Fast Magnetic Resonance Imaging via Adaptive Broadband Encoding of the MR Signal Content

Dimitris Mitsouras, Frank J. Rybicki, Alan Edelman, Gary P. Zientara

***Abstract*— Our goal is to increase the time-efficiency of continuous data acquisition in Magnetic Resonance Imaging. In order to realize this increase in efficiency, we follow the proposition of extending the applicability of the physics that MR imaging is based on. That is, the MR signal may be partially encoded at the excitation step. In so doing, we create a novel application of algebraic matrix factorization technologies. This work presents the underlying technologies that enable this non-Fourier broadband encoding path to offer time efficiencies compared to present day fast MRI methods. Finally, we present an implementation combining these principles in a proof-of-concept practical system. The system works in concert with a commercial MRI scanner, requires minor modification to the scanner, and, using readily available computing resources successfully increases imaging efficiency by up to one order of magnitude for 2D dynamic MRI.**

***Index Terms*— Dynamic MRI, Adaptive MRI, Fast MRI, Minimum data MR image reconstruction, Broadband MRI.**

I. Introduction

Magnetic Resonance Imaging (MRI) demonstrates excellent soft tissue contrast and can be used for a wide variety of clinical applications. However, with respect to a diagnostic scan for routine clinical practice, MRI is inherently slow, for example in comparison with state-of-the-art multirow computed tomography. On one hand, the acquired MR signal is radiated from the sample and decays over time both via the return of the sample's magnetization to thermal equilibrium (T_1 longitudinal relaxation), as well as the loss of phase coherence of spins still in their excited state (T_2^* transverse relaxation). On the other hand, in order to achieve image reconstruction (i.e. solve the inverse problem) enough encoding steps of the signal must be performed in order to localize it in each spatial dimension at the desired resolution. With few exceptions, present-day clinical MR imaging relies on Fourier encoding of the MR signal [1].

Fourier Encoded MRI employs gradient magnetic field-induced Fourier basis functions on the sample's transverse magnetization for signal localization. The acquired MR signal is the discretized value of the voltage induced on surrounding receiver coil(s) by the ensemble of excited spins in the sample that are rotating about and proportionally to the strength of the magnetic field. Manipulation of constant field gradients yields the k-space signal matrix, i.e., the Discrete Fourier Transform (DFT) of the excited MR spin density. Practically then, most sampling schemes can fill in only a small portion of the k-space matrix per excitation, since the received signal decays after sample excitation. Relaxation times of the order of 1sec in human tissue make the process of filling in the entire matrix time consuming if performed naïvely.

Fast MRI is necessary for dynamic studies where a continuous series of images is required, for example in contrast bolus tracking or the monitoring of interventional procedures. Many methods have been devised to increase MRI efficiency without sacrificing spatial resolution. Generally, these fall into roughly one of three categories. The most popular, *Signal Squeezing* methods, acquire more Fourier coefficients per unit time (for a review see [2]). This can be accomplished by employing optimized trajectories through k-space, such as echo planar (EPI) [3] and spirals [4], or, by acquiring multiple spin echoes after a single excitation (e.g. RARE [5]). Alternatively, steady states of the magnetization throughout the imaging experiment, enforced by fast, periodic, RF excitation, allow one to retain constant signal strength throughout the experiment so that the k-space matrix can be filled in a streaming fashion (see e.g. [2], or [6]). *Parallel Imaging* methods (e.g. [7], [8], [9]) use the inherent receiver coil sensitivity to location of radiating nuclei in order to replace some of the gradient-induced signal encoding. When an array of coils with known independent spatial sensitivities, acquire the MR signal in parallel, the ensemble of spatial encoding information can be used for partial signal localization.

Another category of fast MRI is *Data-Selective* methods that aim to use *a priori* assumptions about the imaged object, or place constraints on image reconstruction. For example, restricted Field-of-View (FOV) imaging of only a selected region of interest maintains resolution with fewer encoding steps (e.g. [10]). Alternatively, simplifying assumptions, such as the nature or limits of the dynamic or quasi-static portions of the image that undergoes dynamic changes, can be used to offer reduced encoding requirements (see e.g. [11], [12]). Data selection can be performed in a space other than Fourier; the RIGR method [13] expands a high resolution reference image as a generalized series of complex exponentials which is subsequently used to extrapolate low resolution dynamically acquired images into high resolution ones. The "feature-recognizing" method [14] uses a large database of similar

D. Mitsouras is with the Department of Electrical Engineering and Computer Science, Lab for Computer Science, Massachusetts Institute of Technology, Cambridge, Massachusetts.

F.J. Rybicki is with the Department of Radiology, Brigham and Women's Hospital and Harvard Medical School, Boston, Massachusetts.

A. Edelman is with the Department of Applied Mathematics, Massachusetts Institute of Technology, Cambridge, Massachusetts.

G.P. Zientara is with the Department of Radiology, Brigham and Women's Hospital and Harvard Medical School, Boston, Massachusetts.

This research was supported in part by NIH NCI R01-CA78299, NIH NCI P01-CA67165, a BWH Research Council Interdisciplinary Seed Grant, an Edgerly Science Partnership Award (A.E. and G.Z.), an Oxygen Alliance grant and equipment loans by Mercury Computer Systems Inc.

anatomy images, hypothesizing that e.g., any new brain image can be expressed as a linear combination of the initial set, and, acquiring the Fourier coefficients that best represent that expansion basis. Wavelet MRI [15] proposed an increase in acquisition efficiency by directly encoding the MR signal using non-Fourier bases. Wavelets with a small support can be excited in an order that maximizes the time between excitation of spatially overlapping wavelets, thus yielding longer effective tissue relaxation times [16]. Rather than k-space Fourier coefficients, wavelet encoding yields wavelet coefficients of the spin density.

All clinical MRI scanners today use spatially selective radio frequency (RF) excitation, typically for selection of the slice of interest through the sample. However, this capacity is not generally used to full advantage. When the superposition principle is applicable [17], a signal generated by a linear combination of any set of elemental regions can be obtained via selective RF excitation. Direct manipulation of these combinations can then yield expansions of the MR signal onto any basis. Coupling this ability of MRI with linear algebraic methods, a more efficient process of MR signal acquisition can be realized. Associated with every matrix M is a column and row space. In many cases, the useful information content may be concentrated in a relatively small part of these vector spaces. In matrix language, if Mv or $v^T M$ is large (small) relative to v, then v contains important (unimportant) information about M. Using spatially selective RF excitation, *any* MR system has the capacity to produce Mv or $v^T M$ for a large class of v.

This work expands on the ability to acquire and reconstruct a good approximation to the FOV in a fraction of typical acquisition times by limiting the subspace spanned by the acquired data [18]. Rather than focusing on the use of a specific set of encoding vectors v, this work aims to reduce the technological and implementation gaps between non-Fourier encoding in MRI and present-day fast dynamic MRI. In summary, we contribute some methods and systems that enable any novel encoding method that does not rely on Fourier basis functions to be applied to fast dynamic MR imaging. Following a brief introduction of the digital non-Fourier encoding principle [19] in Section II-B, the first goal is largely achieved in Section III where we show how to control the efficiency provided by non-Fourier encoding while minimizing subspace truncation errors. Furthermore, we show how non-Fourier encoding can coexist with or take advantage of current fast MRI approaches such as parallel MRI and multi-echo acquisitions. The second goal is achieved in Section IV, where we present the design and implementation of an adaptive "pipeline" system that works in tandem with a clinical MR scanner to perform continuous, adaptive, non-Fourier encoding. Finally, in Section V we present some preliminary results obtained with this system.

II. 2D Non-Fourier Encoded MRI

Fourier MRI produces a discretized form of the Fourier coefficients of the magnetization distribution of a sample, and, via subsequent DFTs the image is reconstructed. Non-Fourier encoded MRI produces expansion coefficients of the magnetization distribution onto an arbitrary basis. The magnetization distribution or its Fourier coefficients are then reconstructed using the basis. The digital non-Fourier encoding principle [19] describes an alternative to Fourier encoding using familiar constructs, such as k-space, and in a way that allows non-Fourier encoding to be described and analyzed on an equal footing to Fourier encoding.

A. Fourier Encoding Basics

A 2D Fourier-basis MRI experiment typically begins with the excitation of the sample's magnetization within the slice $z_0 \pm \Delta z$ by means of a selective RF excitation. During application of this RF pulse an associated gradient imposes a linear dependence of magnetic field strength (i.e. resonance frequency) to location in the sample. An RF pulse that contains a smaller band of frequencies than the sample from edge to edge excites only a slice of the sample's magnetization.

The pulse sequence (PSD) for a *Spin Echo Spin-Warp Experiment* [1] that produces one row of the k-space of a sample per repetition is shown in Fig. 1. The diagram describes the application of the gradients in the three physical axes ($G_{x,y,z}$), the excitatory RF pulses and the locations at which samples from the receiver coils are obtained. Time starts on the left of the diagram and flows to the right. Application is initiated and completed typically in a few tens of milliseconds. A dead time follows to complete the experiment repetition time, TR at which time the sequence is repeated to obtain the next row. The term "spin echo" refers to the use of a 180° flip angle RF pulse in the middle of the experiment in order to partially account for transverse relaxation due to external magnetic field inhomogeneities. Transverse magnetization partially rephases symmetrically in time about the 180° RF pulse.

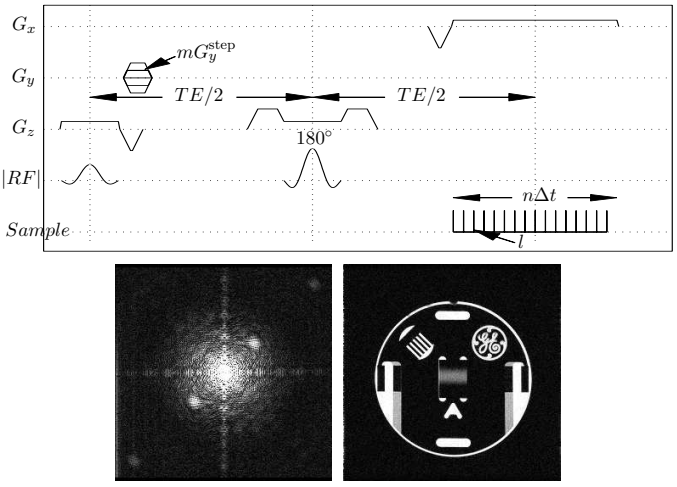

Fig. 1. A simple 2D Fourier encoded spin echo PSD and image produced using 256 repetitions each with 256 samples. Imaging parameters are TE = 60ms, TR = 400ms, 16cm FOV, 10mm slice thickness and 15° flip angle. Bottom left: raw MR data (magnitude image of Fourier coefficients) collected; k_x spans the horizontal and k_y the vertical axis. Bottom right: image after 2D DFT of the raw k-space.

The MR signal is *frequency encoded* along the x axis (arbitrarily) by sampling it at Δt-equidistant points under a G_x constant gradient. The signal equation for the lth sample, obtained at time $t = l\Delta t$ after sampling begins, and assuming

a fixed gradient, $G_y = mG_y^{\text{step}}$ for some m, has been applied, is given by

$$S(t, G_y, z_0) = \iiint \rho(x, y, z)P(z)\times e^{-i2\pi(k_x x + k_y y)}dxdydz, \quad (1)$$

where $\rho(x, y, z)$ is the excited MR spin density weighted by the various imaging factors (e.g. relaxation terms, hardware characteristics etc.), $P(z)$ is the spatial transformation of the RF pulse envelope played under the gradient G_z, and is approximately unity in the interval $[z_0 - \Delta z, z_0 + \Delta z]$ and zero elsewhere, and

$$k_x = \gamma G_x \left(l - 1/2n\right) \Delta t, \quad (2)$$

$$k_y = \gamma G_y \tau, \quad (3)$$

where τ is the fixed duration of the G_y gradient, and γ is the gyromagnetic ratio. The $1/2n\Delta t$ factor in k_x, due to the negative lobe of the G_x gradient applied prior to sampling, is used to center the samples in k-space.

The signal produced by the excited slice of the sample is localized in the remaining two dimensions by manipulation of the magnetic field gradients; the G_x gradient during sampling (readout) gives rise to the factor $\exp(-i2\pi k_x x)$ in Eq. (1) thus producing the Fourier coefficients of the projection of the excited MR spin density onto x. The desired FOV size,

$$\text{FOV}_x = 1/(\gamma G_x \Delta t) \equiv 1/\Delta k_x, \quad (4)$$

in turn determines sampling frequency (i.e. Δt) and gradient strength. When n samples, $l = 1 \ldots n$, are acquired, the resulting spatial resolution along x is FOV_x/n.

Localization along y is achieved via repeating the sequence while varying the *phase encoding*, i.e., the factor $\exp(-i2\pi k_y y)$. The G_y gradient chosen in each repetition of the sequence produces one row of the k-space matrix. Using the gradient $G_y^{(m)} = (1/2n - m)G_y^{\text{step}}$ in the mth repetition, the FOV along y is given by $\Delta k_y = \gamma G_y^{\text{step}} \tau$. The Fourier coefficients are then obtained from positive to negative order of spatial frequency, since

$$k_y = \gamma G_y^{(m)} \tau. \quad (5)$$

A 2D DFT of The k-space matrix, formed by placing the lth sample of the mth experiment in a matrix, $F^{\mathcal{F}}_{m,l} = S(l\Delta t, G_y^{(m)}, z_0)$ (where $\mathcal{F}$ stands for Fourier encoded), produces the MR image.

The k-space matrix, formed by the discretized spin echoes produced by this MR experiment, intrinsically represents a *system response*. Non-Fourier encoding focuses its attention on how to efficiently and accurately sample the FOV to reconstruct this system response.

B. Non-Fourier Encoding Basics

If the initial RF pulse excitation in the Fourier encoded experiment is replaced with an RF pulse defined by the envelope

$$p(t) = \sum_{m=1}^{n} p_m \hat{\delta}(t - (m-1)\Delta t), \quad (6)$$

where $\hat{\delta}(t)$ is unity in the interval $0 \le t < \Delta t$ and zero elsewhere, the resulting PSD can be used to acquire practically the same results as the Fourier encoded sequence of Fig. 1, but it can also encode it using any arbitrary basis [19]. The intent is to use the RF pulse to replace the phase encoding gradient, hence the RF pulse is "played out" under a G_y spatially selective gradient of duration $n\Delta t$, as shown in Fig. 8.

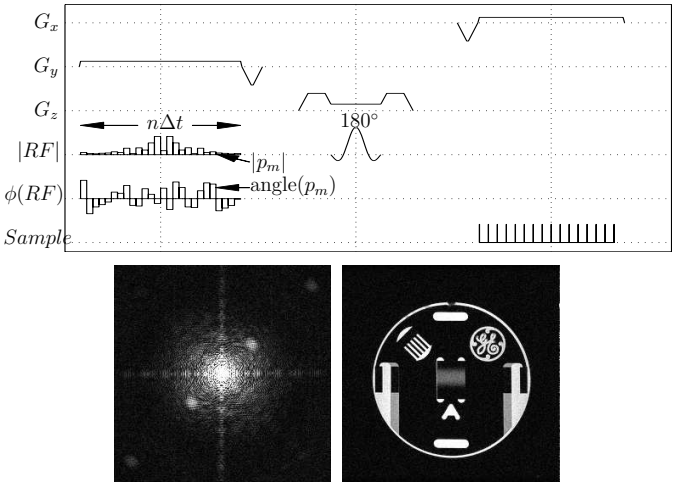

Fig. 2. A simple 1D-spatial/1D-Fourier encoded spin echo PSD and image produced using same parameters as in Fig. 1 but using each row of the identity matrix as the RF pulse of each repetition. Bottom left: raw MR data collected. Bottom right: corresponding image after 2D DFT.

With Δt sufficiently short, each hard pulse (i.e., narrow boxcar $p_m\hat{\delta}(t - t_m) \in \mathbb{C}$) can be approximated by an impulse at time $t_m = m\Delta t$. An impulse (with infinite bandwidth) instantaneously flips all magnetization throughout the FOV to the transverse plane by a flip angle proportional to its power. If impulses are low flip, the resulting transverse magnetization remains undisturbed by subsequent impulses and undergoes free precession under the influence of the remaining gradients only. Since the y-axis gradient is followed by a rephasing lobe of half area, magnetization excited by the impulse at $m\Delta t$ then acquires a phase offset proportional to $((1/2n - m)G_y\Delta t)$.

Linearity and superposition allow us to use the simple Fourier modes of the sample's magnetization as "building blocks" in order to produce arbitrary complex 1D spatial excitation profiles and obtain the response of these profiles of spin density. The necessary condition for superposition is that the entire RF pulse is low flip. If this flip angle, θ, is small (e.g. $\theta < 30°$ [20]), the RF pulse $p(t)$, expressed by the vector $\vec{p}$ containing the p_m, produces a signal response that is the linear combination of the n constituent hard pulse responses [19]:

$$\begin{aligned} S_{\mathcal{SE}}(t, \vec{p}, z_0) &= \int_{z_0-\Delta z}^{z_0+\Delta z} \int_y \int_x \rho(x, y, z)\times \\ &\left(\sum_m p_m e^{-i2\pi k_m y}\right) e^{-i2\pi k_x x} dxdydz \quad (7) \\ &= \sum_m \left\{ p_m S_{\mathcal{SE}}(t, \vec{e}^{(m)}, z_0) \right\}, \end{aligned}$$

where k_x as in Eq. 2, $k_m = \gamma G_y(1/2n - m)\Delta t$ and where

the 180° RF pulse was used to refocus the signal only within the slice $z \in [z_0 - \Delta z, z_0 + \Delta z]$. It is assumed that all other magnetization in the sample is significantly dephased and does not contribute any observable signal at sampling time. Each impulse at m produces a single Fourier mode of the spin density along y, but scaled by the complex value p_m. Relaxation effects aside, using the mth row of the identity matrix, $\vec{e}^{(m)}$, for $\vec{p}$, we obtain precisely the same echo as the mth repetition of the Fourier encoded experiment of Eq. (1). Repeating the non-Fourier experiment using each of the rows of the identity matrix thus produces the system response matrix $F^{\mathcal{F}}$. For arbitrary RF excitations, the term in parentheses of Eq. (7) describes a spatially selective profile that is excited: it is the Fourier transform of the hard pulse train:

$$p(y) = \sum_{k=-1/2n}^{1/2n} \hat{p}_k e^{i(2\pi/\text{FOV})ky}, \tag{8}$$

where $\hat{p}_k = \vec{p}(1/2n - k)$. The spatial profile of magnetization that generates the signal is built by the superposition of the ensemble of Fourier modes excited – or "activated" – by the scaled impulses in the RF train.

Spatial Profile $\equiv$ *Encoding Function.* The linear system expression of Eq. (7) [21], [22],

$$s = p^T F^{\mathcal{F}}, \tag{9}$$

where s is the sampled response, p is the digital RF excitation, and $F^{\mathcal{F}}$ is the sample's MR system response matrix, shows that effectively any MR imager has the ability to compute $v^T M$ if the sample's Fourier coefficient matrix is M and the vector v describes a low flip excitation. Exchanging the x and y gradient axes, Mv can be similarly obtained from any MR system. Efficient non-Fourier encoding aims to acquire this response only for those spatial modes which have the strongest contributions in the imaged object, and possibly adapting these modes as necessary.

Via shaped RF excitation, the digital linear system response model of Eq. (9) enables any arbitrary *boradband* combination of Fourier modes to generate each signal response, rather than the single Fourier mode of the sample per response that is generated in Fourier imaging. Using the rows of an invertible matrix P as excitations in the repetitions of the spatially encoded experiment and gathering the resulting responses in a matrix, we can express the MR imaging process as

$$S = PF, \tag{10}$$

by stacking each sampled response in the rows of the matrix S. Inversion of Eq. (10) yields the k-space matrix F as

$$F = P^\dagger S. \tag{11}$$

When the matrix P is square invertible, the full k-space matrix F is obtained, and, via Fourier transformation, the desired image. Based on Eq. (11), Hadamard, Wavelet and Singular Value Decomposition (SVD) bases have been studied and used in MRI [23], [15], [21]. For example, Hadamard encoding derives the matrix P from Hadamard shaped spatial profiles. In the low flip angle regime this is done by taking the Fourier transform of the rows of the Hadamard matrix.

III. Strategies for Fast Non-Fourier MRI Encoding

More generally, Eq. (11) provides the necessary tools to compress the MR image at the excitation step resulting in the necessity for fewer signal acquisitions. The columns of the reconstructed k-space matrix lie in the subspace spanned by the columns of the encoding matrix P. Careful selection of the subspace spanned by the encoding matrix P yields an acquisition and subsequent reconstruction that can retain most pertinent image information. Furthermore, careful selection of the encoding matrix itself yields a short-wide matrix that, when used for acquisition, requires fewer encoding steps than phase encodes for the desired image resolution. For example, SVD encoded MRI [21] uses a reduced set of basis functions derived from the left singular vectors of the k-space matrix or an estimate to it. When only the first few singular vectors are used for acquisition, the power of F that is captured by the reduced set of encoding steps is maximized while it is expected that this maximization is equivalent to capturing pertinent image information [24]. The DATUM approach [25] proposed continually adapting the small subspace used for each prospective image acquisition in a dynamic series by encoding a linear prediction of the change in F.

A. Adaptive Spatial Encoding – Efficiency Control

However an encoding basis is generated, compression in MRI via control of the encoding matrix cannot sacrifice salient image information. Since any coherent MR signal detected is always contaminated by white, zero mean, Gaussian noise due to the thermal radiation generated in the receiver by the sample [26], any encoding basis can be separated into significant and insignificant subspaces for a given F. Most MR images obtained today have an SNR of approximately between 10 (e.g. functional MRI studies) to 120 (e.g. morphological imaging). SNR is defined as the ratio of the signal energy to the standard deviation of noise (which we denote as σ_{noise}). For some unit vectors v we have that $\|vF\|_2 \leq \sigma_{\text{noise}}$; the projection of the columns of F along the direction of v is smaller than can be explained by the noise component alone and is therefore statistically insignificant.

According to Fig. 3.1, using 84 randomly selected MR images (256 by 256 matrix size) from a clinical MRI database, assuming an SNR of 100, a 179-dimensional subspace contains information less significant than noise. However, 155 Fourier basis functions produce signal that contains information more significant than noise. Similarly, for an SNR of 30, 218 dimensions are insignificant. Using these assumptions, typical Fourier encoded MRI experiments expend 60-85% of scanning time acquiring data that when encoded using a basis other than the Fourier one, is indistinguishable from noise in the imaging equipment.

Approaching the subspace truncation problem in relation to image noise, we can always ignore any basis vectors that form statistically insignificant subspaces of k-space, hence increasing imaging efficiency. The SVD example chooses the smallest k s.t. $\Sigma(k,k) \geq \sigma_{\text{noise}}$. For typical MR images, this translates to speedups of up to 8 over Fourier encoding. An

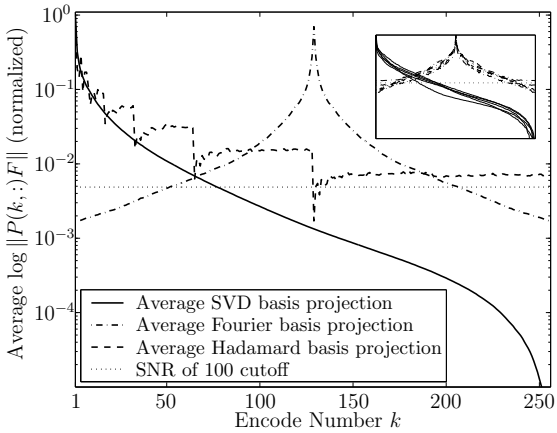

3.1: Log plot of average singular values and phase encode and Hadamard encode norms of 84 clinical MR images, and cutoff for SNR of 150.

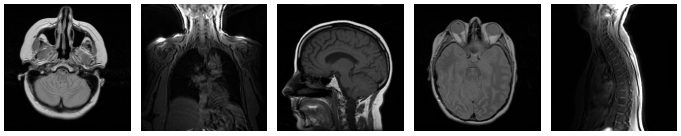

3.2: Five of the 84 images used to generate the plot in Fig. 3.1. Their individual projections appear in the inset plot in Fig. 3.1.

Fig. 3. Statistical significance of encoding basis projections with respect to noise is established via the SNR cutoff.

experimental result of truncated SVD MRI is shown in Fig. 4. Naturally, there are advantages to using fixed encoding bases that are oblivious to FOV content, such as the Fourier one. When acquisition of a complete basis for the column space of F is desired, oblivious bases do not require any computational overhead. Furthermore, in reduced basis acquisitions, the trade-offs can be more easily identified and quantified for oblivious bases. For example, a reduced set of Fourier basis functions results in lower spatial resolution and Gibbs ringing.

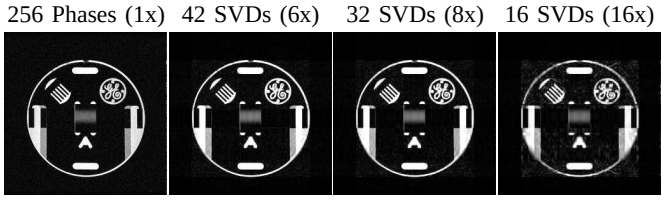

Fig. 4. Images acquired on a 1.5T MR scanner using full phase encoding (1x speedup) and spatially selective encoding using the truncated SVD basis at 6, 8 and 16x speedups (same parameters as in Fig. 1). The SNR of the phase encoded image approached the 40th singular value.

In contrast, the trade-offs of subspace projection in compressed non-Fourier imaging are not as easily assessed. Separating the time-varying F into an initial representation plus a time-dependent change, $F(t) = F(t_0) + \Delta F(t)$, when a truncated encoding basis described by the orthogonal matrix $U(t)$ is used to acquire the response $S(t)$ of the MR system's state at time t, the k-space estimate reconstructed is

$$\tilde{F}(t) = U(t)S(t) = \tilde{F}(t_0) + \left\{U(t)U^T(t)\right\}\Delta F(t). \quad (12)$$

Choosing the subspace and its dimensionality based on the SNR ensures than $\tilde{F}(t_0)$ contains all significant information that $F(t_0)$ does. However, changes are only acquired to the extent that they have a component in the subspace spanned by the given choice of encoding vectors. The principal angles analysis of [24] asserts that using an encoding basis computed to optimally encode the FOV contents at one time step to acquire an image at a subsequent time step is a well-defined proposition. Alternatively, the encoding subspace can be designed specifically to minimize such errors as the DATUM method [25] demonstrates. Applicability is further suggested by our experimental results in section V.

B. Spatial Encoding and SMASH Parallel Imaging

Parallel imaging methods use the intrinsic spatial encoding induced by sampling the signal using multiple coils with independent spatial sensitivities in order to achieve some signal localization that typically requires the use of gradients. However, the intrinsic sensitivity encoding of the coils is present regardless of how one encodes the signal content. By combining non-Fourier encoding with parallel imaging we show that the column space *dimensionality* of k-space that needs to be encoded via RF pulses can be reduced. Thus, in a parallel non-Fourier acquisition only the dimensions of the space that are not naturally encoded via the coils need to be encoded via RF excitation. A related advantage, not encountered in Fourier parallel imaging, is that the digital RF excitation length is reduced precisely because of this dimensionality reduction.

When the signal emanating from the spin distribution is weighted by the spatial sensitivity $C_l(y)$ of a coil, Eq. (1) becomes

$$S_l(k_x, k_y^{(m)}) = \iint \rho(x,y)C_l(y)e^{-i2\pi(k_y^{(m)}y + k_x x)}dxdy, \quad (13)$$

where we have simplified the spin density to a 2D function for brevity. In a typical SMASH experiment [7] a regular subsampling factor along the direction of phase encoding is enforced, yielding insufficient data to reconstruct an alias-free image. Skipping k-space samples reduces the FOV and the reconstructed image is aliased. When L coils with independent spatial sensitivities $C_l(y)$, $l = 1, \ldots, L$, are used to acquire the MR signal, the skipped k-space samples can be reconstructed by taking linear combinations of the raw coil samples of Eq. (13) that mimic the effect of the skipped gradient encoding steps.

Suppose, without loss of generality, the y direction is phase encoded. A 1D DFT of the raw data from each coil along the readout direction allows the use of coils with sensitivities that vary in both image plane axes [9]. Given the phase encoding gradient step $\Delta k_y = \gamma G_y^{\text{step}} \tau$ necessary to yield the desired FOV (from Eq. (4)) SMASH computes a set of scalars c_l^s, for $s = 0, \ldots, g-1$, where $g < L$, such that

$$\sum_{l=1}^{L} c_l^s C_l(y) = e^{-i2\pi(s\Delta k_y)y}. \quad (14)$$

The sum in Eq. (14) is precisely the complex exponential that separates harmonic (i.e. phase encode step) $k_y^{(m)}$ from $k_y^{(m+s)}$. Experimentally measured coil sensitivities for a four-element cardiac phased array coil and a least squares fit of

these particular sensitivities, yielding the 0th ($s = 0$) and 1st ($s = 1$) harmonic, are shown in Figs. 5 and 6 respectively.

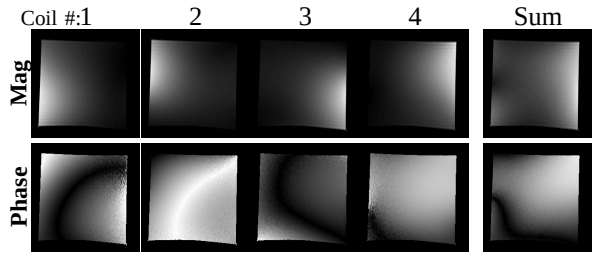

Fig. 5. Measured coil sensitivities (magnitude and phase) of a four-element cardiac phased array receiver. The complex-valued sensitivity profiles were measured by full acquisition of a relatively homogeneous doped water phantom (e.g. $\rho(x, y) \simeq$ const) placed in the FOV.

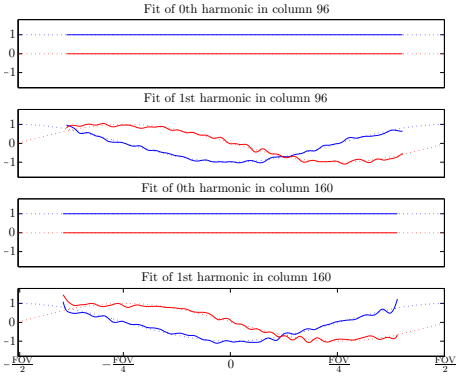

Fig. 6. SMASH fits for two harmonics, yielding a speedup factor of 2 in a SMASH-enabled acquisition experiment. True harmonics shown in dotted and fitted shown in solid (real components in blue, imaginary components in red). Harmonics were not fitted outside the region of support of the measured sensitivity profiles.

When a set of harmonics $\{k_y^{(m)}\}$ is induced via gradients, the complex exponentials generated by the linear combinations of the coil sensitivities in Eq. (14) can be used to obtain the remaining harmonics. In particular, the data acquired from the coils for the gradient-induced k-space step $k_y^{(m)}$, can be recombined via the c_l^s to generate the sample's $k_y^{(m)} + s\Delta k_y = k_y^{(m+s)}$ harmonic:

$$\sum_{l=1}^{L} c_l^s S_l(k_x, k_y^{(m)}) = \iint \rho(x,y) \left[\sum_{l=1}^{L} c_l^s C_l(y)\right] e^{-i2\pi k_y^{(m)} y} e^{-i2\pi k_x x} dxdy = \iint \rho(x,y) e^{-i2\pi k_y^{(m+s)} y} e^{-i2\pi k_x x} dxdy. \quad (15)$$

SMASH experiments skip g phase encode steps in each of n/g experiment repetitions. The signal obtained at the $m = 1, \ldots, n/g$ repetition of the experiment uses gradients to acquire the harmonic $k_y^{(g(m-1)+1)}$. Combining raw coil signals from each acquired harmonic $m_p = g(m-1)+1$ for each s generates the ensemble of harmonics $\{M\} = m_p + \{s\} = \{g(m-1)+1, \ldots, gm\}$. These recombined signals complete the full set of encodes $k_y^{(M)}$, $M = 1, \ldots, n$ that produce the desired FOV and alias-free image.

Now consider an experiment using non-Fourier excitation with the spatial profile $p(y)$ corresponding to the Fourier coefficients p_k that we excite with the RF pulse (as in Eq. (8)). Each coil produces the signal

$$S_l(k_x, p(y)) = \int\int \rho(x,y) p(y) C_l(y) e^{-i2\pi k_x x} dxdy. \quad (16)$$

Combining these signals according to the same c_l^s as in the Fourier SMASH experiment, similarly yields

$$\sum_{l=1}^{L} c_l^s S_l(k_x, p(y)) = \iint \rho(x,y) \left(\sum_m p_m e^{-i2\pi(k_{m+s})y}\right) e^{-i2\pi k_x x} dxdy, \quad (17)$$

after substitution of the Fourier coefficient formulation of $p(y)$, and where $k_m \equiv k_y^{(m)}$ is the harmonic induced by the remainder spatially selective gradient from the location of p_m in the RF pulse (as in Eq. (7)).

Equation (17) reveals that a scalar p_m at the mth position in the RF pulse can be used to produce a number of signal responses in each of which it scales a different set of harmonics (Fourier modes k_{m+s}) depending on which combination of raw signals is used. Effectively, we can think of a given RF encoding vector as capable of producing s *separate* responses, each one offset to begin the inner product in Eq. (9) at the sth row of k-space. Now, suppose that our RF vector is made to contain impulses at only every gth position. Then, using the appropriate linear combination of acquired signals, the entire RF vector $\vec{p}$ produces one linear response with the subsampled-by-g rows of k-space starting from the top row (when combined with c_l^0). It also produces another linear response, this time with the subsampled-by-g rows of k-space starting from the second row (when combined with c_l^1), and so on, for each offset row $1, \ldots, g$. It is therefore sufficient to choose the encoding vectors $\vec{p}$ so that they represent an optimal encoding of all the subsampled-by-g matrices of k-space *separately and simultaneously*. The physical analogy here is that a spatial excitation containing impulses only at every gth position is in itself aliased, and hence encodes each superimposed alias of the SMASH aliased image identically.

The *NF-SMASH* method we present here is a modified SMASH method that uses non-Fourier spatial encoding in conjunction with a post acquisition SMASH reconstruction. The original k-space matrix F of dimensions n by n is split into g matrices, each of size n/g by n, containing the subsampled-by-g rows of F starting from each of the first to the gth row. That is, let $F^f = F(s\!:\!g\!:\!n, :)$ (in MATLAB notation) for $f = 1, \ldots, g$. We then concatenate these matrices side by side, i.e.

$$F^{\mathcal{PSE}} = [F^1, \ldots, F^g], \quad (18)$$

where the superscript $\mathcal{PSE}$ stands for parallel spatially encoded. It is this matrix that we need to non-Fourier encode; for the SVD example we compute $F^{\mathcal{PSE}} = U\Sigma V^*$. The gn columns of $F^{\mathcal{PSE}}$ live in n/g dimensional space which can be fully encoded using at most n/g encoding functions. The matrix U has n/g columns corresponding to non-zero singular values and each column only contains n/g entries, reflecting the dimensionality reduction of k-space. The prospect of shorter digital RF excitation pulses in NF-SMASH is clear

given their n/g length. Only the k-space distance between the Fourier modes excited by each hard pulse in these short RF pulses needs to be increased by g. In some cases, this can be simply accomplished by increasing the strength of the gradient G_y by g, or increasing the dwell interval Δt of the impulses by g, or any combination of the two.

NF-SMASH uses knowledge of the parallel reconstruction in order to shift the encoding burden from the spatial RF encoding (P) to the acquisition portion (the matrix S in Eq. (10)). In the SVD example this manifests in that the matrix ΣV^* has gn columns. Image reconstruction given the acquired short-wide matrix S is straightforward; the non-Fourier encoding is inverted via Eq. (11), yielding the concatenated subsampled k-space matrices, which, after reshaping, produce the full k-space and subsequently image. In order to produce S given the raw coil responses to each shaped encoding function, the SMASH combinations are taken: responses to each encoding function received by the coils are combined according to the c_l^s for each s. Each combination yields the sth length-n portion of the acquisition matrix S (e.g. the ΣV^* estimate for SVD).

The same configuration of the standard four-element cardiac phased array receiver with the sensitivities shown in Fig. 5 was used to achieve a SMASH speedup of two. The Fourier SMASH experiment produced the reconstructed image shown in Fig. 7.1. A row by row reconstruction was performed using Eq. (15) by fitting the spatial sensitivities at that particular x_0 to pre-conditioned harmonics [27]. Example fits for two locations x_0 are shown in Fig. 6. Following the NF-SMASH

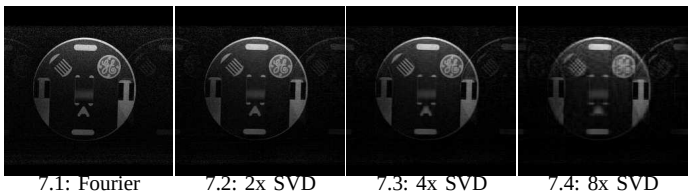

Fig. 7. Experimental NF-SMASH results reconstructed using a four-element array coil and 2x SMASH speedup. The horizontal dimension is parallel coil and non-Fourier encoded.

algorithm, the matrix $F^{\mathcal{PSE}}$ was decomposed via the SVD and the resulting $n/2$-dimensional vectors were used to acquire and reconstruct the sample. These vectors represented a two-fold aliased spatial excitation profile crafted so as to simultaneously best encode each of the two image aliases that are superimposed by the SMASH downsampling. The reconstructions at various speedup factors attained by truncating the acquisition basis dimensionality are shown in Fig. 7.

The NF-SMASH algorithm is a simple method that combines these two spatial encoding approaches for the particular case of SMASH reconstruction. Comparison of the parallel imaging results with non-Fourier and the standard Fourier encoding approaches demonstrates that the incorporation of parallel and non-Fourier spatial encoding approaches not only allows scan time reduction, but also reduces the burden of excitation of the digital encoding RF pulses.

C. Increasing Spatial Encoding Throughput via Multi-Echo "Fuzzy" Pulse Sequences

Through application of linear algebraic techniques, an temporal efficiency increase of almost an order of magnitude in acquisition time may be obtained in 2D MRI compared to phase encoding. The method presented in the previous section allows this efficiency increase even when employing parallel imaging techniques. Many practical Fourier encoding techniques obtain more than one row of k-space per repetition interval thereby increasing temporal resolution. It is desirable to have the same functionality available in the non-Fourier encoding toolbox. In this section we show that it is possible to acquire more than one system response per repetition in a non-Fourier spatially encoded MRI experiment as well [28].

First, any longitudinal magnetization remaining after a low flip excitation can be used by subsequent excitations within the same TR. Second, immediately following excitation, the emanating signal decays due to T_2^* dephasing, which is partially reversed by the refocuser in a spin echo sequence. This refocusing pulse precisely pinpoints when in time transverse magnetization will be in phase again and thus produce the spin echo. This can be used to multiplex multiple spatial RF excitations in a single TR by allowing sufficient time in between encoding RF pulses so that at the moment when magnetization due to one RF pulse refocuses, magnetization due to all other RF pulses is adequately dephased. Lastly, spatially selective profiles are composed by an ensemble of Fourier modes "excited" by the hard pulses. These modes in turn depend on the rephasing of the spatially selective gradients. Unbalanced gradients can selectively force certain modes to correspond to low energy outskirts of k-space, and can therefore be used as another form of "demultiplexing."

Combining these facts allows us to use multiple spatial excitations in a single experiment repetition to acquire and reconstruct images that closely resemble those obtained with a single excitation and subsequent response sampling per repetition. Consider a hard pulse train $p_{j,m}$ of $j = 1, \ldots, k$ RF pulses, $\vec{p}^{(j)}$, each composed of $m = 1, \ldots, n$ constituent hard pulses and excited serially as shown in Fig. 8. Transverse

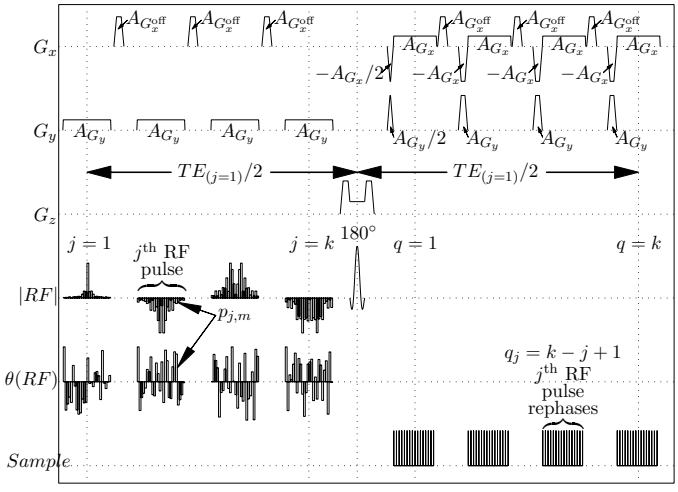

Fig. 8. A multiple excitation, multiple echo 1D-spatial/1D-Fourier encoded spin echo pulse sequence.

magnetization resulting from the constituent hard pulses of the jth RF pulse rephases precisely during the sampling interval $q_j = k - j + 1$. At that interval we wish to receive a response from the transverse magnetization due to $\vec{p}^{(j)}$ and ensure that no significant response from transverse magnetization due other RF pulses ($\vec{p}^{(l)}$, $l \neq j$) is received. We accomplish this by relying both on dephasing, as well as unbalanced gradients to force unwanted magnetization to the edges of k-space so that their contribution is indistinguishable from noise.

For each sampling interval q_j, transverse magnetization due to the hard pulses in RF pulses other than j (i.e. $p_{l,m}, l \neq j$) are dephased proportionally to the time difference between the time of their excitation to the refocuser and the time of the refocuser to the sampling interval. The amount of dephasing of the signal due to transverse magnetization from pulse j at some sampling interval q' is proportional to the time distance between q' and q_j, $d = q_j - q' = k - j + 1 - q'$, which we call "rephasing to sampling" distance. The time t_d, equal to d times the echo spacing, is an indicator of the suppression of the signal due to $\vec{p}^{(j)}$ at sampling interval q' . At q_j, $d = 0$ for the jth RF pulse, so that this pulse is rephased producing most signal, while $d \neq 0$ for all other RF pulses. Since T_2^* and T_2 are intrinsic to the sample being imaged and can not be controlled, additional suppression would necessitate increasing the echo spacing. However, because of T_2 relaxation, doing so is impractical, and we thus turn to the use of gradient spoilers to provide additional suppression.

With insufficient dephasing, at any sampling interval $q_j = k - j + 1$, some of the RF pulses $l \neq j$ still account for some of the coherent magnetization. Setting dephasing aside, equation (19) describes the signal sampled at location r of the qth sampling interval, that is, at time $t = t_q + r\Delta t$, where t_q is the time corresponding to the beginning of that sampling interval:

$$S_{\mathcal{MSE}}(q, r, \{\vec{p}^{(1)}, \ldots, \vec{p}^{(k)}\}, z_0) = \int_{z_0-\Delta z}^{z_0+\Delta z} \iint \rho(x, y, z) \times \sum_{j=1}^{k} \sum_{m=1}^{n} p_{j,m} e^{-i2\pi k_y(j,m,q)y} e^{-2\pi k_x(j,r,q)x} dx dy, \tag{19}$$

where the Fourier modes $k_y(j, m, q)$ and $k_x(j, r, q)$ can be written, using d, as:

$$k_y(j, m, q) = \gamma G_y(1/2n - m)\Delta t + \gamma G_y dn\Delta t, \tag{20}$$

$$k_x(j, r, q) = \gamma G_x(r - 1/2n)\Delta t + \gamma G_x^{\text{off}} d\tau_{\text{off}}. \tag{21}$$

At the sampling interval in which $\vec{p}^{(j)}$ rephases ($d = 0$), the signal obtained contains the correct sum of modes for its constituent hard pulses. It also contains a sum of modes from the hard pulses of the other RF pulses, but they are offset by $\gamma G_y dn\Delta t$ (i.e. d times the desired k-space extent) along k_y and by $\gamma G_x^{\text{off}} d\tau_{\text{off}}$ along k_x. Since $k_{x,y}$ define the wavelengths of the complex exponentials in the integrals of Eq. (19), these wavelengths can be made small enough so that the ensemble of Fourier modes produced by the unwanted RF pulses, i.e. $d \neq 0$, represent information about oscillations of the sample's properties on insignificantly small spatial scales, making them indistinguishable from thermal noise oscillations. We found no need for experimental G_x^{off} offset gradients. The natural offset added in k_y due to the excitatory gradients appears enough to crush any undephased components for typical SNRs. This is perhaps expected given the low energy in the higher spatial frequencies found in Fig. 3.1.

Experimental results indicate that the natural refocusing and excitation gradient spoiler mechanisms alone provide sufficient demultiplexing of the responses. Images reconstructed using the identity matrix for RF encoding, producing Fourier basis imaging, shown in Fig. 9, demonstrate that image quality is not compromised when multiple excitations are multiplexed in each TR, with the exception of gradual loss of SNR as expected. Similarly, a broadband experiment using SVD

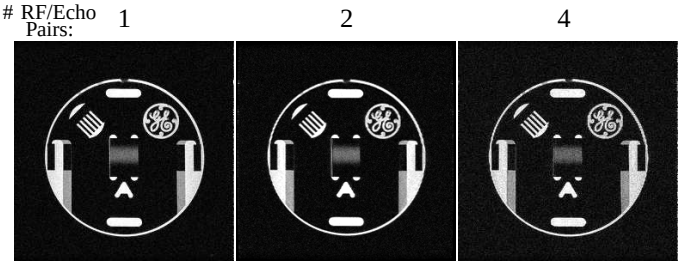

Fig. 9. Images acquired using the PSD of Fig. 8 for variable number of encoding RF pulses per repetition. The rows of the identity matrix were used for the RF pulses. Imaging parameters where Eff.TE $= 40$ms, TR $= 400$ms, 16cm FOV, 10mm Slice Thickness and 15° flip angle.

encoding demonstrates large scan time reduction without significant loss of image quality. The rightmost image in Fig. 10 was acquired in only four TRs. We note that the effective

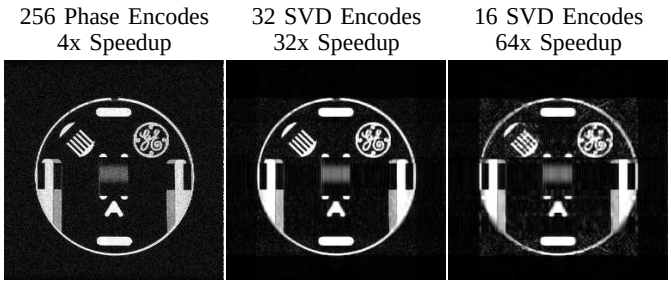

Fig. 10. Images acquired using the multi-echo sequence of Fig. 8 with 4 RF/Echo pairs per TR. Full phase encoding via the identity matrix (left) and non-Fourier encoding using the truncated SVD basis (middle and right). Same imaging parameters as Fig. 9.

echo time in this multi-echo acquisition is controlled via the placement of the lowest spatial frequency encoding pulses in the RF pulse train. The term "fuzzy" is chosen since, even in the absence of noise, the sequence does not produce the exact linear response expected from the Fourier coefficients of the sample that we are trying to reconstruct, but rather, it includes contributions from Fourier coefficients outside the spatial frequency region of interest.

IV. A Dynamic Adaptive Non-Fourier MRI Platform

This section describes a hardware and software platform capable of performing real-time adaptive non-Fourier spatial

encoding in concert with a commercial MR imager. In the next section we demonstrate that this currently operational system can be used in dynamic imaging cases where it is necessary to observe a sample at high temporal as well as spatial resolution.

The system we have developed [29] employs non-Fourier encoding methods to acquire images in a dynamic adaptive cycle. Each iteration of the imaging cycle produces a new, high spatial resolution, MR image. At a high level, the steps that are performed at each iteration are shown in Fig. 11. This

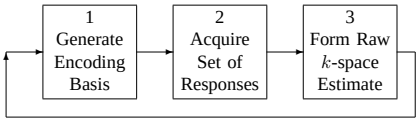

Fig. 11. High-level view of a dynamic adaptive non-Fourier encoded imaging cycle.

simple high level imaging cycle is capable of capturing many non-Fourier encoding methods, thus the system is designed to be easily extensible and/or modifiable in order to handle new methods, whether or not they adapt to the sample over time.

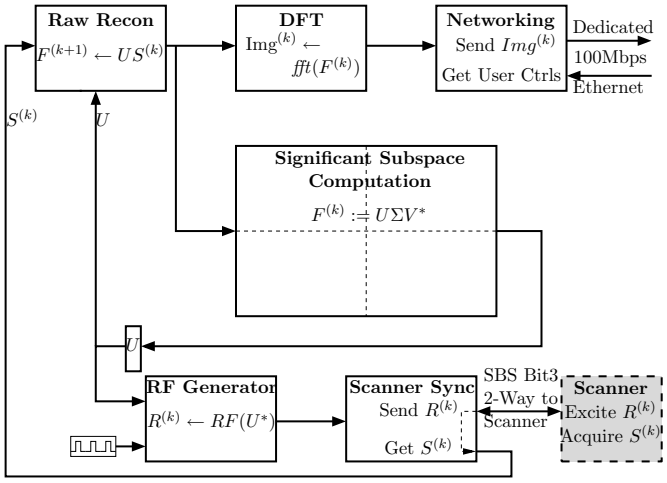

Fig. 12. The dynamic computation pipeline ring. Each solid rectangle represents a pipeline task and is executed on separate processors. Computation modules encapsulate separate tasks, as indicated in boldface and may have multiple pipeline stages or execute on multiple processors. The imaging pipeline interfaces via a dedicated Ethernet connection to the display system and via an SBS 413-1 VMEBus Bit3 adapter to the MR scanner.

A. System Implementation – Design

Accordingly, a modular pipeline platform was designed to implement the imaging cycle of Fig. 11. The pipeline ring is depicted in Fig. 12, and is *asynchronous* and *locally timed* (also called self-timed). Each logical computation task, or module, performs its task on the input provided to it and, once completed (asynchrony), it distributes the result to all consumer units along with a notification (local timing). Pipelining the imaging cycle allows overlapping the execution of different imaging cycles. Provided sufficient computing resources all stages of the pipeline execute in parallel. For example, while the samples received from one acquisition are being used to form the current k-space estimate, that of the previous acquisition is being Fourier transformed to produce the image, while the image from two acquisitions ago is being sent to the display. Parallel execution of the tasks allows us to maximize throughput for every dynamic experiment, yielding the maximal temporal resolution that can be achieved with the underlying computing hardware or as limited by operator-prescribed acquisition parameters such as TR, number of excitations per TR, subspace dimensionality etc. The latter may be dictated by desired image characteristics such as contrast and SNR.

The pipeline platform design offers a number of advantages to adaptive spatial encoding, related to extensibility and performance. Maximum throughput is maintained regardless of which pipeline stage is the slowest. Although in some cases this is limited by operator-prescribed scan acquisition parameters, in other cases, the slowest stage may be a novel encoding or reconstruction method that requires additional computing, such as spiral acquisitions for 3D imaging or DATUM encoding. In such cases, the pipeline design allows easy replacement of existing modules with improved software module implementations or with implementations that use additional computing hardware. Furthermore, new encoding methods that require additional encoding or reconstruction tasks can likewise be accommodated, while replacement modules with additional pipeline depth allow some coarse-grained control of throughput. Finally new encoding methods that may require additional information, for example, receiver coil sensitivity calibration for parallel imaging experiments, can be easily implemented. New pipeline modules that produce and distribute information as needed, can be inserted into the pipeline and the pipeline can be rewired with minimal programmer effort.

Our current implementation of the pipeline, shown in Fig. 12, is targeted to non-Fourier encoding methods such as Hadamard, SVD and DATUM. To present this implementation, we analyze the steps involved in completing one imaging cycle. The RF Generator and Scanner Synchronization modules handle the generic tasks of transforming any set of basis vectors to excitations and acquiring the sample responses from the MR scanner. Given an encoding basis, the RF Generator produces RF pulses of the desired flip angle that conform to the scanner's waveform generators' data format. In the low flip angle regime this requires computing vector magnitude and phase in the k-space domain. Once formed, RF pulses are sent to the scanner's RF waveform generators via the Synchronization module. This module uploads the RF pulses to the scanner and instructs the pulse sequence active on the scanner to apply the shaped RF waveforms in order to produce the spatial profiles of excited magnetization (Eq. (8)). Apart from the MR scanner hardware (e.g. magnet, coils, receivers etc) the Scanner module encapsulates the non-Fourier pulse sequence. The sequence is responsible for acquiring the sample response for each RF pulse. Like any sequence provided with the scanner, the non-Fourier sequence enters a real time scanning loop for an appropriate number of repetitions, depending on, for example, whether it is the multi-echo sequence of the previous section. The MR scanner places each sampled result (as requested by the pulse sequence) in its raw data memory. The Scanner Synchronization module then

downloads these acquired responses from the scanner's raw data memory directly. Currently, the set of responses (without any post-processing such as filtering) are combined with the basis vectors (Eq. (11)) by the Raw Reconstruction module to produce the current k-space estimate of the FOV contents. Non-Fourier Keyhole [30], [31] and DATUM reconstructions can also be performed by this module. The k-space estimate is subsequently Fourier transformed in the DFT module to produce the image which is sent to the display and archiver process by the Networking module. In order to initiate the next imaging cycle, the Raw Reconstruction module also distributes the current k-space estimate to the Significant Subspace Computation module which, via an operator-selected rank-revealing linear algebraic decomposition, produces a new near-optimal encoding basis. For non-oblivious basis encoding methods, this final step is almost always necessary due to small arbitrary rotations of the most significant subspaces stemming both from noise in the measurements and excited basis profiles as well as new data from changes in the FOV captured by those noise components. The first imaging cycle begins by acquiring a full estimate of the current contents of the FOV, usually by encoding with an orthogonal matrix of full rank, such as the identity (i.e. producing phase encoded responses) or Hadamard matrix. The Significant Subspace module currently executes on up to four processors (as shown by the dashed boxes in Fig. 12) in order to minimize adaptivity latency.

In order for the pipeline to fill, one set of response data must be available at the frequency of the slowest pipeline stage. This amounts to beginning acquisition of the next set of responses as soon as the acquisition of the current set of responses is completed. The RF Generator module therefore differs from other modules in that it starts computations for the next imaging cycle as soon as it has completed the current imaging cycle rather than waiting for notification (represented by the clock wave in Fig. 12). The RF Generator module drives the pipeline ring at a period as low as 25-50ms, using the latest encoding basis for multiple acquisitions if that is necessary.

B. System Implementation – Hardware

The computation pipeline has been implemented on a Mercury multi-processor computer (Mercury Computer Systems Inc, Chelmsford MA) connected to a GE 1.5T Signa LX EchoSpeed MR scanner (GE Medical Systems, Milwaukee, WI). The three-generation old VME bus-based multi-processor computer contains a number of boards: first, a FORCE dual-slot board (Force Computers Inc, Fremont, CA) based on a μSparc processor (sun4m architecture, 110MHz, 32MB RAM). This board provides a network interface to the display and control unit via a fast Ethernet controller. Second, two Mercury MCH6 boards each holding 4 PowerPC 750 processors (each 375MHz, 32MB RAM). Each module in the pipeline executes on one or more of these processors, as shown in Fig. 12, except for the Networking module, which executes on the μSparc, and the Scanner module, which operates on the scanner hardware. Each board contains a fast proprietary interconnect (called ILK) which routes data in between the processors at a high rate. Finally, an SBS 413-1 adapter board (SBS Technologies Inc., Carlsbad, CA) is used to interface the multi-processor with the scanner's computing hardware. It provides access to the raw memory board and waveform memory of the scanner at sustained DMA rates of up to 26MB/sec.

The choice of a VME bus architecture for the external system is supported by the VME bus architecture of our GE scanner's data acquisition system. It is possible, if desired, to insert the boards used by our external system directly into the scanner thereby completely integrating the dynamic system with the scanner. With few exceptions (Scanner Synchronization module, waveform format and pulse sequence implementation), the design is portable to any other MR scanner. The scanner-specific modules may be ported so long as the scanner manufacturer allows access to their hardware and provides sufficient information for programming it.

C. System Performance

Most modules in the pipeline are optimized to a moderate degree in order to reduce latency. Typical module execution times are given in Table I. A number of methods for subspace truncation are currently implemented in order to accommodate different adaptive time-resolution requirements and provide some versatility in encoding choice. Although many different algorithms may be used to identify the encoding subspace, our method of choice is a gradient descent (CG) algorithm that searches for the vector associated with the highest singular value by maximizing the Rayleigh quotient [32]. The matrix is iteratively deflated before searching for the next largest singular vector. Computing time is reduced using *warm restarts* to speed up subsequent decompositions. That is, singular vectors computed for the previous cycle are used to re-initiate calls to the algorithm. The parallel implementation of the algorithm uses 4 processors to achieve a moderate speedup of 2, partially due to the small size of the matrix in 2D imaging (typically 256 by 256) which makes communications overhead a significant portion of running time.

Other subspace identification methods currently offered include a randomized approximate version of the SVD [33] and a spatial-frequency ordered Krylov subspace method [34]. These methods represent trade-offs between subspace accuracy and computational cost. The main trade-off for MR imaging purposes is the number of signal responses, or TRs, necessary to encode all important image information. We have found that using such methods one needs to retain a much larger subspace to keep the imaging error small relative to noise (e.g. typically about 80 encoding vectors are needed in comparison to 32 of the singular vectors). Reducing the approximation error of these methods, e.g. using more columns of the matrix in the approximate SVD computation or using more iterations of Lanczos tridiagonalization, the computational expense reaches that of CG SVD, making their use impractical at this time.

V. Dynamic Imaging Results

In this section we present some results obtained with the dynamic system. Via doped water and animal tissue phantom experiments we try to cover a range of cases that may be

Raw Recon	FFT	Network	Rank-32 Subspace Computation						RF Generate	Scanner Sync
			L	R	C	CW	CP	F		
50ms	25ms	55ms	600ms	225ms	3s	1.2s	600ms	13.6s	25ms	75ms

TABLE I

MODULE TIMES FOR ONE IMAGING CYCLE. SUBSPACE METHODS ARE: LANCZOS (L), RANDOMIZED SVD (R), CONJUGATE GRADIENT SVD (C), CONJUGATE GRADIENT SVD WITH WARM RESTARTS (CW), CW PARALLELIZED ON 4 PROCESSORS (CP) AND FULL SVD (F).

encountered in a dynamic MR imaging session. In all images presented, the direction of the spatially selective excitation is the horizontal, while the vertical is frequency encoded. The warm restart CG SVD algorithm was used for subspace identification.

First Experiment: In this study, a small syringe filled with contrast is continuously moved around a large static bottle filled with doped water, to assess adaptation of encoding to movement in the FOV. Imaging parameters for this sequence were 56ms Eff.TE, 115ms TR, 256×256 image matrix, 24cm FOV and the multiple echo spin echo pulse sequence described in Section III-C was used, with 4 excitation/echo pairs. The 256-dimensional subspace was truncated to the 40 primary SVD vectors. This leads to a total acquisition time of 1.15s per image, compared to 29.4s for a single echo and 7.4s for a four echo Fourier encoded pulse sequence. A small portion of this series is shown in Fig. 13. Notably, the smaller object is in constant motion, while responses that will be used for a single image reconstruction are being acquired. The ability to image this motion properly is inherent because a single image requires only 10 TRs to be formed. Furthermore, the first k SVD encodes always account for the greatest percentage of energy in k-space and encode the lowest spatial frequencies, which determine salient features [18].

Second Experiment: In an experiment to assess the ability to capture changes within an object during dynamic imaging, a 22G biopsy needle was continuously inserted into and subsequently removed from an animal tissue phantom. Part of the removal series is presented in Fig. 14. Imaging parameters were 20ms TE, 180 ms TR, 256×256 matrix, 24cm FOV, 32 primary singular vectors yielding 5.8s per image, a speedup of 8 over the equivalent single-echo Fourier acquisition. Spatial resolution is not sacrificed in non-Fourier imaging as it is with reduced phase encoded Fourier MRI (e.g. low resolution scans). In the results shown in Fig. 15, the surface of the tissue does indeed appear to be pulled upward as the needle is being removed. This surface tension detail is captured even at speedup of 8 over Fourier imaging. The equivalent 32-coefficient Fourier acquisition would possess a 7.5mm pixel size, versus 0.9375mm in this non-Fourier acquisition, which would have undoubtedly caused loss of this detail.

Third Experiment: Since SVD imaging becomes more problematic when large changes in the FOV yield large changes in the basis vectors, the adaptivity of such cases is demonstrated in Fig. 16 where the animal tissue phantom was deformed. A human hand entered the FOV and applied pressure to the phantom. Same imaging parameters were used as in the last experiment. The 32-dimensional subspace contains most of the useful information in the image and manages to capture those changes in the FOV, as demonstrated by the results. Although the individual maximal response basis vectors may change significantly between two acquisitions, the subspace they form may remain fairly constant.

VI. CONCLUSIONS

Although non-Fourier encoding of the MR signal has been presented [23], [15], [18], used [35], [36], analyzed [37], [24], and compared to reduced encoding Fourier imaging (such as RIGR and Keyhole) [25], it has not come to routine use. This is perhaps in part attributed to the lack of non-Fourier MR imaging technologies that are competitive to present-day Fourier MRI. This paper attempts to produce some technologies that are necessary to obtain fast non-Fourier MRI.

The conclusion of this work is that accelerated fast MRI via non-Fourier signal encoding is a viable solution for dynamic MRI and that non-Fourier imaging is able to offer many of the advantages that are used, and are necessary, in the fast Fourier encoding arena. The prototype multi-echo pulse sequence developed in section III-C shows that non-Fourier encoding can achieve MR signal compression efficiency benefits even when compared to fast multi-echo Fourier imaging. The theory of non-Fourier SMASH parallel imaging developed in section III-B shows that further imaging efficiency can be obtained by combining MR signal compression with parallel imaging methods. Finally, the implementation of a complete, currently operational system, described in section IV, brings together some of these components, enabling non-Fourier encoding to be an alternative to a dynamic imaging session in a commercial clinical MR scanner. Yet, many questions remain unanswered, and many more technologies must be developed, as has been done for Fourier imaging.

We summarize some of the issues we believe are central to the future of broadband compressed MR imaging. First, fast, optimal pulse sequences must be developed. On this issue we can note that stimulated and indirect echoes [38] do not immediately appear to be useful in non-Fourier encoding. These echoes contain encoding information that can not be easily removed, for example, via gradient reversal in Fourier MRI. Their usefulness may lie in new encoding methods that are yet to be discovered. For the time being, magnetization from previous encoding steps must be spoiled. On the other hand, implementation of fast spoiled gradient recalled sequences is certainly well within reach, as is implementation of fast, low flip angle, short-TR spin echo sequences. These pose a further complication related to the effect of the refocusing RF pulse on longitudinal magnetization, which can be overcome by applying a second refocusing pulse as soon as the sampling interval is completed [2], [39].

Second, the issue of compression must be more closely studied in the context of adaptivity to a changing FOV. Implementation of useful fixed basis imaging, such as Wavelet encoding to e.g. obtain high resolution in a given region of interest, is now possible with the pipeline system we have

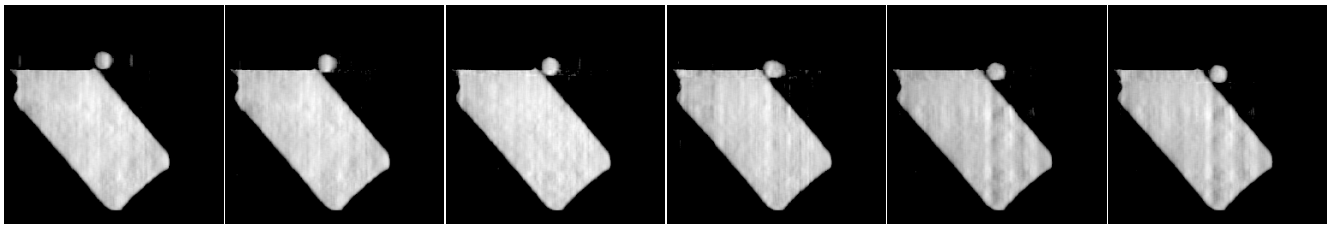

Fig. 13. Syringe filled with contrast agent, moving around a bottle filled with doped water. Time flows from left to right. Images were acquired at a rate of 1.2sec per image, a speedup of 6.4 over an equivalent 4 echo train length Fourier encoded acquisition, and 25.6 speedup over a single echo Fourier encoded acquisition. Note that the syringe was moving *constantly* throughout the experiment, therefore contaminating the acquired responses with inter-response motion artifact, as in any Fourier encoded image.

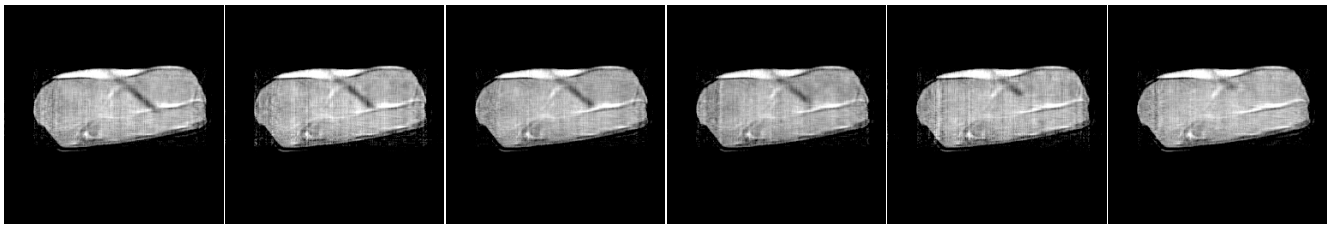

Fig. 14. A 22G biopsy needle removed from animal tissue phantom. Imaging at a speedup of 8 over equivalent single-echo Fourier encoded imaging.

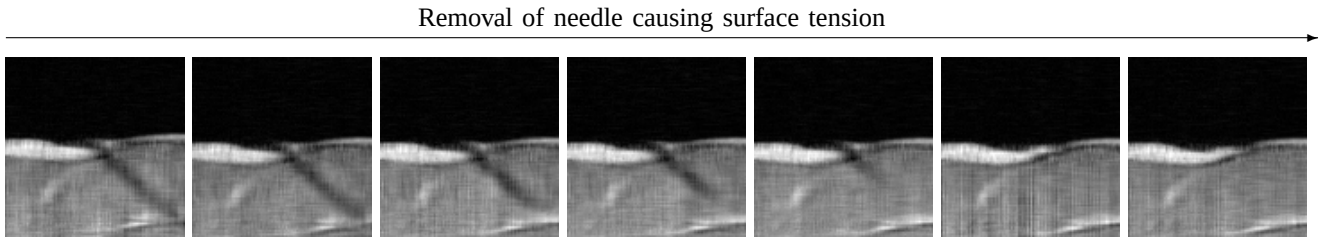

Fig. 15. Animal tissue phantom detail captured by truncated SVD encoding at speedup of 8. Close inspection shows that the tissue appears to move upward at the exit point of the needle.

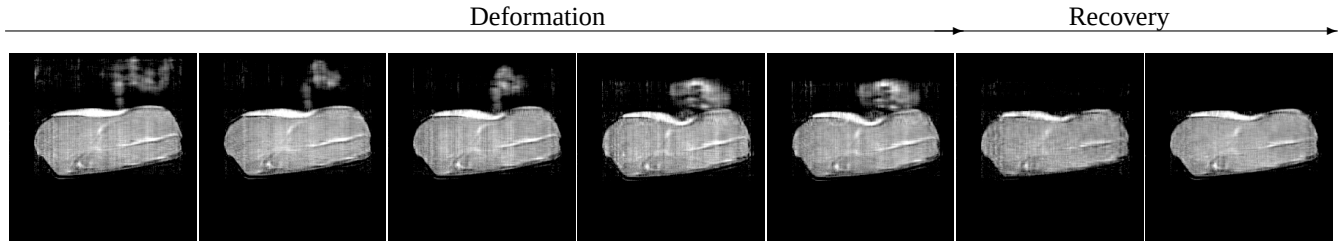

Fig. 16. Animal tissue phantom being deformed. Same imaging parameters as in Fig. 14

developed. However, all adaptive signal compression methods pose the risk of improperly encoding important changes in the FOV. On this subject, we note that it is possible to obtain an estimate of the component of $\Delta F(t)$ that lies in the subspace complementary to the one acquired by switching the readout and non-Fourier encoded directions. One can obtain Fe, with e a column of the identity matrix. The complementary subspace projection for that column of $\Delta F(t)$ can be used to either restart with a full space acquisition, or update the acquisition subspace. Furthermore, navigator acquisitions can be used much like in Fourier imaging in order to assess and correct for sample movement in the spatial encoding profiles. Finally, new methods, in the direction of the DATUM approach [25], that continually adapt the subspace, must be developed.

Last, but not least, non-Fourier encoding research is necessarily intimately linked to the low level details of MR scanners. These details concerning RF subsystems, acquisition and reconstruction technologies etc, are known to scanner manufacturers and can greatly enhance or hinder non-Fourier imaging quality. In this sense, our experimental results are in essence "worst-case" imaging, due to the very limited knowledge of the implementation of proprietary MR scanners available to us.

In closing, we hope that this research serves two purposes: Firstly, to re-introduce the ideas and theory of non-Fourier encoding, while at the same time extending it to compete with present-day fast MRI techniques. Secondly, we hope this work will instigate further research on this exciting and promising subject.

REFERENCES

[1] A. Kumar, D. Welti, and R. R. Ernst, "NMR Fourier Zeugmatography," *Journal of Magnetic Resonance*, vol. 18, pp. 69–83, 1975.

[2] F. W. Wehrli, "Fast-Scan Magnetic Resonance: Principles and Applications," *Magnetic Resonance Quarterly*, vol. 6, pp. 165–236, 1990.

[3] P. Masnfield, "Real-Time Echo-Planar Imaging by NMR," *British Medical Bulletin*, vol. 40, no. 2, pp. 187–190, 1984.

[4] G. H. Glover, "Simple Analytic Spiral K-Space Algorithm," *Magnetic Resonance in Medicine*, vol. 42, pp. 412–415, 1999.

[5] J. Hennig, A. Naureth, and H. Friedburg, "RARE Imaging: A Fast Imaging Method for Clinical MR," *Magnetic Resonance in Medicine*, vol. 3, pp. 823–833, 1986.

[6] J. Hennig, "Echoes – How to Generate, Recognize, Use or Avoid Them in MR-Imaging Sequences
Part II: Echoes in Imaging Sequences," *Concepts in Magnetic Resonance*, vol. 3, pp. 179–192, 1991.

[7] D. K. Sodickson and W. J. Manning, "Simultaneous Acquisition of Spatial Harmonics (SMASH): Fast Imaging with Radiofrequency Coil Arrays," *Magnetic Resonance in Medicine*, vol. 38, pp. 591–603, 1997.

[8] K. P. Pruessman, M. Weiger, M. B. Scheidegger, and P. Boesiger, "SENSE: Sensitivity Encoding for Fast MRI," *Magnetic Resonance in Medicine*, vol. 42, pp. 952–962, 1999.

[9] W. E. Kyriakos, L. P. Panych, S. F. Kacher, C.-F. Westin, S. M. Bao, R. V. Mulkern, and F. A. Jolesz, "Sensitivity Profiles from and Aray of Coils for Encoding and Reconstruction in Parallel (SPACERIP)," *Magnetic Resonance in Medicine*, vol. 44, pp. 301–308, 2000.

[10] D. A. Feinberg, J. C. Hoenninger, L. E. Crooks, L. Kaufman, J. C. Watts, and M. Arakawa, "Inner Volume MR Imaging: Technical Concepts And Their Application," *Radiology*, vol. 156, pp. 743–747, 1985.

[11] X. Hu and T. Parrish, "Reduction of the Field-of-View for Dynamic Imaging," *Magnetic Resonance in Medicine*, vol. 31, pp. 691–694, 1994.

[12] B. Madore, G. H. Glover, and N. J. Pelc, "Unaliasing by Fourier-Encoding the Overlaps Using the Temporal Dimension (UNFOLD), Applied to Cardiac Imaging and fMRI," *Magnetic Resonance in Medicine*, vol. 42, pp. 813–828, 1999.

[13] Z.-P. Liang and P. C. Lauterbur, "An Efficient Method for Dynamic Magnetic Resonance Imaging," *IEEE Transactions on Medical Imaging*, vol. 13, pp. 677–686, 1994.

[14] Y. Cao and D. N. Levin, "Feature-Recognizing MRI," *Magnetic Resonance in Medicine*, vol. 30, pp. 305–317, 1993.

[15] J. B. Weaver, D. M. Healy, and J. R. Driscoll, "Wavelet Encoded MR Imaging," *Magnetic Resonance Medicine*, vol. 24, pp. 275–287, 1992.

[16] D. M. Healy and J. B. Weaver, "Joint Best Basis for Fast Encoding in Magnetic Resonance Imaging," *ICASSP-95, International Conference on Acoustics, Speech, and Signal Processing*, vol. 5, pp. 2923–2926, 1995.

[17] J. Pauly, D. Nishimura, and A. Macovski, "A k-space Analysis of Small Tip Angle Excitation," *Journal of Magentic Resonance*, vol. 81, pp. 43–56, 1989.

[18] G. P. Zientara, L. P. Panych, and F. A. Jolesz, "Near-Optimal Encoding for Dynamically Adaptive MRI: Mathematical Principles and Computational Methods," *International Journal of Imaging Systems and Technology*, vol. 10, pp. 151–165, 1999.

[19] L. P. Panych, G. P. Zientara, and F. A. Jolesz, "MR Image Encoding by Spatially Selective RF Excitation: An Analysis Using Linear Response Models," *International Journal of Imaging Systems and Technology*, vol. 10, pp. 143–150, 1999.

[20] W. S. Hinshaw and A. H. Lent, "An Introduction to NMR Imaging: From the Bloch Equation to the Imaging Equation," in *Proceedings of the IEEE*, vol. 71, no. 3, Mar. 1983, pp. 338–350.

[21] G. P. Zientara, L. P. Panych, and F. A. Jolesz, "Dynamically Adaptive MRI with Encoding by Singular Value Decomposition," *Magnetic Resonance in Medicine*, vol. 32, pp. 268–274, 1994.

[22] L. P. Panych, C. Oesterle, G. P. Zientara, and J. Hennig, "Implementation of a Fast Gradient-Echo SVD Encoding Technique for Dynamic Imaging," *Magnetic Resonance in Medicine*, vol. 35, pp. 554–562, 1996.

[23] C. Oh, H. W. Park, and Z. H. Cho, "Line-integral projection reconstruction (LPR) with slice encoding techniques: Multislice regional imaging in NMR tomography," *IEEE Trans Med Imag*, vol. MI-3, pp. 170–178, 1984.

[24] G. P. Zientara, L. P. Panych, and F. A. Jolesz, "Applicability and Efficiency of Near-Optimal Spatial Encoding for Dynamically Adaptive MRI," *Magnetic Resonance in Medicine*, vol. 39, pp. 204–213, 1998.

[25] W. S. Hoge, E. L. Miller, H. Lev-Ari, D. H. Brooks, and L. P. Panych, "A Doubly Adaptive Approach to Dynamic MRI Sequence Estimation," *IEEE Transactions Image Processing*, vol. 11, no. 11, pp. 1168–1178, Oct. 2002.

[26] G. A. Wright, "Magnetic Resonance Imaging," *IEEE Signal Processing Magazine*, vol. 14, pp. 56–66, 1997.

[27] D. K. Sodickson, "Tailored SMASH Image Reconstructions for Robust In Vivo Parallel MR Imaging," *Magnetic Resonance in Medicine*, vol. 44, pp. 243–251, 2000.

[28] D. Mitsouras, A. S. Edelman, L. P. Panych, F. A. Jolesz, and G. P. Zientara, "Fast Multiple-Excitation Multiple-Echo Spin-Echo Pulse Sequence for Non-Fourier Spatial Encoding," in *Proceedings of the* 9^{th} *ISMRM*, Glasgow, Scotland, Apr. 2001, p. 1808.

[29] ——, "Sub-Second Continuous 2D non-Fourier Dynamic Adaptive MRI Using Near-Optimal Spatial Encoding," in *Proceedings of the* 9^{th} *ISMRM*, Glasgow, Scotland, Apr. 2001, p. 1779.

[30] J. J. van Vaals, M. E. Brummer, , W. T. Dixon, H. H. Tuithof, H. Engels, R. C. Nelson, B. M. Gerety, J. L. Chezmar, and J. A. den Boer, "Keyhole Method for Accelerating Imaging of Contrast Agent Uptake," *Journal of Magnetic Resonance Imaging*, vol. 3, pp. 671–675, 1993.

[31] G. P. Zientara, L. P. Panych, and F. A. Jolesz, "Keyhole SVD Encoded MRI," in *Proceedings* 2^{nd} *SMR*, San Francisco, USA, 1994, p. 778.

[32] R. Haimi-Cohen and A. Cohen, "Gradient-type Algorithms for Partial Singular Value Decomposition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 9, pp. 137–142, 1987.

[33] A. Frieze, R. Kannan, and S. Vempala, "Fast Monte-Carlo Algorithms for Finding Low-Rank Approximations," in *Proceedings of the 39th Annual Symposium on Foundations of Computer Science*, 1998.

[34] G. P. Zientara, L. P. Panych, and F. A. Jolesz, "Lanczos Spatial Encodings for Dynamically Adaptive MRI," in *Proceedings of the* 3^{rd} *ISMRM*, Nice, France, Aug. 1995, p. 664.

[35] C. H. Cunningham, G. A. Wright, and M. L. Wood, "High-Order Multiband Encoding in the Heart," *Magnetic Resonance in Medicine*, vol. 48, pp. 689–698, 2002.

[36] L. P. Panych, L. Zhao, F. A. Jolesz, and R. V. Mulkern, "Dynamic Imaging with Multiple Resolutions Along Phase-Encode and Slice-Select Dimensions," *Magnetic Resonance in Medicine*, vol. 45, no. 6, pp. 940–947, 2001.

[37] L. P. Panych, "Theoretical Comparison of Fourier and Wavelet Encoding in Magnetic Resonance Imaging," *IEEE Transactions in Medical Imaging*, vol. 15, no. 2, pp. 141–153, 1996.

[38] J. Hennig, "Echoes – How to Generate, Recognize, Use or Avoid Them in MR-Imaging Sequences Part I: Fundamental and Not So Fundamental Properties of Spin Echoes," *Concepts in Magnetic Resonance*, vol. 3, pp. 125–143, 1991.

[39] J. A. Tkach and E. M. Haacke, "A Comparison of Fast Spin Echo and Gradient Field Echo Sequences," *Magnetic Resonance Imaging*, vol. 6, pp. 373–389, 1988.