

Compressive sampling with prior information in remotely detected MRI of microfluidic devices

Thomas Z. Teisseyre^{a,b}, Jeffrey L. Paulsen^{b,c}, Vikram S. Bajaj^{b,c,*}, Nicholas W. Halpern-Manners^{b,c}, Alexander Pines^{a,b,c}

^a Graduate Program in Bioengineering, University of California, Berkeley and University of California, San Francisco, CA, USA

^b Materials Sciences Division, E.O. Lawrence Berkeley National Lab, Berkeley, CA, USA

^c Department of Chemistry, University of California, Berkeley, CA, USA

ARTICLE INFO

Article history:

Received 27 June 2011

Revised 29 September 2011

Available online 19 October 2011

Keywords:

Compressed sensing

MRI

Remote detection

Microfluidics

ABSTRACT

The design and operation of microfluidic analytical devices depends critically on tools to probe micro-scale chemistry and flow dynamics. Magnetic resonance imaging (MRI) seems ideally suited to this task, but its sensitivity is compromised because the fluid-containing channels in “lab on a chip” devices occupy only a small fraction of the enclosing detector’s volume; as a result, the few microfluidic applications of NMR have required custom-designed chips harboring many detectors at specific points of interest. To overcome this limitation, we have developed remotely detected microfluidic MRI, in which an MR image is stored in the phase and intensity of each analyte’s NMR signal and sensitively detected by a single, volume-matched detector at the device outflow, and combined it with compressed sensing for rapid image acquisition. Here, we build upon our previous work and introduce a method that incorporates our prior knowledge of the microfluidic device geometry to further decrease acquisition times. We demonstrate its use in multidimensional velocimetric imaging of a microfluidic mixer, acquiring microscopically detailed images 128 times faster than is possible with conventional sampling. This prior information also informs our choice of sampling schedule, resulting in a scheme that is optimized for a specific flow geometry. Finally, we test our approach in synthetic data and explore potential reconstruction errors as a function of optimization and reconstruction parameters.

© 2011 Elsevier Inc. All rights reserved.

1. Introduction

A multitude of analytical techniques are currently used to probe microscopic fluid flow, chemistry, and chemical dynamics. With few exceptions, the most generic of these employ miniaturized analogs of laboratory-sized analytical devices based on mass spectrometry [1,2] or optical spectroscopy. For chemical analysis, the most frequent incarnations of the latter depend either on fluorescence-detected immunoassays or covalent modification of the detected analyte with a fluorescent chromophore. Velocimetry for flow design and optimization most frequently employs particle imaging velocimetry (PIV), which also requires exogenous chromophores, to understand microscale fluid dynamics [3–5]. While several optical spectroscopic techniques, including attenuated total reflection-fourier transform infra-red (ATR–FTIR) [6] spectroscopy and surface enhanced Raman spectroscopy (SERS) [7,8], promise to deliver chemical information from the analytes directly, they are difficult to apply to complex analytes (e.g. biological mac-

romolecules, emulsions), difficult to integrate in three-dimensional imaging schemes without confocal optics, and insensitive to fluid motion.

We have therefore sought a generic tool that is capable of probing both chemistry and fluid dynamics without perturbing the analyte or dynamics that are the subjects of our measurements. In macroscopic applications, nuclear magnetic resonance (NMR) and magnetic resonance imaging (MRI) can peer deep within opaque materials to directly measure their complex constituents without covalent labeling. NMR is intrinsically sensitive to fluid motion [9]. It has been used but not limited to studying complex fluid mechanics in porous media [10–12], living organisms [13–15], and mixing in microchannels [16–18]. Due to its successful use in similar applications, there have been significant efforts to integrate NMR spectroscopy and MRI onto lab on a chip devices [19–22].

Unfortunately, these microfluidic applications of MRI are compromised by the conventional geometry of the experiment, in which an inductive radiofrequency coil encloses the entire sample. Since the sensitivity of an inductive detector is proportional to the magnetic flux it encloses, the experiment is inherently insensitive for microfluidic and microporous structures whose fluid channels

* Corresponding author at: Graduate Program in Bioengineering, University of California, Berkeley, CA, USA. Fax: +1 510 366 3768.

E-mail address: vikbajaj@gmail.com (V.S. Bajaj).

occupy only a small fraction of the detector volume. Worse, magnetic susceptibility anisotropies at the interface of liquid samples and solid channel boundaries causes line broadening that further reduces sensitivity and complicates the interpretation of spectra. As a result, microfluidic applications of MRI and NMR usually involve arrays of microfabricated inductive detectors placed on a chip at specific points of interest [23,24].

As a favorable alternative to this approach, we have previously demonstrated the application of remotely detected MRI [25,26] in microfluidic devices [27–31]. In this method, a conventional inductive coil encloses the entire microfluidic chip and is used to encode spatial, chemical, and velocimetric information. This information is then stored longitudinally and transported with the fluid to a single, optimized detector located at the device outflow. Since the volume of this detector is matched with the volume of the microfluidic features of interest, we are able to detect the information with orders of magnitude greater sensitivity than with the enclosing volume coil [32]. Fourier transformation of these data then yields the multidimensional image and a correlated time of flight parameter that reflects transport of fluid from the encoding region to the detector. Our results indicate that MRI spatial resolutions of $\sim 10 \mu\text{m}$ and temporal resolutions of $\sim 20 \text{ms}$ are possible in generic microfluidic devices [27,30]. The time-varying flow is resolved here only in the case of steady-state flow, and the time resolution refers to the resolution of the time of flight dimension, not the speed at which images can be acquired. The resolution of the time of flight dimension is limited by the stroboscopic detection speed, which is dictated principally by the residence time of the fluid in the microcoil detector.

There are still considerable opportunities to accelerate the acquisition of remotely detected microfluidic MRI. In our initial study, we noted that microscale systems in which remote detection is most likely to be useful are sparse because their interesting features occupy only a small fraction of the image field of view. We thus employed compressed sensing [33,34], now also being introduced in clinical MRI to acquire and reconstruct images from a fraction of the data ordinarily recorded in an exhaustively sampled experiment. These algorithms rely on a transform to a domain in which the image is sparse (i.e. the wavelet domain) and a nonlinear reconstruction whose objective maximizes the sparsity of the image such that it agrees with experimental data. Previously, we have achieved subsampling factors of 8–16 with phase-sensitive velocimetric data, and up to 32 with absolute-valued intensity data.

Here, we introduce an improvement to this method which uses our *a priori* knowledge of the flow geometry to improve the robustness of the reconstruction process with even fewer data. Similar to the approach in [35], our reconstruction minimizes the l_1 -norm of the compressed representation of the reconstruction to optimize for sparsity while both maintaining consistency with data and enforcing the *a priori* constraint on the image. We expand this technique and apply it in remotely detected MRI of microfluidic devices, achieving subsampling levels of $128\times$. We incorporate prior knowledge in the optimization (a) through a soft outline mask outside of which any image intensity is penalized in the constraint function (b) effectively enforcing an irregular field of view:

$$\min \|\psi m\|_1 \quad (\text{a})$$

$$\text{s.t. } \|Fm - k\|_2 + s\|M_r m\|_2 \quad (\text{b})$$

Here, ε is set to a fraction of the noise level to force agreement with the data, M_r is the image space *a priori* mask, s is the image-mask compliance multiplier. This image-mask compliance multiplier, s , is a weighting factor that biases the reconstruction towards the real space *a priori* mask. The compressive transformation is the wavelet transform, ψ , and the data is manipulated in Fourier space, F , minimizing the reconstructed data, m , while ensuring agreement

with the acquired data, k . The relative weighting of the two constraints in (b) is adjustable through the image mask compliance multiplier s , which is empirically optimized to minimize the error. Furthermore, reducing ε improves agreement with the data while increasing reconstruction time. We also show that prior information about the microfluidic chip geometry alters the optimal k -space sampling scheme, resulting in one that is specific to the features of the image. Specifically, the distribution of sampled points in k -space can be modified to minimize reconstruction errors and tailor the acquisition based on prior information. While our mask contains no details of the internal structure of the fluid channels, we show in simulations and experiments that its inclusion in the objective function and in the k -space sampling optimization improves the fidelity with which these internal details are reconstructed.

Using this method, we are able to accurately reconstruct six-dimensional MRI images (three spatial and three vector velocity dimensions) of a serpentine microfluidic mixer from an MRI data set that is undersampled by a factor of 128. Since knowledge of flow pathways is a generic feature of all fabricated microfluidic devices, we anticipate that this method will be useful in any MRI experiment in a lab on a chip device.

2. Methods

2.1. Remote detection data acquisition

NMR experiments were performed on a 7 T Oxford Instruments superconducting magnet mated to a Varian console probe and gradient coils. Imaging experiments were conducted using magnetic field gradients with nominal maximum values of 100 Gauss/cm along all axes, also produced by Varian, and encoding was performed using a 40 mm Varian volume imaging probe.

A complete description of the remote detection experiment and hardware appear elsewhere [27]. Briefly, water flows from a regulated pressure-driven flow apparatus to the chip, which is enclosed by a volume imaging coil through which RF pulses are applied to accomplish spatial and velocity encoding. The encoded spin information is then stored as long-lived longitudinal magnetization, and it travels to a microsolenoid NMR probe for detection. The microfluidic serpentine mixing chip is etched in glass and consists of three channels that terminate in two inputs and one output. The glass microfluidic chip was fabricated by etching a rectangular channel, $\sim 150 \pm 10 \mu\text{m}$ in width and $90\text{--}120 \mu\text{m}$ in depth, limited by manufacturing tolerances. While scanning confocal microscopy would have been capable of resolving the precise shape of this channel, no such experiment was attempted here, and hence, information about the cross sectional, or through-plane, dimension was not included in the mask which constrains the image to the known geometry of the chip. The chip outlet was connected to a microcapillary ($150 \mu\text{m}$) through which its contents reached the microcoil detector, a 16-turn microsolenoid wound around a tube that encloses the capillary. Pressure-driven flow of water, doped with 2% isopropanol to prevent bubble formation, was maintained with a driving pressure of 10 psi for the duration of these experiments.

The NMR pulse sequence for the remote detection experiment is illustrated in Fig. 1. Following slice-selective excitation of only the region of interest, phase encoding of spatial and velocity information is performed in the context of a spin echo, as detailed in [27]. We employ gradient moment nulling to cancel the phase accrued due to the motion of spins during their spatial encoding. The information is then stored as long-lived longitudinal magnetization by a terminal $\pi/2$ pulse and detected stroboscopically after flow to the detector. During transport to the detector, a gradient pulse is applied to dephase any residual transverse magnetization that

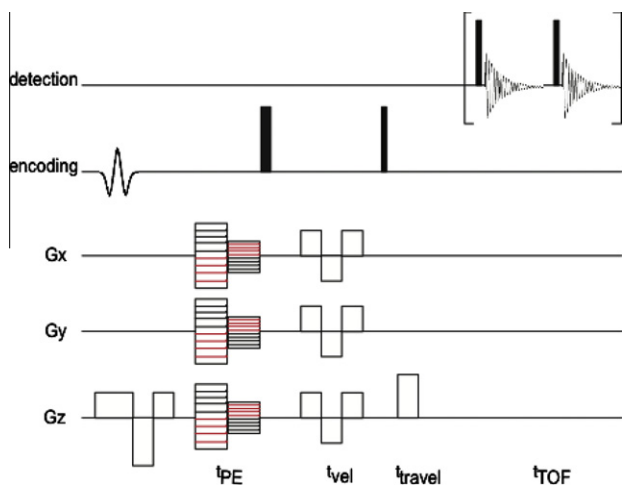


Fig. 1. Pulse sequence for remotely detected MRI with velocity encoding. Radio-frequency pulses for encoding and detection are applied in two separate encoding and detection coils, as indicated. Magnetic field gradients (G_x , G_y , G_z) are used for phase and velocity encoding in each of three spatial dimensions. Transverse magnetization is first created in a region of the sample by a frequency-selective pulse applied concomitantly with a slice selection gradient (G_z) in the longitudinal direction. Next, phase (PE) and velocity (vel) encoding gradients are applied in the context of a spin echo refocusing period. Encoded information is then stored as long-lived longitudinal magnetization and detected in a separate coil stroboscopically with a series of pulses.

persists due to imperfections in the storage pulse, and the pulses are phase cycled to eliminate the contribution of any unencoded signal.

In the following experiment, a three-dimensional velocity encoded image was acquired with 128×128 points in the plane of the chip and 16 points perpendicular to the plane of the chip. However, because the image acquisition was undersampled by a factor of 128 (vide infra), we recorded only 4096 points instead of the 524,288 points that would be required for each exhaustively sampled hyper-complex image. These images were encoded with velocity information in a separate dimension in which the velocity encoding gradient was switched between positive and negative phase to generate a phase contrast proportional to the velocity. For each point, the total time for the encoded volume to flow through the detection region was 2 s, with each of the 50 stroboscopically acquired FIDs lasting

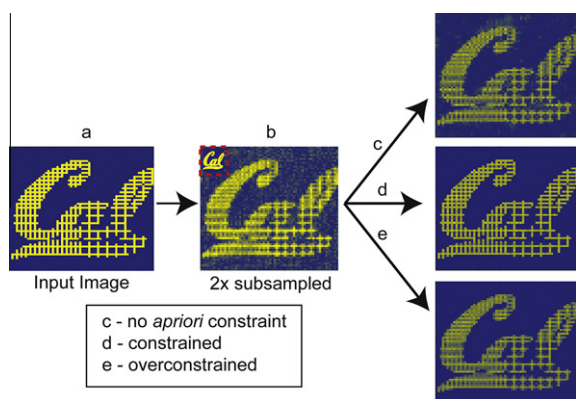


Fig. 2. Illustrative example of sample data reconstruction scheme. A 64×64 Cal logo containing a grid and noise (1%, see text) (a) was Fourier transformed and 2x subsampled (b). The subsampled data are then reconstructed using a soft mask (upper left of Fourier Transformed image) to incorporate *a priori* constraints on the geometry. The mask is applied with different weights during the reconstruction. The illustration shows the effect of increased mask scaling starting with no prior knowledge (c), constrained with $s = 2$ (d), and overconstrained with $s = 32$ (e).

40 ms. A separate set of 3D images were acquired for each of the three velocity components: x , y , and z .

2.2. Image reconstruction

Images were reconstructed using software implemented in MATLAB and the l_1 -minimization was implemented with an iterative optimizer [36], called from within MATLAB. Following their apodization with a decaying exponential, we transformed and integrated the primary data around the position of the water resonance to yield 50 FIDs for each k space point, constituting a time of flight curve weighted by the k -space interferogram. The data were then arranged into a multidimensional matrix of $128 \times 128 \times 16$ k -space points for each of two velocity-encoded dimensions. The points in the conjugate space that were not sampled were initially replaced with zeros.

Details of our reconstruction scheme for sparse data in the absence of prior information constraints have been described elsewhere [31] and will only be treated briefly here. First, sampling tables were determined by choosing an optimal distribution from among many randomly generated k -space distributions, each weighted toward the center of k space. The reconstruction uses an iterative algorithm that constrains the difference between the image and the measured data and maximizes its sparsity in the transform (wavelet) domain by minimizing the l_1 norm as discussed earlier. It was implemented in Matlab using the Wavelab v.8.02 [37] package and spgl1 v.1.7 [36] for l_1 -norm optimization. Geometric data about the flow geometry were used as added constraints to this minimization, expressed weakly (with the image-mask compliance multiplier) in the objective function by an outline mask outside which image pixel intensities are constrained to lie below a threshold value. By “outline mask,” we mean that the mask is a geometrically weak constraint; in that, it is completely homogeneous with respect to the internal structure of the fluid channels and only serves to constrain the edge topology of those channels within the field of view.

2.3. Synthetic data reconstruction

Reconstruction fidelity is dictated by the strength of this image-mask compliance multiplier. We therefore explore the possible regimes of scaling *in silico* by constructing a model image that is internally structured on several length scales: a 64×64 image of the University of California, Berkeley, logo, divided into an intersecting line grid, and subsampled in the Fourier domain by a factor of two. We simulated experimental noise in the test image by adding normally distributed noise set to 1% of the maximum image intensity. Next, the mask corresponding to this was a similar logo that was homogeneous in its internal structure. During the reconstruction, the image-mask compliance multiplier (s) was applied as noted in the aforesaid equation and varied for eight different values ranging from 0 to 32.

2.4. Experimental data reconstruction and processing

Experimental data in the microfluidic mixing chip were reconstructed separately for each time of flight point. Boundary constraints that incorporate *a priori* information about the chip geometry were imposed in the form of an image space mask generated from a processed photograph of the chip (Fig 3a) and aligned to a low-resolution scout image. Contrast between the channel and the glass in the photograph was created by injecting dye into the channel. The chip geometry boundaries were determined from the photograph in MATLAB. The channel was identified by thresholded pixels above the average glass pixel value, and assigning these a value of 1. Values below the threshold were

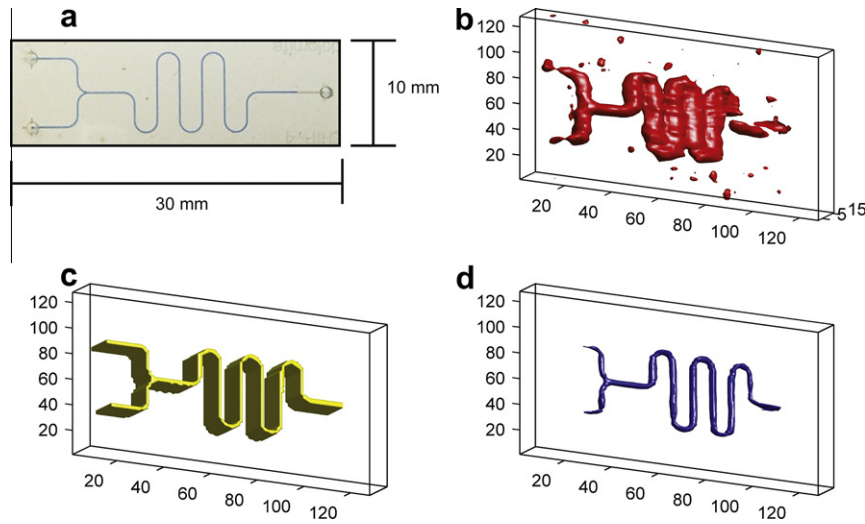


Fig. 3. Comparative figure illustrating benefits of *a priori* masking. Image of the microfluidic serpentine mixer (a), image resulting from direct Fourier transformation of subsampled data (b), applied mask generated from serpentine chip photograph (c), and reconstructed surface of the chip geometry in which the mask was used during the reconstruction (d). All surfaces were cropped to include only the first 90 points along the longitudinal direction to avoid image wrapping from the outlet capillary.

assigned to zero. Lastly, the image was Gaussian filtered in k -space to accommodate for registration imperfections. Artifacts were manually removed from the mask photograph, and the 2D section was replicated across all 16 slices of the acquired image dimension. The reconstruction was carried out at moderate image-mask compliance value, $s = 2$, to control relative weighting of compliance to the *a priori* mask and data. The inequality constraint for the l_2 -norm constraint was determined by taking a fraction of the average noise value from an early TOF point at which no encoded signal had yet reached the detector. A schematic illustration of this procedure is shown in the [Supporting information](#).

We next extracted vector velocity information from the reconstructed data. For each component of the velocity, the phase difference between positive and negative gradient acquisitions was determined for each voxel. The data were corrected by subtracting a linear phase in the Z dimension that persists even in a static phantom. Finally, the phase data were converted to velocity units and plotted on a 3D surface as velocity vectors ([Fig. 5](#)).




Fig. 4. Images of the serpentine mixing chip $128 \times$ subsampled and reconstructed without (a) and with prior information (b). Prior information was a mask during the reconstruction with an optimized scaling parameter. The three-dimensional image is displayed here as a projection along the orthogonal dimension, not illustrated.

This same experimental data were used for comparison with equivalent reconstructions conducted without the benefit of flow information. Two-dimensional projections from the three (spatial)-dimensional data were acquired by summing the intensity information across the through plane dimension. These 2D projections images were displayed for an image-mask compliance multiplier reconstruction of $s = 0$ (no mask) and $s = 2$ ([Fig. 4](#)).

2.5. Reconstruction and k -space sampling optimization

Lastly, we produced an optimized sampling schedule by varying both the random Gaussian sampling distribution and image-mask compliance (s) multiplier. The Gaussian sampling distribution, used to determine the k -space gradient values in an experimental subsampling acquisition, was varied by its relative weighting toward the center of k -space. Ideally, k -space sampling captures




Fig. 5. Three-dimensional, vector velocity field superimposed on an image of the chip surface. All planes of the velocity field are displayed for this perspective of the 3D volume. The three-dimensional surface model, reconstructed from intensity data, is shown in red. Arrows correspond to velocity vectors and are shown in blue. The spacing of the velocity vectors is dictated by a grid corresponding to the acquired voxels. There is asymmetric flow in each inlet due to differing lengths of capillary tubing that supply the chip from a source at common pressure (a). The velocity direction changes at each turn in the chip (b). The average linear velocity for the straight segments of flow in the microfluidic chip was 19.04 cm/s with a standard deviation of 1.82 cm/s .

both high- and low-frequency components. However, with significantly fewer sampled points, the relative distribution of the points should be optimized to yield minimal reconstruction error and try to match the statistics of the data. Additionally, *a priori* information about the reconstructed geometry changes the optimal sampling scheme by providing additional information in the reconstruction. Second, the image-mask compliance multiplier (s), which dictates the relative adherence of the data to the *a priori* mask, was also optimized to generate minimal reconstruction error. Both variables were tested in a sample reconstruction of a synthetic data set that resembles the experimental microfluidic chip. First, a synthetic three-dimensional data set was generated from the experimental data as described earlier for generating the mask. Voxels inside the channel were given a value of 1 and voxels outside the channel were given a value of zero. The resulting synthetic data resulted in a three-dimensional image that was, like the experimental data $128 \times 16 \times 128$ pixels. Unlike the *a priori* mask, the synthetic sample data only occupied four slices of the 16 slice three-dimensional data set. Normally distributed noise was then added to the data set, as described previously. The data were Fourier transformed and subsampled with five different $128 \times$ subsampled randomly generated k -space distributions with various weightings toward the center of k -space. The errors from the best and worst reconstructions were analyzed in detail. The two corresponding sampling schemes, which result in the lowest and highest reconstruction errors, are illustrated in the left column of Fig. 6 as two-dimensional slices of a three-dimensional k -space sampling scheme. The k -space sam-

pling scheme at the upper panel of this figure, which corresponds to the reconstructions in Fig. 6a and b, was more weighted toward the center of k -space. The scheme illustrated in the bottom panel is a more diffuse distribution with more points weighted to sample higher-resolution points in k -space. Each sampling table was restricted to fully sample the centermost points of k -space. The reconstructions were calculated with the same *a priori* mask which outlines the channel geometry as in the experimental reconstruction. An intermediate value of image-mask compliance (s) was used. Each of the reconstructed data sets is quantitatively analyzed to determine the error resulting from $128 \times$ subsampling and reconstruction. Reconstruction errors were determined by comparing the reconstruction to noiseless input synthetic data. Because our reconstruction algorithm scales the data due to a Fourier transform function convention, the reconstructed image was first normalized, rendering it on same scale as the input data. The errors were calculated by subtracting the normalized reconstructed image from the noiseless fully sampled input synthetic data. For percentage calculations, this error was normalized to the noiseless fully sampled input synthetic data. The sampling scheme that yielded the smallest average reconstruction error in the channel was determined to be the optimal sampling schedule. We then investigated the errors resulting from varying the image-mask compliance variable(s) for a data set subsampled with this optimized sampling scheme. As with the Cal logo, reconstructions were calculated for eight different image-mask compliance multipliers (s), ranging from 0 to 32. Reconstruction errors were determined

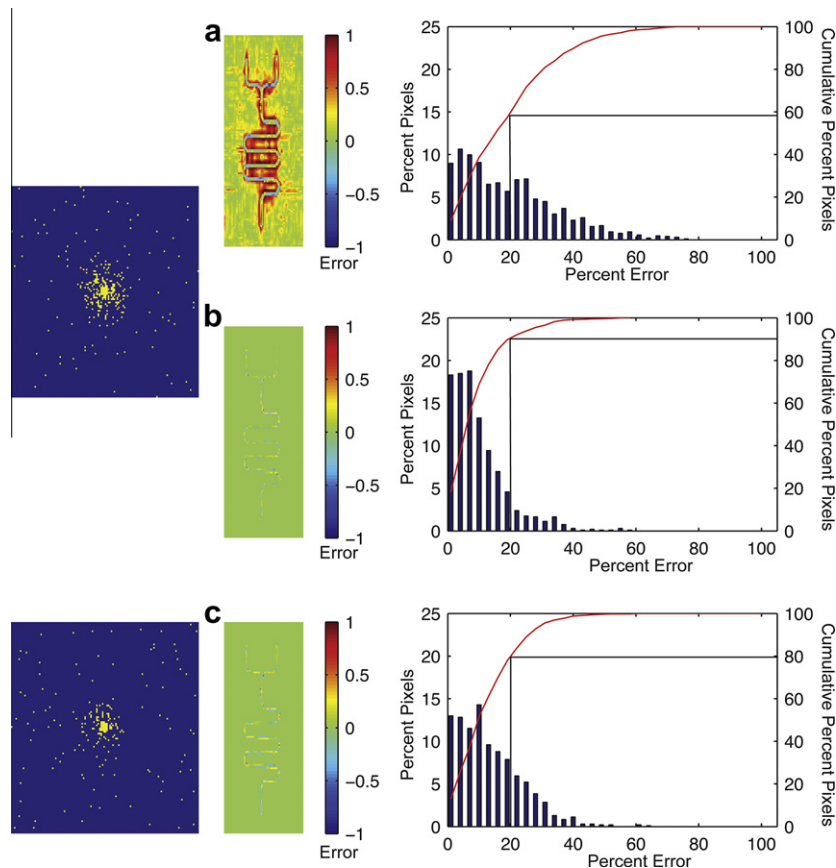


Fig. 6. Subsampling mask optimization and error calculations by reconstruction of synthetic three-dimensional $128 \times$ subsampled data. Reconstructions were done without (a) and with (b) prior information for an optimized sampling schedule. Sampling optimizations were done by varying the weighting toward the center of k -space (c). Errors are illustrated in one slice of real space images (middle column) for the entire region. A quantitative error assessment of the interior of the channel is shown in histograms (left column). Percentage error is binned and shown on the vertical axis of each histogram. The percentage of pixels within each error bin is shown on the left vertical axis and the cumulative percentage is shown with a red line referenced to the right vertical axis. Each histogram also contains a reference line showing the cumulative percentage of pixels that fall within 20% error.




Fig. 7. Iterative reconstruction with prior information of a $64 \times 64 \times$ subsampled University of California, Berkeley logo. Each of the three images illustrates iterations in the thresholding and masking algorithm which eventually converges to a final solution.

as described earlier, and the optimal image-mask compliance value was selected to be the reconstruction with least average error within the channel.

2.6. Iterative thresholding of synthetic data

Because this kind of reconstruction is computationally intensive, we also implemented an iterative soft thresholding scheme with *a priori* information. This scheme has been introduced previously and shown to be equivalent to l1-norm minimization, but can potentially be much faster to compute [38–40].

In our implementation of this technique, the data are multiplied by the image space *a priori* mask to enforce the prior knowledge constraint. In the sparse, wavelet domain, the data is then soft thresholded by a threshold value τ , here empirically determined, as defined by the following function:

$$(\psi m)_\tau = \begin{cases} \psi m - \frac{\tau(\psi m)}{|\psi m|}, & |\psi m| > \tau \\ 0, & |\psi m| < \tau \end{cases} \quad (2)$$

Finally, the image is transformed into the conjugate k space, and any data points that were removed by thresholding but which were determined experimentally are replaced by their measured values. This sequence is then repeated until the reconstruction converges. For reconstruction of the University of California logo, we used 50 iterations to complete the reconstruction. These results are shown in Fig. 7.

3. Results and discussion

3.1. Sample data reconstruction

To investigate the effects of varying the image-mask compliance multiplier on the fidelity of reconstruction, we used a synthetic image consisting of an internally structured feature placed upon a blank image field, as described in Fig. 2. Direct Fourier transformation of the subsampled k -space representation of this image (Fig. 2a) resulted in unacceptable artifacts (Fig. 2b). This linear transformation was not expected to be an acceptable solution of the inherently ill-posed inverse problem, and indeed the resulting artifacts distort the internal structure at all spatial wavelengths and similarly obscure the background.

Next, utilizing a compressed sensing reconstruction without *a priori* information (Fig. 2c) improves the reconstruction quality significantly. Both the vertical and horizontal high-density grid lines become more apparent, but the upper left grid is unclear. The image also exhibits characteristic wavelet artifacts and significant noise in the image background.

The imposition of geometrical constraints in the reconstruction dramatically improves the result, at least within an appropriate range of weighting parameters (Fig. 2b). Near the optimal

image-mask compliance scaling, Fig. 2c, middle, the entire grid is accurately reconstructed at all wavelengths, and the background is also correctly rendered. For high mask-image compliance scaling multipliers, however (Fig. 2c bottom), we note a degradation in reconstruction quality. In the case, the reconstruction overemphasizes agreement with the mask relative to the acquired data. In particular, incoherent artifacts (noise) are folded into the region of interest and distort its internal structure.

3.2. Serpentine mixing chip imaging

To probe the limits of our technique, we applied it to multidimensional velocimetric imaging of a serpentine microfluidic mixer (Fig. 3a). We have acquired velocity-encoded images and, from them, derived intensity images by summing the absolute values of these velocity-encoded data. Fig. 3 illustrates the directly Fourier transformed three-dimensional image (b), corresponding three-dimensional mask (c), the image reconstructed with our algorithm, including the imposition of a known geometrical constraint (d). The reconstructed three-dimensional image, illustrated as a surface layer, is an accurate qualitative depiction of the overall chip geometry, including the proportions of the rectangular channels. The image has relatively few artifacts considering the degree of subsampling involved. This reconstruction is also illustrated as a two-dimensional projection along the orthogonal direction (Fig. 4). The two-dimensional image is shown for both the reconstruction performed without any mask biasing the results towards a known geometry (Fig. 4a) and with the optimal mask and image-mask compliance multiplier (Fig. 4b). The image reconstructed with the mask clearly shows reduced artifacts and an improved representation of the chip geometry.

3.3. Three-dimensional velocimetry in three spatial dimensions

Fig. 5 shows the three-dimensional surface with a three-dimensional velocity vector field superimposed. The input volumetric flowrate is biased more heavily to the lower branch (Fig. 5a), a consequence of an asymmetry in the flow pathway leading to the microfluidic chip. In this image, the overall direction of flow is consistent with our expectations, including at turns and channel boundaries (Fig. 5b). The mean linear velocity within the straight region of the channel is 19.04 cm/s with a standard deviation of 1.82 cm/s. This closely matches previously published results [31], as well as the expected linear velocity for a rectangular channel $125 \mu\text{m} \times 150 \mu\text{m}$ with a volumetric flow rate of 250 $\mu\text{L}/\text{min}$. This is the expected volumetric flow rate at 10 psi for this system.

3.4. Reconstruction error and sampling mask optimization

The manner in which a fixed number of data points are distributed in the conjugate space significantly impacts the reconstruct-

tion quality. We therefore optimized the sampling distribution with a three-dimensional synthetic data set resembling the microfluidic chip. Reconstruction fidelity was calculated quantitatively by comparing the reconstruction with the synthetic input data. In Fig. 6, we show the results of image reconstruction in a three-dimensional synthetic data set subsampled by a factor of 128 in the Fourier domain, illustrating only the center slice (128×128) of the full data set ($128 \times 16 \times 128$). Synthetic data were created using five different 128-fold k -space subsampling schemes, two of which are shown in the left column. We display both differences in intensity as a function of reconstruction parameters (with and without noise) and quantify these errors by means of histograms, illustrating that reconstructions without (Fig. 6a) and with prior information (Fig. 6b) are dramatically different. Reconstruction at this high subsampling ratio yields images with significant artifacts outside the channel and large errors within the channel, unless prior information constraints are used (Fig. 6b). In this case, the majority of pixels are included within 10% error and over 90% of pixels fall within 20% error. This is in contrast to Fig. 6a, where less than 60% of pixels are within 20% error.

The sampling table optimization is also illustrated in Fig. 6b and c, which differ qualitatively in the degree to which the sampling is biased toward the center of k -space. Biasing the sampling toward the center of k -space increases reconstruction error. In this case, the high-resolution components, which are important to sparse microfluidic acquisitions, are lost. While our optimization here is empirical, our results illustrate the potential of achieving higher levels of subsampling by rigorous optimization of the sampling scheme to take advantage of the prior information about microfluidic geometry.

3.5. Iterative thresholding

Lastly, we demonstrate in Fig. 7 the results of the iterative thresholding reconstruction scheme with *a priori* masking, which was used as an alternative but substantially equivalent reconstruction scheme. The results in Fig. 7 illustrate equivalent convergence after 50 iterations of the algorithm.

4. Conclusions

We have demonstrated a method to acquire and reconstruct MRI images of microfluidic devices for which there is some prior knowledge of the flow geometry. Logical extensions of this technique include the encoding of chemical information in microfluidic chemical assays. In those cases, prior knowledge about the NMR spectrum may be similarly integrated into the reconstruction to allow for a higher degree of subsampling. In all cases, we have not yet exploited this prior information in the optimal design of sampling schedules or the devices themselves, and this may result in further savings in acquisition time.

Acknowledgments

This work was supported by the Director, Office of Science, Office of Basic Energy Sciences, Materials Sciences and Engineering Division, of the US Department of Energy under Contract No. DE-AC02-05CH11231. T.Z.T. gratefully acknowledges support from the National Science Foundation Graduate Research Fellowship. The authors thank Professor Michael Lustig for helpful discussions and acknowledge Schlumberger-Doll Research, the Agilent Foundation, and Chevron for their generous and unrestricted support of our research.

Appendix A. Supplementary material

Supplementary data associated with this article can be found, in the online version, at doi:10.1016/j.jmr.2011.10.001.

References

- [1] M.-H. Fortier, E. Bonneil, P. Goodley, P. Thibault, Integrated microfluidic device for mass spectrometry-based proteomics and its application to biomarker discovery programs, *Journal/Analytical Chemistry* 77 (2005) 1631–1640.
- [2] A.A. Dawoud, H.A. Sarvaiya, I.M. Lazar, Microfluidic platform with mass spectrometry detection for the analysis of phosphoproteins, *Journal/Electrophoresis* 28 (2007) 4645–4660.
- [3] Y. Sato, G. Irisawa, M. Ishizuka, K. Hishida, M. Maeda, Visualization of convective mixing in microchannel by fluorescence imaging, *Journal/Measurement Science and Technology* 14 (2003) 114–121.
- [4] N.C.H. Le, R. Yokokawa, D.V. Dao, T.D. Nguyen, J.C. Wells, S. Sugiyama, Versatile microfluidic total internal reflection (TIR)-based devices: application to microbeads velocity measurement and single molecule detection with upright and inverted microscope, *Journal/Lab on a Chip* 9 (2009) 244–250.
- [5] J.G. Santiago, S.T. Wereley, C.D. Meinhart, D.J. Beebe, R.J. Adrian, A particle image velocimetry system for microfluidics, *Journal/Experiments in Fluids* 25 (1998) 316–319.
- [6] K.L.A. Chan, S. Gulati, J.B. Edel, A.J. de Mello, S.G. Kazarian, Chemical imaging of microfluidic flows using ATR-FTIR spectroscopy, *Journal/Lab on a Chip* 9 (2009) 2909–2913.
- [7] D. Choi, T. Kang, H. Cho, Y. Choi, L.P. Lee, Additional amplifications of SERS via an optofluidic CD-based platform, *Journal/Lab on a Chip* 9 (2009) 239–243.
- [8] H. Cho, B. Lee, G.L. Liu, A. Agarwal, L.P. Lee, Label-free and highly sensitive biomolecular detection using SERS and electrokinetic preconcentration, *Journal/Lab on a Chip* 9 (2009) 3360–3363.
- [9] E. Fukushima, Nuclear magnetic resonance as a tool to study flow, *Journal/Annual Review of Fluid Mechanics* 31 (1999) 95–123.
- [10] A. Sederman, Magnetic resonance imaging of liquid flow and pore structure within packed beds, *Journal/Chemical Engineering Science* 52 (1997) 2239–2250.
- [11] a. Sederman, Structure-flow correlations in packed beds, *Journal/Chemical Engineering Science* 53 (1998) 2117–2128.
- [12] Y.E. Kutsovsky, L.E. Scriven, H.T. Davis, B.E. Hammer, NMR imaging of velocity profiles and velocity distributions in bead packs, *Journal/Physics of Fluids* 8 (1996) 863.
- [13] D.H. Szolar, H. Sakuma, C.B. Higgins, Cardiovascular applications of magnetic resonance flow and velocity measurements, *Journal/Journal of Magnetic Resonance Imaging* 6 (1996) 78–89.
- [14] E. Be'eri, S.E. Maier, M.J. Landzberg, T. Chung, T. Geva, In vivo evaluation of Fontan pathway flow dynamics by multidimensional phase-velocity magnetic resonance imaging, *Journal/Circulation* 98 (1998) 2873–2882.
- [15] S. Sharma, A.E. Ensley, K. Hopkins, G.P. Chatzimavroudis, T.M. Healy, V.K. Tam, K.R. Kanter, A.P. Yoganathan, In vivo flow dynamics of the total cavopulmonary connection from three-dimensional multislice magnetic resonance imaging, *Journal/The Annals of Thoracic Surgery* 71 (2001) 889–898.
- [16] X. Zhang, A.G. Webb, Magnetic resonance microimaging and numerical simulations of velocity fields inside enlarged flow cells used for coupled NMR microseparations, *Journal/Analytical Chemistry* 77 (2005) 1338–1344.
- [17] S. Ahola, F. Casanova, J. Perlo, K. Münnemann, B. Blümich, S. Stapf, Monitoring of fluid motion in a micromixer by dynamic NMR microscopy, *Journal/Lab on a Chip* 6 (2006) 90–95.
- [18] L.G. Raguin, L. Cioabanu, Multiple echo NMR velocimetry: fast and localized measurements of steady and pulsatile flows in small channels, *Journal/Journal of Magnetic Resonance (San Diego, Calif.: 1997)* 184 (2007) 337–343.
- [19] Y. Maguire, I.L. Chuang, S. Zhang, N. Gershenfeld, Ultra-small-sample molecular structure detection using microslot waveguide nuclear spin resonance, *Journal/Proceedings of the National Academy of Sciences of the United States of America* 104 (2007) 9198–9203.
- [20] H.G. Krojanski, J. Lambert, Y. Gerikalan, D. Suter, R. Hergenröder, Microslot NMR probe for metabolomics studies, *Journal/Analytical Chemistry* 80 (2008) 8668–8672.
- [21] J. Bart, A.P.M. Kentgens, J.W.G. Janssen, P.J.M. van Bentum, J.G.E. Gardeniers, Optimization of stripline-based microfluidic chips for high-resolution NMR, *Journal/Journal of Magnetic Resonance* 201 (2009) 175–185.
- [22] J. Bart, A.J. Kolkman, A.J. {Oosthoek-de Vries}, K. Koch, P.J. Nieuwland, H.J.W.G. Janssen, J.P.J.M. van Bentum, K.a.M. Ampt, F.P.J.T. Rutjes, S.S. Wijmenga, H.J.G.E. Gardeniers, A.P.M. Kentgens, A microfluidic high-resolution NMR flow probe, *Journal/Journal of the American Chemical Society* 131 (2009) 5014–5015.
- [23] J.D. Trumbull, I.K. Glasgow, D.J. Beebe, R.L. Magin, Integrating microfabricated fluidic systems and NMR spectroscopy, *Journal/IEEE Transactions on Biomedical Engineering* 47 (2000) 3–7.
- [24] C. Massin, Planar microcoil-based microfluidic NMR probes, *Journal/Journal of Magnetic Resonance* 164 (2003) 242–255.
- [25] A.J. Moulé, M.M. Spence, S.-I. Han, J.A. Seeley, K.L. Pierce, S. Saxena, A. Pines, Amplification of xenon NMR and MRI by remote detection, *Journal/Proceedings of the National Academy of Sciences of the United States of America* 100 (2003) 9122–9127.

- [26] J. Granwehr, E. Harel, S. Han, S. Garcia, A. Pines, Time-of-flight flow imaging using NMR remote detection, *Journal/Physical Review Letters* 95 (2005).
- [27] V.S. Bajaj, J. Paulsen, E. Harel, A. Pines, Zooming in on microscopic flow by remotely detected MRI, *Journal/Science (New York, NY)* (2010).
- [28] C. Hilty, E.E. McDonnell, J. Granwehr, K.L. Pierce, S.-I. Han, A. Pines, Microfluidic gas-flow profiling using remote-detection NMR, *Journal/Proceedings of the National Academy of Sciences of the United States of America* 102 (2005) 14960–14963.
- [29] E.E. McDonnell, S. Han, C. Hilty, K.L. Pierce, A. Pines, NMR analysis on microfluidic devices by remote detection, *Journal/A analytical Chemistry* 77 (2005) 8109–8114.
- [30] E. Harel, A. Pines, Spectrally resolved flow imaging of fluids inside a microfluidic chip with ultrahigh time resolution, *Journal/Journal of Magnetic Resonance* 193 (2008) 199–206.
- [31] J. Paulsen, V.S. Bajaj, A. Pines, Compressed sensing of remotely detected MRI velocimetry in microfluidics, *Journal/Journal of Magnetic Resonance (San Diego, Calif.: 1997)* 205 (2010) 196–201.
- [32] J.A. Seeley, J. Granwehr, Sensitivity quantification of remote detection NMR and MRI, *Journal/Journal of Magnetic Resonance* 179 (2006) 280–289.
- [33] D. Donoho, Compressed sensing, *Journal/IEEE Transactions on Information Theory* 52 (4) (2006) 1289–1306.
- [34] E.J. Candes, J. Romberg, T. Tao, Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information, *IEEE Trans. inf. Theory* 52 (2006) 489–509.
- [35] D.J. Holland, D.M. Malioutov, A. Blake, A.J. Sederman, L.F. Gladden, Reducing data acquisition times in phase-encoded velocity imaging using compressed sensing, *Journal/Journal of Magnetic Resonance* 203 (2010) 236–246.
- [36] E.v.D. Berg, M.P. Friedlander, Probing the frontier for basis pursuit solutions, *Journal/SIAM Journal of Scientific Computing* 31 (2008) 890–912.
- [37] J. Buckheit, S. Chen, D. Donoho, I. Johnstone, J. Scargle, About Wavelab, *Journal/Stanford University Tech. Report*, 2005.
- [38] K. Bredies, D.A. Lorenz, Linear convergence of iterative soft-thresholding, *Journal/Journal of Fourier Analysis and Applications* 14 (2008) 813–837.
- [39] A.S. Stern, D.L. Donoho, J.C. Hoch, NMR data processing using iterative thresholding and minimum $l(1)$ -norm reconstruction, *Journal/Journal of Magnetic Resonance (San Diego, Calif.: 1997)* 188 (2007) 295–300.
- [40] I. Drori, Fast $l(1)$ minimization by iterative thresholding for multidimensional NMR spectroscopy, *Journal/EURASIP Journal on Advances in Signal Processing* 2007 (2007) 1–11.