

Reducing data acquisition times in phase-encoded velocity imaging using compressed sensing

D.J. Holland^{a,*}, D.M. Malioutov^b, A. Blake^b, A.J. Sederman^a, L.F. Gladden^a

^a Department of Chemical Engineering and Biotechnology, University of Cambridge, New Museums Site, Pembroke Street, Cambridge CB2 3RA, United Kingdom

^b Microsoft Research Cambridge, Roger Needham Building, 7 J. J. Thomson Avenue, Cambridge CB3 0FB, United Kingdom

ARTICLE INFO

Article history:

Received 14 September 2009

Revised 23 December 2009

Available online 7 January 2010

Keywords:

Magnetic resonance

Velocity imaging

Compressed sensing

Gas

Liquid

Porous media

ABSTRACT

We present a method for accelerating the acquisition of phase-encoded velocity images by the use of compressed sensing (CS), a technique that exploits the observation that an under-sampled signal can be accurately reconstructed by utilising the prior knowledge that it is sparse or compressible. We present results of both simulated and experimental measurements of liquid flow through a packed bed of spherical glass beads. For this system, the best image reconstruction used a spatial finite-differences transform. The reconstruction was further improved by utilising prior knowledge of the liquid distribution within the image. Using this approach, we demonstrate that for a sampling fraction of $\sim 30\%$ of the full \mathbf{k} -space data set, the velocity can be recovered with a relative error of 11%, which is below the visually detectable limit. Furthermore, the error in the total flow measured using the CS reconstruction is $< 3\%$ for sampling fractions $\geq 30\%$. Thus, quantitative velocity images were obtained in a third of the acquisition time required using conventional imaging. The reduction in data acquisition time can also be exploited in acquiring images at a higher spatial resolution, which increases the accuracy of the measurements by reducing errors arising from partial volume effects. To illustrate this, the CS algorithm was used to reconstruct gas-phase velocity images at a spatial resolution of $230 \mu\text{m} \times 230 \mu\text{m}$. Images at this spatial resolution are prohibitively time-consuming to acquire using full \mathbf{k} -space sampling techniques.

© 2010 Elsevier Inc. All rights reserved.

1. Introduction

Magnetic Resonance (MR) imaging is a powerful tool for imaging the velocity distribution in single- and multi-phase flows. However, the time taken to acquire a velocity image can be prohibitively long, particularly when the velocity distribution varies with time or the signal-to-noise ratio (SNR) is low. Compressed sensing (CS), a technique now attracting considerable interest in the signal processing field, has recently been demonstrated to offer the opportunity to significantly reduce the acquisition time of MR images [1]. In the present paper a CS algorithm is developed and implemented to measure the velocity distribution in gas and liquid flows. This work reports the first practical demonstration of CS to phase-encoded, as opposed to magnitude, MR data.

Phase-encoded MR velocity imaging is used in medical imaging to study the distribution and variation in flow in blood vessels and around the heart, particularly with reference to congenital heart disease and the heart valves [2]. MR velocity imaging has also been widely used in the physical sciences [3,4]. Some examples of systems studied include the rheology of complex fluids [5], the variation in the velocity of liquids flowing through packed beds [6], and

granular flows [7,8]. The main advantage of MR for studying flow is that it is possible to non-invasively image systems without the use of a tracer. However, a major drawback of the technique, in both medical and non-medical applications, is the acquisition time of the measurement. In the medical field, long acquisition times (\sim min) necessitate breath hold or triggered acquisition techniques that, although impressive, are susceptible to artefacts and may be impractical in certain cases. In non-medical applications, long acquisition times require very stable systems and can prohibit the study of certain features, e.g. fine vortices in turbulent flow. To overcome these limitations many studies have explored methods to increase the temporal resolution of velocity-encoded imaging (e.g. [9–13]). These “ultra-fast” (< 1 s) techniques have all been demonstrated to provide a substantial improvement in the temporal resolution. However, despite these advances, each has limitations in terms of the systems that can be studied and the trade-off between spatial and temporal resolution.

CS provides a method of reducing the data acquisition times characteristic of many imaging techniques, including the ultra-fast MR measurements described above. CS utilises the fact that images are highly compressible, and certain linear transforms render them sparse, that is, they can be accurately represented by only a few non-zero elements. Sparsity in a transform domain allows good reconstruction from an under-sampled set of measurements in

* Corresponding author.

E-mail address: djh79@cam.ac.uk (D.J. Holland).

k-space. CS has been demonstrated for spin density imaging in both medical [1] and non-medical [14] applications and to Fourier encoded velocity imaging [15], but has not previously been applied to phase-encoded imaging of, for example, velocity.

Our interest in implementing CS is twofold. First, we seek to increase the temporal resolution with which we can study dynamic systems. Second, in systems in which an increase in spatial resolution is required, we are able to exploit the reduction in data acquisition times to achieve greater signal averaging, thereby decreasing the voxel size from which we can acquire a desired SNR. To illustrate this, we report a case study of velocity imaging of gas flow within a packed bed of spheres. Recently, we reported the ability of MR to acquire maps of the local velocity distribution in both gas and liquid flow through packed beds [16]. However, it was noted in this work that the low SNR associated with data acquisition of the gas phase limited the resolution that was achievable in a practical image acquisition time to $714 \mu\text{m} \times 714 \mu\text{m}$, compared with $179 \mu\text{m} \times 179 \mu\text{m}$ for the liquid phase for a slice thickness of 1.5 mm. Whilst this work demonstrated how gas and liquid velocity varied within the packed bed, greater spatial resolution of the gas-phase velocity image is needed for determination of the gas and liquid velocities at the gas–liquid interface. This is of particular interest in identifying the appropriate closure relationships required to improve the accuracy of computational fluid dynamics simulations of gas–liquid flows. Implementation of CS to increase spatial resolution in gas velocity imaging represents a significant step forward in acquiring the data we need to improve the implementation of these numerical codes.

The paper is structured as follows. First, we build on a CS algorithm of Lustig et al. [1] for spin density imaging and develop it for use with phase-encoded velocity imaging. The CS reconstruction is tested on model data sets for water flow through a packed bed. We use the model data sets to study the effects of various parameters of the reconstruction and set them to get the best reconstruction for a given SNR. This approach was then implemented on a real system and the CS algorithm was used to reconstruct images of liquid-phase flow through a packed bed. The flow rate from the liquid-phase measurements obtained using both full **k**-space sampling and reduced **k**-space sampling with CS reconstruction was compared with the measured macroscopic flow rate to confirm that the measurements were quantitative. Finally, the technique was extended to the application of velocity imaging of gas flows at a spatial resolution that was not practical using conventional MR imaging techniques.

2. Theory – compressed sensing

It is well-established that acquisition times of MR imaging experiments can be decreased by using sparse **k**-space sampling schemes [17–21]. These methods become preferable when the dynamics of the system are too fast to capture using conventional, full **k**-space techniques or when the SNR is low, thus requiring many signal averages. When **k**-space is not fully sampled, linear recovery methods will lead to lossy reconstruction and large artefacts in the resulting image. CS provides a way to achieve performance beyond the limitations of linear methods, and reduces the artefacts arising from sparse sampling to recover the underlying image.

CS allows accurate (and possibly even perfect) reconstructions of dramatically under-sampled **k**-space data sets by utilising the principles behind image compression [22,23]. Typical image compression algorithms take advantage of sparsity by first sampling the full image, computing a transform, e.g. the discrete cosine transform or the wavelet transform, and then discarding the small coefficients. The resulting transform domain data will contain few non-zero coefficients. Thus, the image is said to have a sparse rep-

resentation in the transform domain. The original image can be reconstructed from this sparse representation with minimal loss of information by applying the inverse transform. These ideas underlie the popular JPEG and JPEG-2000 image compression algorithms [24]. CS takes this idea a step further and suggests that near-perfect reconstructions may be possible from a much smaller number of ‘incoherent’ initial samples – which capture just enough information to reconstruct the image. In this context, ‘incoherent’ means that the under-sampling causes incoherent artefacts, or more formally that the sampling operator must not be easily (sparsely) represented in the transform domain [22]. These ‘incoherent samples’ are fused with the prior knowledge that the image is sparse in a transform domain, and the image is then recovered using a non-linear reconstruction method (e.g. [23,25–27]). A variety of transforms exist in which images can be sparsely represented and in which the under-sampling leads to incoherent artefacts. Examples of these transforms include spatial finite-differences (i.e. computing differences of neighbouring pixels), wavelet and curvelet [28] transforms. Each of these transforms will be appropriate for different types of images; all three are investigated in this work.

The image reconstruction algorithm that we have used is a variant of basis pursuit [1,23], which uses the ℓ_1 -norm as a surrogate for sparsity. The ℓ_1 -norm formulation is attractive not only because it leads to a convex optimisation problem which can be solved by efficient algorithms, but moreover, it has allowed strong theoretical analysis of the quality of reconstruction [23,29]. In the following paragraphs the optimisation problem required in the CS reconstruction is stated and details particular to our implementation for this application to phase-encoded velocity imaging are summarised.

Consider the case that the image to be reconstructed is stacked as a vector **x**, Ψ is the operator that compresses the image from pixel representation to a sparse representation (e.g. the wavelet transform), **F** is the under-sampled Fourier transform mapping the image domain to **k**-space and **y** is a vector containing all the **k**-space measurements. The reconstruction is then obtained by solving the following constrained optimisation problem:

$$\begin{aligned} \min \quad & \|\Psi\mathbf{x}\|_1 \\ \text{s.t.} \quad & \|\mathbf{F}\mathbf{x} - \mathbf{y}\|_2 \leq \varepsilon, \end{aligned} \quad (1)$$

where ε is a threshold that can be set to the expected noise level. The ℓ_1 -norm, $\|\mathbf{x}\|_1 = \sum_i |x_i|$, acts as a proxy for sparsity – i.e. minimising the above objective produces an image which has the sparsest representation in the transform domain whilst remaining consistent with the acquired measurements. A variety of methods have been developed to solve Eq. (1), or its unconstrained Lagrangian form

$$\arg \min_{\mathbf{x}} \|\mathbf{F}\mathbf{x} - \mathbf{y}\|_2 + \lambda \|\Psi\mathbf{x}\|_1, \quad (2)$$

including interior-point methods [25], iterative thresholding [26] and Bregman iterations [27]. We follow the approach of Lustig et al. [1] which uses projected conjugate gradients to solve Eq. (2). Since the objective is convex, the algorithm is guaranteed to find the global minimum of the function.

2.1. CS for velocity imaging

The basic CS algorithm described above can be improved for velocity imaging by incorporating additional prior knowledge to the reconstruction. For example, it has previously been suggested that the reconstruction can be improved by including the spatial finite-differences transform in the reconstruction along with another sparse transform [30]. The optimisation problem could then be expressed as:

$$\begin{aligned} \min \quad & \|\Psi_a \mathbf{x}\|_1 + \alpha \|\Psi_b \mathbf{x}\|_1 \\ \text{s.t.} \quad & \|\mathbf{F}\mathbf{x} - \mathbf{y}\|_2 \leq \varepsilon, \end{aligned} \quad (3)$$

where Ψ_a is the operator to transform the image into the domain in which the image is sparse, Ψ_b is the spatial finite-differences or total variation transform, and α is a weighting parameter that trades sparsity in each of the two domains. However, in the case of phase-encoded velocity imaging, as studied here, more specific prior knowledge is available. When measuring the velocity of a liquid or gas, the spatial distribution of the liquid or gas is often known prior to obtaining the image. This prior knowledge can be incorporated through an additional constraint on the reconstruction. If the spatial distribution of the signal is given by a binary image mask, \mathbf{M} , which is a vector with a value of 1 where signal is expected and 0 where only noise is expected, then the reconstruction can proceed by solving a modified optimisation problem:

$$\begin{aligned} \min \quad & \|\Psi \mathbf{x}\|_1 \\ \text{s.t.} \quad & \|\mathbf{F}\mathbf{x} - \mathbf{y}\|_2 \leq \varepsilon \\ & \|(1 - \mathbf{M}) \cdot * \mathbf{x}\|_2 \leq \varepsilon_2, \end{aligned} \quad (4)$$

where the “ \cdot ” operator is the element by element vector multiplication. The second constraint of this optimisation ensures the reconstruction minimises the signal intensity that is incorrectly allocated to the region outside the mask. As with Eq. (1), we solve Eq. (4) in its unconstrained Lagrangian form:

$$\arg \min_{\mathbf{x}} \|\mathbf{F}\mathbf{x} - \mathbf{y}\|_2 + \lambda_1 \|\Psi \mathbf{x}\|_1 + \lambda_2 \|(1 - \mathbf{M}) \cdot * \mathbf{x}\|_2. \quad (5)$$

The relative weight of each term in the reconstruction is controlled by adjusting the parameters λ_1 and λ_2 ; the values of these parameters being determined by the characteristics and SNR of the data. The improvement in the reconstruction achieved by utilising this mask will be examined in Section 4.

The final point, specific to our implementation of CS for acquisition of phase-encoded velocity images, addresses how the detail of the CS method as described by Eqs. (1)–(5) is modified to handle phase, as opposed to magnitude, data sets. As is well established, velocity measurements are achieved by encoding the velocity of the fluid, V , in the phase of the observed signal, ϕ , over an observation time, Δ , using a pair of magnetic field gradients applied with a strength g and for a duration δ :

$$\phi = \gamma g \delta \Delta V, \quad (6)$$

where γ is the gyromagnetic ratio of the observed nuclei. To produce an image, this measurement is combined with imaging gradients such that the signal in any given pixel of an image, $S_{m,n}$, is described by:

$$S_{m,n} = \rho_{m,n} \exp(-i\phi_{m,n}) = \rho_{m,n} \exp(-i\gamma g \delta \Delta V_{m,n}), \quad (7)$$

where $\rho_{m,n}$ is the signal magnitude in the pixel m, n . The velocity is then calculated from the phase of the complex signal in each pixel of an image. The approach described by Eqs. (1)–(5) above is equally applicable to both real and complex images. However, it is difficult to apply CS to the phase data of a complex image directly. Instead the real and imaginary intensity images are reconstructed using CS and the phase of the signal is calculated after reconstruction. A similar approach has previously been used in the denoising of MR images [e.g. 31] and has been suggested for use with CS [32]. In this approach the sparse transform is calculated for the real (\mathbf{x}_r) and imaginary (\mathbf{x}_i) components of the image separately and the minimisation is performed over the magnitude of the signal in the transform domain:

$$\arg \min_{\mathbf{x}} \|\mathbf{F}\mathbf{x} - \mathbf{y}\|_2 + \lambda_1 (\|\Psi \mathbf{x}_r + i\Psi \mathbf{x}_i\|_1) + \lambda_2 \|(1 - \mathbf{M}) \cdot * \mathbf{x}\|_2. \quad (8)$$

The result of this optimisation is a complex image containing the intensity of the real and imaginary components in each pixel. Reconstructing the real and imaginary intensity images in this

manner requires a piecewise-smooth variation in signal intensity in both images. Depending on the Fourier transform algorithm employed, this may require a rearrangement of the data prior to Fourier transformation according to the shift theorem.

3. Experimental details

All experiments were performed on a Bruker DMX 200 spectrometer with a vertical 4.7 T superconducting magnet. The experiments were performed using a 64 mm diameter radiofrequency (r.f.) coil operating at 199.7 MHz for proton (^1H) and 188.3 MHz for fluorine (^{19}F). The magnet was equipped with a three-axis shielded gradient producing a maximum gradient strength of 0.136 T m^{-1} in the x, y , and z directions.

3.1. Development and optimisation of the CS algorithm

To optimise the parameters and transform domain for the CS reconstruction, a velocity map was obtained from a lattice-Boltzmann (LB) simulation of water flowing through a 39 mm diameter column that was randomly packed with 3 mm diameter spheres. The geometry for the simulation was derived from a three-dimensional spin-echo MR image of the bed that was obtained at a resolution of $164 \mu\text{m} \times 164 \mu\text{m} \times 164 \mu\text{m}$. This data set was converted to a binary image and used as input to the LB simulation. The LB method models the hydrodynamics of a fluid on a mesoscopic scale, where fluid kinetics are described by probability distribution functions of the location of particles. The LB code used has been described and validated previously [33,34]. The time step for the simulations was $4.5 \times 10^{-3} \text{ s}$, such that the dimensionless relaxation was 1, which optimises the simulation accuracy and stability [33]. The velocity map resulting from the LB simulation was converted to a phase-encoded image suitable for testing the CS algorithm by selecting a single axial (z) slice from the centre of the image. This velocity image was converted to a phase map by assuming that the maximum velocity in the image corresponded to a phase shift of 1.8π radians, thus ensuring that the total range of velocities was less than 2π radians. The phase map and the binary signal intensity map were then converted to a complex image, thus giving real and imaginary intensity data.

The \mathbf{k} -space sampling strategy followed the Monte Carlo approach of Lustig et al. [1] on a Cartesian grid. Firstly, a \mathbf{k} -space grid was chosen based on the desired resolution. This grid was under-sampled in the dimension corresponding to the phase-encoding direction. The under-sampling was controlled by constructing a probability density function (pdf) and randomly drawing indices from that density. The incoherence of this sampling pattern was measured by calculating the maximum intensity of the side lobes of the point spread function in the transform domain. This procedure was repeated 50 times and the sampling pattern with the lowest intensity interference was chosen. The same procedure was repeated for each desired image resolution. The sampling pattern was biased to the origin of \mathbf{k} -space by using a suitable choice of pdf: the pdf was calculated using a polynomial decay:

$$\text{pdf} = (1 - r)^p, \quad (9)$$

where $r = k/k_{\text{max}}$ is the normalised distance from the origin of \mathbf{k} -space. This is beneficial as most of the intensity of real images is concentrated near the \mathbf{k} -space origin [1]. Typically, the pdf was calculated with $p = 5/2$.

The choice of the parameters λ_1 and λ_2 in Eq. (5) is critical to obtaining a good quality reconstruction using CS. Previously it has been suggested that the value of the parameters λ_i can be determined by solving Eq. (5) for different values, and then choosing λ_i such that $\|\mathbf{F}\mathbf{x} - \mathbf{y}\|_2 \approx \varepsilon$ [1]. In practice, it was found that the

quality of the reconstruction was not strongly influenced by the value chosen for λ_i above a given threshold. Furthermore, this threshold was approximately the same for all the images in this paper at a given SNR. Therefore, it was possible to optimise the values of λ_1 and λ_2 on one data set and use these values for subsequent reconstructions at the same SNR.

3.2. Acquisition of velocity images

Velocity images of single-phase liquid and gas flow were acquired using the same pulse sequence, and using the same packed bed. The bed used was a cylindrical column of inner diameter 27 mm, randomly packed with 5 mm diameter glass spheres. The column was of length 1 m. Velocity images were obtained using a spin-echo sequence that was designed to minimise the total echo time [8,16]. The velocity was encoded using half a period of a sine-shaped gradient waveform. Two images were acquired, each with flow-encoding gradient strengths of the same magnitude but opposite direction.

The experimental set-up for acquisition of the liquid-phase velocity images was as follows. Deionised water was pumped in a closed circuit by a Verder VG330-10 gear pump and controlled by a PC-operated Bronkhorst Cori-flow (model M55C4-AAD-11-K-C) mass flow controller; the flow rate of water was varied from 0 to 60 kg h⁻¹. The T₁ relaxation time constant of the water was reduced to 50 ms by adding gadolinium chloride to the water at a concentration of 0.49 g L⁻¹. The flow-encoding gradients were applied for duration $\delta = 1.19$ ms and separated by an observation time (Δ) of 3.3 ms. The gradient strength used was varied for different water flow rates to optimise the dynamic range of the velocity measurement. First, a full **k**-space acquisition was performed. Velocity images were obtained with a field-of-view of 30 mm \times 30 mm at an in-plane resolution of 178 μ m \times 178 μ m with a 1.5 mm slice thickness. The repetition time of the experiment was 300 ms and a 4 step phase-cycle was used, giving a total acquisition time of about 7 min. CS measurements were then performed using the same pulse sequence, with the phase-encoding gradient set to only acquire the desired lines of **k**-space; in these acquisitions only 28% of the **k**-space raster was sampled. The sampling distribution was determined as described in Section 3.1.

Gas-phase velocity images were acquired on the same packed bed as used for the liquid-phase studies. Sulphur hexafluoride (SF₆) gas was supplied in a closed circuit using a DILO Piccolo compressor (model B022R01). The gas from the compressor was stored in a pressure vessel of volume 6 L. Gas was drawn from the pressure vessel and the flow rate was measured using a rotameter (Brooks Sho-rate 1357/D2B5D1B00000). A bypass line ensured excess gas was returned to the feed of the compressor. The pressure in the column was regulated using a back pressure regulator set to 4.9 bar absolute. As for the liquid velocity measurements, in implementing the pulse sequence the flow-encoding gradient strength was adjusted to optimise the dynamic range of the measurement. A flow-encoding gradient duration and observation time of 0.5 ms and 1.9 ms, respectively, were used.

Gas velocity images were acquired using different acquisition times and at different spatial resolutions. In all cases, the field-of-view was 29.4 mm \times 29.4 mm and the slice thickness was 1.5 mm. The acquisitions were as follows:

- (i) A full **k**-space raster of 64 \times 64 data points was acquired giving a spatial resolution of 460 μ m \times 460 μ m; 512 scans were acquired with a repetition time of 34 ms, giving a total acquisition time of 37 min.
- (ii) A CS data acquisition was acquired in a total acquisition time of 37 min, but with an array of 84 \times 84 data points, giving a spatial resolution of 350 μ m \times 350 μ m. For the CS recon-

struction, only 28 points were sampled in the phase-encoding direction and the repetition time was 26 ms. Thus whilst the same acquisition time was used as in (i), 1536 scans were acquired compared to the 512 scans when full **k**-space sampling was employed. The data points sampled were chosen using the approach outlined in Section 3.1.

- (iii) High resolution velocity maps were obtained using a full **k**-space raster of 128 \times 128 data points, giving a spatial resolution of 230 μ m \times 230 μ m. The repetition time was 34 ms and 608 scans were acquired, giving a total imaging time of 90 min. This was chosen as the practical upper limit for the duration of a gas-phase velocity imaging pulse sequence.
- (iv) CS velocity mapping was performed at a spatial resolution of 230 μ m \times 230 μ m but only sampling 39 data points in the phase-encoding direction. Thus, 2048 scans were acquired in a total imaging time of 90 min.

4. Results

The results are reported in two sections. First, a numerical study is presented which explores the influence of the detail of the reconstruction algorithm, the number and location of **k**-space points selected and the noise level on the performance of the CS reconstruction. Second, CS is applied to reduce the acquisition time of velocity imaging of single-phase flow of both liquid and gas, whilst maintaining the accuracy of the total flow measurement to within $\pm 5\%$ of that recorded with a mass flow controller (liquid) or rotameter (gas). The measurements on gas-phase flow are presented at a resolution that was not practicable without the use of CS.

4.1. Development and optimisation of the CS algorithm on simulated data

4.1.1. Demonstration of CS

Fig. 1a shows a velocity map obtained from a LB simulation [33] of water flowing through a packed bed of 3 mm diameter glass spheres. Since the map was derived from a numerical simulation, it is essentially free from variations in the signal due to noise and thus provides an ideal starting point for developing a CS technique for velocity imaging. This velocity map was converted to a phase map by assuming that the velocity range in the image corresponded to a phase shift of 1.8π . This was then used to produce a complex image of the bead pack in which the velocity was encoded in the phase of the signal. This image was Fourier transformed to obtain the **k**-space distribution shown in Fig. 1b, which was then used to develop the CS technique.

Fig. 2 illustrates the CS reconstruction of the velocity image shown in Fig. 1. Fig. 2a shows the **k**-space data points used for the CS reconstruction. This corresponds to using only 28% of the full **k**-space data. The **k**-space data were under-sampled by removing entire lines of **k**-space, consistent with the use of a frequency-encoded imaging sequence. Fig. 2b shows the 2D velocity image reconstructed using the CS algorithm described in Section 2 with the **k**-space data shown in Fig. 2a. The typical image reconstruction time for a velocity image such as that in Fig. 2b was ~ 5 min on a standard desktop PC. Although clearly a subjective measure of the error, the image in Fig. 2b shows no visually perceptible differences from the original velocity image shown in Fig. 1a.

The absolute error introduced by under-sampling the **k**-space data is shown in Fig. 3. The absolute error is defined as the difference between the velocity map obtained from the LB simulations and the velocity map reconstructed using the under-sampled **k**-space data. Fig. 3a shows the absolute error that would be obtained

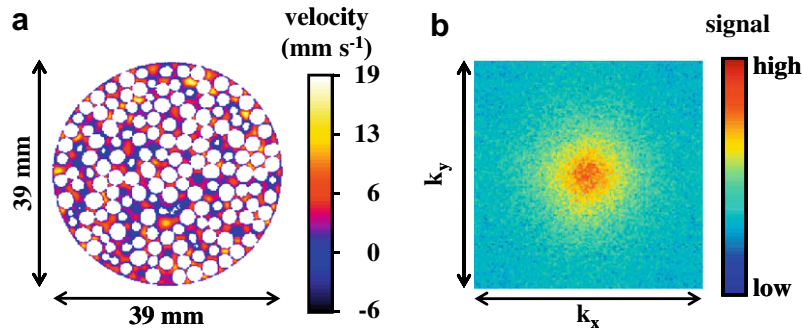


Fig. 1. (a) Velocity map obtained from a LB simulation of flow through a packed bed of 3 mm diameter spheres. The velocity map was converted to a phase map for development of the compressed sensing reconstruction. The phase map was converted to a complex image, which was subsequently Fourier transformed to obtain the k -space data set shown in (b).

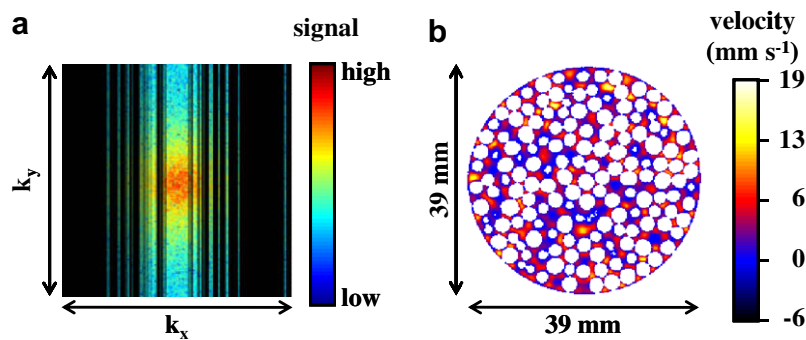


Fig. 2. Demonstration of a CS reconstruction from under-sampled k -space data. (a) Shows the k -space data points that were selected from the data in Fig. 1b. The data selected correspond to a sampling fraction of 28%. These data were used to reconstruct the velocity map shown in (b). The velocity map shows excellent agreement with the original image shown in Fig. 1a.

if the missing k -space data were replaced by zeros prior to Fourier transformation. The zero-filled reconstruction shows alternate positive and negative errors in the phase-encoded (horizontal) direction. This is expected as under-sampling the k -space data will increase the width of the point spread function. Thus, regions of high signal intensity in the image domain will tend to decrease in intensity and neighbouring regions of low signal intensity will tend to increase in intensity, leading to the alternating positive and negative errors in the reconstructed signal. The error obtained from the CS reconstruction is shown in Fig. 3b. The error shows only slight systematic variations in the phase-encoded direction. Errors in the reconstruction will now be considered in more detail.

The proportional reconstruction error within the voxels of the image associated with flow is:

$$\ell_2 - \text{error} = \frac{\sqrt{\sum_i |M_i (V_i^{\text{CS}} - V_i)|^2}}{\sqrt{\sum_i |M_i (V_i)|^2}}, \quad (10)$$

where V_i is the velocity in the i th pixel of the original, fully-sampled, image, V_i^{CS} is the velocity in the i th pixel of the image reconstructed from the under-sampled k -space data using CS, and M_i takes the value of 0 or 1 depending on whether the pixel is solid or water, respectively. The ℓ_2 -error obtained from Eq. (10) using the $(V^{\text{CS}} - V)$ data shown in Fig. 3b yields a value of 0.11% or 11%, despite only sampling 28% of the full k -space data set. By comparison, when the data are zero-filled before Fourier transformation (Fig. 3a), the ℓ_2 -




Fig. 3. Images showing the absolute error in the velocity map reconstructed from the k -space data in Fig. 2a. The error is shown for (a) a zero-filled reconstruction and (b) CS reconstruction. The error in the CS reconstruction is markedly reduced from that shown in the zero-filled reconstruction.

error is 23%. Thus, the CS reconstruction has reduced the error by a factor of 2.

The ℓ_2 -error provides a simple quantitative description of the accuracy of the reconstruction and therefore will be used throughout this paper. However, it should be noted that other factors, for example the range and bias of the error, must also be considered when defining the accuracy of the reconstruction. These will now be examined for the data in Fig. 3.

A major concern with applying a non-linear reconstruction algorithm such as that described in Section 2, is that the error becomes biased. For example, the velocity in regions of high velocity might be systematically under-estimated. Fig. 4 shows a plot of the error in the velocity obtained from the CS reconstruction as a function of the fully-sampled velocity. The CS reconstruction shows a slight bias toward positive error (i.e. to over-estimate the velocity) for very low velocities, but is essentially randomly scattered about zero. This indicates that although the CS reconstruction introduces noise into the velocity distribution, it introduces little bias. By contrast, the zero-filled reconstruction shows both a broader distribution of errors and a significant bias, with low velocities over-estimated and high velocities under-estimated. It was observed that if the sampling fraction using CS is decreased further or the parameters λ_1 and λ_2 used in the reconstruction (Eq. (5)) are incorrectly set, the CS error may become biased. This effect was most pronounced in the fine details of the velocity distribution, for

example in the constrictions between local elements of void space in the packed bed. However, provided that the images being considered were similar, e.g. were of different slices in the packed bed, the optimised sampling fraction and values of λ_1 and λ_2 were approximately the same.

In summary, Figs. 2–4 show that CS can accurately reconstruct velocity imaging data from as little as 28% of the full \mathbf{k} -space data set. These results demonstrate the significant time reductions that can be achieved using CS for velocity-encoded imaging. The next sections will investigate which sparse transforms to use, modifications to the reconstruction algorithm and how robust the reconstruction is when the measurement contains noise.

4.1.2. Sparsity in velocity images

The reconstruction of an image using CS utilises the fact that the image can be represented sparsely in some transform domain. In the present work, four domains have been considered in which the image may be sparsely represented. Fig. 5a shows a plot of the rank ordered intensity (energy) of the image in the image domain, and three common transform domains: the spatial finite-differences domain, a wavelet domain and a curvelet [28] domain. The intensity in each case was calculated as the absolute value of the transform of the real component of the complex image derived from the LB simulation shown in Fig. 1a. The total number of pixels in the image was 65,536 (256 pixels \times 256 pixels). The rapid




Fig. 4. Plots of the absolute error associated with a given pixel, obtained from Fig. 3, as a function of the local velocity in that pixel as obtained from Fig. 1a. The errors shown are for (a) the zero-filled reconstruction and (b) the CS reconstruction.

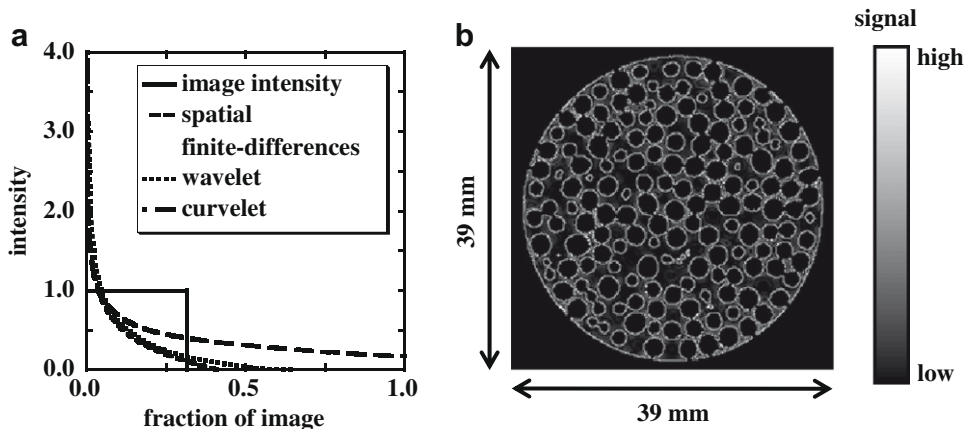


Fig. 5. (a) Plot of the rank ordered intensity for the raw image data, the spatial finite-differences transform, a wavelet transform and a curvelet transform of the complex image derived from the data in Fig. 1a. The intensity in each case was calculated as the absolute value of the transform of the real component of the complex image derived from the LB simulation shown in Fig. 1a. These data indicate that most of the energy in the transform domains is recorded in only a small number of voxels in the image. (b) The absolute value of the spatial finite-differences transform of the real component of the complex image derived from the velocity image shown in Fig. 1a.

decrease in the intensity indicates that all four domains show a high degree of sparsity, with approximately 70% of pixels containing negligible signal. The image intensity domain contains the fewest pixels with negligible signal, and thus is not as sparse a representation as the other domains. All three transform domains show a rapid initial decrease followed by an extended tail in the distribution. An example of the spatial finite-differences transform domain is given in Fig. 5b. Most of the image is dark, confirming that the image is sparsely represented in this domain. Given the similarity of the results for the transform domains in Fig. 5a, it is not obvious which domain will yield the best sparse representation. However, the tail in the curvelet domain is the most pronounced and this is the slowest transform to calculate. Therefore, these results suggest that the most promising sparse domains to consider are the spatial finite-differences and the wavelet domain.

Previous reports have indicated that in order to accurately reconstruct the under-sampled image, the number of \mathbf{k} -space points sampled should be between 2 and 5 times the number of sparse coefficients [1]. For the sparse transforms calculated here this would suggest that the number of \mathbf{k} -space samples required would be in the region of 50% of the full data set. It was found that the reconstruction error with the spatial finite-differences transform was significantly less than that obtained with the wavelet transform. Fig. 6 shows a plot of the ℓ_2 -error for reconstructions of the LB velocity map in Fig. 1a using the spatial finite-differences transform, and from 15% through to 75% of the full \mathbf{k} -space data set. The reconstruction error increases as the fraction of \mathbf{k} -space sampled decreases. The error increases rapidly when less than 30% of \mathbf{k} -space is sampled.

4.1.3. Modifications to the CS algorithm

Two modifications to the CS algorithm were examined: the addition of a second sparse transform and the use of a mask of the image to guide the reconstruction. The effect of each of these modifications on the accuracy of the reconstruction will now be described.

The use of two sparse transform domains was detailed in Eq. (3). Two additional domains were considered: the image domain and a wavelet transform domain. It was found that incorporating either of these second sparse transforms led to a reduction of the reconstruction ℓ_2 -error of <0.5%. This was accompanied by an increase in the reconstruction time of about 50% for the image domain and about 100% for the wavelet transform. Therefore, as the decrease in error was not significant, a second transform domain was not considered to be beneficial.




Fig. 6. The effect of decreasing the sampling fraction on the accuracy of the image reconstructed using CS. The data were taken from the LB simulation of water flow through a packed bed of 3 mm diameter spheres. The reconstructed error follows closely an exponential decay (dotted line) for sampling fractions between 75% and 15%.

The incorporation of the prior knowledge of the signal distribution was found to significantly improve the CS reconstruction. This is achieved straightforwardly by acquiring a high SNR image of the system prior to obtaining the velocity-encoded images; a mask data set is readily derived from this image which is used to bias the reconstruction such that signal only appears where it is expected. This helps to eliminate the misregistration of the signal resulting from the under-sampling of \mathbf{k} -space. For the LB velocity map shown in Fig. 1a, the reconstruction ℓ_2 -error can be reduced from 14% without utilising the mask to 10% by utilising the mask. This is a substantial improvement in the reconstruction that effectively allows a reduction in the sampling fraction of $\sim 10\%$, depending on the desired final reconstruction error. The additional mask required in this approach is easily obtained from a liquid- or gas-phase image of a stationary bead pack and significantly increases the accuracy of the reconstruction. Therefore, the use of a mask data set makes a significant impact on the data acquisition time or spatial resolution of the resulting gas- or liquid-phase velocity images without a significant complication to the reconstruction.

4.1.4. Effect of noise on the reconstruction

The SNR of the acquired data will influence the accuracy of the reconstruction. This effect was investigated by systematically adding Gaussian noise to the LB simulated flow field. Fig. 7 shows the accuracy of the reconstruction as a function of increasing standard deviation of the noise distribution for sampling fractions of between 21% and 75%. The signal was normalised to a value of 1 such that the SNR is simply the inverse of the standard deviation of the noise. The reconstruction error is seen to increase with decreasing sampling fraction and to increase with increasing noise, as would be expected.

The increase in the error shown in Fig. 7 with increasing noise is slower than would be expected for a fully-sampled image reconstructed using a conventional Fourier reconstruction. Thus, for a typical data acquisition in which the SNR might be ~ 10 , the ℓ_2 -error in a data set that is reconstructed using CS but sampling 30% of \mathbf{k} -space is 15%, compared with an ℓ_2 -error of 8% for the fully-sampled data set. Furthermore, for higher sampling fractions the reconstructed error using a CS approach is less than the error in a fully-sampled image that is reconstructed with a conventional Fourier transformation. For example, if we consider the data with a 75% sampling of \mathbf{k} -space the reconstructed error is less than the error arising in a fully-sampled data set if the noise standard deviation exceeds ~ 0.1 , which corresponds to a SNR of ~ 10 . It should be noted that the decrease in error is not due to the




Fig. 7. The effect of increased noise on the accuracy of the CS reconstruction is shown for sampling fractions of (\blacktriangledown) 75%, (\bullet) 50%, (\blacksquare) 39%, (\blacklozenge) 30% and (\blacktriangle) 21%. The solid line indicates the error arising from a fully-sampled image at the given noise level. The noise is given by a Gaussian function with the standard deviation indicated. The signal intensity in the image was normalised to a maximum value of 1.

under-sampling directly, but rather due to the requirement for sparsity in the reconstruction. These results indicate that a CS approach can be used to achieve a desired accuracy in the reconstructed image in a given time by optimising the combination of the sampling fraction and the signal averaging used.

4.2. Experiments

The CS algorithm developed in the previous section is now used to reconstruct images of the velocity for single-phase gas and liquid flow through a packed bed of 5 mm glass spheres, using real experimental data. The CS algorithm was first applied to the acquisition of a liquid-phase velocity image. The aim of the experiment was to compare a velocity image reconstructed with CS using an under-sampled \mathbf{k} -space data set, directly with a fully-sampled \mathbf{k} -space image acquired at the same spatial resolution. The CS reconstruction was then used to acquire gas-phase velocity images at a resolution that was greater than that which could be obtained using fully-sampled \mathbf{k} -space data.

4.2.1. CS velocity imaging of the liquid phase

The CS reconstruction algorithm developed in Section 4.1 was applied to velocity measurements of liquid-phase flow. Fig. 8a shows an image of the velocity of water flowing through the

bead pack of 5 mm diameter spheres at a flow rate of 5.3 ml s^{-1} . This corresponds to a Reynolds number of 46, based on the superficial velocity of the water and the diameter of the beads. At this Reynolds number the flow will be stable. Fig. 8a was obtained from a full \mathbf{k} -space acquisition. Fig. 8b shows an equivalent reconstruction of the velocity distribution using the CS algorithm defined by Eq. (5), when only 28% of the \mathbf{k} -space data were acquired. Applying Eq. (10) to these data gives an ℓ_2 -error of 11%. The SNR in these images was 58, which is equivalent to $\sigma_{\text{noise}} = 0.02$, and thus this estimate of the error is consistent with that shown in Fig. 7.

Fig. 9a shows the rank-ordered distribution of the intensity for the image domain, the spatial finite-differences, a wavelet transform and a curvelet transform of the data used to obtain Fig. 8a. As with the data derived from the LB simulation, the plot decays rapidly for all three transform domains, indicating that few pixels in these domains contain non-zero values. This confirms that the image can be represented sparsely in all three domains; the spatial finite-differences domain will be used for the reconstruction as this was found to be most suitable with the simulated data. An example of the spatial finite-differences domain image is shown in Fig. 9b. As with Fig. 5b, Fig. 9b is mostly dark confirming that the image can be sparsely represented in the spatial finite-differences domain.

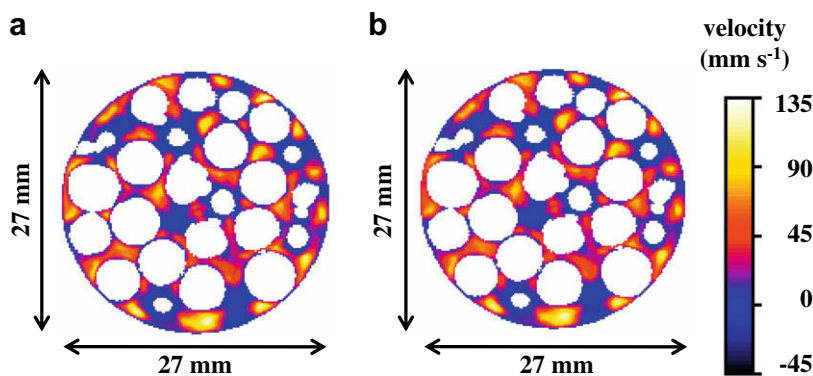


Fig. 8. Example of the use of CS to image the velocity of liquid flowing through a packed bed of 5 mm diameter spheres. (a) Fully-sampled image and (b) CS reconstruction of an image with only 28% sampling. The velocity was measured in the axial direction (i.e. coming out of the plane) and is given by the colour bar. The image resolution was $178 \mu\text{m} \times 178 \mu\text{m}$ and the slice thickness was 1 mm. The total imaging time was (a) 430 s for the fully-sampled image and (b) 120 s for the under-sampled image. The ℓ_2 -error of the velocity in the under-sampled image was 11% and the error on the total flow was 0.9%. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

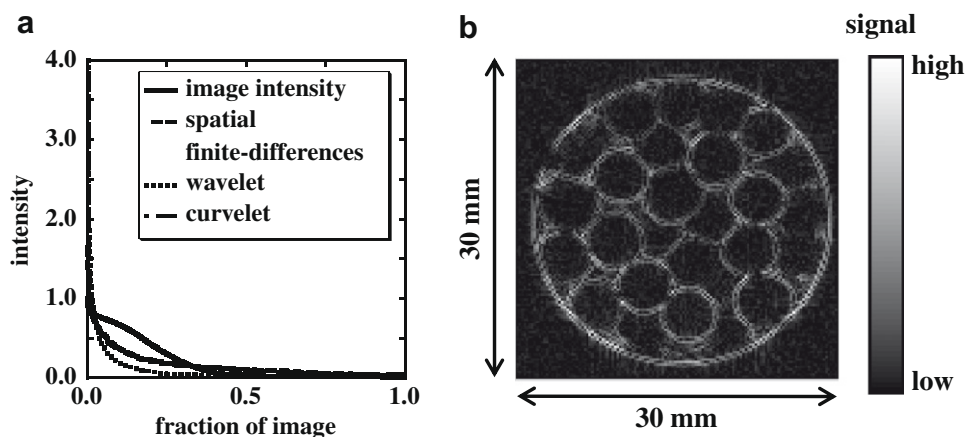


Fig. 9. Analysis of the sparsity in an experimental phase-encoded velocity image. (a) Plot of the ranked intensity for the raw image data, the spatial finite-differences transform, a wavelet transform and a curvelet transform. The intensity in each case was calculated as the absolute value of the transform of the real component of the complex image. The image was obtained for water flowing through a packed bed of 5 mm diameter spheres. The image was obtained with a data array of 168×168 pixels at a resolution of $178 \mu\text{m} \times 178 \mu\text{m}$. (b) Example of the sparsity in the spatial finite-differences transform of the velocity image in Fig. 8a, which was used to derive the data in (a).




Fig. 10. Comparison of the flow rate measured by MR and the flow rate measured using a mass flow metre. The MR data shown are for both (●) fully-sampled k -space and (▲) CS reconstructions of the velocity. The solid line $y = x$ is shown.

To check the validity of the CS reconstruction, the measured velocity distribution was used to calculate the total flow rate through the bead pack for the fully-sampled and CS reconstructions. Fig. 10 shows a comparison of the flow rates determined from the MR images for three different image slices with the flow rate of water measured by the mass flow controller. The root mean squared error in the flow rate was 1.7 kg h^{-1} , or 3%, for both the fully-sampled k -space data and the CS reconstruction. These values

compare favourably with the error in the flow rate of the mass flow controller of 1%. These results suggest that the CS reconstruction has negligible effect on the quantitative reconstruction of the velocity distribution and that the CS algorithm has been implemented correctly on the spectrometer.

4.2.2. CS velocity imaging of the gas phase

The CS reconstruction algorithm developed and validated in Sections 4.1 and 4.2.1 was applied to velocity measurements in the gas phase. Fig. 11 shows gas-phase velocity measurements for SF_6 flow through a packed bed of 5 mm diameter spheres using a conventional full k -space acquisition and a CS reconstruction where only 33% (28 points in the phase-encoding direction) of k -space was sampled. The flow rate of SF_6 was $(15 \pm 1) \times 10^{-6} \text{ m}^3 \text{ s}^{-1}$ at a pressure of 4.9 bar absolute. This corresponds to a Reynolds number of 260, based on the superficial velocity of the gas and diameter of the beads. Fig. 11a shows an image of a full k -space acquisition at a resolution of $460 \mu\text{m} \times 460 \mu\text{m}$. This image was acquired in 19 min. Fig. 11b shows a CS reconstruction at a resolution of $350 \mu\text{m} \times 350 \mu\text{m}$, acquired in the same total acquisition time. The total flow calculated from the fully-sampled data set was $18 \times 10^{-6} \text{ m}^3 \text{ s}^{-1}$ and from the CS reconstruction was $16 \times 10^{-6} \text{ m}^3 \text{ s}^{-1}$. These measured flow rates are within 13% and 5% of the flow rate measured with the rotameter for the full k -space acquisition and the CS reconstruction, respectively. Thus, the error in the total flow rate measured using CS was within the




Fig. 11. Images of the velocity distribution for SF_6 flow through a packed bed of 5 mm diameter spheres. (a) Velocity image obtained with a fully-sampled k -space data set and reconstructed using a conventional Fourier transform. The resolution is $460 \mu\text{m} \times 460 \mu\text{m} \times 1.5 \text{ mm}$. (b) Velocity image reconstructed using CS at a resolution of $350 \mu\text{m} \times 350 \mu\text{m} \times 1.5 \text{ mm}$. The acquisition time for both (a) and (b) was 37 min. The flow rate of SF_6 was $(15 \pm 1) \times 10^{-6} \text{ m}^3 \text{ s}^{-1}$. The CS reconstruction allows images at higher spatial resolution to be acquired with no loss of SNR.




Fig. 12. High resolution images of the velocity distribution for SF_6 flow through a packed bed of 5 mm diameter spheres. (a) Conventional velocity image obtained with a fully-sampled k -space data set and reconstructed using a conventional Fourier transform. (b) Velocity image reconstructed using CS. The resolution in both (a) and (b) was $230 \mu\text{m} \times 230 \mu\text{m} \times 1.5 \text{ mm}$ and the acquisition time was 90 min. The flow rate of SF_6 was $(15 \pm 1) \times 10^{-6} \text{ m}^3 \text{ s}^{-1}$. The CS reconstruction allows significantly greater data averaging for a given acquisition time and spatial resolution.

uncertainty of the flow rate measured with the rotameter ($15 \pm 1 \times 10^{-6} \text{ m}^3 \text{ s}^{-1}$). The reduction in the error in the CS reconstruction may be due to the higher resolution that is obtained, which reduces partial volume effects in the reconstruction. These results demonstrate that utilising a CS reconstruction can improve the spatial resolution achievable without compromising the quantitative reconstruction of velocity measurements.

The CS image reconstruction approach was then used to demonstrate the improved spatial resolution that can be achieved, as compared to full \mathbf{k} -space sampling, for an acquisition time of 90 min. Fig. 12 shows a comparison of (a) a fully-sampled velocity image obtained at a resolution of $230 \mu\text{m} \times 230 \mu\text{m}$ with 928 scans and (b) a CS reconstruction obtained at the same resolution but from only 30% of the full \mathbf{k} -space data set, thereby allowing signal averaging over 3072 scans. The image reconstructed using CS (Fig. 12b) shows significantly less noise than the fully-sampled image (Fig. 12a) whilst retaining detailed measurements of the high resolution features of the flow field. The image obtained from the CS reconstruction gives an 11-fold decrease in the voxel size of the image compared with the highest resolution gas-phase velocity images we have reported [16]. By under-sampling the \mathbf{k} -space data, the CS reconstruction allows an increase in the SNR of the resulting velocity image and hence an increase in the accuracy of the reconstruction. The SNR in the fully-sampled image was 5.2, compared with 21 in the under-sampled image reconstructed using CS. The ℓ_2 -error of the velocity images, given the sampling fraction and SNR, can be estimated from Fig. 7 to be 15% and 10% for the fully-sampled and CS image, respectively. Furthermore, the error in the total flow measurement decreased from 8% to 3%, which is within the experimental uncertainty of the rotameter measurement.

5. Conclusions

In this paper we have presented a CS reconstruction for phase-encoded velocity imaging. The reconstruction algorithm was demonstrated for simulated flow data and liquid-phase imaging where comparison with the full \mathbf{k} -space acquisition was possible. These results demonstrate that for a sampling fraction of $\sim 30\%$ of the full \mathbf{k} -space data set the velocity image can be recovered with an ℓ_2 -error of 11%, which is below the visually detectable limit. The total flow rate calculated from the liquid-phase velocity measurements was found to be within 3% of the independently measured flow rate (using a mass flow meter) for both the conventional and CS velocity images. Furthermore, the absolute error in the CS reconstruction was demonstrated to be unbiased, confirming that the reconstruction artefacts are negligible.

The CS algorithm was then used to obtain gas-phase velocity images at a spatial resolution that was previously prohibitively time-consuming to acquire. The final image resolution was $230 \mu\text{m} \times 230 \mu\text{m}$ with a 1.5 mm slice thickness. To the best of our knowledge, the voxel size in these images is an order of magnitude less than in any previously reported gas-phase velocity images acquired using MR imaging. The total flow calculated from these CS velocity images was found to be within 5% of the flow rate determined using a rotameter, which is within the experimental uncertainty of the rotameter measurement.

The CS image reconstruction with the smallest ℓ_2 -error for a given sampling fraction was found using the spatial finite-differences as the sparse transform. The reconstruction was further improved by utilising prior knowledge of where the signal should be located. This was achieved by using a mask of the image obtained prior to commencing the flow experiments. The use of a second sparse transform, such as the wavelet transform or the image itself, was also considered. This was found to only margin-

ally reduce the reconstructed error at the cost of an increased reconstruction time, and was therefore not employed.

This paper has demonstrated the potential for CS to improve the reconstruction of sparse phase-encoded velocity imaging. The approach is simple to implement and is applicable to any MR technique for velocity imaging. It should also be possible to further improve the reconstruction by incorporating more specific prior knowledge of the system, for example by using sparse coding techniques to design a dictionary specifically for a given class of velocity measurements [35,36]. A further possibility is to use more detailed signal models, such as the embedded zero trees of the wavelet transform [37,38] to further reduce the number of coefficients required to reconstruct the image.

Acknowledgments

The authors would like to acknowledge Thomas Pintelon for the lattice-Boltzmann simulations and the EPSRC (EP/F047991/1) and Microsoft Research for funding.

References

- [1] M. Lustig, D. Donoho, J.M. Pauly, Sparse MRI: the application of compressed sensing for rapid MR imaging, *Magn. Reson. Med.* 58 (2007) 1182–1195.
- [2] P.D. Gatehouse, J. Keegan, L.A. Crowe, S. Masood, R.H. Mohiaddin, K.-F. Kreitner, D.N. Firmin, Applications of phase-contrast flow and velocity imaging in cardiovascular MRI, *Eur. Radiol.* 15 (2005) 2172–2184.
- [3] E. Fukushima, Nuclear magnetic resonance as a tool to study flow, *Ann. Rev. Fluid Mech.* 31 (1999) 95–123.
- [4] M.D. Mantle, A.J. Sederman, Dynamic MRI in chemical process and reaction engineering, *Prog. Nucl. Mag. Res. Spectrosc.* 43 (2003) 3–60.
- [5] P.T. Callaghan, Rheo-NMR: nuclear magnetic resonance and the rheology of complex fluids, *Rep. Prog. Phys.* 62 (1999) 599–670.
- [6] A.J. Sederman, M.L. Johns, P. Alexander, L.F. Gladden, Structure-flow correlations in packed beds, *Chem. Eng. Sci.* 53 (1998) 2117–2128.
- [7] E. Fukushima, Granular flows, in: S. Stapf, S.-I. Han (Eds.), *NMR in Chemical Engineering*, Wiley-VCH, Weinheim, 2006.
- [8] D.J. Holland, C.R. Müller, J.S. Dennis, L.F. Gladden, A.J. Sederman, Spatially resolved measurement of anisotropic granular temperature in gas-fluidized beds, *Powder Tech.* 182 (2008) 171–181.
- [9] K. Kose, Instantaneous flow-distribution measurements of the equilibrium turbulent region in a circular pipe using ultrafast NMR imaging, *Phys. Rev. A* 44 (1991) 2495–2504.
- [10] K. Kose, One-shot velocity mapping using multiple spin-echo EPI and its application to turbulent flow, *J. Magn. Reson.* 92 (1991) 631–635.
- [11] A.J. Sederman, M.D. Mantle, C. Buckley, L.F. Gladden, MRI technique for measurement of velocity vectors, acceleration and autocorrelation functions in turbulent flow, *J. Magn. Reson.* 166 (2004) 182–189.
- [12] P. Galvosas, P.T. Callaghan, Fast magnetic resonance imaging and velocimetry for liquids under high flow rates, *J. Magn. Reson.* 181 (2006) 119–125.
- [13] B. Jung, M. Honal, P. Ullmann, J. Hennig, M. Markl, Highly $\mathbf{k-t}$ -space-accelerated phase-contrast MRI, *Magn. Reson. Med.* 60 (2008) 1169–1177.
- [14] P. Parasoglou, D. Malioutov, A.J. Sederman, J. Rasburn, H. Powell, L.F. Gladden, A. Blake, M.L. Johns, Quantitative single point imaging with compressed sensing, *J. Magn. Reson.* 201 (2009) 72–80.
- [15] U. Gamber, P. Boesiger, S. Kozerke, Compressed sensing in dynamic MRI, *Magn. Reson. Med.* 59 (2008) 365–373.
- [16] M.H. Sankey, D.J. Holland, A.J. Sederman, L.F. Gladden, Magnetic resonance velocity imaging of liquid and gas two-phase flow in packed beds, *J. Magn. Reson.* 196 (2009) 142–148.
- [17] K. Scheffler, J. Hennig, Reduced circular field-of-view imaging, *Magn. Reson. Med.* 40 (1998) 474–480.
- [18] D.C. Peters, F.R. Korosec, T.M. Grist, W.F. Block, J.E. Holden, K.K. Vigen, C.A. Mistretta, Undersampled projection reconstruction applied to MR angiography, *Magn. Reson. Med.* 43 (2000) 91–101.
- [19] C.-M. Tsai, D.G. Nishimura, Reduced aliasing artifacts using variable-density \mathbf{k} -space sampling trajectories, *Magn. Reson. Med.* 43 (2000) 452–458.
- [20] G. McGibney, M.R. Smith, S.T. Nichols, A. Crawley, Quantitative evaluation of several partial Fourier reconstruction algorithms used in MRI, *Magn. Reson. Med.* 30 (1993) 51–59.
- [21] D.K. Sodickson, W.J. Manning, Simultaneous acquisition of spatial harmonics (SMASH): fast imaging with radiofrequency coil arrays, *Magn. Reson. Med.* 38 (1997) 591–603.
- [22] E.J. Candès, J. Romberg, T. Tao, Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information, *IEEE Trans. Inform. Theor.* 52 (2006) 489–509.
- [23] D. Donoho, Compressed sensing, *IEEE Trans. Inform. Theor.* 52 (2006) 1289–1306.
- [24] D.S. Taubman, M.W. Marcellin, JPEG2000: standard for interactive imaging, *Proc. IEEE* 90 (2002) 1336–1357.

- [25] S.-J. Kim, K. Koh, M. Lustig, S. Boyd, D. Gorinevsky, A method for large-scale ℓ_1 -regularized least squares, *IEEE J. Sel. Top. Signal. 1* (2007) 606–617.
- [26] M. Fornasier, H. Rauhut, Iterative thresholding algorithms, *Appl. Comput. Harmon. Anal.* 25 (2008) 187–208.
- [27] W. Yin, S. Osher, D. Goldfarb, J. Darbon, Bregman iterative algorithms for ℓ_1 -minimization with applications to compressed sensing, *SIAM J. Imaging Sci.* 1 (2008) 143–168.
- [28] E.J. Candès, F. Guo, New multiscale transforms, minimum total variation synthesis: applications to edge-preserving image reconstruction, *Signal Process.* 82 (2002) 1519–1543.
- [29] E.J. Candès, T. Tao, Near optimal signal recovery from random projections: universal encoding strategies?, *IEEE Trans Inform. Theor.* 52 (2006) 5406–5425.
- [30] M. Lustig, Sparse MRI, Ph.D. thesis, Stanford University, 2008.
- [31] J. Wood, K. Johnson, Wavelet-packet denoising of magnetic resonance images: importance of Rician statistics at low SNR, *Magn. Reson. Med.* 41 (1999) 631–635.
- [32] J. Trzasko, A. Manduca, Highly undersampled magnetic resonance image reconstruction via homotopic ℓ_0 -minimization, *IEEE Trans. Med. Imaging* 28 (2009) 106–121.
- [33] B. Manz, L.F. Gladden, P.B. Warren, Flow and dispersion in porous media: lattice-Boltzmann and NMR studies, *AIChE J.* 45 (1999) 1845–1854.
- [34] M.D. Mantle, A.J. Sederman, L.F. Gladden, Single- and two-phase flow in fixed-bed reactors: MRI flow visualisation and lattice-Boltzmann simulations, *Chem. Eng. Sci.* 56 (2001) 523–529.
- [35] B.A. Olshausen, D.J. Field, Natural image statistics and efficient coding, *Network Comp. Neural Syst.* 7 (1996) 333–339.
- [36] M.S. Lewicki, T.J. Sejnowski, Learning overcomplete representations, *Neural Comp.* 12 (2000) 337–365.
- [37] J.M. Shapiro, Embedded image coding using zerotrees of wavelet coefficients, *IEEE Trans. Signal Process.* 41 (1993) 3445–3462.
- [38] R.G. Baraniuk, V. Cevher, M. Duarte, C. Hegde, Model-based compressive sensing, *IEEE Trans. Inform. Theor.*, submitted for publication, arXiv:0808.3572v5.