation of this PSF is covered in Appendix 13.8).

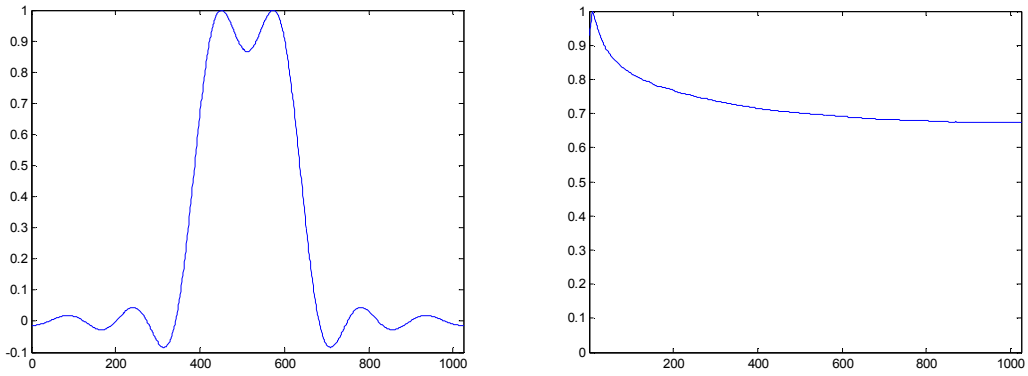


(a) (b)

Figure 13.3 – The simulated data used to test the deconvolution algorithms; (a) the "true signal", a sinc function; (b) the "PSF function", a propagator for laminar flow in a square geometry.

When we convolve these two functions together, we obtain the simulated "observed signal", which is shown in Figure 13.4.

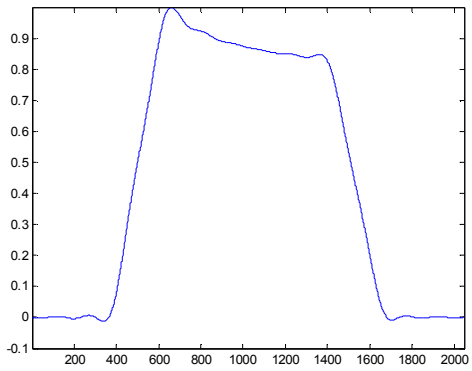


Figure 13.4 – The result of the convolution of the true signal and PSF function from Figure 13.3.

### 13.6.1. Fourier deconvolution

On using a simple inverse Fourier transform to deconvolve the known PSF from the convolution, we obtain the original function, as seen in Figure 13.5.

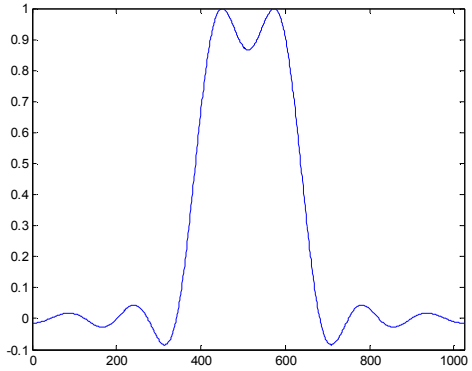


Figure 13.5 – The result of using Fourier deconvolution to remove the effect of the PSF (Figure 13.3b) from the convolved data (Figure 13.4).

So, for the noiseless case, the Fourier deconvolution method is seen to work perfectly. If we now attempt the same deconvolution, but this time add in 1% additive noise to the convolved signals (see Figure 13.6a), on deconvolution our result is swamped by noise (Figure 13.6b). This clearly demonstrates the problems encountered in the

presence of even small amounts of noise, and highlights the need for filtering / smoothing functions.

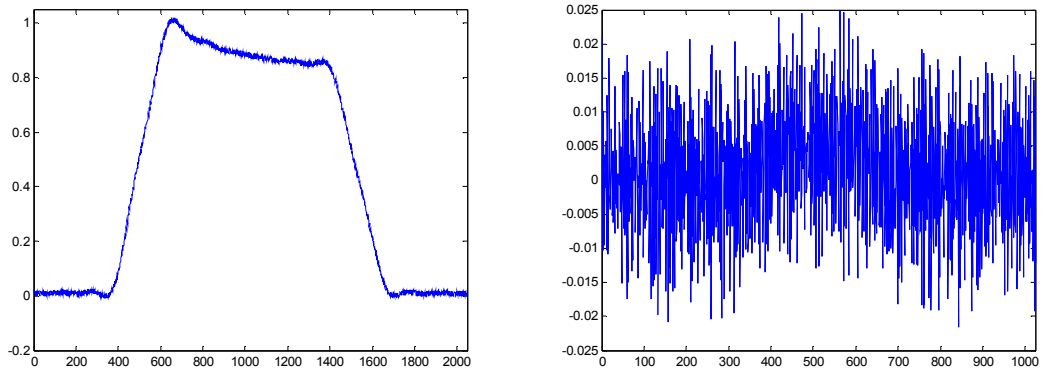


(a) (b)

Figure 13.6 – Illustration of the noise-sensitivity of the Fourier deconvolution method; (a) shows the data of Figure 13.4 with 1% additive noise; (b) shows the result of performing the Fourier deconvolution of (a).

If we examine the Fourier dimension, we can observe the cause of the problem. We see in Figure 13.7 the Fourier transform of the obtained estimate of the "true signal" signal for the noiseless case and for 1% additive noise.

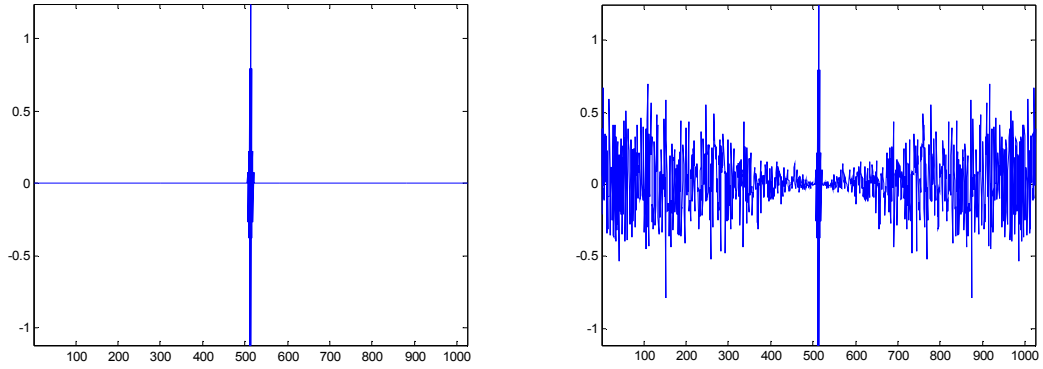


(a) (b)

Figure 13.7 – Fourier spectra of the convolved test signals; (a) is the noiseless case of Figure 13.5; (b) is the noisy case of Figure 13.6b.

Clearly, during the deconvolution process the noise has been greatly amplified at the edges of Fourier space, while the main signal peak is largely unaffected. A solution to this is to employ a smoothing pass, such as a Gaussian function. The results of

multiplying the deconvolved signal by a Gaussian in Fourier space, and then transforming back to real space are shown in Figure 13.8.

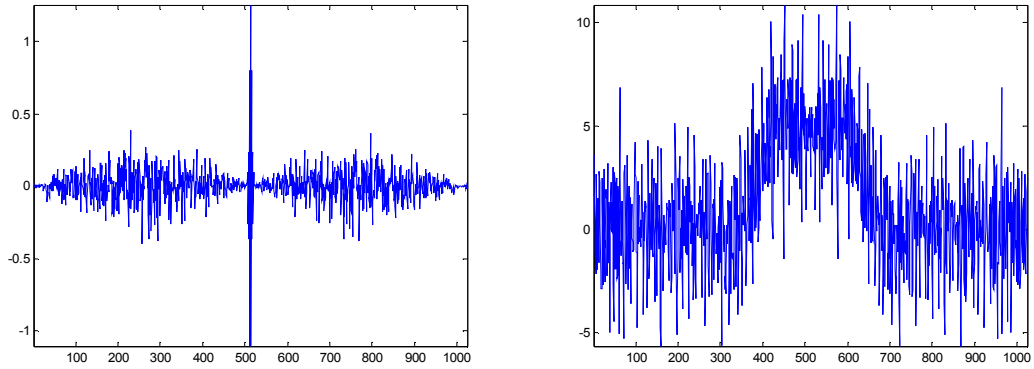


(a) (b)

Figure 13.8 – The results of a single smoothing pass; (a) shows the Fourier spectra after the multiplication of Figure 13.7b with a Gaussian function (Section 7.2.3); (b) shows the Fourier transform of (a) in real space.

It is apparent that the results are improved, but still very noisy. Further smoothing passes yield the data of Figure 13.9, and now we can clearly see the original sinc function of Figure 13.3a.

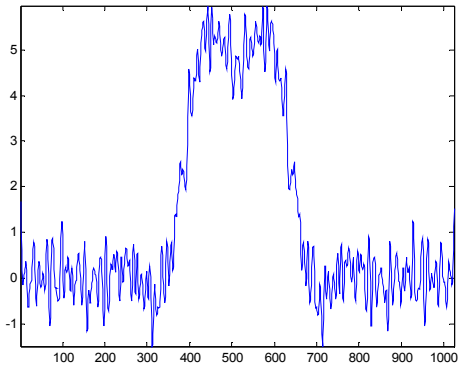


Figure 13.9 – The result of 10 smoothing passes with the Gaussian smoothing function

This investigation has shown that the Fourier deconvolution method is not an ideal process since we have lost information due in the presence of a relatively small level of noise. We will now investigate the Weiner deconvolution technique to see if it fares better, as has been claimed in the literature.

## 13.6.2. Weiner deconvolution

The Weiner deconvolution algorithm was applied to the same noisy data, with varying values of $K$, the parameter which controls the degree of smoothing / sharpness. Results are shown in Figure 13.10 for $K$ = 1, 10 and 100. In this case, with a $K$ factor of 100, we have recovered a good representation of the original function.

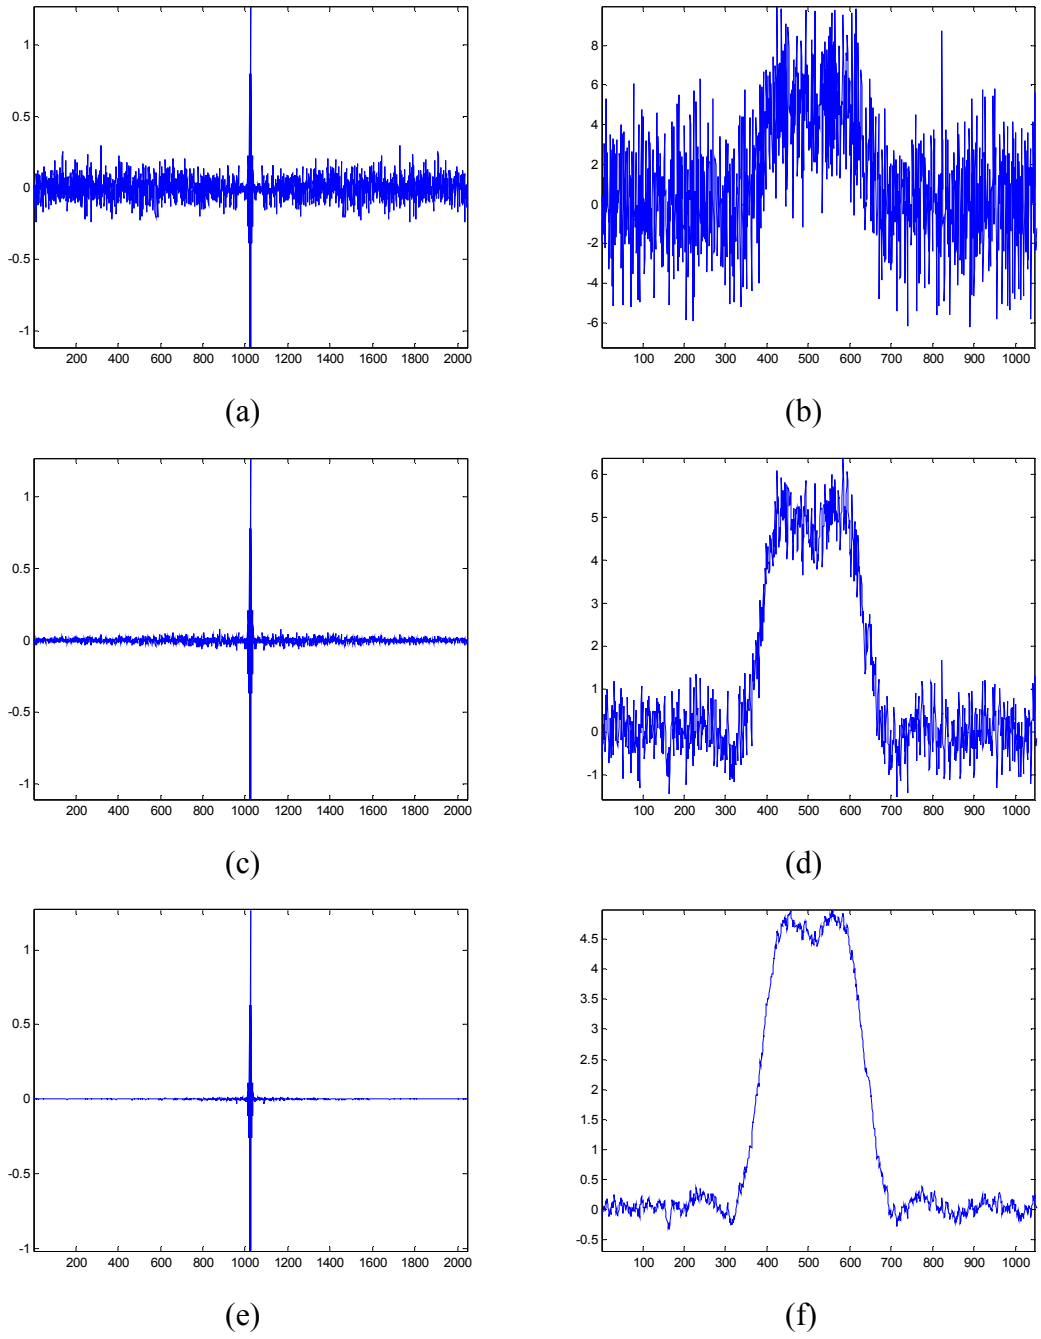


Figure 13.10 – The results of Weiner deconvolution in the presence of 1% additive noise. In each case, the Fourier spectrum of the obtained "true signal" and the real space "true signal" are shown, for $K$ values of (a-b) 1; (c-d) 10, and (e-f) 100.

### 13.6.3. Lucy-Richardson

The Lucy-Richardson deconvolution algorithm was next applied, with 5, 50 and 500 iterations, and the results are shown in Figure 13.11.

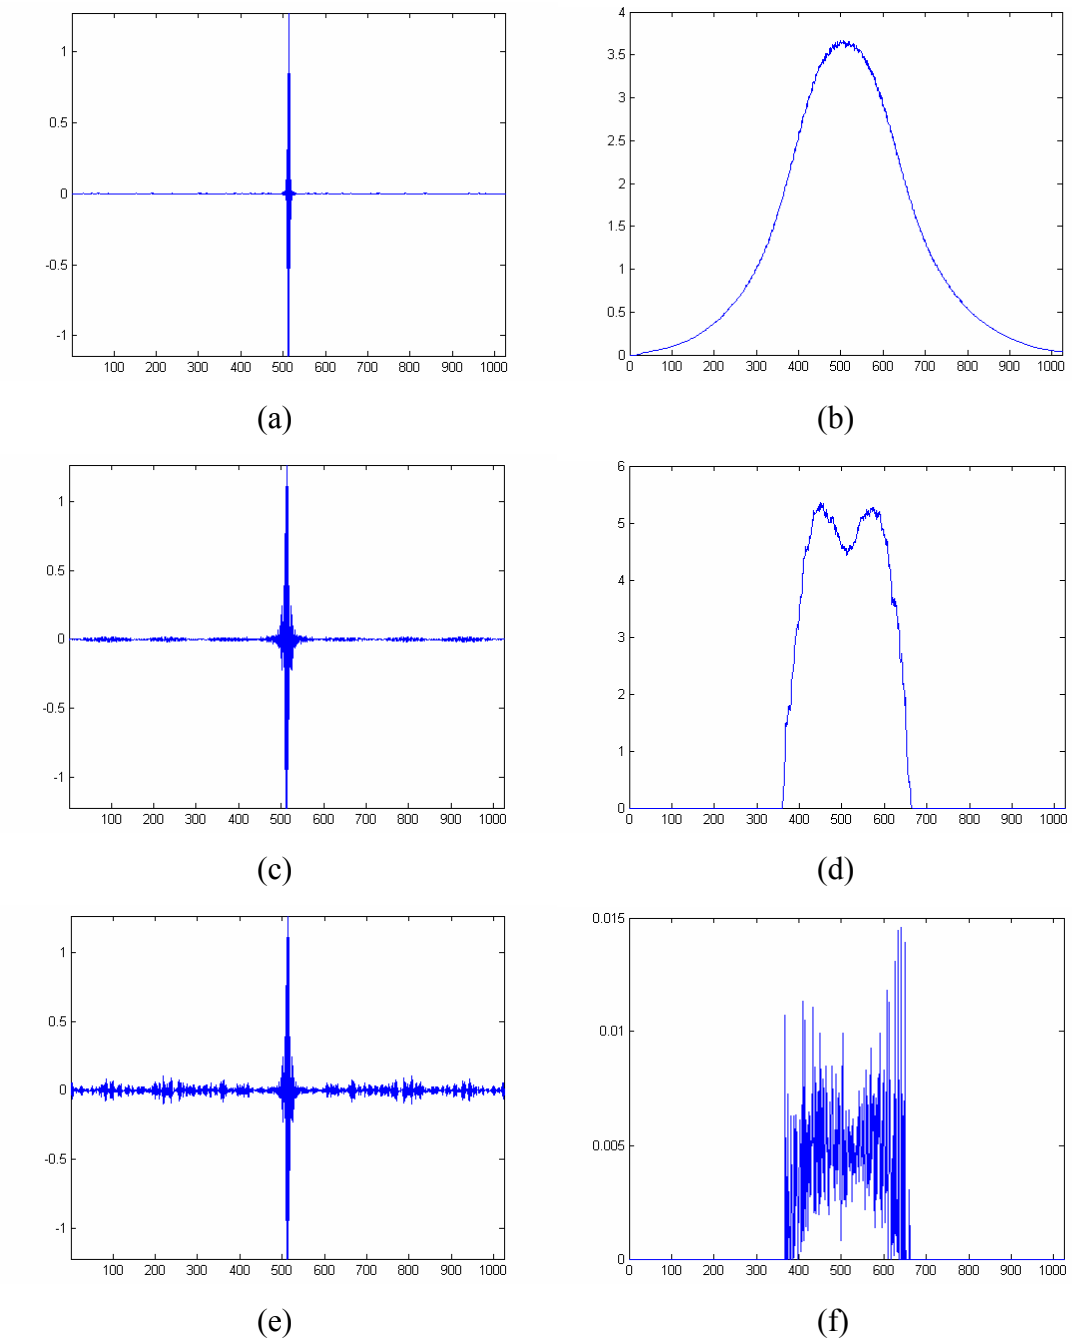


Figure 13.11 – The results of Lucy-Richardson deconvolution in the presence of 1% additive noise. In each case, the Fourier spectrum of the obtained "true signal" and the real space "true signal" are shown, for iterations of (a-b) 5 steps; (c-d) 50 steps, and (e-f) 500 steps.

With this filter, it was seen that several iterations are required to approximate the original function, however using too many iteration steps appears to then deviate from the ideal solution. An optimum was found around 50 steps. However, the sensitivity of this algorithm to the number of iteration steps is a potential problem when we consider the robustness of the algorithm to be important. For this reason, the Lucy-Richardson algorithm was considered to be less practical than the Weiner deconvolution method.

### 13.6.4. Iterative Blind Deconvolution

For the specific case of motion-blurring within the channels of the monolith reactor, which is the aim of this work, the degree of blurring varies according to the respective propagators within each channel. Therefore, if we are to avoid assumptions on the form of the propagator, we must use a blind deconvolution scheme to estimate the shape of the propagator involved.

For the blind deconvolution techniques, it was decided to use an estimate for the PSF, rather than use the known function since there would then be no need for blind deconvolution. A block pulse was used as an estimate for the PSF (shown in Figure 13.12).

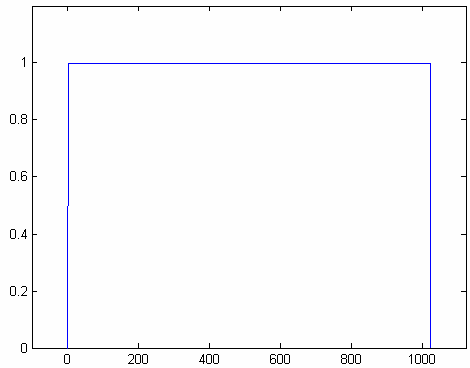


Figure 13.12 – The test PSF function, used as an input to the blind deconvolution algorithms. This was used instead of the known PSF.

The IBD algorithm was implemented and executed, and in each case the obtained true signal and obtained PSF are shown. The algorithm was run for 5, 50 and 100 steps,

and the results are shown in Figure 13.13. The algorithm produces reasonable results considering the conditions under which it must operate, however it is likely that it can be further improved by the addition of an incremental Weiner filter within the IBD algorithm.

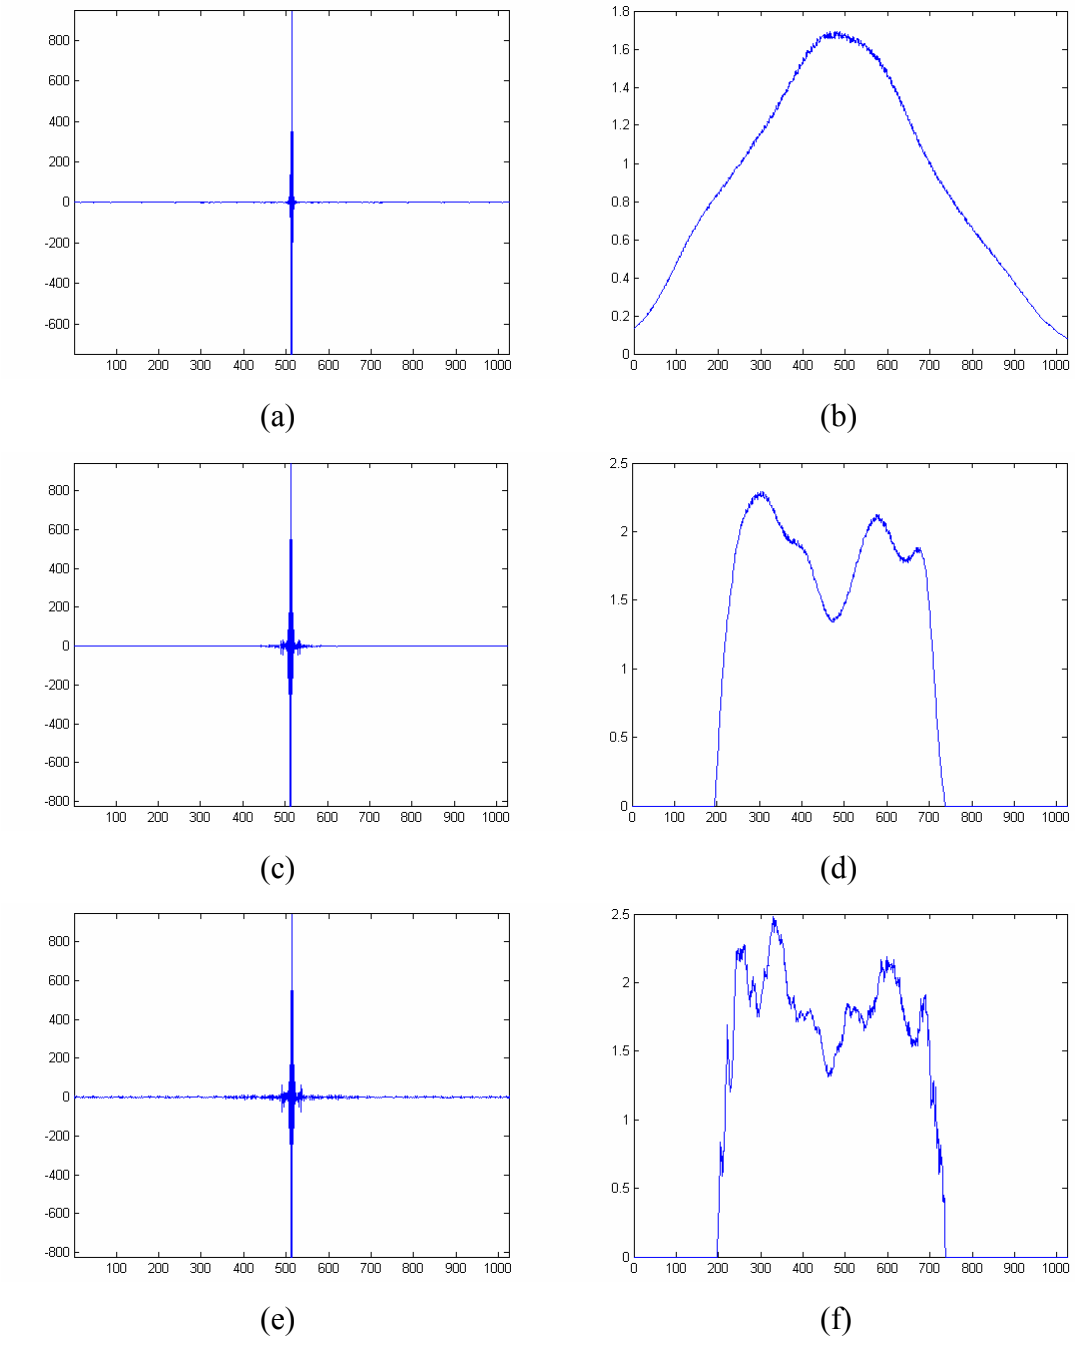


Figure 13.13 – The results of the IBD deconvolution algorithm in the presence of 1% additive noise. In each case, the Fourier spectrum of the obtained "true signal" and the real space "true signal" are shown, for iterations of (a-b) 5 steps; (c-d) 50 steps, and (e-f) 100 steps.

### 13.6.5. Simulated Annealing

The Simulated Annealing algorithm was implemented, however it was found to be very sensitive to the input parameters and took an unfeasibly long time to process. The results did not always appear to be converging satisfactorily. It was decided to abandon this method in favour of other, more robust techniques. This technique may benefit from more controlled convergence formulations, perhaps by use of information measures described in Section 7.4.

## 13.7. Conclusions

In summary, we have introduced the main theory behind deconvolution, and reviewed the primary techniques used in signal processing. Some trial experiments on simulated data proved the Weiner filter to be the most promising, due to its reliability and insensitivity to the input parameters of the algorithm. The IBD algorithm performed the best of the blind deconvolution methods investigated, and is likely to improve with the inclusion of an incremental Weiner filter into the algorithm. This has now been implemented, and its application to real MRI data is described in Chapter 14.

## 13.8. Appendix - calculation of propagator for flow in a square pipe

This appendix will address the calculation of a propagator representing flow in a square pipe, which is the geometry associated with the monolith channels.

The velocity profile for flow in a square pipe is more complex to solve than for that of circular geometry. The Navier-Stokes equations can only be solved analytically for a few simple cases with basic geometries, and the solution to the second-order differential equation is not so straightforward for square geometry. However, the velocity profile can be decomposed into an infinite sum of Fourier series (Ericson, 2004; Verschelde, 2005), or approximated numerically using the method of Galerkin (http://www2.math.uic.edu/~jan/mcs494/Lec38/galerkin1.html), shown in Eq. 13.33.

$$\begin{aligned} trial := {} & \alpha_{1,1}(x-x^2)(y-y^2) + \alpha_{1,2}(x-x^2)(y^2-y^3) + \alpha_{1,3}(x-x^2)(y^3-y^4) + \alpha_{1,4}(x-x^2)(y^4-y^5) \\ & + \alpha_{2,1}(x^2-x^3)(y-y^2) + \alpha_{2,2}(x^2-x^3)(y^2-y^3) + \alpha_{2,3}(x^2-x^3)(y^3-y^4) + \alpha_{2,4}(x^2-x^3)(y^4-y^5) \\ & + \alpha_{3,1}(x^3-x^4)(y-y^2) + \alpha_{3,2}(x^3-x^4)(y^2-y^3) + \alpha_{3,3}(x^3-x^4)(y^3-y^4) + \alpha_{3,4}(x^3-x^4)(y^4-y^5) \\ & + \alpha_{4,1}(x^4-x^5)(y-y^2) + \alpha_{4,2}(x^4-x^5)(y^2-y^3) + \alpha_{4,3}(x^4-x^5)(y^3-y^4) + \alpha_{4,4}(x^4-x^5)(y^4-y^5) \end{aligned} \quad \textbf{(Eq. 13.33)}$$

A velocity plot of the velocity profile described by Eq. 13.33 is shown in Figure 13.14. The profile is quasi-parabolic, and the velocity falls to zero at the four edges to fit the square geometry.

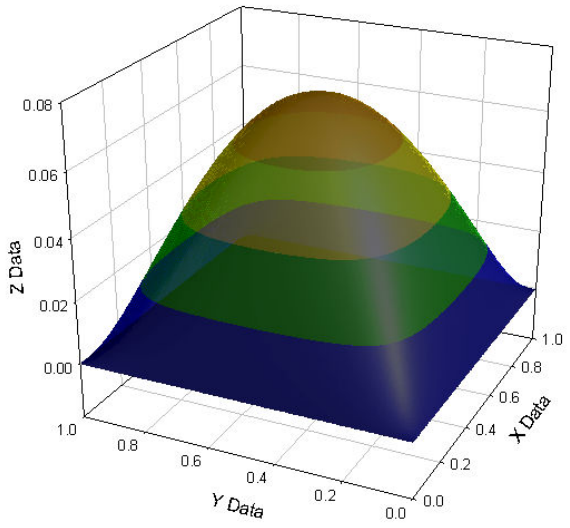


Figure 13.14 – Velocity plot of the numerical approximation to laminar flow in a square pipe, described by Eq. 13.33.

We can use this data to calculate the corresponding propagator (the probability of a fluid element having a particular displacement), and we obtain the result shown in Figure 13.15.

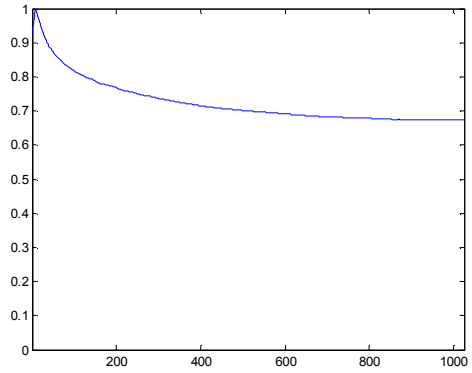


Figure 13.15 – Calculated propagator for laminar flow in a square pipe.

In comparison, the propagator for flow in a circular pipe is flat, i.e. the probability of a fluid moving any particular distance is uniform. We see that for flow in a square pipe there are more of the slower-moving fluid elements, which is to be expected for a geometry which has more boundary edges, and therefore more zero-velocity, and low-velocity zones.

## 13.9. Nomenclature

| Symbols | Definition |
|---|---|
| $a$ | Scalar factor |
| $a_k$ | Scale of perturbations |
| $c_l$ | Motion-blurring parameter |
| $C(u)$ | Convolution of true signal and PSF |
| $d_x$, $d_y$ | Linear offset in $x$, and $y$ direction |
| $E(\cdot)$ | Expectation of |
| $f_0$, $h_0$ | Initial guess of true signal, PSF |
| $\tilde{f}(x,y)$ | Estimate of true signal |
| $\tilde{F}(u)$, $\tilde{G}(u)$ | Estimate of Fourier transforms of true signal, and PSF |
| $F(\cdot)$, $F^{-1}(\cdot)$ | Fourier transform, inverse Fourier transform |
| $f(x,y)$, $g(x,y)$, $h(x,y)$ | Functions representing the true signal, the convolved signal, and the PSF in real space $(x,y)$ |
| $F(u,v)$, $G(u,v)$, $H(u,v)$ | Functions representing the true signal, the convolved signal, and the PSF in Fourier space $(u,v)$ |
| $F_{old}$, $F_{new}$ | Old and new estimates of the true function (in Fourier space) |
| $i$ | Imaginary number, or iteration number (depending on context) |
| $H^*$ | Complex conjugate of $H$ |
| $J_1$ | Bessel function |
| $k$ | Scalar factor |
| $K$ | Weiner filtering parameter |
| $M,N$ | Dimensions of a 2D dataset |
| $n(x,y)$ | Additive noise term |
| $n_c$, $N_c$ | Cycle counter, max number of cycles |
| $n_s$, $N_s$ | Scan counter, max number of scans |
| $P(x)$, $P(x\|y)$ | Probability of event $x$, conditional probability of $x$, given $y$ |
| $Q$, $dQ$ | Cost function, change in cost function |
| $S(u,v)$ | Error array for iterative algorithms |
| $T$ | Total exposure (acquisition) time |
| $T_k$ | Temperature parameter (SA) |
| $V_x$, $V_y$ | Components of linear velocity in $x$ and $y$ directions |
| $W_{nn}$, $W_{xx}$ | Power spectra of the noise, and of the object |
| | |
| $\beta$ | Weighting parameter for IBD algorithm |
| $\varepsilon$ | Additive noise term, entropy measure (depending on context) |
| $\sigma^2$ | Variance |
| $\sigma_{image}$, $\sigma_{noise}$ | Standard deviation of the signal strength of image pixels, noise pixels |
| $*$ | Convolution symbol |

| Abbreviations | Definition |
|---|---|
| FFT | Fast Fourier transform |
| IBD | Iterative blind deconvolution |
| MEM | Maximum entropy method |
| MRI | Magnetic resonance imaging |
| MSE | Percentage mean squared error |
| NMR | Nuclear magnetic resonance |
| PSF | Point spread function |
| SA | Simulated annealing |
| SNR | Signal to noise ratio |

## 13.10. References


Adam, D. and O. Michailovich (2002). "Blind Deconvolution of Ultrasound Sequences using Non-parametric Local Polynomial Estimates of the Pulse." IEEE transactions on Biomedical Engineering **42**(2): 118-131.

Ayers, G. R. and J. C. Dainty (1988). "Iterative Blind Deconvolution Method and Its Applications." Optics Letters **13**(7): 547-549.

Bertero, M. and P. Boccacci (2000). "Image Restoration Methods for the Large Binocular Telescope." Astron Astrophys Suppl Ser **147**: 323-333.

Cooley, J. W. and J. W. Turkey (1965). "An Algorithm for the Machine Calculation of Complex Fourier Series." Mathematics of Computation **19**: 297-301.

Cornwell, T. and A. Bridle (1996). "Deconvolution Tutorial." http://www.cv.nrao.edu/~abridle/deconvol/deconvol.html.

Ericson, S. (2004). "Square Duct Verification Equations." http://www.jics.utk.edu/~sericson/misc/Verification/square_channel/.

Fish, D. A., A. M. Brinicombe*, et al.* (1995). "Blind Deconvolution by Means of the Richardson-Lucy Algorithm." Journal of the Optical Society of America a-Optics Image Science and Vision **12**(1): 58-65.

Gonzales, R. C. and R. E. Woods (1992). Digital Image Processing. Reading, MA, Addison-Weasley.

Gull, S. F. and G. J. Daniell (1978). "Image Reconstruction from Incomplete and Noisy Data." Nature **272**: 686-690.

Jefferies, S. M. and J. C. Christou (1993). "Restoration of Astronomical Images by Iterative Blind Deconvolution." Astrophysical Journal **415**(2): 862-874.

Katsaggelos, A. K. and K. T. Lay (1991). Identification and Restoration of Images Using the Expectation Maximisation Algorithm. ch 6, pp143-176, Springer-Verlag.

Kundur, D. and D. Hatzinakos (1996). "Blind image deconvolution." Ieee Signal Processing Magazine **13**(3): 43-64.

Lam, E. Y. and J. W. Goodman (2000). "Iterative statistical approach to blind image deconvolution." Journal of the Optical Society of America a-Optics Image Science and Vision **17**(7): 1177-1184.

Lane, R. G. and R. H. T. Bates (1987). "Automatic Multidimensional Deconvolution." Journal of the Optical Society of America a-Optics Image Science and Vision **4**(1): 180-188.

Lucy, L. B. (1974). "An Iterative Technique for the Rectification of Observed Distributions." Astrophysical Journal **79**: 745.

Markham, J. and J. A. Conchello (1999). "Parametric blind deconvolution: a robust method for the simultaneous estimation of image and blur." Journal of the Optical Society of America a-Optics Image Science and Vision **16**(10): 2377-2391.

McCallum, B. (1990). "Blind Deconvolition by Simulated Annealing." Optics Communications **75**(2): 101-105.

Mouyan, Z. and R. Unbehauen (1995). "A Deconvolution Method for Spectroscopy." Measurement Science & Technology **6**(5): 482-487.

Muller, J. P. (1988). Digital Image Processing in Remote Sensing. Philadelphia, Taylor and Francis.

Oppenheim, A. V., R. W. Schafer, *et al.* (1968). "Nonlinear Filtering of Multiplied and Convolved Signals." IEEE Proceedings **56**: 1264-1291.

Press, W., B. P. Flannery, *et al.* (1986). Numerical Recipies - The Art of Scientific Computing, Cambridge University Press.

Richardson, W. (1972). "Bayesian-based Iterative Method of Image Restoration." Journal of the Optical Society of America **62**: 55-59.

Russ, J. C. (2002). The Image Processing Handbook: Fourth Edition, CRC Press.

Stockham, T. G., T. M. Cannon, *et al.* (1975). "Blind Deconvolution Through Digital Signal Processing." IEEE Proceedings **63**(4): 678-692.

Tikhonov and Arsenin (1977). Solutions of Ill-Posed Problems. Washington, Winston & Sons.

van Kempen, G. M. P., H. T. M. van der Voort, *et al.* (1996). "Comparing Maximum Likelihood Estimation and Constrained Tikhonov-Miller Restoration." Ieee Eng Med Biol **15**: 76-83.

Verschelde, J. (2005). "Galerkin's method for ODEs and PDEs." http://www2.math.uic.edu/~jan/mcs494/Lec38/galerkin1.html.

Weiner, N. (1949). Extrapolation, Interpolation, and Smoothing of Stationary Time Series. Cambridge (MA), Wiley and Sons.

Wilson, T. and S. J. Hewlett (1991). "Imaging Strategies in Three-dimensional Confocal Microscopy." Proc SPIE **1245**: 35-45.

# Chapter 14 – Propagator deconvolution

This chapter will demonstrate a novel data processing method, which is designed to extract additional flow information from existing MRI data. It will be shown that by post-processing the line-excitation data shown in Chapters 6 and 11, we can estimate the propagators within a row of monolith channels in around 150 ms. The method utilises deconvolution techniques described in Chapter 13, and allows extra information to be extracted from an existing MRI measurement.

**CHAPTER 14 – PROPAGATOR DECONVOLUTION**

## 14.1. Introduction

MRI has been shown to be a fast progressing area of research with rapid development in new and faster acquisition techniques, putting more demand on hardware and data analysis procedure. The aim of the approach presented in this chapter is to use the line-excitation time-of-flight technique (Section 5.3) with novel data analysis, in order to extract the maximum amount of information possible.

It will be shown that, for a suitable system, it is possible to extract propagator information in as little as 150 ms, which is an order of magnitude less than for comparable traditional NMR techniques (Section 2.4.2). The technique presented here has some limitations, but nevertheless is very powerful and gives a rapid measurement, allowing data acquisition in dynamically changing systems where traditional techniques would break down completely.

The theory behind the data processing technique will first be outlined, and then the development of an analytical model for processing the data will be developed, first for pipe flow (Section 14.3), then for single-phase flow in a monolith (Section 14.4), and finally applied to two-phase flow systems (Section 14.5) for discussion and analysis.

## 14.2. Theory

The technique developed in this work is based upon a modification of the line-excitation time-of-flight technique previously reported in Heras *et al.* (2005), applied to monolith operation. The principle of the approach is as follows. The technique has been further optimised, and the associated data analysis is also presented. The line-excitation imaging method tags spins in a thin line, and then tracks the movement of these spins over the acquisition time. By examining the spread of the spins relative to the original excitation line, we can extract a superficial velocity since the distance moved by spins is proportional to the velocity of the fluid. By more careful analysis we can estimate the propagator which characterises the spreading of these spins. This therefore gives a rapid measure of the range of velocities present within a monolith channel, all from an MRI technique which originally appeared to be merely a time-of-flight tagging method.

Two main factors need to be taken into account in this analysis. Firstly, the shape of the initial excitation slice needs to be deconvolved from the acquired data, i.e. if we used a Gaussian excitation pulse, then the spins would have begun with a Gaussian intensity distribution which then spreads according to the flow within the channel, and this effect needs to be removed. Secondly, the line-excitation acquisition is not an instantaneous imaging technique – it is very rapid at 150 ms, but the signal is acquired over the entire acquisition time, meaning that there is a degree of motion-blurring which needs to be compensated for. Since we can assume that the flow is steady over this time (and we therefore have a linear motion-blur) we can evaluate the effect of this blurring, and deconvolve the spreading function to obtain a measure of the effective profile at the mid-point of the acquisition. The resultant spread of intensity within the channel should then be representative of the propagator for that channel.

## 14.3.Single-phase pipe flow

This section will show the development of the procedure and its application to single-phase flow in a pipe, which is the simplest single-phase flow system to analyse. This test system will show whether the technique is fundamentally reliable.

### 14.3.1. Experimental

MRI data were acquired using the SEMI-RARE sequence (Section 5.2) for single-phase flow in a 10 mm i.d. Perspex pipe. The driving force for the flow was provided by gravity to ensure steady-state conditions. The liquid flow rate was varied between zero and 58 cc/min (superficial velocities of zero to 12.3 mm/s). Images are taken with a field-of-view of 11 mm × 11 mm, a resolution of 86 μm × 172 μm, an excitation slice thickness of 1 mm, and a refocusing slice thickness of 1 mm. The experiments were performed 64 times and averaged to improve the signal-to-noise ratio. A recovery time of 5 seconds was allowed between experiments. An example of a data set is shown in Figure 14.1.

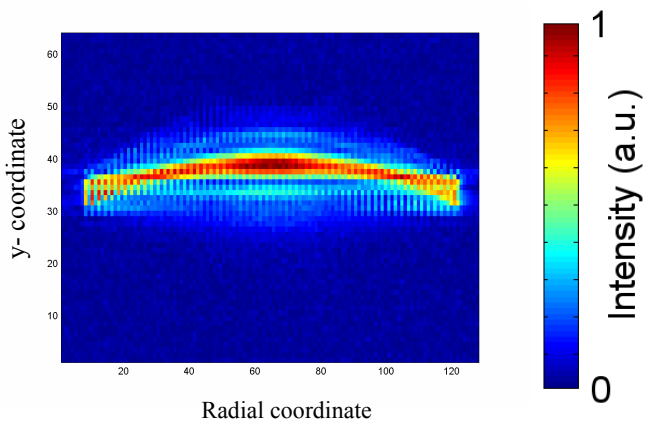


Figure 14.1 – An example of a SEMI-RARE acquisition for laminar flow in a 10 mm i.d. circular pipe at a flow rate of 17.1 cc/min of water. The parabolic nature of the flow can be seen, as well as the motion-blurring associated with the acquisition.

The processing procedure is outlined in Figure 14.2, showing the blind deconvolution of the linear motion-blurring, and the deconvolution of the initial excitation pulse shape. The inputs to the system are (i) the raw data profile which is the spread of intensity from excited spins at any given $x$-coordinate in the pipe, (ii) the corresponding intensity profile at zero flow, which gives a measure of the initial spread of intensity caused by the shape of the excitation pulse, and (iii) an initial guess for the motion-blurring PSF. The deconvolution step can be performed with the Weiner filter (with a $K$ value on the order of 100) if the initial excitation shape is known, otherwise an Iterative Blind Deconvolution (IBD) step can be used. There is no choice for the motion-deblurring step, an IBD process must be used since we should not assume anything about the shape of the propagator. Constraints may be added to the IBD loop, such as enforcing non-negativity, or imposing a particular shape such as a linear motion-blur.

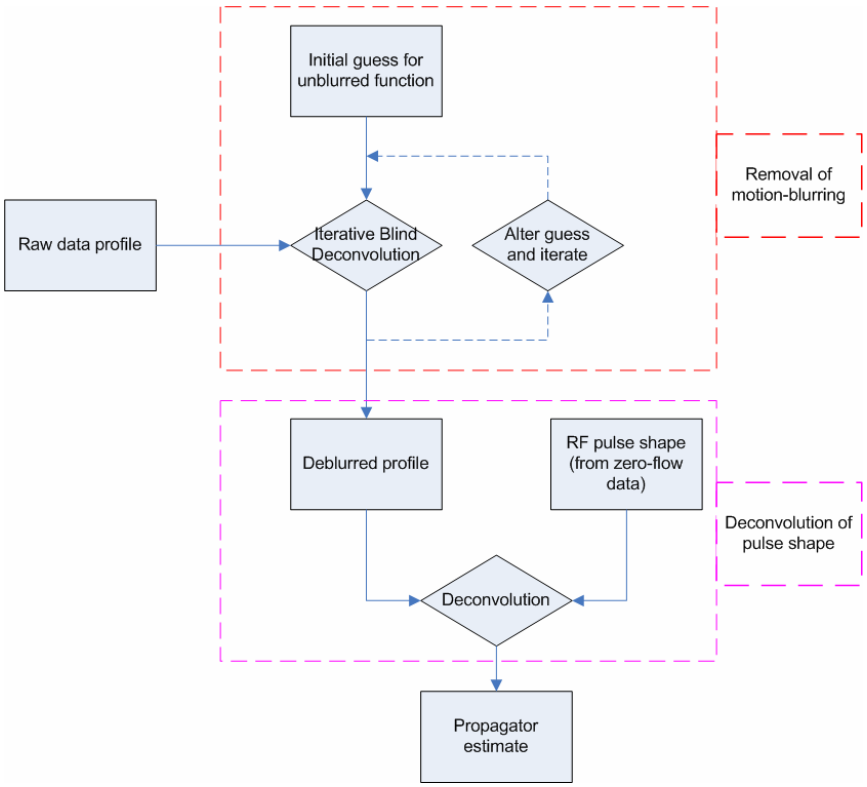


Figure 14.2 – Flow sheet of the propagator deconvolution procedure, which will be described in this chapter.

## 14.3.2. Algorithm development

The acquired MRI profile consists of the convolution of the propagator with the initial excitation slice, which is then blurred over the acquisition time. Therefore, to obtain the original propagator measurement, we must reverse the process, i.e. we must remove the motion-blurring first, and then remove the initial slice shape[1](1).

The motion-blurring can be assumed to be linear since the velocity is steady; however the PSF associated with this blurring cannot be assumed, and hence we require a blind deconvolution technique. This was performed with an IBD algorithm (Section 13.4.1), which was modified to utilise a Weiner deconvolution filter (Section 13.3.2) during the deconvolution steps, in order to improve the algorithm's performance in the presence of noise. The number of iterations was set at 100, and a value for the weighting parameter, $\beta$, of 0.9 was chosen, based on recommendations

[1] In fact, the order of the deconvolution steps shouldn't matter. However, later in this chapter we will encounter some additional complications when addressing flow in the monolith reactor, and it makes sense to perform the deconvolutions in this particular order.

in the literature. The initial guess for the blurring PSF was taken as a typical linear motion-blur shape, and iterated from there.

One of the disadvantages of using the IBD technique is that the iterated solution does not necessarily contain the correct scale and offset (as described in Section 13.4). The resultant deconvolved profiles must therefore be fitted to the initial data to correct for this. In general, this is done by a least-squares fitting algorithm, and the corrected results are shown in Figure 14.3. It can be seen that our profiles are "sharpened" or "narrowed" compared to the original spread of intensity, which is expected from a deconvolution procedure such as the removal of motion-blurring.

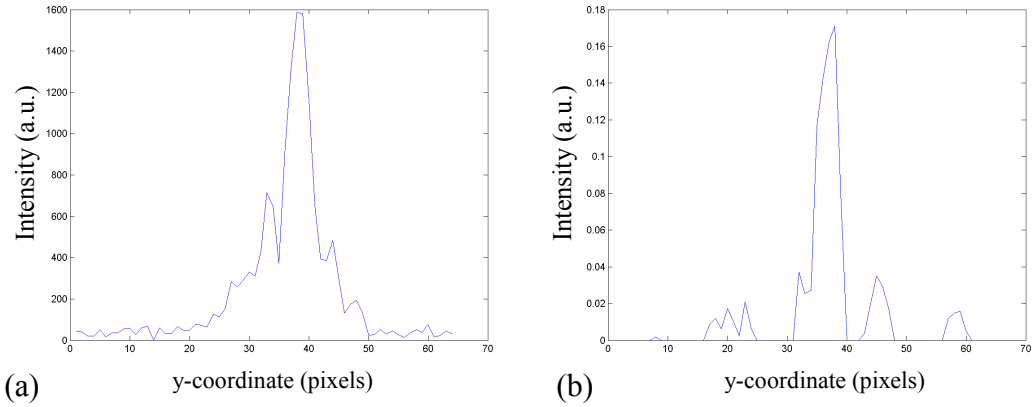


Figure 14.3 – Result of blind deconvolution on a selected profile; (a) shows the raw data from the central slice of the pipe-flow data; (b) shows the result of applying the blind deconvolution technique to remove motion-blurring.

The full 2D data set of pipe flow is shown in Figure 14.4, showing the results of removing the motion-blurring from the pipe flow dataset.

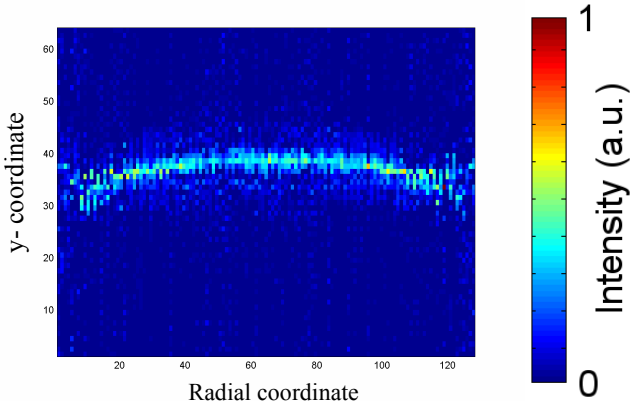


Figure 14.4 – The result of applying the blind deconvolution technique to remove motion-blurring from the data of Figure 14.1.

The next step is to remove the bias contributed from the initial excitation shape (a sinc function) from the profiles. Since the shape of the excitation pulse is known (and is easily measured in the zero-flow case) this can be done by a direct deconvolution technique, such as the Weiner deconvolution filter (Section 13.3.2). The Weiner filter is based on the Fourier deconvolution method – $F(f * g) = F(f) \times F(g)$ – but has far superior performance in the presence of noise.

However, we encounter some problems in this procedure. Errors associated with the blind deconvolution technique may mean that we have lost some of the distinctiveness of the sinc profile in Figure 14.3b, and also relaxation effects will have affected the observed profile, such that the sinc shape may not be completely recognisable. Therefore, the correct deconvolution is not necessarily to use the precise sinc shape. In practice, this step was performed with a second blind deconvolution step, but using the shape of the sinc pulse as an initial guess. The results of this are shown for a sample profile in Figure 14.5, and for the full 2D dataset in Figure 14.6.

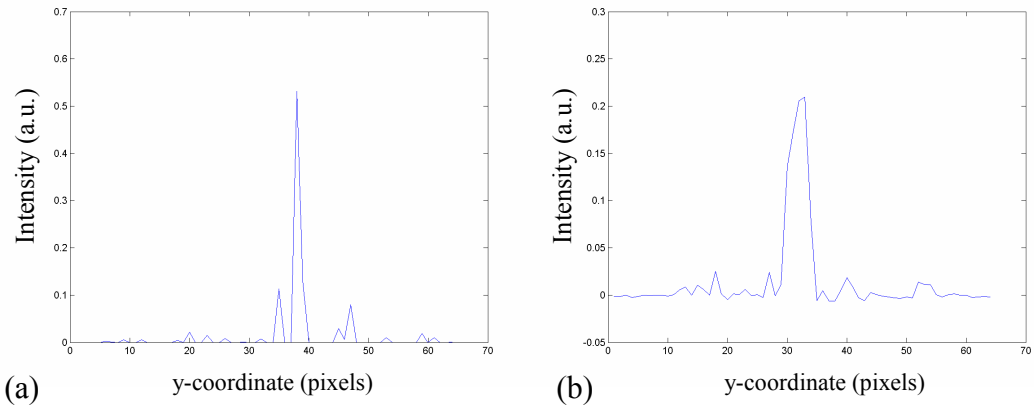


Figure 14.5 – Result of the blind deconvolution of the initial excitation shape on the profile of Figure 14.3a; (a) shows the result of applying the blind deconvolution technique to remove the excitation pulse shape; (b) shows the obtained PSF, which closely resembles the sinc function we were expecting.

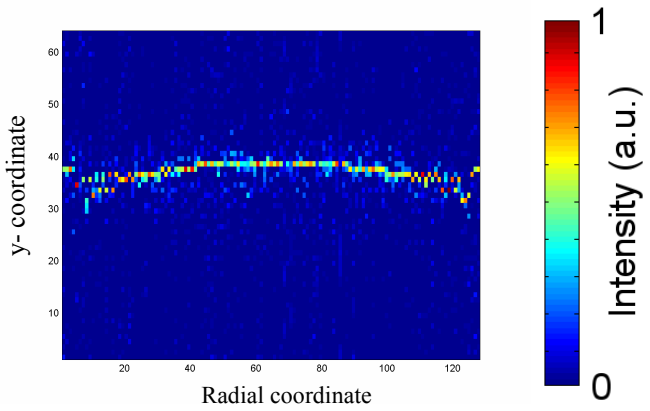


Figure 14.6 – Result of blind deconvolution for the data of Figure 14.4b, showing the result of deconvolving the excitation pulse shape.

As we can see from Figure 14.5, we have obtained a very sharp, thin line. This is consistent with the expected result for laminar pipe flow, since there is only a single velocity at each radius in the pipe and we would expect a delta function for a propagator.

## 14.3.3. Results for pipe flow

We may validate the results of Figure 14.6 by projecting it left-to-right, thereby obtaining a measure of the overall propagator for this central slice of pipe flow. The theoretical result is shown in Figure 14.7. This propagator has been calculated analytically in Matlab. The real data are shown in Figure 14.8, and it can be seen that it compares favourably with the simulated profile.

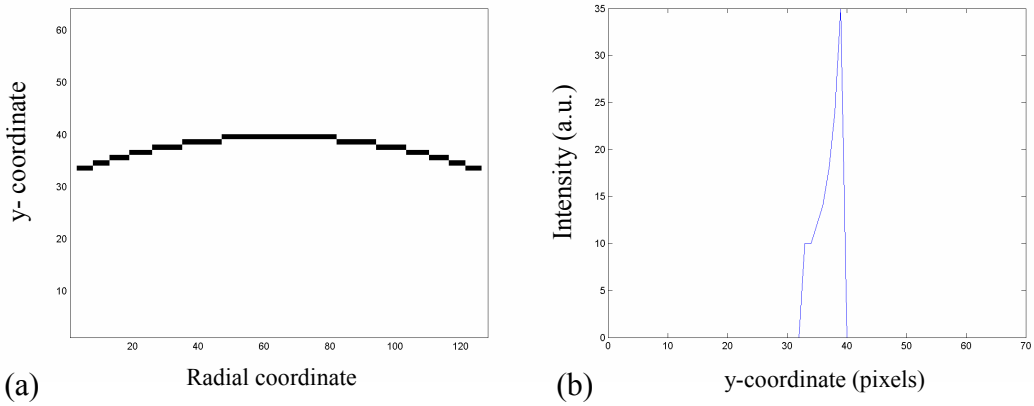


Figure 14.7 – Simulated propagator for laminar flow in a pipe; (a) shows an idealised measurement of the velocity at each radius in the pipe; (b) shows the projected propagator of the data in (a).

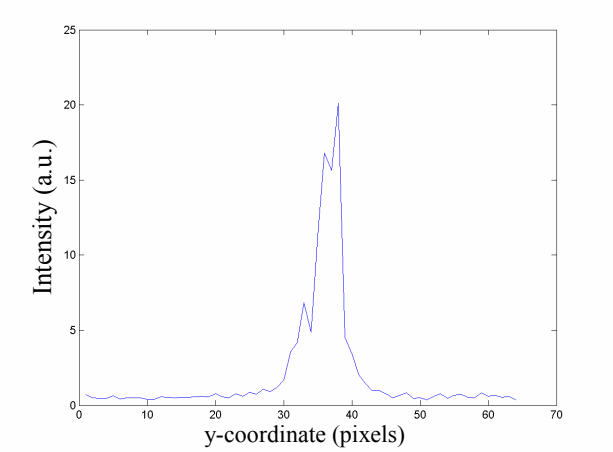


Figure 14.8 – Real propagator, obtained by projecting the data of Figure 14.6.

Since we have now shown that the analysis procedure works, we can calculate propagators for a range of liquid flow rates, and these are shown in Figure 14.9(a-f). It can be seen that the propagator is best defined at low flow rates; at higher flow rates, the intensity is less well-defined, and this causes a loss of signal at higher velocities, thereby skewing the obtained propagator.

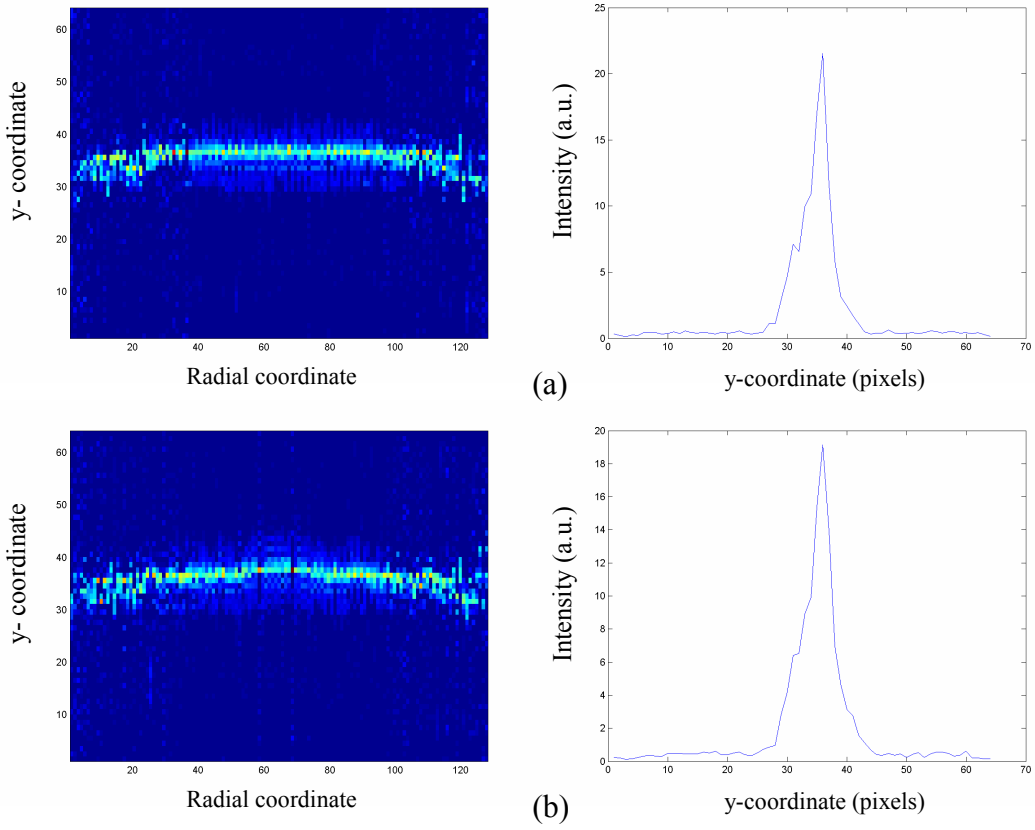

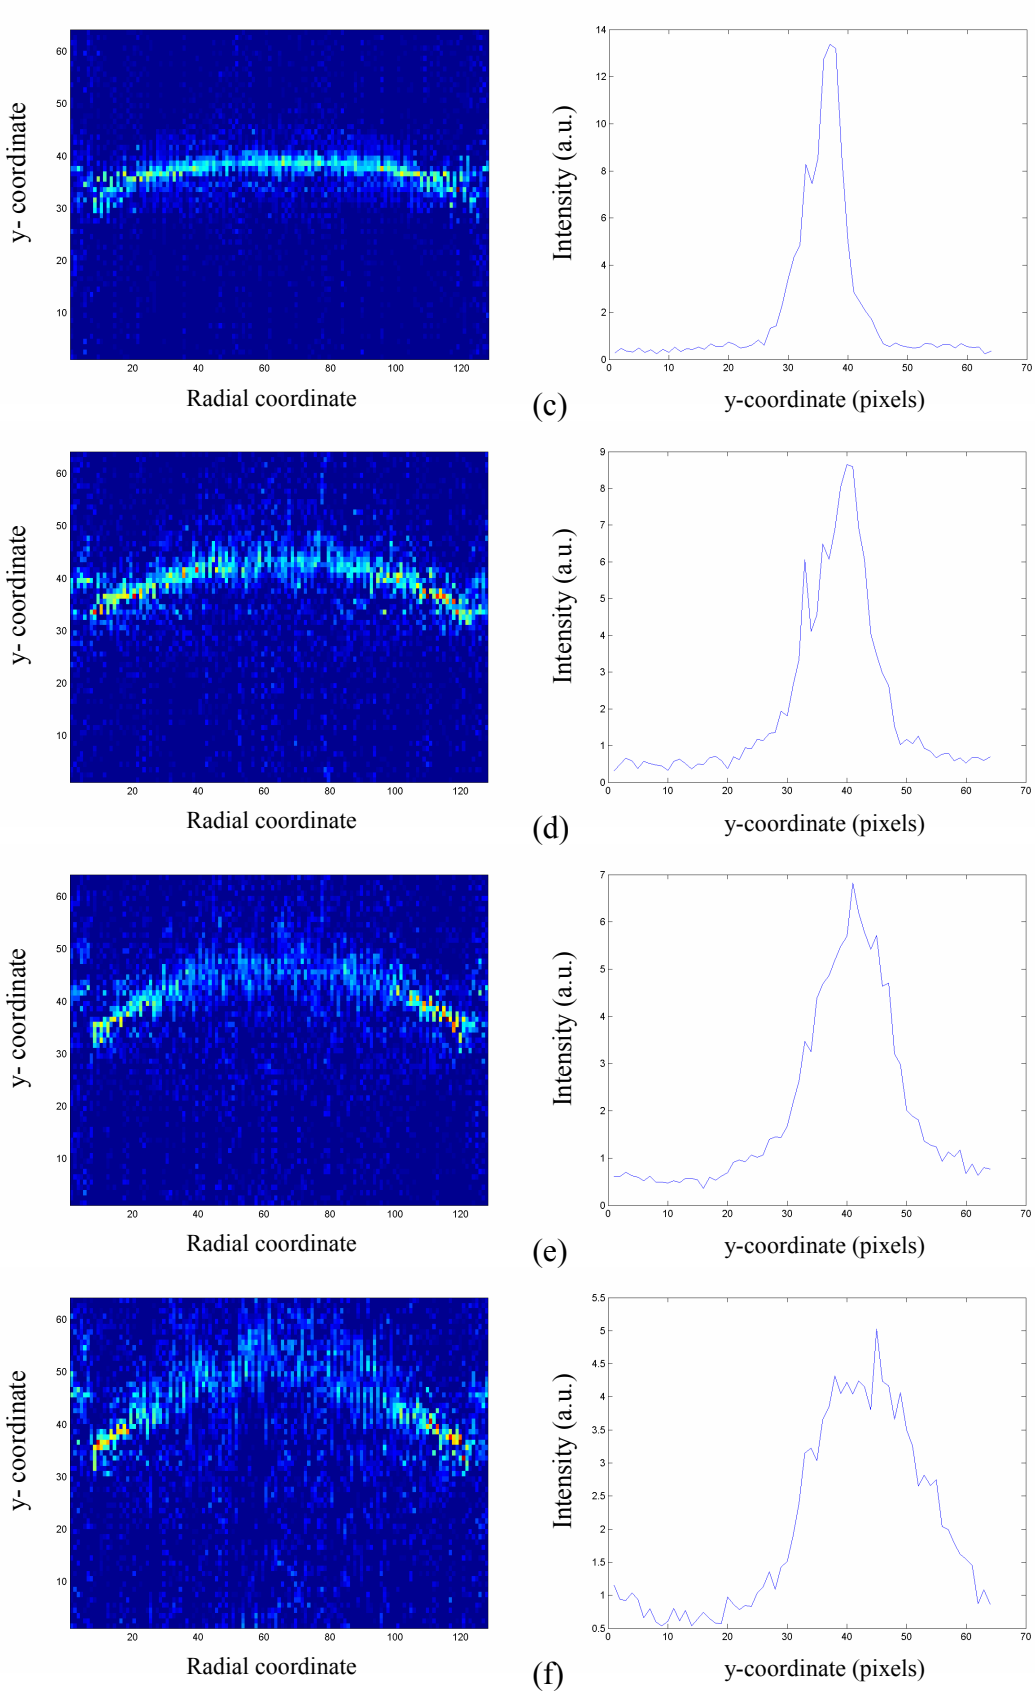


Figure 14.9 – Results of processing the MRI data. 2D datasets and projected propagators are shown for flow rates of (a) 10.8 cc/min, (b) 11.1 cc/min, (c) 17.1 cc/min, (d) 32.1 cc/min, (e) 44.8 cc/min, and (f) 57.7 cc/min.

The following section will consider the case of the monolith reactor, and further develop the processing method for application to this system.

## 14.4. Single-phase flow in a monolith

The geometry of the monolith reactor adds several complications, which need to be taken into account in our method. First, we will discuss the effect of limited resolution, and then investigate the motion-blurring more closely, developing a model for linear motion-blurring and adapting this to better describe the monolith reactor. Relaxation effects are also important, and these will be addressed. This will allow us to develop a deconstructive algorithm to obtain the propagator for flow in a monolith system.

### 14.4.1. The effect of limited spatial resolution

In the case of single-phase flow in a pipe, we have been able to assume that the blurring within the data is a result of a linear motion-blur over the acquisition time, i.e. a spin packet moving at velocity, $v$, will move a distance of $v \times t_a$ over an acquisition time, $t_a$, with the acquisition being a linear combination of this motion. However, if we consider the case of limited resolution we encounter some problems with this assumption. For example, if we assume that we have only one voxel over the width of our pipe then this voxel will contain a range of velocities, not just a single velocity. This means that slow moving spins within the voxel are displaced less than fast moving spins, and the effect is to *stretch* the propagator over the acquisition time, rather than to simply apply a linear blur (which shifts it over the acquisition time). In short, we can no longer assume the motion-blurring is linear. Instead, we must try a different approach to simulate the blurring effect of voxels which contain a range of velocities.

This resolution effect is minimal for the pipe-flow example of Section 14.3, since voxels can be assumed to contain only spins of a certain velocity (or more accurately each voxel has a very narrow velocity distribution). However, this effect will be an important one when we consider flow in a monolith where our resolution is

insufficient to allow the above assumption. Our processing method must be adapted to allow for non-linear blurring, and this is discussed in the following section.

### 14.4.2. Motion-blurring simulation

The correction for motion-blurring will be discussed in this section. First, linear motion-blurring will be considered, followed by the more complicated case of non-linear motion-blurring. In both cases, we must consider how the spins will spread out, according to their local velocities within the channels.

### 14.4.3. Linear motion-blurring

A simple linear motion-blurring may be thought of either as a convolution procedure in Fourier space, or a discretised operation in real space. There are several ways of performing such a motion-blur, each of which are mathematically equivalent:

(i) The convolution procedure consists of convolving together the original image with the PSF associated with linear motion-blur. This gives the desired result directly.

(ii) The point-by-point sampling of Fourier space, which is analogous to the MRI acquisition procedure. We generate a series of profiles which evolve over the acquisition time and Fourier transform them. We then create a composite Fourier profile by taking the first point from the first Fourier profile, *etc*., and then transform this profile back to real space.

(iii) The complimentary way to think about convolution is by the summation of a series of linear shifts. For example, in the case of linear motion-blur we may think of it in this way – for every point in the original image, we substitute in a copy of the PSF (multiplied by the intensity of the original image point) at this position. When we sum together the series of generated PSFs, we obtain the same result as through convolution.

(iv) We may also generate a series of profiles over the acquisition time, and sum the results. This is related to the method of (ii), which is simply a linear averaging of the profiles.

Method (iv) is very easy to implement in the case of non-linear blurring. The direct convolution methods are non-trivial in the case of non-linear blurs, and have not been considered. The discretised operation will be considered in this work, and a brief illustration of the method, and the equivalence to the direct convolution method will now be presented.

### 14.4.4. A comparison of discretised and direct convolution methods

Let us suppose that the profiles of Figure 14.10 represent (a) our true signal, and (b) a PSF which represents a linear motion-blur.

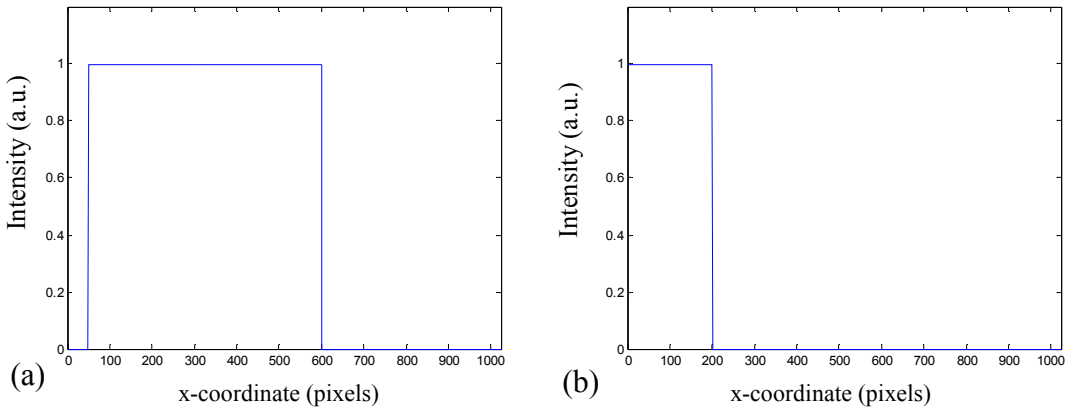


Figure 14.10 – Simulated data consisting of (a) the original signal, and (b) the PSF.

So, for the case of direct convolution, we obtain the result shown in Figure 14.11.

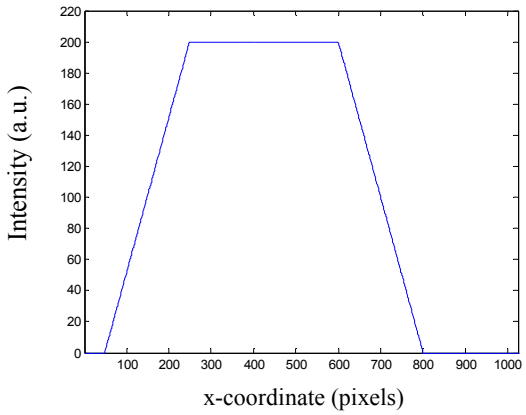


Figure 14.11 – The result of directly convolving the data and PSF of Figure 14.10.

To simulate this acquisition by the discretised method, we generate a series of profiles at increasing time steps, which correspond to the initial signal, all the way up to the final profile which has moved an amount designated by the width of the PSF. This series of profiles are them summed. Some of the profiles in the series are shown in Figure 14.12a, and the final result is shown in Figure 14.12b. It can be seen that this is identical to the profile obtained by direct convolution (Figure 14.11).

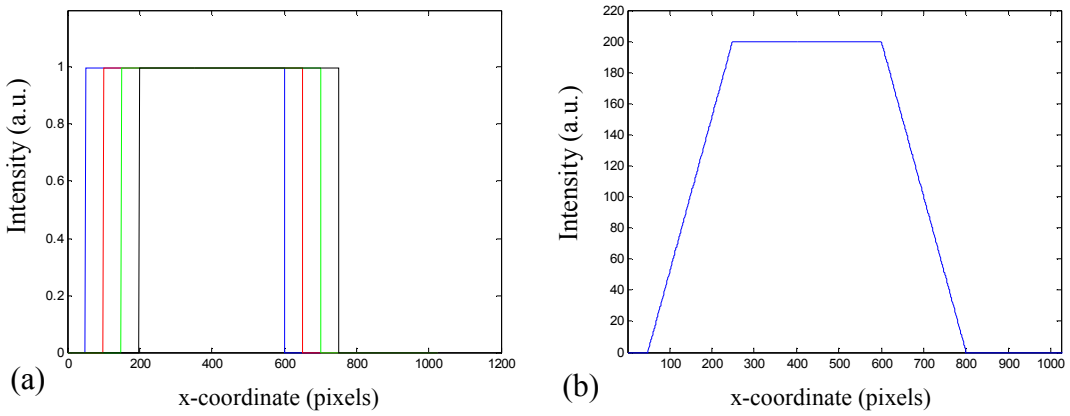


Figure 14.12 – Illustration of the discretised motion-blurring method; (a) shows some of a series of profiles generated to correspond to the position of the true signal over the observation time – the movement between the first and last profiles is equivalent to the distance denoted by the width of the PSF; (b) shows the result of summing the series of profiles. This data is equivalent to the result obtained by direct convolution of the true signal and the PSF.

## 14.4.5. Non-linear motion-blurring

The simulation described above is for a linear motion-blur, which is analogous to plug flow. However, in the case of the monolith we do not have plug flow; instead the propagator stretches with time, rather than shifting with time. This is because fast spins move further than slow spins (by a proportionate amount), hence the effect is to elongate the propagator in a linear manner with time. Since the number of spins is conserved, the height of the propagator must reduce to give the same overall area under the shape (barring relaxation effects which will be considered later).

If we begin with the propagator for flow in a square pipe, we obtain a series of profiles described in Figure 14.13. The result of summing this series of profiles results in the "observed" signal shown in Figure 14.14.

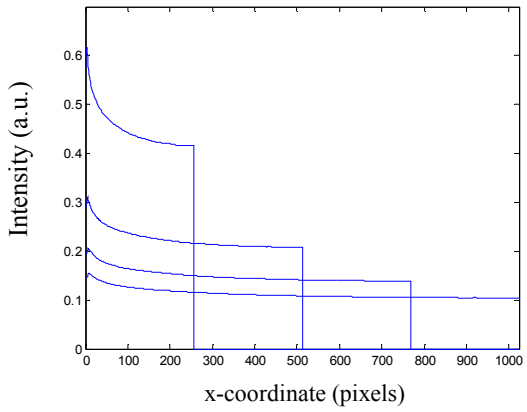


Figure 14.13 – The evolution of a propagator (for flow in a square pipe) as a function of time. The propagator stretches, rather than being shifted as was the case for linear blurring.

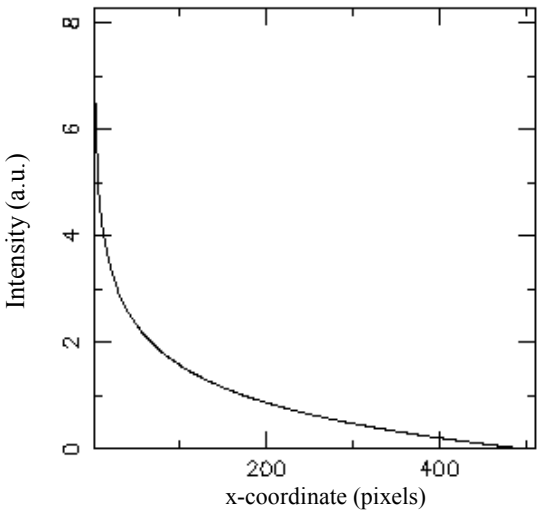


Figure 14.14 – The result of summing the series of profiles described in Figure 14.13. This is the simulated "observed signal" associated with flow in a square pipe.

This process of stretching and summing the propagator is not obviously described by a convolution step – it could possibly be a non-linear convolution, but this is beyond the scope of this project. Hence we must seek another method to reverse the process and obtain the original propagator from the observed, composite signal; i.e. to reconstruct the original propagator from the composite data of Figure 14.14.

It is proposed that an iterative method is implemented – the chosen method is a simple brute force approach, similar in nature to the Simulated Annealing method of blind deconvolution (Section 13.4.5), and is detailed in the following section.

## 14.4.6. The iterative de-blurring algorithm

In order to reconstruct the original propagator from the blurred result, we may begin by guessing an initial starting function for the propagator. We then randomly perturb it, apply the known non-linear blurring process, and then evaluate the error in the output compared to the actual observation. We may then iterate and approach the correct propagator by choosing those results producing the least error. It is a computationally intensive approach, however it is by far the simplest and most robust way to obtain the solution we are seeking, and should perform well even in the presence of noise.

There is however, a caveat. This procedure will be rather insensitive and we should not expect to obtain a high degree of precision in the solution. To illustrate this, a comparison between the simulated observed signal for flow in a square pipe is shown with that of flow in a circular pipe in Figure 14.15. It can be seen that the results are very similar, and hence it is unlikely that we would be able to distinguish the subtle differences such as these in the original propagators.

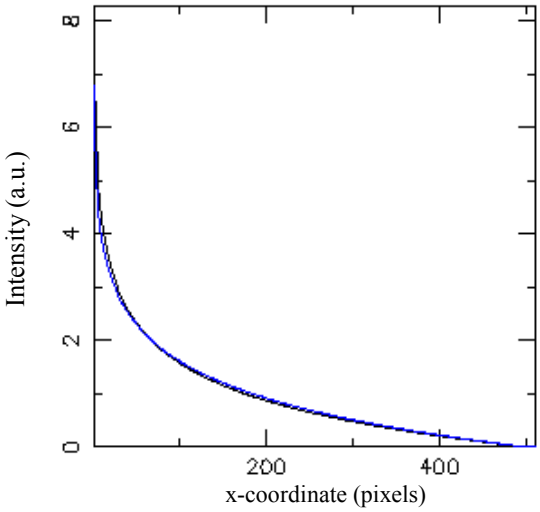


Figure 14.15 – A comparison of the results of the simulation procedure; the blue line shows the observed signal for laminar flow in a circular pipe, and the black line shows the observed signal for laminar flow in a square pipe.

The Simulated Annealing algorithm (Section 13.4.5) applies a uniform perturbation to the input, however this is not the fastest way to approach a solution. It is proposed that we begin by making coarse changes to the guessed input in order to approximate the rough shape, and then reduce the width of our perturbations to refine and add detail to the curve. The amplitude of the perturbations are reduced, and repeated until no more perturbations are accepted. We continue iterating until we have a sufficiently precise solution.

Some robustness tests were performed with simulated data. The test input was created using the propagator for laminar flow in a square pipe and this was transformed with the simulated blurring, as described in Section 14.4.5. The input signal was 1024 pixels in length, and the following iterations were performed; 10,000 iterations with perturbation width of 50 pixels; 10,000 iterations with perturbation width of 25 pixels; 10,000 iterations with perturbation width of 5 pixels. The initial input to the iterative algorithm was a signal whose value was 1 throughout. The results are shown in Figure 14.16, with the result overlaid on the true solution.

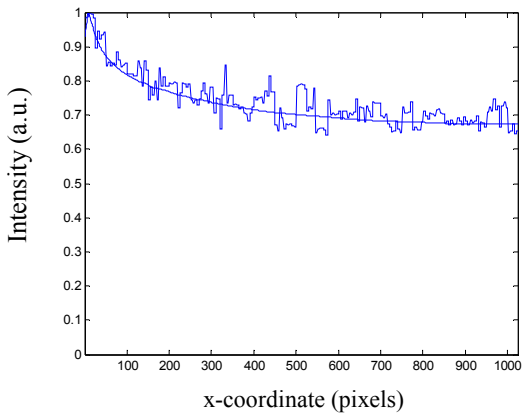


Figure 14.16 – The result of the initial iterative algorithm, overlaid on the true solution to the system (depicted by the smooth line).

However, the iteration became stuck in a local minimum, and no further perturbations were accepted by the algorithm. One reason why we may become trapped in a local minimum is that when, for example, we make a perturbation such that one of the under-estimated areas is increased, the overall area under the graph is therefore increased, and during the simulation process the function is normalised. This means

that every other point on the curve will be reduced slightly, and hence the overall error may in fact increase. No matter how many iterations we perform of even the smallest magnitude, the error will always be unfavourable, and we remain trapped in this local minimum.

There are several possibilities for improving this convergence:

(i) We could accept some of the unfavourable perturbations in order to push convergence past such local minima. This is similar to the Simulated Annealing process, and our result illustrates the reason for its inclusion in their algorithm. The adjustment of what proportion of unfavourable perturbations to be accepted needs investigation, since if we accept too many the rate of convergence will be extremely slow.

(ii) Rather than making just one perturbation, we could make two at random positions – one positive and one negative. Hence the overall sum will remain the same, and the problem of normalisation will be irrelevant. The convergence time may increase.

(iii) We could add smoothing steps. Since we know that the true solution is likely to be a relatively smooth function (rather than a function with jagged peaks) we could add a smoothing function every few thousand iterations to remove the unnecessary detail, to help average out these errors and accelerate the rate of convergence. Then the algorithm can further improve the convergence from the smoothed approximation.

Option (ii) was chosen because it is simple to implement, and will not adversely affect the end result. The convergence will be slower as a consequence, since we must find pairs of perturbations which lower the overall error. This is less likely than finding one useful perturbation, and so the convergence will be slower, especially towards the end of the iteration where many perturbations are rejected.

Applying the new algorithm results in the profile of Figure 14.17a. The error associated with this solution is $7.3\times10^{-8}$. We can see that the results are much closer to the true solution than for the previous algorithm (Figure 14.16). A further 10,000 steps of the smallest width were performed, yielding the result of Figure 14.17b. The

error associated with this solution is $5.5\times10^{-8}$. The error in this second iteration is 25% lower than that of the previous result, however it is nearly identical in appearance. This indicates that the problem may be ill-conditioned, and we may in fact require some smoothing passes (option iii).

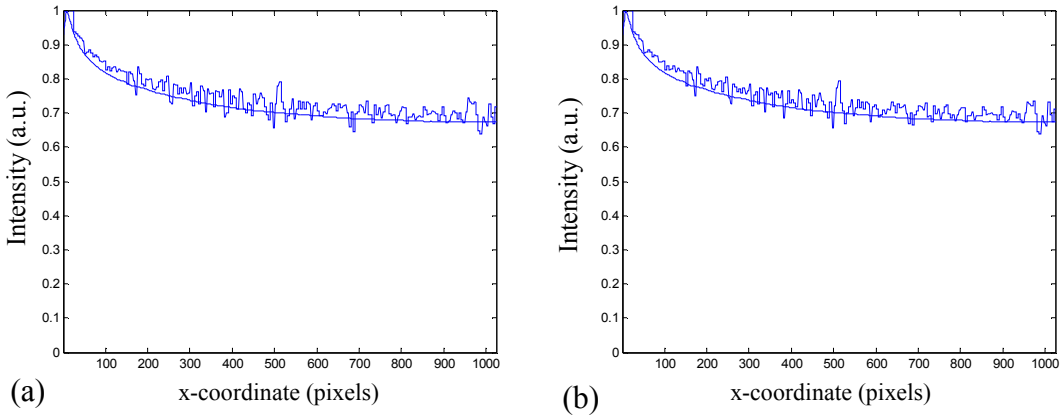


Figure 14.17 – The result of the improved iterative algorithm, (a) after 10,000 steps, and (b) after 20,000 steps. The profiles are overlaid on the true solution (depicted by the smooth line).

A quadratic Savitsky-Golay[2] smoothing pass (with a filter width of 101 pixels) was applied to the solution of Figure 14.17b, and then further iterated. The result is shown in Figure 14.18 (note the different scale on the y-axis), and the solution appears much improved to the eye. The error associated with this solution is $9.0\times10^{-6}$. This error is larger than the previous case, however the apparent fit is clearly much better and our result is very close to that of the original signal.

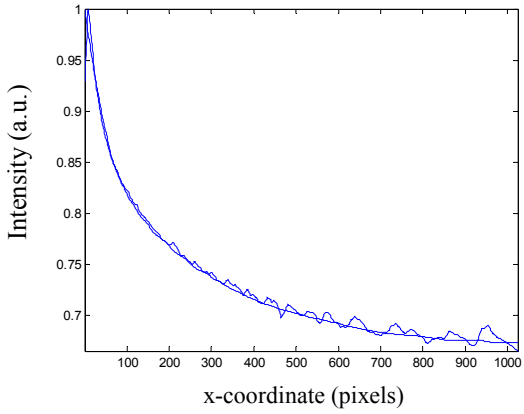


Figure 14.18 – The result of applying a smoothing pass to the data of Figure 14.17b.

[2] This is a standard function, built into the Matlab toolbox. Further information is available at http://www.mathworks.com/access/helpdesk/help/toolbox/signal/sgolay.html

Further iterative steps lead to the result of Figure 14.19a (error of $1.8\times10^{-8}$). Even though the error is lower, this is clearly a poorer fit than the previous solution. Hence, it would appear that there is a natural tendency for noise to appear in the solution, and a smoothing pass is necessary to retain a physically realistic solution. The error is also not necessarily representative of the quality of the fit. The results of a second smoothing pass are shown in Figure 14.19b. The result is now an excellent approximation for the original input function.

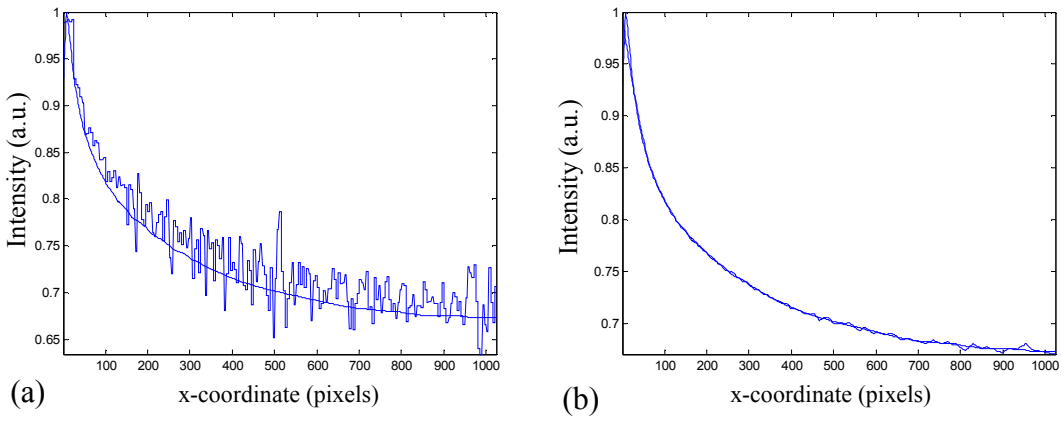


Figure 14.19 – Plots showing the result of (a) 20,000 more iterative steps on the profile of Figure 14.18, and (b) shows the result of a further smoothing pass on (a).

Hence, it has been shown that this method can, through a combination of iteration steps and smoothing passes, reverse-engineer the original signal from an observation of the blurred output. It should be noted that this simulated data has a far higher resolution than we require the code to run on – typically MRI profiles will be 64 or 128 pixels in length rather than 1024 pixels. The high resolution examples are simply to show the proof of concept and to test the performance of the algorithm.

## 14.4.7. Inclusion of the excitation convolution

It is relatively straightforward to include the convolution of the sinc profile (the initial spread of the distribution, caused by the shape of the excitation pulse) in this simulation, rather than to try to perform the simulation and later deconvolve the sinc contribution.

To include it in the simulation, we simply convolve together the series of profiles with the initial pulse shape (as acquired during zero flow conditions), and we have our combined simulation algorithm. The iterative procedure is then able to simulate both steps (convolution of the propagator with the initial slice shape, and then non-linear motion-blurring).

The sinc function used in this test is shown in Figure 14.20a, and the resulting series of profiles generated by the simulation are shown in Figure 14.20b. The output (containing both the convolution of the sinc shape, and motion-blurring) is shown in Figure 14.20c.

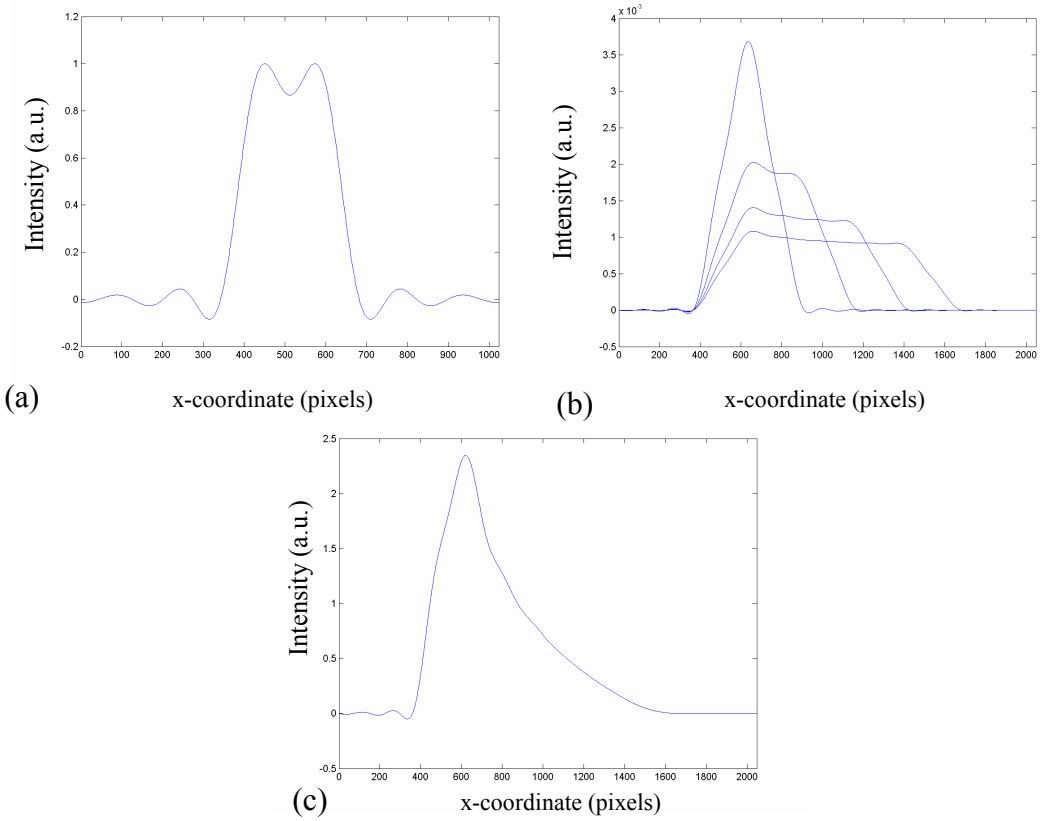


Figure 14.20 – Plot of (a) the sinc function used; (b) the result of convolution of (a) with the evolving series of profiles of the simulation (Figure 14.13); and (c) is the output of the simulation procedure.

Running the new iterative algorithm with inputs of just the observed data of Figure 14.20c and the known sinc function, produces the predicted profile of Figure 14.21 which is in very close agreement with the original input, shown as the smoother line in the figure.

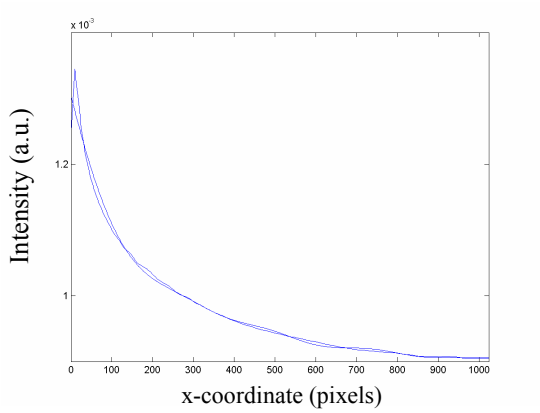


Figure 14.21 – Plot of the results of the iterative algorithm, overlaid on the original function.

We have now developed a robust algorithm to simulate the non-linear blurring of a propagator, and an iterative procedure to reconstruct the original propagator from the observed profile and knowledge of the initial excitation shape. The algorithm is illustrated in Figure 14.22. It has been shown to work well for noise-less simulated data. The next step is to apply the process to real MRI data and evaluate its performance.

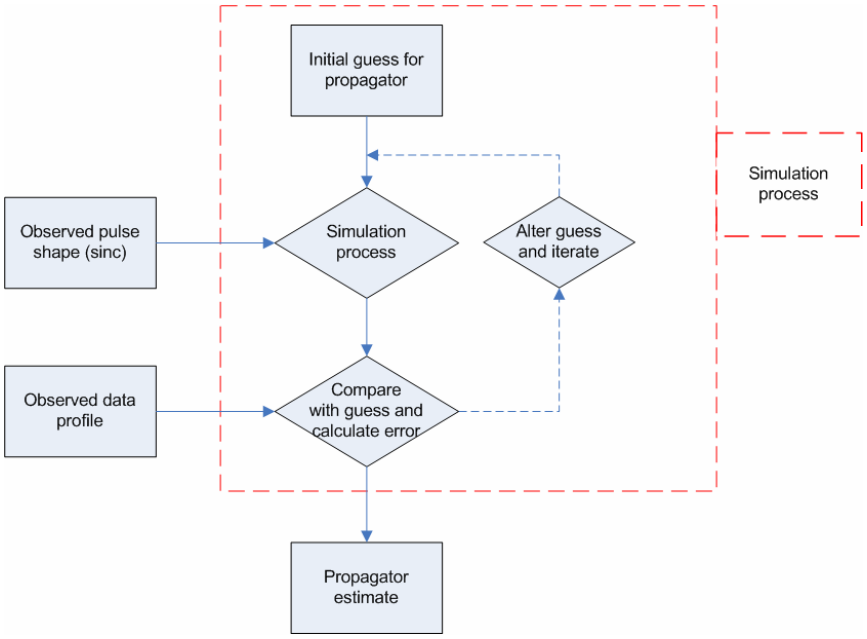


Figure 14.22 – Flowsheet of the iterative algorithm, for the simulation of non-linear motion-blurring. The inputs to the system are the observed, blurred MRI profile, and the known pulse shape (acquired from a zero-flow observation).

## 14.4.8. Real MRI data

MRI data were acquired for single-phase flow in a 200 cpsi monolith reactor at a range of liquid flow rates (setup in Section 6.1). Sample data are shown in Figure 14.23(a-b) for zero flow and a single-phase liquid flow rate of 29 l/hr. The channels are processed individually, averaging together voxels associated with each. Sample profiles are shown in (c-d). These profiles were then interpolated to higher resolution and cleaned to remove noise. This was done by lowering the baseline, and regions where no flow existed were nulled, as shown in (e-f).

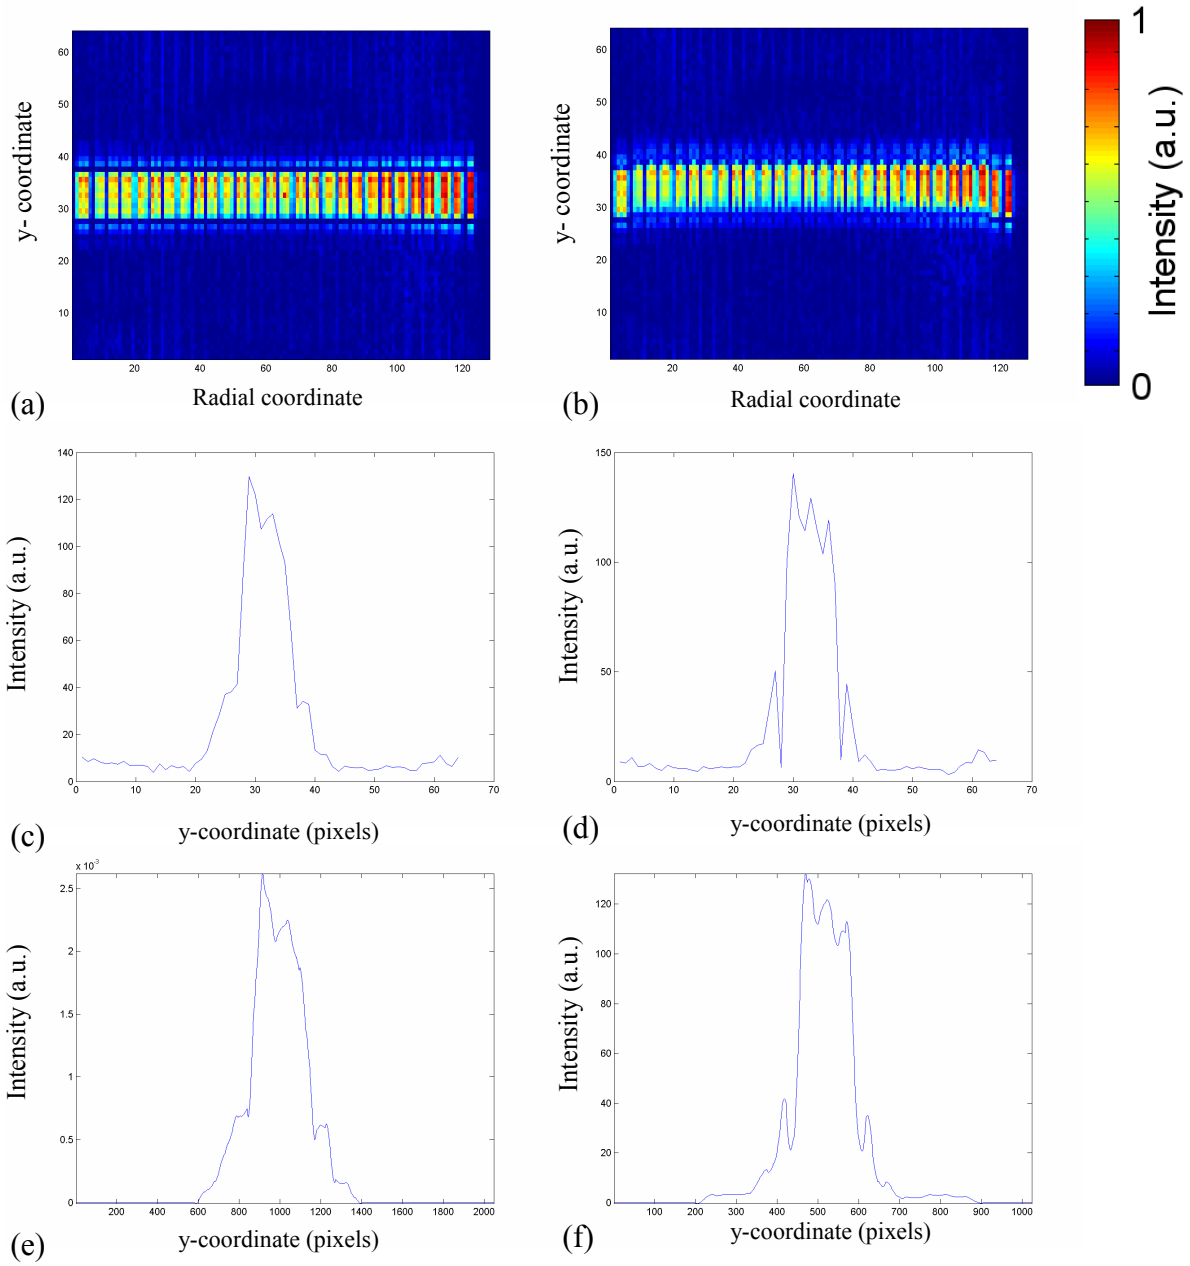


Figure 14.23 – Sample MRI data acquired at (a) zero and (b) 29 l/hr of water, in the 200 cpsi monolith; (c-d) show sample profiles; (e-f) show profiles cleaned to remove noise.

The iterative algorithm was then applied to predict the original propagator, taking as inputs the convolved and blurred observation (Figure 14.23e) and the known initial excitation shape (Figure 14.23f). The calculated propagator is shown in Figure 14.24a. However, it is not a physically realistic propagator; there are no spins at low displacement, and a sharp spike at high displacement. If we use this propagator to reconstruct the observed signal, we obtain the profile of Figure 14.24b. It can be seen that this does not compare very favourably with the actual observed signal (Figure 14.23e).

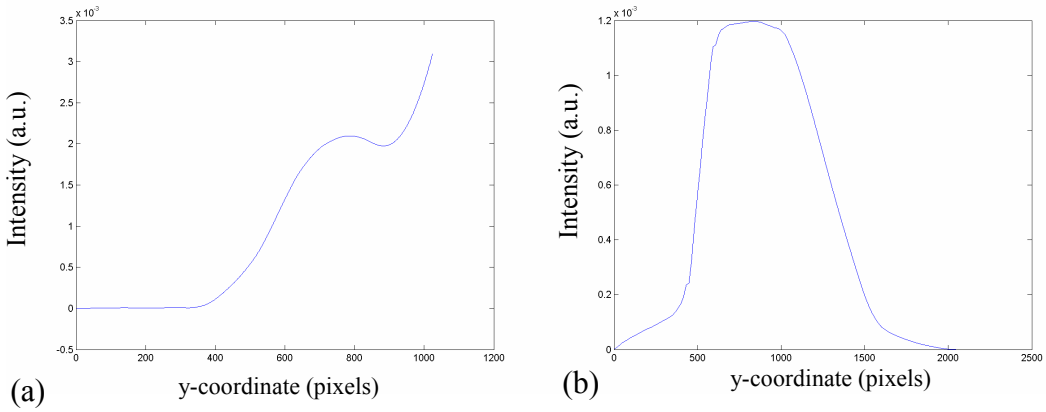


Figure 14.24 – The results of applying the iterative deconvolution procedure to real MRI data; (a) shows the obtained propagator, and (b) shows the "observed signal" this propagator would produce.

This deviation from satisfactory behaviour is believed to be due to nuclear spin relaxation effects within the monolith, which have not been accounted for in the model. In reality, spins will decay at a rate characterised by their $T_1$ value. This means that profiles at later time steps should be given less weighting in the final simulated output. Also, spins at small displacement (near the walls) will experience faster relaxation than the bulk due to interaction with the walls, and higher shearing. Fast-moving spins will also appear to relax faster than the bulk due to diffusive attenuation. Therefore, we will actually lose signal from both the left- and right-hand sides of the propagator, and this needs to be built into our model.

An explanation of the unusual shape of the obtained propagator profile is given by the following. The relaxation effects mean that the MRI experiment has very low signal at low velocities, whereas in the model we *always* have high signal at low velocities.

The iterative model recognises that the intensity at low velocities are low, and the only way it can generate a profile to fit this, is to put a large spike at very *high* velocities. In the simulation process, where we sum together profiles at different time steps (as shown in Figure 14.25) we can see that this spike has a bigger impact at low displacements (on the left of the plot, where many profiles are averaged together) and a low impact at high displacements (on the right of the plot, where only a handful of plots may be averaged together). Therefore, by producing a propagator with a spike at high velocities, the algorithm has succeeded in following the rules of its construction, but has skewed the result because the construction is flawed and does not adequately represent the real system.

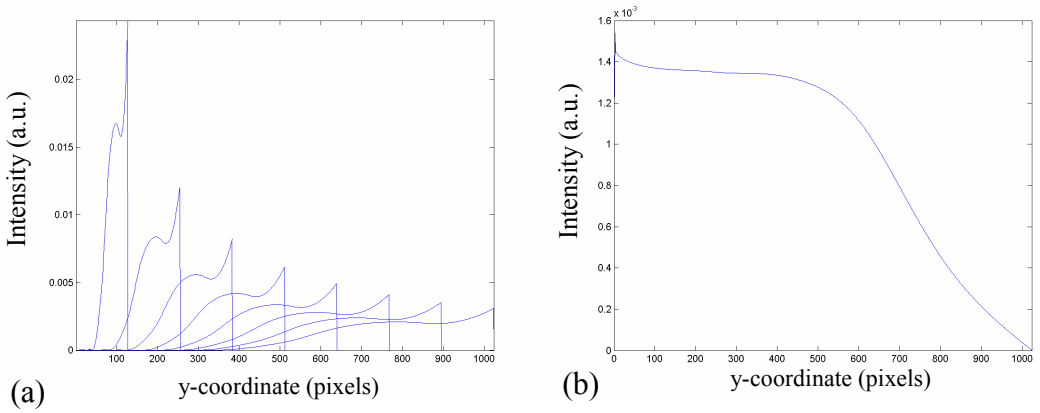


Figure 14.25 – Illustration of the simulation procedure using the profile of Figure 14.24a. Profiles for the spin distribution at increasing time steps are shown in (a). When summed, these profiles give the plot of (b). When the profiles are convolved with the sinc function, they result in the plot of Figure 14.24b.

Hence, we must now incorporate the relaxation effects into the model in order to produce any sort of meaningful result. The following section will address the means by which we can do this.

## 14.4.9. Relaxation effects

As mentioned in the previous section, we need to take into account:

(i) $T_1$ relaxation,

(ii) $T_2$ relaxation, and

(iii) velocity attenuation effects.

Each of these effects will be considered in turn in this section.

Relaxation effects are a little complex in the monolith system, since the relaxation time varies with the local environment – i.e. spins near the wall of the channel relax more quickly than the bulk fluid, giving a range of inherent relaxation times. This will be exaggerated further in the presence of flow, since the spins near the walls will experience greater shearing, and will relax even more quickly. These two relaxation effects are rather hard to quantify, but can be measured by a range of techniques.

### 14.4.9.1. $T_1$ calculation

We know that each successive echo in the acquisition is acquired 2.44 ms apart, and we know that the measured $T_1$ for the system is approx 1.5 seconds. Therefore, we can weight each successive profile in our simulation by the appropriate relaxation factor, shown in Eq. 14.1. In fact, this is a relatively small effect, and even at the end of the acquisition the signal has only decayed to 90% of the original value.

$$T_1 \text{ Relaxation Factor } (t) = \exp\left(-\frac{t}{T_1}\right) \qquad \textbf{(Eq. 14.1)}$$

where $t$ is the time associated with the data, and $T_1$ is the characteristic spin-lattice relaxation time. This factor is a function of time, and each profile is simply multiplied by the appropriate factor. This means that later profiles will have less of a contribution to the overall observed signal.

### 14.4.9.2. $T_2$ and flow-attenuation

The other relaxation properties of the monolith system could be measured by a spatially-resolved $T_2$ map. However, this would not take into account the flow attenuation present in the system. A method was developed to measure both types of relaxation effects simultaneously, utilising the SEMI-RARE pulse sequence. The RARE echo train inherently measures the $T_2$ properties of the system since it is $T_2$-weighted. By acquiring RARE images with a RARE factor of 1, we obtain an image where every point in **k**-space corresponds to the same observation time after the initial excitation (rather than in a one-shot RARE image, where the points in **k**-space are acquired at increasing times after the initial excitation). If we acquire a train of images with the SEMI-RARE sequence, we will have a series of images

corresponding to increasing observation times; e.g. if the echo time corresponding to the RARE sequence is 5 ms, then we acquire an image where every pixel corresponds to an acquisition time of 5 ms, another image at 10 ms, the third at 15 ms, and so on. By comparing the contrast of pixels from one image to the next, we may then fit a relaxation curve, and extract an apparent characteristic relaxation time. This time will combine both the $T_2$ relaxation and flow attenuation effects.

For clarity, an illustration of the relaxation fitting process is shown in Figure 14.26. This shows the decrease in intensity of a given voxel, over the course of the SEMI-RARE acquisition. The intensity of the voxel decreases with each image in the SEMI-RARE train, since the echo time of the observation is increasing. By fitting an exponential decay to this data, we may calculate an apparent relaxation time constant. This is then repeated for each voxel in the data. Relaxation maps for the real data are shown in Figure 14.27, for zero flow rate and a flow rate of 95 l/hr.

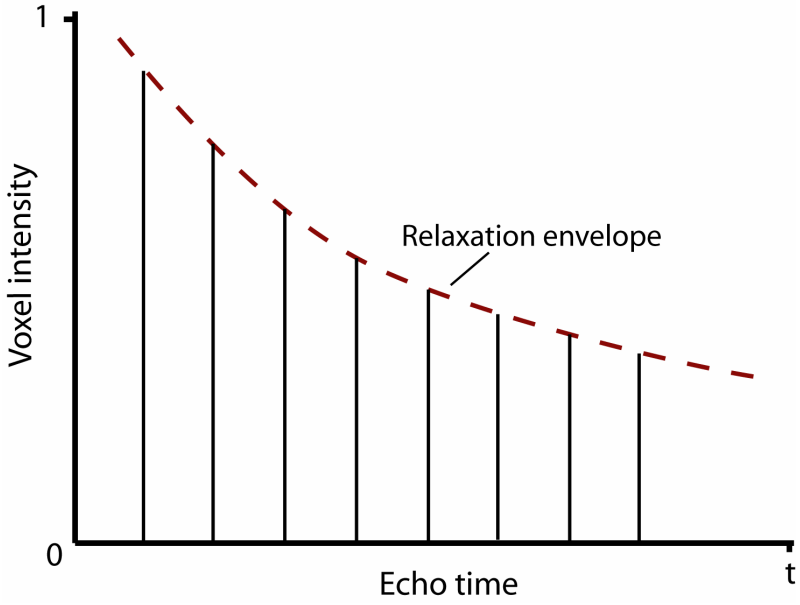


Figure 14.26 – Illustration of the fitting procedure to calculate an apparent relaxation rate constant. Sample data for a single voxel are shown, at increasing observation time (increasing echo time).

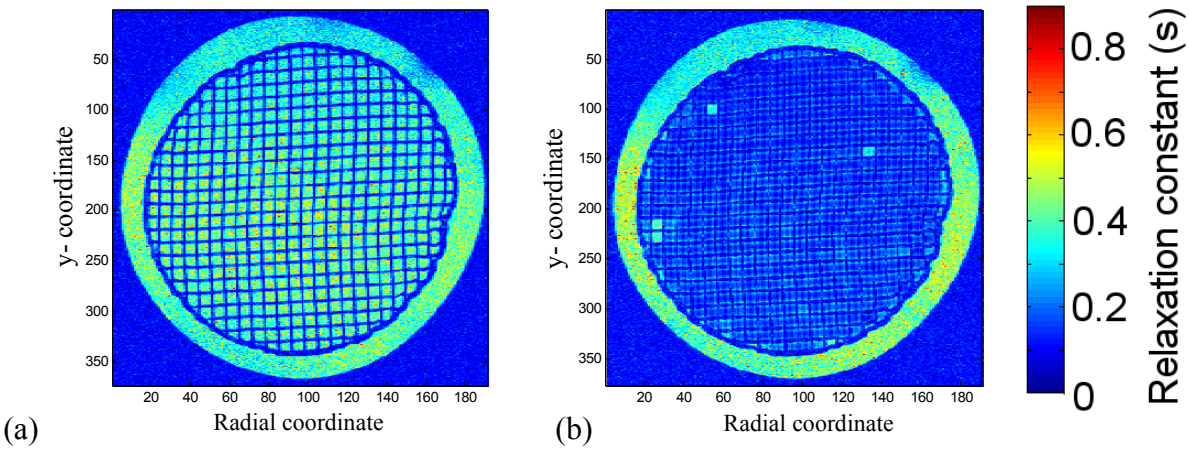


Figure 14.27 – High-resolution flow-relaxation maps. Computed from a SEMI-RARE measurement, with a RARE factor of 1 and an echo time of 5.278 ms. Intensity scale is in seconds; (a) shows the zero-flow case, and (b) is the flowing case at 95 l/hr.

If we examine the data of Figure 14.27a, we can see that the relaxation rate of the monolith walls and the fluid near the walls is much faster than for the bulk fluid. The relaxation times are in general short compared to the bulk properties of water, which is around 3 seconds. Figure 14.27b shows the effect of flow attenuation. It is seen that even though the shearing at the walls has an effect, the biggest effect is to the bulk fluid – this is due to diffusive attenuation.

It is stressed that this does not represent a quantitative $T_2$ analysis. The reasons for this are two-fold; firstly, diffusive attenuation, and hence additional contrast within the images arises due to the application of a magnetic field gradient; secondly, the imperfect refocusing of the nominally 180° RF refocusing pulse will introduce a proportion of $T_1$ relaxation while the magnetisation is in *z*-storage, thereby increasing the observed decay time constant. Sederman *et al.* (2003) reported that the $T_2$ values obtained in zero flow are 10% higher than those taken with a CPMG train.

The information within these data sets was used by correlating the reduction in intensity observed with the local velocity of fluid within each voxel. This was done by sorting the apparent relaxation times of the data pixels in Figure 14.27b into order from lowest to highest characteristic relaxation time. It is proposed that short relaxation times correspond to areas of low velocity, and *vice versa*. For the purposes of this approximation, we will linearly map the velocity distribution to the relaxation

distribution; i.e. if the velocity of a spin corresponds to the 25$^{th}$ percentile (between the min and max velocities present in the monolith), then we will use the relaxation look-up table, and choose the value associated with the value of the 25$^{th}$ percentile (between the min and max values).

The relaxation look-up table for the combined effects of $T_2$ and flow-attenuation is shown in Figure 14.28a. A look-up table for diffusive attenuation was also constructed in a similar manner. RARE acquisitions for zero flow and for a flow rate of 95 l/hr were compared. The apparent relaxation factor was calculated for each pixel, according to Eq. 14.2.

$$\text{Relaxation Factor} = \frac{I_F}{I_Z} \qquad \textbf{(Eq. 14.2)}$$

where $I_Z$ is the intensity in the zero flow data set, and $I_F$ is the intensity in the flowing data set. Calculating this array of values, and ordering them as before gives the look-up table of Figure 14.28b.

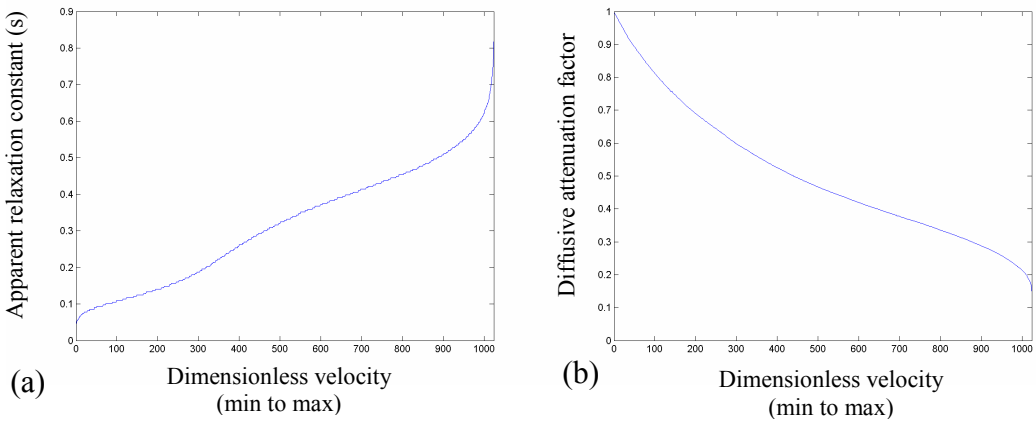


Figure 14.28 – Plots of the look-up tables for (a) combined $T_2$ / flow-attenuation characteristic relaxation times; and (b) diffusive attenuation factors. The *x*-axes in both cases correspond to the velocity of the trial propagators (from min to max), and the plots are used as look-up tables to obtain the appropriate correction factors. The effect of (a) is to reduce the contribution from slow-moving spins (where the time constant is short), and the effect of (b) is to reduce the contribution from fast-moving spins (where the diffusive attenuation factor is small).

As with the $T_1$ correction factor (Section 14.4.9.1), we know the time of the associated profiles we are simulating. The respective relaxation factors for $T_2$ / Flow attenuation and for diffusive attenuation are given in Eq. 14.3 and Eq. 14.4. These are a function both of time and velocity, since the relaxation factors depend on the local velocity of the spins.

$$T_2 \text{ and Flow-attenuation Factor } (t,v) = \exp\left(-\frac{t}{T_A(v)}\right) \qquad \textbf{(Eq. 14.3)}$$

$$\text{Diffusive Attenuation Factor } (v) = F_D(v) \qquad \textbf{(Eq. 14.4)}$$

where $T_A$ is the apparent characteristic relaxation time associated with $T_2$ and flow attenuation, and $F_D$ is the factor obtained from the diffusive attenuation look-up table. These two values are obtained from the look-up tables of Figure 14.28. The total relaxation factor is therefore given by Eq. 14.5.

$$\begin{aligned}\text{Relaxation Factor } (t,v) &= \exp\left(-\frac{t}{T_1}\right)\times\exp\left(-\frac{t}{T_A(v)}\right)\times F_D(v) \\ &= F_D(v)\exp\left[-t\left(\frac{1}{T_1}+\frac{1}{T_A(v)}\right)\right]\end{aligned} \qquad \textbf{(Eq. 14.5)}$$

These relaxation factors were then built into the iterative algorithm, and the results are given in the following section.

## 14.4.9.3. Results for single-phase flow in a monolith

The results of applying the new iterative deconvolution method to the real MRI data (Figure 14.23b) results in the data of Figure 14.29. Plot (a) shows the resultant propagators for 26 channels (all acquired from a single line-excitation image), and (b) shows a sample propagator from (a). It can be seen that there is some variation in the propagators, and this is believed to be due to processing artefacts such as local minima in the iterative process.

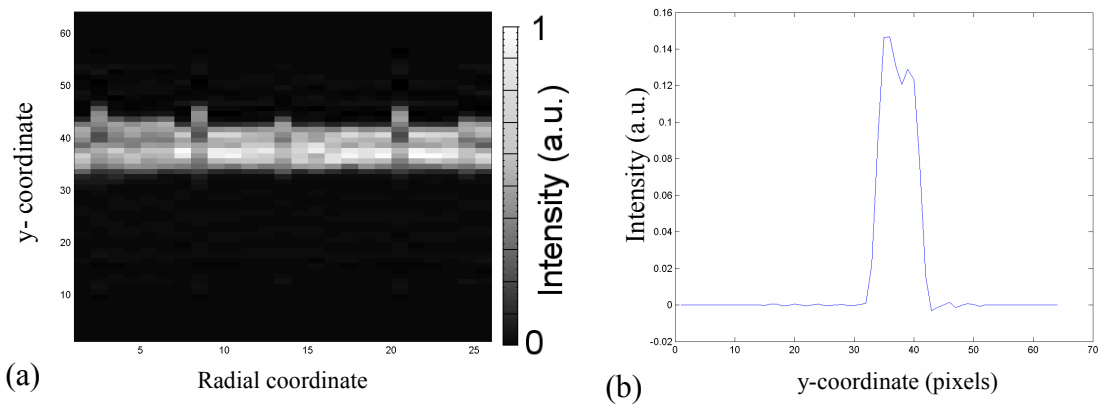


Figure 14.29 – Result of the new analysis applied to single-phase liquid flow at 86 l/hr; (a) shows a series of 26 propagators, corresponding to a row of 26 channels, acquired from a single RARE image in 156 ms; (b) shows a sample propagator from (a).

Applying the process to a range of flow rates results in the propagators of Figure 14.30. In each of these, 26 channels and 32 repetitions have been averaged together to get a representative propagator for the whole system. (However, if individual channel behaviour is of interest, the data may be treated individually).

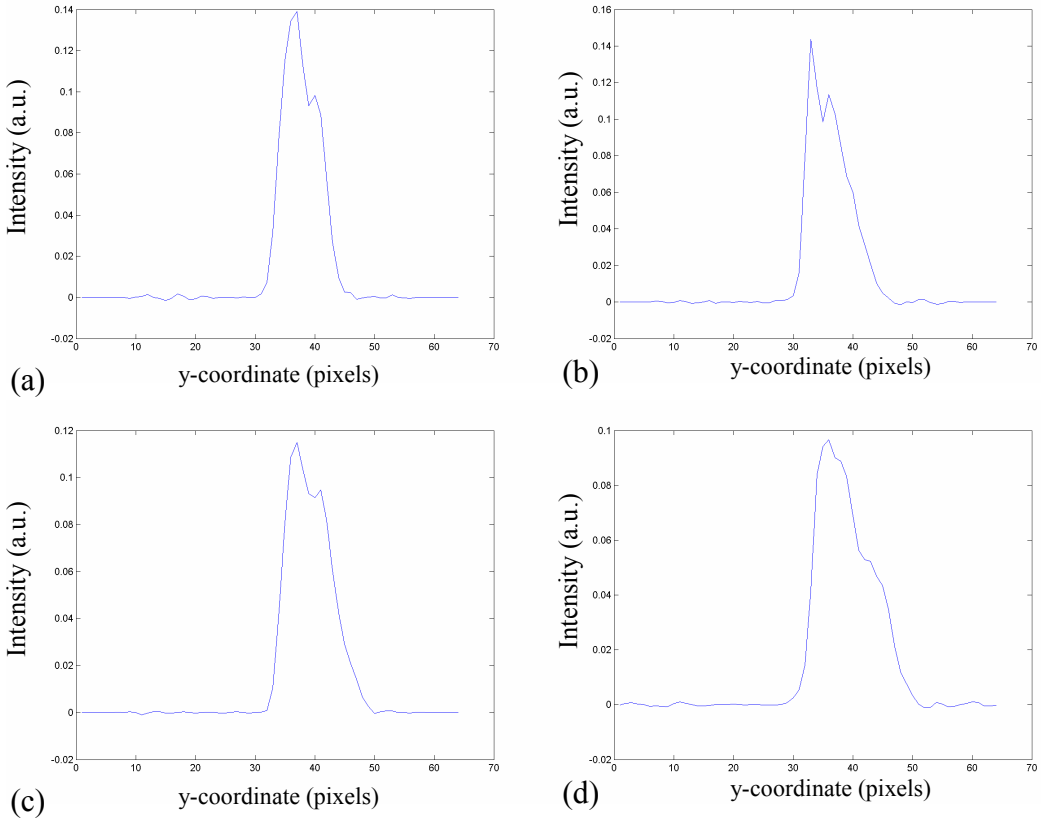


Figure 14.30 – Average propagators for flow rates of (a) 29 l/hr, (b) 58 l/hr, (c) 86 l/hr, and (d) 115 l/hr. In each plot, 26 channels and 32 repetitions were averaged together to obtain the plots shown.

## 14.5. Two-phase flow

This section will cover some simple two-phase flow systems, namely the bubble sparger and co-current flow setups (see setup in Section 8.2.1). MRI data were acquired as for the single-phase flow case, and the calculated propagators are averages of 26 channels, and 32 repetitions.

### 14.5.1. Bubble sparger

The results for the bubble sparger system at a range of gas flow rates is shown in Figure 14.31.

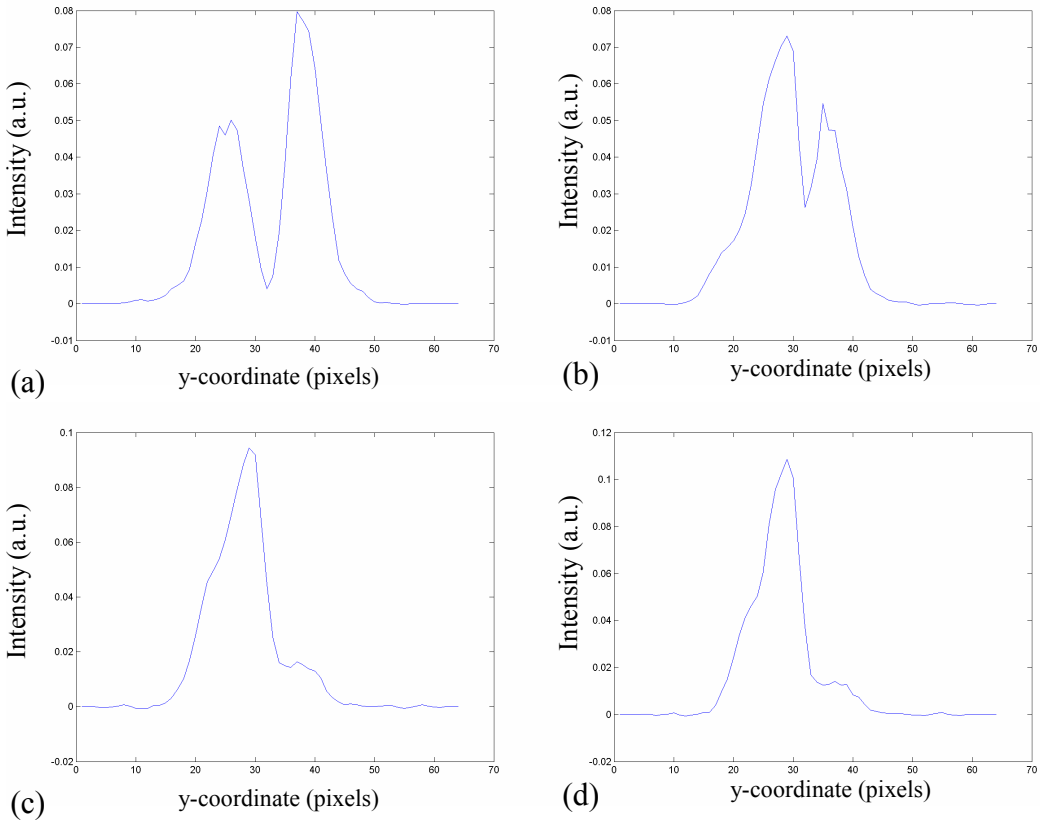


Figure 14.31 – Resultant propagators for the bubble sparger system operating at gas flow rates of (a) 100, (b) 200, (c) 300, and (d) 500 cc/min, through stagnant fluid.

## 14.5.2. Co-current flow

The results for the co-current flow system at a range of liquid flow rates and gas flow rates of 100 cc/min is shown in Figure 14.32, and at a gas flow rate of 1000 cc/min is shown in Figure 14.33. The plots show that on increasing the gas flow, the histograms spread to higher velocities, both positive and negative, with fewer channels stagnant with near-zero velocities.

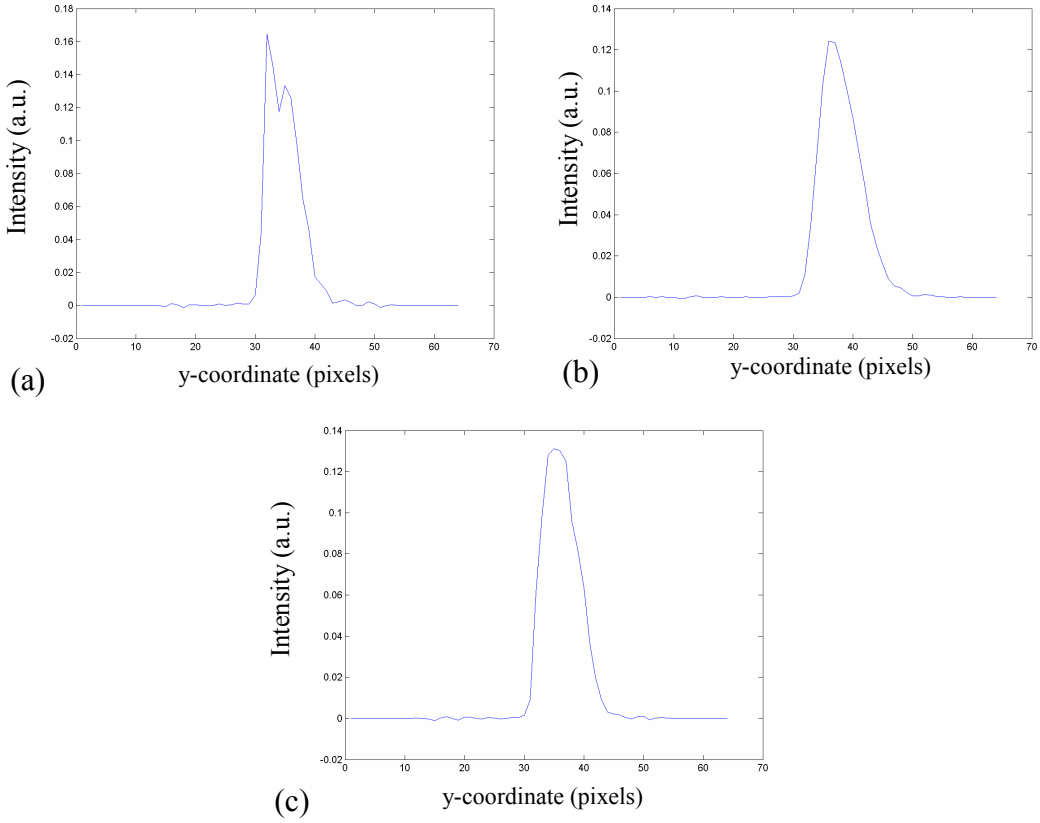


Figure 14.32 – Obtained propagators for the co-current flow system at a gas flow rate of 100 cc/min, and liquid flow rates of (a) 29 l/hr, (b) 78 l/hr, and (c) 106 l/hr.

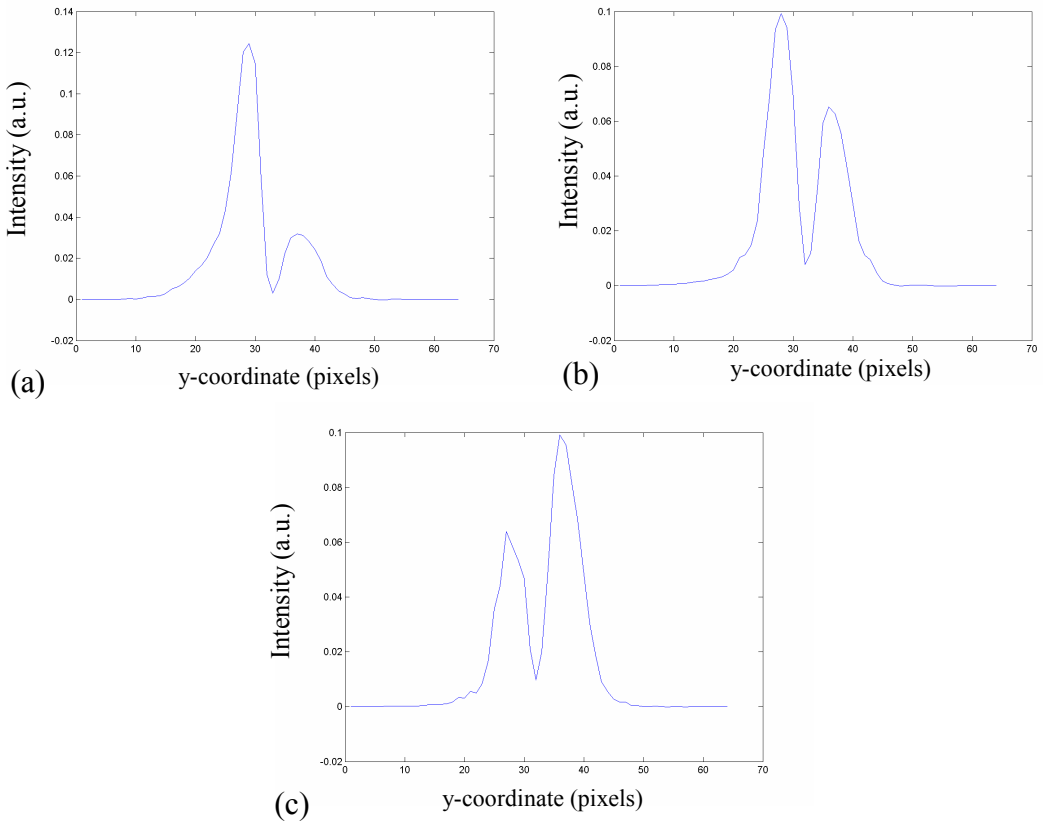


Figure 14.33 – Obtained propagators for the co-current flow system at a gas flow rate of 1000 cc/min, and liquid flow rates of (a) 29 l/hr, (b) 78 l/hr, and (c) 106 l/hr.

## 14.6.Conclusions

A novel analysis method has been developed and presented, which is capable of deblurring MRI data acquired with the line-excitation method in such a way that we can calculate propagator information for the monolith reactor system. The line-excitation sequence acquires this information in just 156 ms, and allows us to extract propagator information for each channel in the row under observation. Application is only possible to systems with suitable geometry (we require axial flow), but is a much faster technique than traditional propagator measurements (Section 2.4.2) and allows the study of dynamic two-phase flow systems. The maximum flow rate that can be imaged depends on the signal to noise ratio, and hence on the channel thicknesses, material properties and imaging parameters. In this case, velocities of up to 5 cm/s have been investigated.

Although the results obtained are not as accurate as traditional propagator measurements, the analysis presented is based on a very rapid acquisition technique and is applicable to dynamic systems where traditional techniques would fail. This post-processing of our data allows us to extract additional information from existing data in a robust and repeatable manner.

## 14.7. Further notes

The line-excitation data were acquired using the SEMI-RARE acquisition procedure, thereby giving four images from one excitation. The first image is by far the clearest, so this was used for the propagator analysis. As a side note, the analysis can be applied to any of the subsequent images (thereby giving the evolution of the propagator with time),

## 14.8. Nomenclature

| **Symbols** | **Definition** |
|---|---|
| $f, g$ | Functions (in real space) |
| $F(\cdot)$ | Fourier transform |
| $F_D$ | Diffusive attenuation factor |
| $I_F, I_Z$ | Intensity from flowing data, and from zero-flow data |
| $\mathbf{k}$ | Reciprocal space vector |
| $K$ | Weiner filtering parameter |
| $t$ | Time |
| $t_a$ | Acquisition time |
| $T_1$ | Spin-lattice relaxation constant |
| $T_2$ | Spin-spin relaxation constant |
| $T_A$ | Apparent (combined) relaxation time constant |
| $v$ | Velocity |
| | |
| $\beta$ | Weighting parameter for IBD algorithm |
| $*$ | Convolution symbol |

| **Abbreviations** | **Definition** |
|---|---|
| CPMG | Carr-Purcell-Meiboom-Gill |
| cpsi | Cells per square inch |
| IBD | Iterative blind deconvolution |
| MRI | Magnetic resonance imaging |
| NMR | Nuclear magnetic resonance |
| PSF | Point spread function |
| RARE | Rapid acquisition with relaxation enhancement |
| RF | Radio frequency |
| SEMI-RARE | Single excitation, multiple image RARE |

## 14.9. References

Heras, J. J., A. J. Sederman*, et al.* (2005). "Ultrafast velocity imaging of single- and two-phase flows in a ceramic monolith." Magnetic Resonance Imaging **23**(2): 387-389.

Sederman, A. J., M. D. Mantle*, et al.* (2003). "Single excitation multiple image RARE (SEMI-RARE): ultra-fast imaging of static and flowing systems." Journal of Magnetic Resonance **161**(1): 15-24.

# Chapter 15 – Conclusions and Future Work

---

**CHAPTER 15 – CONCLUSIONS AND FUTURE WORK**



---

## 15.1. Conclusions

The development and quantification of novel nuclear magnetic resonance (NMR) and magnetic resonance imaging (MRI) techniques have been presented. These were applied to the various process configurations of the monolith reactor, with the aim of characterising the flow phenomena in such porous systems.

Parameters such as holdup, bubble and slug sizes, fluid and velocity maldistribution, stability criterion, and dynamic velocity data have been measured in these challenging, dynamic, porous, two-phase systems. These operational parameters were acquired as a function of process configuration, flow rate and design geometry. Certain MRI techniques were not applicable to the most high flow rate configurations, and the limitations of each technique were discussed. However, the COTTAGE technique performed beyond expectations, and allowed the observation of dynamic fluctuations due to its very high temporal resolution.

The development and implementation of novel image processing techniques allowed the robust processing of MRI data. This allowed the interpretation of textural features within images, where manual judgements had previously been necessary. This allowed the quantification and ranking of design parameters, and allowed the optimisation of monolith process configurations, in particular for the selection of an appropriate method of fluid distribution. By considering the way in which MRI data were acquired, it was shown to be possible to further process existing data to extract information such as propagator data from a simple time-of-flight technique.

## 15.2. Further work

An obvious extension to this work would be to include reaction in the various flow configurations, in order to study mass transfer effects and the kinetics associated with such systems. Due to the chemically-specific nature of NMR, spectra may be obtained in a spatially-resolved way, giving a quantitative measurement of chemical conversion throughout the reactor space. This would allow an investigation into the effects of flow maldistribution on local conversion, as well as the effect on selectivity. The choice of reaction is crucial to ensure that products (and by-products) can be distinguished from reactants; the line-width of spectra in the monolith system is unfavourably broad (due to magnetic susceptibility effects) and it may be necessary to use $^{13}C$ imaging to identify individual peaks, since it does not suffer so badly from line-broadening. The lower signal strength of $^{13}C$ would mean long imaging times, and systems would have to be operated in steady-state. It has been shown that bubble and slug lengths can be observed in a non-invasive manner, and this could be compared to reaction extent, allowing the in-situ measure of the sensitivity of mass-transfer and conversion to the flow phenomena within individual channels. Different catalysts and loadings could also be investigated to evaluate their effect.

The effect of surface tension has been briefly mentioned in this work – it was noted that in bubbly upflow, bubbles could enter channels more easily and channels were less likely to be stagnant when the surface tension of the fluid was decreased. Other parameters of monolith design have not been investigated in this work, such as the effect of wall thickness, monolith length, and material of construction. The choice of channel geometry is also important, since the square channels will behave quite differently to circular channels, particularly in upflow where there will be significant downflow in the corners, giving unwanted axial mixing. It would also be of interest to examine the top and bottom of the monolith to visualise entrance and exit effects. For example, in the bubble sparger system (upflow), bubbles may coalesce and deform before entering a channel, and the diameter of the bubbles will affect how easily they may enter. These are all important factors in the design of a functioning monolith reactor, and deserve consideration. If these parameters can be well understood, the modelling of two-phase flow in monolith reactors will be possible, and we may use MRI to aid the development and validation of such models. This will

enable the optimal design of monolith reactors, with a minimum of empirical experiments.

The novel MRI pulse sequences used in this work are primarily based on tagging methods which measure displacement (the exception being GERVAIS which is a phase-encoded method). As a result, they are only applicable to uni-axial flow systems. However, they can still be used in more complex systems, even if quantification is not straightforward. For example, it might be possible to visualise recirculation patterns within a single slug – by using the line-excitation sequence and visualising how the spins travel, we will observe the circular motion within the slug. However, the rate of recirculation is slower than the superficial velocity of the slug through the channel, which will present some hurdles to be overcome. It might be possible to tailor the MRI pulse sequence to compensate for the linear velocity in a similar manner to rotationally-compensated RARE, developed by Sederman *et al.* (2004), or to use some motion-deblurring techniques to clean up MRI data. Or the capillary could be moved at a constant speed against the direction of flow, such that the slug is stationary in space.

The COTTAGE pulse sequence has been shown to observe dynamic fluctuations in the monolith reactor on a timescale of 10s of milliseconds, and it would be interesting to further investigate the relationship between the local flow pattern and velocity within a channel. For example, we could investigate how the velocity changes when bubbles enter a monolith channel. It is not possible to use COTTAGE at the same time as visualising the slug profiles with RARE; however a traditional measurement such as a conductivity probe could measure the passing bubbles, and then combine this information with the dynamic velocity observations from COTTAGE. This would be possible even in high flow rate systems. There are also many observations of flow behaviour in the literature, and some of the correlations could be investigated, particularly in the effect of scale-up in going from a single capillary system to a multiple-channel system such as the monolith reactor.

Since NMR is not a feasible option for inclusion in an industrial process pipeline, it is desirable to try to link observations made by traditional chemical engineering techniques to the knowledge gained through MRI. For example, if we observe a

particular sort of pressure fluctuation, what can this tell us about the flow processes occurring within the reactor? The acquisition of pressure drop data with simultaneous MRI acquisition would help to shed some light on such correlations, and to allow the optimisation of process conditions during industrial operation, where controllers do not have the luxury of live MRI data.

The image processing algorithms developed in this work have only been used to a fraction of their potential. They could be applied to a wide range of systems – to process and quantify features within the acquired data, and to amplify trends that may be hard to distinguish otherwise. A possible application of the Hough transform is to identify pellet locations in MRI structural data. There is interest in 3D packing algorithms, and validation with real systems is difficult. MRI data are generally not acquired at higher resolutions than 256×256×256 voxels, and it can be hard to pinpoint the exact position of the pellets and their surface, particularly when the sensitivity of the RF coil varies over the 3D volume and complicates binary gating (Section 7.2.1). The Hough transform could be a reliable method of identifying pellets (based on parameters such as rotation, position and shape), and this could then be compared with packing algorithms in the literature. Such structural information could also be input to CFD or Lattice-Boltzmann flow simulations. At present, raw MRI structural data has been used in such flow simulations, however the limited resolution is a problem, particularly with non-Newtonian flow which is far more sensitive to the geometry of the pore space.

The autocorrelation techniques which compare an image with itself can be very simply extended to compare two different images instead. This would allow subtle differences between images to be identified, and interesting structural features and fluctuations may become apparent.

## 15.3. Nomenclature

| Abbreviations | Definition |
|---|---|
| CFD | Computational fluid dynamics |
| COTTAGE | Continuous translation tracking with alternating gradient encode |
| GERVAIS | Gradient echo rapid velocity and acceleration imaging sequence |
| NMR | Nuclear magnetic resonance |
| MRI | Magnetic resonance imaging |
| RF | Radio frequency |

## 15.4. References

Sederman, A. J., K. G. Hollingsworth*, et al.* (2004). "Development and application of rotationally compensated RARE." Journal of Magnetic Resonance **171**(1): 118-123.