

Comparison of Three Algorithms for Solving Linearized Systems of Parallel Excitation RF Waveform Design Equations: Experiments on an Eight-Channel System at 3 Tesla

ADAM C. ZELINSKI,¹ LAWRENCE L. WALD,^{2,3} KAWIN SETSOMPOP,¹
VIJAYANAND ALAGAPPAN,³ BORJAN A. GAGOSKI,¹ VIVEK K GOYAL,¹
FRANZ HEBRANK,⁴ ULRICH FONTIUS,⁴ FRANZ SCHMITT,⁴ ELFAR ADALSTEINSSON^{1,2}

¹ *Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA*

² *Harvard-MIT Division of Health Sciences and Technology, Massachusetts Institute of Technology, Cambridge, MA*

³ *A.A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital, Harvard Medical School, Charlestown, MA*

⁴ *Siemens Medical Solutions, Erlangen, Germany*

ABSTRACT: Three algorithms for solving linearized systems of RF waveform design equations for calculating accelerated spatially-tailored excitations on parallel excitation MRI systems are presented. Their artifact levels, computational speed, and RF peak and root-mean-square (RMS) voltages are analyzed. An SVD-based inversion method is compared with conjugate gradient least squares (CGLS) and least squares QR (LSQR), two iterative algorithms designed to solve large linear systems. The excitation pulses calculated using these methods are used in both Bloch simulations and imaging experiments on an actual eight-channel parallel excitation coil array implemented on a 3T human scanner. Specifically, RF waveforms are designed for accelerated 2D spiral k -space trajectories to produce a variety of 2D target excitations and for a 3D spokes trajectory to produce a uniform thin-slice excitation. Overall, these experiments show that waveforms designed using LSQR and CGLS have significantly lower peak and RMS waveform voltages and produce excitations with fewer artifacts than those generated by the SVD-based method. © 2007 Wiley Periodicals, Inc. Concepts Magn Reson Part B (Magn Reson Engineering) 31B: 176–190, 2007

KEY WORDS: parallel excitation; cancellation artifacts; RF acceleration; RF coil array; conjugate-gradient methods

Received 9 April 2007; accepted 14 May 2007

Correspondence to: Adam C. Zelinski; E-mail: zelinski@mit.edu or Elfar Adalsteinsson; E-mail: elfar@mit.edu

Concepts in Magnetic Resonance Part B (Magnetic Resonance Engineering), Vol. 31B(3) 176–190 (2007)

Published online in Wiley InterScience (www.interscience.wiley.com). DOI 10.1002/cmr.b.20093

© 2007 Wiley Periodicals, Inc.

INTRODUCTION

Parallel RF excitation in the presence of gradient trajectories offers a flexible means for spatially-tailoring excitation patterns for inner-volume excitation (I) and addressing increased B_1 inhomogeneity observed at high field strengths (2–6). These pulses are useful because they may be tailored to impose an arbitrary spatial pattern on the transverse magnetization's magnitude and phase, subject to power deposition limits and constraints on RF and gradient hardware.

Implementations of multi-channel parallel excitation systems were first shown by Ullmann et al. (7) and Zhu et al. (8). More recently, researchers have designed eight-channel parallel excitation systems and demonstrated slice-selective and 2D spatial excitations (5, 9).

A parallel excitation system consists of a coil array with multiple elements capable of independent, simultaneous RF transmission. Assuming the set of gradient waveforms is fixed (i.e., the k -space trajectory is predetermined), the spatial coil profiles (B_1 maps) of the coil elements are known, and that a complex-valued target excitation pattern is chosen, it remains necessary to design a set of RF waveforms for these array elements to perform the excitation. The primary limitation of any excitation pulse is its duration. A parallel excitation system allows one to shorten this overall pulse duration by exploiting variations among the spatial excitation profiles of the coil array elements during the design stage (7, 8, 10–14), i.e., one may “accelerate” a given k -space trajectory via undersampling to reduce pulse duration. For example, an acceleration factor of R for a 2D spiral trajectory means the radial separation between spiral samples is increased R -fold relative to a Nyquist-sampled design. Acceleration is possible due to the extra degrees of freedom provided by the system’s multiple excitation elements and is analogous to acceleration in parallel reception.

This paper investigates three algorithms for solving linearized systems of parallel excitation RF waveform design equations by conducting several experiments. In Experiment 1 (E1), pulses are designed using each algorithm to produce a square target pattern for $R = 1, 4, 6,$ and 8 spiral trajectories. For each algorithm and R value, the mean-square error (MSE) between the resulting Bloch-simulated excitation and target is calculated, which quantifies each method’s excitation artifacts. Each waveform’s peak voltage, V_{peak} , and root-mean-square (RMS) voltage, V_{RMS} , are also determined. In Experiment 2 (E2), the trajectories are again spirals and the target is a text logo. In addition to Bloch-simulation analyses, the designed waveforms are played through a fully-implemented eight-channel parallel excitation system at 3T, and actual excitations are analyzed. In Experiment 3 (E3), the trajectory is a fixed set of “spokes” in k_z that sample the (k_x, k_y) -plane to achieve slice selection in z (3–5, 15), and the in-plane target is a uniform pattern. Finally, in Experiment 4 (E4), the trajectory is an $R = 8$ spiral and the target is again the text logo. Thousands of pulses are designed by looping over each method’s primary control parameter, which provides extensive

empirical data that shows how well each method trades off excitation quality with V_{peak} and V_{RMS} .

For each experiment, after fixing the target and trajectory, the Bloch equations relating the RF waveforms and target excitation are first linearized using Grissom et al.’s formulation (10). Other valid linearization approaches are those of Katscher et al. (12) and Zhu et al. (8, 14). After linearizing the system, each design method is used to generate a set of pulses. The methods have different regularizations and implementations that influence their optimization criteria and finite-precision arithmetic effects, which in turn strongly affect the resulting pulses; thus each method’s designed waveform produces a unique excitation.

The first of these methods uses an approximate pseudoinverse generated via singular value decomposition (SVD), a popular approach for least-squares problems whose use is analytically justifiable (16, 17). The other methods are conjugate gradient least-squares (CGLS) and least-squares QR (LSQR), iterative CG optimization algorithms for solving large linear systems (18–20). An early use of a CG method was the reconstruction of sensitivity encoded (SENSE) data by Kannengießer et al. (21) and Pruessmann et al. (22, 23). More recently, a CG method has been used by Yip et al. to design pulses for a single-coil system (24) and by Grissom et al. to design pulses for an emulated parallel excitation system (10).

THEORY

Parallel Excitation RF Waveform Design

Linearization. The interaction between the RF waveforms driving each array element of the array and the resulting excitation is nonlinear, so to simplify the RF design process it is useful to linearize the system. To do this, we use Grissom et al.’s formulation (10), an extension of Pauly’s small-tip angle approximation (25) to parallel systems. This reduces the design problem to solving a set of linear equations that contain a Fourier-like integral, which may be reformulated to contain a Fourier transform, leading to the following system:

$$m(\mathbf{r}) = i\gamma \sum_{p=1}^P S_p(\mathbf{r}) \int_0^T b_{1,p}(t) e^{i\mathbf{r}\cdot\mathbf{k}(t)} dt, \quad [1]$$

where \mathbf{r} is a spatial variable, $m(\mathbf{r})$ the approximate relative transverse magnetization pattern resulting from the Bloch equation, γ the gyromagnetic ratio, $S_p(\mathbf{r})$ the sensitivity profile of the p th coil, $b_{1,p}(t)$ the wave-

form played along the p th coil, and T the duration of each waveform. $\mathbf{k}(t)$ is the excitation k -space trajectory, equal to $-\gamma \int_t^T \mathbf{G}(\tau) d\tau$, where $\mathbf{G}(t)$ is a multi-dimensional gradient waveform of duration T . Discretizing Eq. [1] reduces it to a discrete linear system:

$$\mathbf{m} = \mathbf{A}\mathbf{b} \quad [2]$$

where \mathbf{m} is $M \times 1$, created by ordering the discretized elements of $m(\mathbf{r})$ within the field of excitation (FOX), \mathbf{b} is a voltage vector of samples of all P RF waveforms, and \mathbf{A} is $M \times N$, incorporating information about the coil profiles and trajectory. M is the number of locations at which the profiles are sampled in the FOX (equal to 1,466 for all our experiments), and N is P times the number of samples in each $b_{1,p}$ waveform. The size of \mathbf{A} and whether Eq. [2] is overdetermined or underdetermined depends on how finely the underlying continuous functions are sampled relative to one another. Note that \mathbf{m} , \mathbf{A} , and \mathbf{b} are complex-valued.

RF Waveform Generation. Choosing a desired magnetization pattern and trajectory implicitly determines \mathbf{d} and \mathbf{A} , where \mathbf{d} is $M \times 1$, formed by ordering the elements of $m(\mathbf{r})$ within the FOX. It is then necessary to find a candidate vector \mathbf{b} that approximately solves $\mathbf{d} = \mathbf{A}\mathbf{b}$. Once \mathbf{b} is found, the P waveforms may be extracted, played through a Bloch simulation or actual system, and an excitation pattern may be recorded, the latter of which will resemble the target if the small-tip angle approximation holds and the SNR is sufficient.

Three Algorithms for Solving a Linear System of Equations

SVD-Based Algorithm. One may solve $\mathbf{d} = \mathbf{A}\mathbf{b}$ via a truncated pseudoinverse generated by an SVD (16, 17), seeking a solution that minimizes $\|\mathbf{d} - \mathbf{A}\mathbf{b}\|_2$. This is accomplished with the Moore–Penrose pseudoinverse of \mathbf{A} , denoted as \mathbf{A}^\dagger , yielding $\mathbf{b} = \mathbf{A}^\dagger \mathbf{d}$. To generate \mathbf{A}^\dagger , an SVD is used to decompose \mathbf{A} into $\mathbf{U}\mathbf{\Sigma}\mathbf{V}^H$, where \mathbf{U} and \mathbf{V} are $M \times M$ and $N \times N$ eigenvector matrices and \mathbf{A}^H the complex transpose. If \mathbf{A} is of rank $J \leq \min(M, N)$, then $\mathbf{\Sigma}$ is $M \times N$ and diagonal, and its diagonal elements $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_J > 0$ are the nonzero singular values (SVs) of \mathbf{A} . Formally,

$$\mathbf{A}^\dagger = \mathbf{V}\mathbf{\Sigma}^+\mathbf{U}^H = \sum_{j=1}^J \sigma_j^{-1} \mathbf{v}_j \mathbf{u}_j^H, \quad [3]$$

where \mathbf{u}_j and \mathbf{v}_j are the j th columns of \mathbf{U} and \mathbf{V} , and $\mathbf{\Sigma}^+$ is implicitly defined. When \mathbf{A} is ill-conditioned, Eq. [3] yields a poor candidate for \mathbf{b} because it uses all non-

zero SVs, even those representing only the system's noise subspace. However, by avoiding the use of smaller SVs and retaining only the first $K < J$, a better-conditioned truncated pseudoinverse is obtained:

$$\mathbf{A}_K^\dagger = \sum_{k=1}^K \sigma_k^{-1} \mathbf{v}_k \mathbf{u}_k^H, \quad [4]$$

allowing one to obtain a better-conditioned estimate $\mathbf{b}_K = \mathbf{A}_K^\dagger \mathbf{d}$. K is therefore this method's control parameter: as it is increased, the error $\|\mathbf{d} - \mathbf{A}\mathbf{b}_K\|_2$ decreases whereas the energy of the solution vector, $\|\mathbf{b}_K\|_2$, increases. One typically applies this method by retaining just enough SVs to yield a solution with acceptably low residual error while keeping $\|\mathbf{b}\|_2$ small. For large matrices this algorithm is slow because it computes an SVD, but for fixed K there exist fast methods to compute \mathbf{A}_K^\dagger directly.

Conjugate Gradient Least-Squares. This algorithm solves the following optimization problem:

$$\min_{\mathbf{b}} \left\| (\mathbf{A}^H \mathbf{A} + \lambda_{\text{CGLS}} \mathbf{I}) \mathbf{b} - \mathbf{A}^H \mathbf{d} \right\|_2, \quad [5]$$

where λ_{CGLS} is a regularization term. One sees from Eq. [5] that as λ_{CGLS} is increased, $\|\mathbf{b}\|_2$ decreases and the residual error $\|\mathbf{r}\|_2 = \|\mathbf{d} - \mathbf{A}\mathbf{b}\|_2$ increases. CGLS does not perform an SVD and requires only $2M + 3N$ complex multiplications per iteration i . When solving Eq. [5], users may restrict the number of iterations or specify a threshold ε such that CGLS halts when $\|\mathbf{s}_i\|_2^2 / \|\mathbf{s}_0\|_2^2 < \varepsilon$. CGLS may also incorporate preconditioning matrices, weighted norms, and initial conditions.

When λ_{CGLS} is 0, CGLS is identical to Hestenes and Stiefel's iterative CG method for least-squares problems (20). Pseudocode for CGLS is provided in the Appendix, illustrating how its sequence of approximations \mathbf{b}_i is generated. Analytically, the \mathbf{b}_i are such that the residual error $\|\mathbf{r}_i\|_2$ decreases monotonically (20).

CGLS is similar to Sutton et al.'s CG method (SCG) (26) that Yip et al. (24) and Grissom et al. (10) use for designing RF pulses, because both CGLS and SCG are based on the Hestenes–Stiefel method. Step 2c of CGLS in the Appendix shows that the numerator of the step size α_i equals $\|\mathbf{s}_{i-1}\|_2^2$, and thus is guaranteed to be nonnegative real (zero if an exact solution is reached). The numerator of SCG's step size, however, is $\mathbf{p}_i^H \mathbf{s}_{i-1} = (\mathbf{s}_{i-1} + \beta_i \mathbf{p}_{i-1})^H \mathbf{s}_{i-1}$, and thus not guaranteed to be positive. We confirmed this numerically by providing SCG with randomly generated inputs and consistently observing complex-valued step sizes.

Least-Squares QR. This algorithm is an implementation of Tikhonov regularization and solves large linear least-squares problems in a numerically attractive manner (18, 19). Its name comes from its use of the QR decomposition (16, 17). The algorithm has one regularization parameter, λ_{LSQR} , and solves the following:

$$\min_{\mathbf{b}} \left(\|\mathbf{d} - \mathbf{A}\mathbf{b}\|_2^2 + \lambda_{\text{LSQR}}^2 \|\mathbf{b}\|_2^2 \right). \quad [6]$$

As λ_{LSQR} is increased, more weight is placed on the energy of \mathbf{b} than on the residual error, causing $\|\mathbf{b}\|_2$ to decrease and $\|\mathbf{d} - \mathbf{A}\mathbf{b}\|_2$ to increase. LSQR also avoids use of an SVD; it requires $3M + 5N$ complex multiplications per iteration. Pseudocode for LSQR when $\lambda_{\text{LSQR}} = 0$ is provided in the Appendix [from Section 4 of (18)].

LSQR, like CGLS, generates \mathbf{b}_i such that $\|\mathbf{r}_i\|_2$ decreases monotonically, but LSQR performs better in practice (18, 27) because of its unique restructuring of the input system [via the Lanczos process (28) and Golub–Kahan bidiagonalization (29)] prior to solving it. Empirical studies conducted by Paige and Saunders (18) and Bjorck and Elfving (27) show that LSQR finds solutions with lower residual error than CGLS when \mathbf{A} is ill-conditioned, and of similar fidelity when \mathbf{A} is well-conditioned. Further description of how LSQR applies the Lanczos process and Golub–Kahan factorization, along with pseudocode when $\lambda_{\text{LSQR}} \neq 0$, is located in (18, 19).

In addition to the above, LSQR’s stopping rules are carefully designed to reflect the data’s accuracy. Relative to CGLS’s stopping rule $\|s_i\|_2^2 / \|s_0\|_2^2 < \epsilon$, LSQR’s ensures that it always shuts down sooner and its corresponding \mathbf{b} estimate is equally acceptable. This advantage becomes more pronounced as \mathbf{A} ’s conditioning worsens (18). (Saunders MA. “CGLS & LSQR: differences in convergence behavior,” Personal Communication 2007).

Note that while LSQR indeed requires $N + 2M$ more complex multiplications per iteration than CGLS, this is mitigated by the fact that LSQR often requires fewer iterations to attain a similar-fidelity solution.

METHODS

Image Quality Evaluation

For each experiment we conduct, we generate an excitation image, $O(x, y)$, either via a Bloch simulation or by performing an excitation on the eight-channel

system. We then evaluate the quality of $O(x, y)$ by using metrics that quantify how closely it matches the target pattern, $T(x, y)$, each of which is explained below.

Mean-Squared Error. This measures how close $O(x, y)$ is to $T(x, y)$ over a chosen region of interest (ROI):

$$\text{MSE}(O, T) = \frac{1}{\text{card}(S)} \sum_{(x,y) \in S} |O(x,y) - T(x,y)|^2, \quad [7]$$

where S is a set of coordinates that implicitly defines the spatial ROI over which the MSE is computed, and $\text{card}(S)$ is the number of elements in S .

Second-Order Statistics. Computing the mean μ and standard deviation σ of $O(x, y)$ in different ROIs quantifies the severity of artifacts and noise present within each, e.g., if the target is uniform in a particular ROI, a small σ value implies that $O(x,y)$ closely matches $T(x, y)$ in that particular region.

Peak Value. The maximum value in an ROI of $O(x, y)$ quantifies the worst-case artifact present, e.g., given two observations of the same target, larger peak values in one indicate it has more artifacts than the other.

Note on Non-MSE Metrics. Recent RF pulse design work uses MSE to evaluate the quality of an excitation (10, 24). As an extension of this methodology, we make use of non-MSE metrics in addition to MSE, because the latter is not always an ideal indicator of excitation quality. For example, Wang et al. provide an example of six images with identical MSE, but three contain significant spike-like noise (30). Using region-by-region peak value and second-order statistics analyses on these images causes the noise-ridden ones to exhibit worse scores, whereas MSE incorrectly indicates all images are of equal quality.

RF Waveform Characterization

The peak voltage of a N -element \mathbf{b} vector, V_{max} or V_{peak} , equals $\max_i |\mathbf{b}_i|$, and the RMS voltage of \mathbf{b} is:

$$V_{\text{RMS}} = \left(\frac{1}{N} \sum_{i=1}^N |\mathbf{b}_i|^2 \right)^{1/2} \quad [8]$$




Figure 1 Magnitudes of the eight coil profiles obtained via B_1 mapping.

For each set of waveforms, knowing V_{\max} allows us to compare pulses' relative peak powers, and since V_{RMS} is proportional to both integrated and average pulse power, comparing V_{RMS} values of different pulses tells how much power they dissipate relative to one another. Note that since \mathbf{b} contains waveforms across all excitation channels, V_{peak} is the peak among all the waveforms, and V_{RMS} is the sum of each individual waveform's RMS voltage.

Experiment Setup

System Configuration. The parallel system is built around a Siemens 3T Tim Trio scanner (Siemens Medical Solutions, Erlangen, Germany). The coil array is composed of eight circular, overlapped, 15-cm diameter, detunable surface coils arranged on a 28-cm diameter acrylic tube (31). All scans are performed in a 17-cm low-dielectric oil phantom. For each RF design, the array's eight independent channels are driven, modulated in magnitude and phase by the pulses. Readouts are performed using a GRE sequence with $\text{TR} = 30$ ms, $\text{TE} = 6$ ms, and $\text{BW} = 400$ Hz/pixel.

Spatial Profiles (B_1 Maps). Spatial profiles within a phantom are obtained by sending a low flip angle pulse through each of the eight coil array elements, one at a time, and receiving on a body coil. B_1 maps are generated by recording a complex-valued image via a GRE sequence ($\text{TR} = 20$ ms, $\text{TE} = 6$ ms, $\text{BW} = 400$ Hz/pixel), yielding 51×51 pixel, 4-mm resolution maps that capture the magnitude and relative

phase of each array element. The magnitudes of the spatial profiles are shown in Fig. 1. Spatial variations in the body coil's reception profile are not removed because the profile is fairly uniform ($<5\%$ variation). Before using the maps to generate the \mathbf{A} matrix, each is scaled by a constant so the largest magnitude across all map pixels equals unity. This scaling makes the maps "qualitative" in the low flip angle domain; they do not convey the exact flip angle achieved.

Target Images. The targets are a square, a Massachusetts Institute of Technology (MIT) logo, and a uniformly-flat excitation. Figure 2 depicts each 51×51 pixel, 4-mm resolution target. The logo has two nonzero intensity levels: the lower part of the "i" is twice the intensity of other letters. The intensities of the square, the uniform target, and the lower part of the "i" equal 0.01; this value is arbitrarily chosen since the B_1 maps are qualitative. This means that the V_{peak} and V_{RMS} of each \mathbf{b} vector designed using these maps and targets are not actually in units of volts, but since 0.01 is used consistently through all experiments, it is possible to make relative comparisons between voltages.

Spiral Trajectories. The 2D k -space spirals are configured to have 4-mm resolution and an 18-cm field-of-view (FOV). Gradient amplitude and slew rate are 35 mT/m and 150 T/m/s, respectively. For $R = 1, 4, 6,$ and 8 , spirals are created that are 9.47, 2.42, 1.64, and 1.26 ms long, undergo 16, 4, 3, and 2 revolutions, and lead to \mathbf{b} vectors of length 15,152;

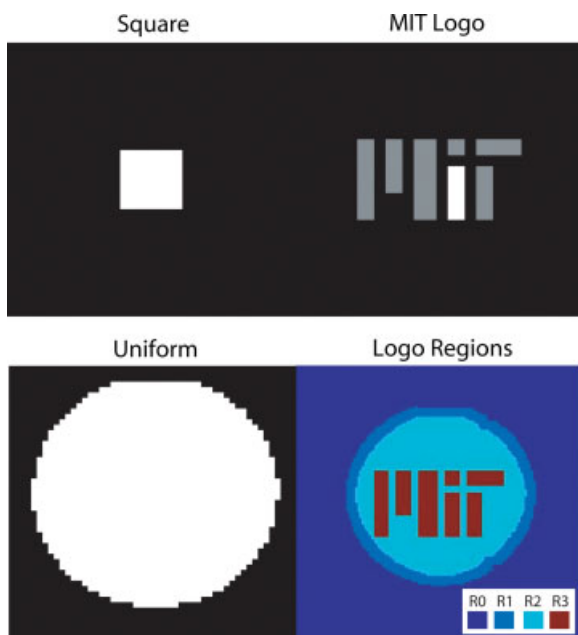


Figure 2 Target excitation patterns and MIT logo regions. The square, MIT logo, and uniformly-flat targets are shown. The MIT logo has two nonzero intensity levels—the lower part of the letter “i” is twice as intense as the others. The region-by-region breakdown of the logo is also shown. Region 0 is where statistical noise occurs; Region 1 is the ring-like edge region where the most glaring artifacts typically occur; Region 2 is the suppression region where the coil profiles interact and attempt to cancel out, but do not do so perfectly in practice; Region 3 is the letters of the target. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

3,872; 2,642; and 2,016, respectively. The $R = 4$ spiral is depicted in Fig. 3.

Spokes Trajectory. The slice-selective trajectory described in (3–5, 15) consists of 10 spokes in k_z placed in the (k_x, k_y) plane to yield an 18-cm FOV, and is shown in Fig. 3. Slice thickness is 1 cm and the center spoke’s time-bandwidth product is 4. Gradient amplitude and slew rate are 30 mT/m and 120 T/m/s, yielding a 5.84-ms pulse. To simplify the design process, we restrict the shape of each waveform to a Hanning-windowed sinc in k_z , which fixes the slice-selectivity of the trajectory and means that each design method only needs to calculate an amplitude and phase for each excitation channel to encode along each spoke. With 10 spokes and eight coil elements, this means \mathbf{A} has 80 columns and \mathbf{b} has 80 elements.

Experiment Summary. Table 1 summarizes the four experiments, listing the key details of each.

Experiments 1, 2, and 3 (Bloch Simulations). In E1 and E2, the targets are the square and text logo, respectively, and the trajectories are $R = 1, 4, 6,$ and 8 spirals. In E3, the in-plane target is the uniform pattern and the spokes trajectory is used. For every fixed target and trajectory, \mathbf{d} and \mathbf{A} are known, and \mathbf{b} vectors are then calculated by solving $\mathbf{d} = \mathbf{A}\mathbf{b}$ using one of the algorithms. Once \mathbf{b} is determined, the pulses are Bloch-simulated as in (5).

For E1, E2, and E3, we first apply the SVD method, retaining enough SVs so the Bloch simulation of the resulting pulse yields an acceptable-looking excitation. Noting the MSE between this excitation and the target, we run LSQR and CGLS, tuning their parameters such that their Bloch-simulated pulses yield equal or lower MSE excitations. We attempt to make CGLS’s MSE close to LSQR’s.

Experiment 2 (Eight-Channel System at 3T). For E2, each waveform designed during the simulation stage is played through the eight-channel excitation system by first scaling each “qualitative” \mathbf{b} vector by a constant so its elements represent actual voltages. This scaling depends on both R and the design algorithm, and is chosen such that the flip angle of the MIT logo excited on the system is approximately constant across all experiments. Each excitation conducted on the system is then stored as a magnitude image, and MSE, voltage, and region-by-region second-order statistics and peak value analyses are

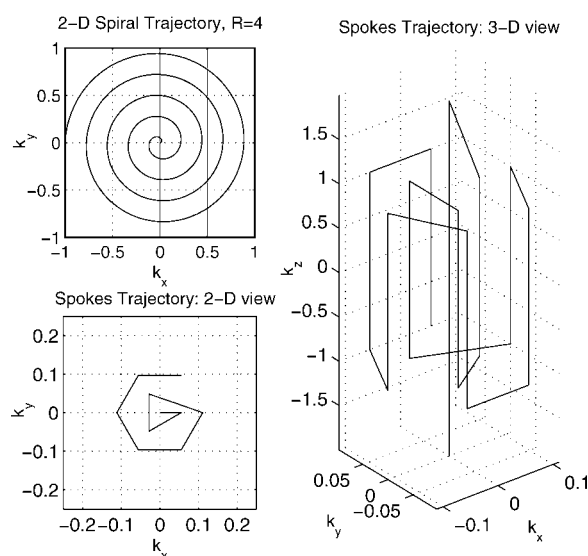


Figure 3 k -space trajectories: the $R = 4$ spiral and 10-spoke slice-selective trajectory.

Table 1 List of Experiments

Label	Target Excitation	Trajectory Type	Methodology
E1	Square	Spirals ($R = 1, 4, 6, 8$)	Bloch simulations
E2	MIT Logo	Spirals ($R = 1, 4, 6, 8$)	Bloch simulations + actual system runs
E3	Uniformly-Flat	Spokes	Bloch simulations
E4	MIT Logo	Spiral ($R = 8$)	Bloch simulations + voltages from actual system runs

conducted. The latter two metrics are calculated over the regions depicted in Fig. 2. Region 0 (R0) is where system noise is present, Region 1 (R1) the edge region where glaring artifacts tend to occur, Region 2 (R2) the suppression region where the profiles are interacting to cancel each other out, and Region 3 (R3) the letters.

Experiment 4. Here the trajectory is the $R = 8$ spiral and the target is the MIT logo. For each design algorithm, we loop over many choices of its control parameter, generating thousands of pulses (\mathbf{b} vectors). Then for each designed pulse, we compute its Bloch-simulated excitation's MSE with respect to the target, along with V_{peak} and V_{RMS} . Since we know from E2 how to properly scale each method's \mathbf{b} vectors to play them on the eight-channel system, we scale them here as well, obtaining the actual voltage characteristics of each pulse in volts. This essentially generates MSE vs. voltage tradeoff curves for each method. Note that extremely low-MSE pulses that yield completely unrealistic voltage values (e.g., $V_{\text{peak}} > 1000$ V) are disregarded.

RESULTS

Experiments 1, 2, and 3 (Bloch Simulations)

For E1, the 12 resulting excitations are shown in Fig. 4. The rows and columns correspond to the $R = 1, 4, 6,$ and 8 spiral trajectories and the three design algorithms, respectively. Each subplot depicts the Bloch-simulated excitation, the MSE between this simulated excitation and target, and ($V_{\text{max}}, V_{\text{RMS}}$). Figure 5 shows the results of E2, and is formatted analogously to Fig. 4. For E3, Bloch-simulated images, MSEs, V_{max} , and V_{RMS} appear in Fig. 8. For E1, E2, and E3, each method's design parameters and runtime are listed in Tables 2–4, respectively. In the interest of space, second-order statistics and peak values are not shown for these experiments.

Experiment 2 (Eight-Channel System at 3T)

Figure 6 shows the system images when the scaled \mathbf{b} vectors are played through the eight-channel system. The region-by-region means, standard deviations, and peak values of each are shown in Fig. 7, where bar graphs of each statistic are shown for R0 through R3. The axes of each such graph are R and algorithm type.

Experiment 4

The upper, middle, and lower plots of Fig. 9 illustrate E4's results. The upper plot shows the iteration-by-iteration MSE performance of LSQR and CGLS when λ_{CGLS} and λ_{LSQR} equal zero. The middle plot shows the MSE vs. V_{peak} tradeoff of a variety of SVD pulses (by retaining different numbers of non-zero SVs) and the tradeoff curve for a large number of LSQR-designed pulses (generated by varying λ_{LSQR} over a wide range). The lower plot is analogous to the middle one, showing the MSE vs. V_{RMS} tradeoffs of the SVD and LSQR methods. CGLS data is not displayed in the middle and lower plots because, at the displayed scale, it is nearly identical to the LSQR data.

DISCUSSION

Mean-Square Errors

Referring to Figs. 4–6, and 8, one sees that each LSQR and CGLS excitation has an equal or lower MSE than the corresponding SVD one, proving that we calibrated the designs in the proposed manner (see Methods). Further, as R increases, the MSEs of the SVD designs grow faster than those of LSQR and CGLS.

Experiments 1, 2, and 3 (Bloch Simulations)

For E1, Fig. 4 makes evident that excitations due to LSQR and CGLS are better than those due to SVD,




Figure 4 Experiment 1’s Bloch-simulated excitations with MSE, V_{\max} , and V_{RMS} overlays. The rows from top to bottom correspond to $R = 1, 4, 6,$ and 8 . From left to right, the columns correspond to RF waveforms designed using the SVD method, LSQR, and CGLS. The MSE between each excitation and the target is shown, along with V_{\max} and V_{RMS} of each designed waveform.

i.e., for fixed R , LSQR and CGLS always result in lower MSE, V_{peak} , and V_{RMS} . Also, regardless of design technique, we see that artifacts always increase rapidly with R . Analyzing Fig. 5, one sees E2 exhibits the same trends, e.g., for $R = 6$, the LSQR image has 1.11 times lower MSE and (1.63, 1.23) times lower voltages than the SVD method. For CGLS, ($V_{\text{peak}}, V_{\text{RMS}}$) are nearly identical to LSQR’s, but CGLS’s MSE is higher, so LSQR outperforms CGLS in excitation quality for the same amount of waveform energy.

In E3, A is 1466×80 in size and thus highly overdetermined, whereas in E1 and E2 it is highly

underdetermined, which means E3 poses a radically different design problem. Yet as Fig. 8 shows, LSQR and CGLS continue to outperform the SVD based method, e.g., the SVD and LSQR images have nearly the same MSE, but LSQR’s V_{RMS} is 1.14 times lower than the SVD pulse’s voltage. We do see that LSQR is not outperforming the SVD method as well as in E1 and E2, and thus conjecture that LSQR and CGLS provide moderately better MSE vs. voltage tradeoffs than the SVD method when the system is small and overdetermined, and significantly better ones when it is large and underdetermined.




Figure 5 Experiment 2's Bloch-simulated excitation images with MSE, V_{\max} , and V_{RMS} overlays, formatted analogously to Fig. 4.

Experiment 2 (Eight-Channel System at 3T)

From Fig. 7, one sees that R_0 's μ and σ are equivalent across all images, confirming the noise floor is equivalent across all trials. For each R , the means of the images in R_3 are nearly equal, proving each excitation achieved nearly the same flip angle. In R_0 through R_2 , if no artifacts at all were present, the μ and σ values in Fig. 7 would equal those of the background noise, because the ideal target has zero intensity there. This means that smaller μ and σ in these regions imply fewer artifacts are present. In R_1 , the ring-like edge region that is particularly artifact-ridden, it is evident that for fixed R , LSQR and CGLS

images exhibit superior second-order statistics than SVD images. For example, in R_1 for $R = 6$, (μ, σ) for the LSQR and CGLS images equal (206, 126) and (203, 127), respectively, whereas SVD's are significantly higher, equal to (260, 175). This same trend occurs in R_2 and R_3 . The peak values also exhibit these trends, e.g., for $R = 4$ in R_2 , the SVD image has a peak of 1,935 whereas the LSQR and CGLS images have much smaller peaks of 1,393 and 1,526.

Comparing the MIT logo simulation results in Fig. 5 with the eight-channel system results in Fig. 6, one sees that LSQR and CGLS's better MSE vs. voltage tradeoffs are present not just in the Bloch simula-

Table 2 Experiment 1 (Square Target, Spiral Trajectories): Algorithm Design Parameters and Runtimes

	SVD			LSQR			CGLS		
	Kept SVs	Total SVs	RT	λ_{LSQR}	Iter.	RT	λ_{CGLS}	Iter.	RT
$R = 1$	976	1,466	476	0.25	500	57	0.0125	1,000	40
$R = 4$	861	1,466	227	0.08	250	29	0.006	1,000	19
$R = 6$	623	1,466	187	0.09	500	16	0.0075	750	12
$R = 8$	391	1,466	159	0.025	500	27	0.0005	750	37

Input parameters to each method used to compute each of the \mathbf{b} vectors are listed. RT stands for runtime in seconds.

tions, but in the system images as well. This proves that the advantage of LSQR and CGLS is due to the fundamental properties of these algorithms exhibited during the simulation stage, and is not due to system hardware effects such as transmission bandwidth.

Experiment 4

From the upper plot of Fig. 9, it is clear that for any of the given numbers of iterations, LSQR slightly outperforms CGLS in terms of MSE. From the middle and lower plots, it is evident that the LSQR method provides superior MSE vs. V_{peak} and MSE vs. V_{RMS} tradeoffs relative to the SVD method. Specifically, for every MSE evaluated, LSQR creates a pulse with lower peak and lower RMS voltage than the SVD method does. In many cases, the performance of LSQR relative to the SVD method is extreme, e.g., the SVD pulse that yields an image quality of $\log_{10}(\text{MSE}) = -5.8$ has $V_{\text{peak}} = 55.3$ V and $V_{\text{RMS}} = 12.3$ V, whereas the LSQR pulse achieves the exact same MSE with V_{peak} and V_{RMS} equal to merely 41.4 and 6.9 V. Note that these superior tradeoffs exhibited by LSQR are not limited to our chosen target and trajectory; they hold across many target-trajectory combinations, but because of space limitations we do not present these results. In conclusion, E4 shows that for the given target, trajectory, and fixed MSE, LSQR will always produce a pulse with better voltage characteristics than the SVD method.

Voltage Characteristics and Specific Absorption Rate

It is clear from E1 through E4 that LSQR and CGLS design pulses with better artifact vs. voltage tradeoffs than the SVD method. Because specific absorption rate (SAR) is heavily influenced by V_{peak} and V_{RMS} , we speculate that the SAR values of LSQR and CGLS waveforms are significantly lower than those of SVD-based pulses. [Note: because V_{peak} and V_{RMS} are not the sole determinants of SAR, it is necessary to quantify the SAR of each pulse to validate this hypothesis, perhaps via the methods in (32–34)].

One noticeable trend across all experiments is the rapid growth of V_{peak} and V_{RMS} with R . For example, LSQR's V_{peak} jumps from 5.54 to 40.0 V when transitioning from $R = 4$ to $R = 6$, and jumps to 95.49 V when $R = 8$. These observations coincide with those of Katscher et al. (35) and Ullmann et al. (36), extending the former's work from a strip-line coil to our circular array, and the latter's from a four-channel to our eight-channel system. This rapid voltage growth poses constraints on in vivo applications and implies the infeasibility of $R \gg 1$ designs due to algorithm development alone. However, since the maximum feasible R with moderate power requirements is strongly linked to the number of transmit elements and the design of the array (35), it may still be possible to achieve $R \gg 1$ pulses by designing arrays to handle higher voltages and have higher efficiencies.

Table 3 Experiment 2 (MIT Logo Target, Spiral Trajectories): Algorithm Design Parameters and Runtimes

	SVD			LSQR			CGLS		
	Kept SVs	Total SVs	RT	λ_{LSQR}	Iter.	RT	λ_{CGLS}	Iter.	RT
$R = 1$	835	1,466	663	0.75	3,000	26	0.5	3,000	10
$R = 4$	674	1,466	220	0.46	250	7	0.22	250	5
$R = 6$	676	1,466	187	0.06	500	22	0.00375	550	18
$R = 8$	767	1,466	158	0.025	500	27	0.0005	1,000	35

Input parameters to each method used to compute each of the \mathbf{b} vectors are listed. RT stands for runtime in seconds.

Table 4 Experiment 3 (Uniformly-Flat Target, Spokes Trajectory): Algorithm Design Parameters and Runtimes

SVD		LSQR		CGLS				
Kept SVs	Total SVs	RT	λ_{LSQR}	Iter.	RT	λ_{CGLS}	Iter.	RT
39	80	0.8	3.0	200	0.15	0.01	100	0.1

Performance of LSQR and CGLS

LSQR and CGLS outperform the SVD method not only because of their numerical properties but because

they directly penalize large \mathbf{b} vectors, whereas the SVD algorithm does not do so.

One sees across all experiments that LSQR and CGLS perform similarly in terms of each metric. This is because the \mathbf{A} matrix is well-conditioned in each case. Because of this, either algorithm may be used for waveform design instead of the SVD-based method. However, although they perform similarly in terms of MSE, voltage values and runtimes, LSQR should be used over CGLS, because of the various empirical studies showing its superior performance, especially in cases where \mathbf{A} is ill-conditioned (18, 27).

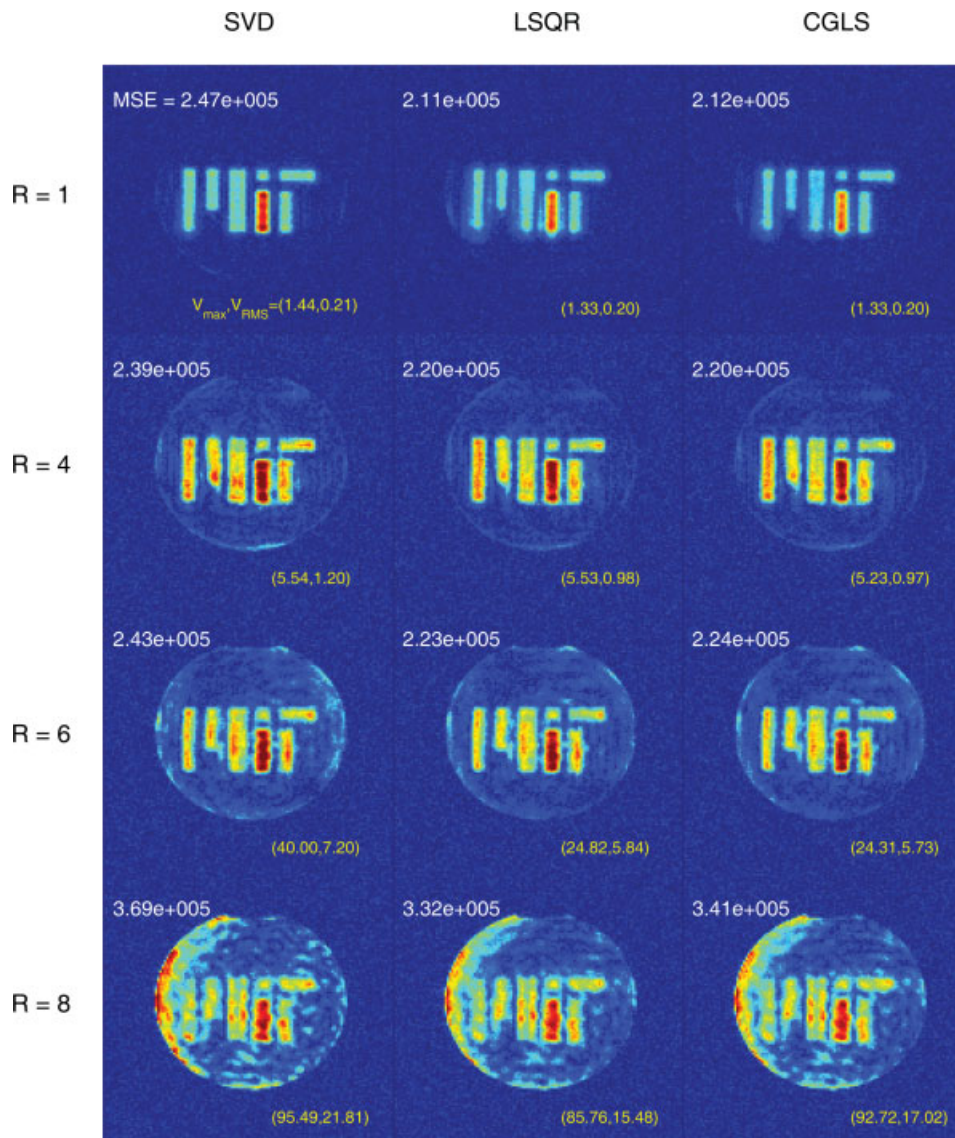


Figure 6 Experiment 2's excitations conducted on the eight-channel system with MSE, V_{max} , and V_{RMS} overlays, formatted analogously to Fig. 4.

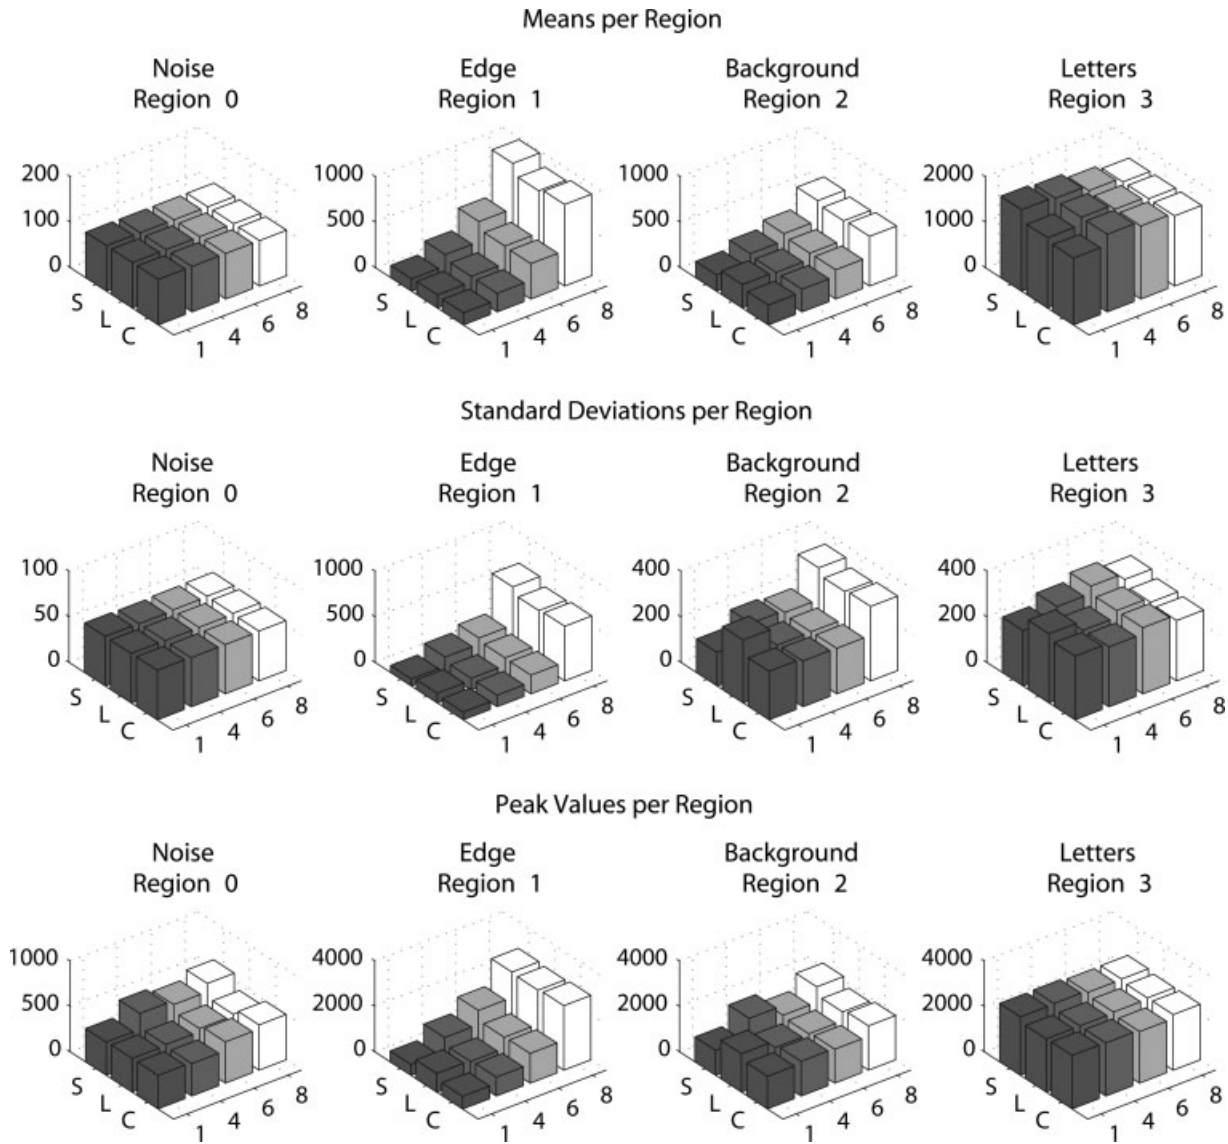


Figure 7 Experiment 2’s excitations conducted on the eight-channel system: means, standard deviations, and peak values per region. For Regions 0 through 3, a bar graph of each statistic is shown. Each graph’s x , y , and z axes denote the acceleration factor, algorithm type, and statistic value, respectively. “S,” “L,” and “C” stand for SVD method, LSQR, and CGLS.




Figure 8 Experiment 3’s Bloch-simulated excitation images with MSE, V_{max} , and V_{RMS} overlays. Each image shows the in-slice excitation achieved by the RF waveform designed with each algorithm. From left to right, the images are due to the SVD method, LSQR, and CGLS.

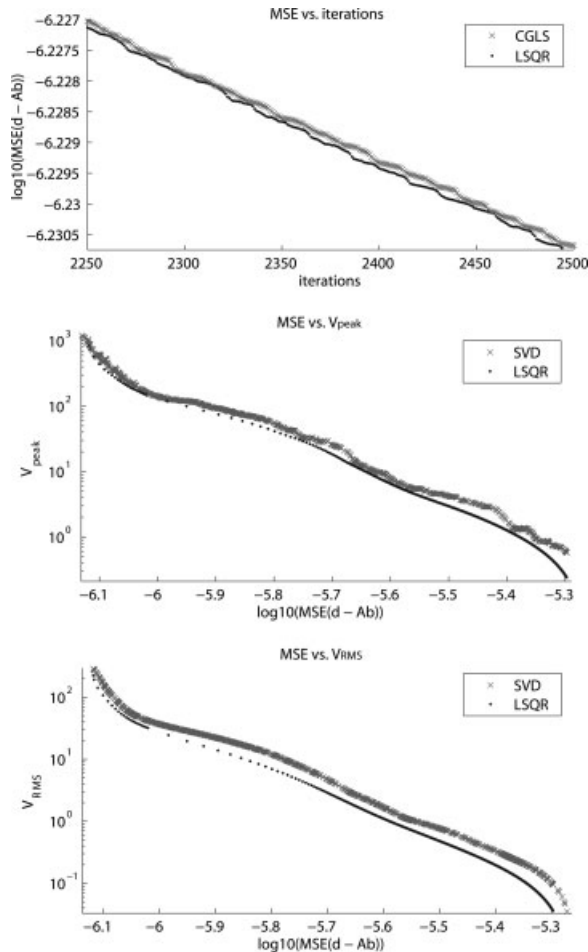


Figure 9 Experiment 4's results. The upper plot shows the iteration-by-iteration MSE performance of LSQR and CGLS when λ_{CGLS} and λ_{LSQR} equal zero. The middle plot shows the MSE vs. V_{peak} tradeoff of a variety of SVD pulses and the tradeoff curve for a large number of LSQR-designed pulses. The lower plot is analogous to the middle one, showing MSE vs. V_{RMS} tradeoffs of the SVD and LSQR methods.

Limitations

One limitation of this work is the qualitative nature of the spatial coil profiles. Obtaining quantitative B_1 maps—those that tell us the exact flip angle at each spatial location generated by each array element per input volt—is currently an open problem. One attempt at solving this problem has been by Cunningham et al. (37).

Another limitation is that during E2's eight-channel runs, several system images are saturated in a small sub-region of R3 because of the system imaging format's limited dynamic range relative to the acquisition parameters. To circumvent this issue, pixels within the more intense part of the letter "i"

are discarded from R3's second-order statistic and peak value computations, preventing error propagation. Note that in images where saturation was not an issue, the intensity level of the "i" was indeed twice that of the other letters.

CONCLUSION

Two iterative CG methods, LSQR and CGLS, have been shown to obtain better or equal quality excitations compared with an SVD-based truncated pseudoinversion method, while consistently depositing significantly less energy into the subject. This was shown to hold across a range of targets, k -space trajectories, and acceleration factors in both Bloch simulations and imaging experiments on an eight-channel system at 3T. Between LSQR and CGLS, the former had equal or superior numerical properties than the latter.

To the best of our knowledge, this work has made the following novel contributions: a) investigating different numerical methods of solving linearized systems of parallel RF waveform equations which have surprising differences in performance, b) validating these new methodologies across three types of targets and two significantly different k -space trajectories, c) introducing new non-MSE metrics that avoid the possible pitfalls of only using MSE to judge excitation quality, d) quantifying large increases in V_{peak} and V_{RMS} voltage as a function of R for an eight-channel system with up to eight-fold trajectory accelerations which let us explore the important issue of pulse energy that influences SAR, e) validating our results by conducting actual accelerated parallel excitations on a realistic eight-channel system at 3T, and f) showing for one target and trajectory, for all realistically-feasible system voltages, LSQR always generates a pulse with lower V_{peak} and V_{RMS} relative to the SVD method when excitation quality is fixed.

ACKNOWLEDGMENTS

This work was supported by Siemens Medical Solutions, the National Institutes of Health P41RR14075, the United States Department of Defense via a National Defense Science and Engineering Graduate Fellowship F49620-02-C-0041, Texas Instruments Leadership University Consortium Program, and the R.J. Shillman Career Development Award. We thank M. Saunders for helpful discussions about LSQR and CGLS.

APPENDIX

Algorithm CGLS (with arbitrary λ_{CGLS})

1. Set $\mathbf{r}_0 = \mathbf{d}$, $\mathbf{s}_0 = \mathbf{A}^H \mathbf{d}$, $\mathbf{p}_1 = \mathbf{s}_0$, $\gamma_0 = \|\mathbf{s}_0\|_2^2$, $\mathbf{b}_0 = \mathbf{0}$
2. For $i = 1, 2, 3, \dots$ repeat the following:
 - a. $\mathbf{q}_i = \mathbf{A} \mathbf{p}_i$ (Hestenes and Stiefel's intermediate vector)
 - b. $\delta_i = \|\mathbf{q}_i\|_2^2 + \lambda_{\text{CGLS}} \|\mathbf{p}_i\|_2^2$ (incorporate the regularization)
 - c. $\alpha_i = \gamma_{i-1} / \delta_i$ (calculate the step size)
 - d. $\mathbf{b}_i = \mathbf{b}_{i-1} + \alpha_i \mathbf{p}_i$ (update the set of RF waveforms)
 - e. $\mathbf{r}_i = \mathbf{r}_{i-1} - \alpha_i \mathbf{q}_i$ (update the residual error vector)
 - f. $\mathbf{s}_i = \mathbf{A}^H \mathbf{r}_i - \lambda_{\text{CGLS}} \mathbf{b}_i$ (incorporate the regularization)
 - g. $\gamma_i = \|\mathbf{s}_i\|_2^2$
 - h. $\beta_i = \gamma_i / \gamma_{i-1}$
 - i. $\mathbf{p}_{i+1} = \mathbf{s}_i + \beta_i \mathbf{p}_i$
 - j. Test for convergence. Exit if a stopping criterion has been met.

Algorithm LSQR (with $\lambda_{\text{LSQR}} = 0$)

1. (Initialize and begin the bidiagonalization.)
Set $\beta_1 \mathbf{u}_1 = \mathbf{d}$, $\alpha_1 \mathbf{v}_1 = \mathbf{A}^H \mathbf{u}_1$, $\mathbf{w}_1 = \mathbf{v}_1$, $\mathbf{b}_0 = \mathbf{0}$, $\bar{\rho} = \beta_1$, $\bar{\rho} = \alpha_1$
2. For $i = 1, 2, 3, \dots$ repeat steps 3 through 6
3. (Continue the bidiagonalization.)
 - a. $\beta_{i+1} \mathbf{u}_{i+1} = \mathbf{A} \mathbf{v}_i - \alpha_i \mathbf{u}_i$
 - b. $\alpha_{i+1} \mathbf{v}_{i+1} = \mathbf{A}^H \mathbf{u}_{i+1} - \beta_{i+1} \mathbf{v}_i$
4. (Construct and apply next orthogonal transformation.)
 - a. $\rho_i = (\bar{\rho}_i^2 + \beta_{i+1}^2)^{1/2}$
 - b. $c_i = \bar{\rho}_i / \rho_i$
 - c. $s_i = \beta_{i+1} / \rho_i$
 - d. $\theta_{i+1} = s_i \alpha_{i+1}$
 - e. $\bar{\rho}_{i+1} = -c_i \alpha_{i+1}$
 - f. $\phi_i = c_i \bar{\phi}_i$
 - g. $\bar{\phi}_{i+1} = s_i \bar{\phi}_i$
5. (Update \mathbf{b} , \mathbf{w} .)
 - a. $\mathbf{b}_i = \mathbf{b}_{i-1} + (\phi_i / \rho_i) \mathbf{w}_i$
 - b. $\mathbf{w}_{i+1} = \mathbf{v}_{i+1} - (\theta_{i+1} / \rho_i) \mathbf{w}_i$
6. Test for convergence. Exit if a stopping criterion has been met.

REFERENCES

1. Spielman D, Pauly J, Macovski A, Enzmann D. 1991. Spectroscopic imaging with multidimensional pulses for excitation: SIMPLE. *Magn Reson Med* 19:67–84.

2. Stenger VA, Saekho S, Zhang Z, Yu S, Boada FE. 2004. B_1 inhomogeneity reduction with transmit SENSE. In: Second International Workshop on Parallel Imaging, Zurich, Switzerland. p 94.
3. Saekho S, Boada FE, Noll DC, Stenger VA. 2005. Small tip angle three-dimensional tailored radiofrequency slab-select pulse for reduced B_1 inhomogeneity at 3T. *Magn Reson Med* 53:479–484.
4. Ulloa J, Hajnal JV. 2005. Exploring 3D RF shimming for slice selective imaging. In: ISMRM, Miami Beach, Florida, USA. p 21.
5. Setsompop K, Wald LL, Alagappan V, Gagoski BA, Hebrank F, Fontius U, et al. 2006. Parallel RF transmission with eight channels at 3 Tesla. *Magn Reson Med* 56:1163–1171.
6. Saekho S, Yip CY, Noll DC, Boada FE, Stenger VA. 2006. Fast- k_z three-dimensional tailored radiofrequency pulse for reduced B_1 inhomogeneity. *Magn Reson Med* 55:719–724.
7. Ullmann P, Junge S, Wick M, Seifert F, Ruhm W, Hennig J. 2005. Experimental analysis of parallel excitation using dedicated coil setups and simultaneous RF transmission on multiple channels. *Magn Reson Med* 54:994–1001.
8. Zhu Y, Watkins R, Gianquinto R, Hardy C, Kenwood G, Mathias S, et al. 2005. Parallel excitation on an eight transmit channel MRI system. In: ISMRM, Miami Beach, Florida, USA. p 14.
9. Graesslin I, Vernickel P, Schmidt J, Findelee C, Roschmann P, Leussler C, et al. 2006. Whole body 3T MRI system with eight parallel RF transmission channels. In: ISMRM, Seattle, Washington. p 129.
10. Grissom WA, Yip CY, Zhang Z, Stenger VA, Fessler JA, Noll DC. 2006. Spatial domain method for the design of RF Pulses in multicoil parallel excitation. *Magn Reson Med* 56:620–629.
11. Katscher U, Bornert P, Leussler C, van den Brink JS. 2002. Theory and experimental verification of transmit SENSE. In: ISMRM, Honolulu, Hawaii, USA. p 189.
12. Katscher U, Bornert P, Leussler C, van den Brink JS. 2003. Transmit SENSE. *Magn Reson Med* 49:144–150.
13. Zhu Y. Acceleration of focused excitation with a transmit coil array. 2002. In: ISMRM, Honolulu, Hawaii, USA. p 190.
14. Zhu Y. 2004. Parallel excitation with an array of transmit coils. *Magn Reson Med* 51:775–784.
15. Blamire AM, Rothman DL, Nixon T. 1996. Dynamic shim updating: a new approach towards optimized whole brain shimming. *Magn Reson Med* 36:159–165.
16. Golub GH, Van Loan CF. 1983. *Matrix Computations*. Baltimore, MD: Johns Hopkins University Press.
17. Strang G. 1993. *Introduction to Linear Algebra*. Wellesley, MA: Wellesley-Cambridge Press.
18. Paige CC, Saunders MA. 1982. LSQR: an algorithm for sparse linear equations and sparse least squares. *ACM Trans Math Software* 8:43–71.

19. Paige CC, Saunders MA. 1982. Algorithm 583; LSQR: sparse linear equations and least-squares problems. *ACM Trans Math Software* 8:195–209.
20. Hestenes MR, Stiefel E. 1952. Methods of conjugate gradients for solving linear systems. *J Res Natl Bur Stand* 49:409–436.
21. Kannengiesser SAR, Brenner AR, Noll TG. 2000. Accelerated image reconstruction for sensitivity encoded imaging with arbitrary k-space trajectories. In: ISMRM, Denver, Colorado. p 155.
22. Pruessmann KP, Weiger M, Bornert P, Boesiger P. 2000. A gridding approach for sensitivity encoding with arbitrary trajectories. In: ISMRM, Denver, Colorado. p 276.
23. Pruessmann KP, Weiger M, Bornert P, Boesiger P. 2001. Advances in sensitivity encoding with arbitrary k-space trajectories. *Magn Reson Med* 46:638–651.
24. Yip CY, Fessler JA, Noll DC. 2005. Iterative RF pulse design for multidimensional, small-tip-angle selective excitation. *Magn Reson Med* 54:908–917.
25. Pauly J, Nishimura D, Macovski A. 1989. A k-space analysis of small-tip angle excitation. *J Magn Reson* 81:43–56.
26. Sutton BP, Noll DC, Fessler JA. 2003. Fast, iterative image reconstruction for MRI in the presence of field inhomogeneities. *IEEE Trans Med Imaging* 22:178–188.
27. Bjorck A, Elfving T. Accelerated projection methods for computing pseudoinverse solutions of systems of linear equations. Linkoping, Sweden: Department of Mathematics, Linkoping University. Res Rep LiTH-MAT-R-1978-5.
28. Lanczos C. 1950. An iteration method for the solution of the eigenvalue problem of linear differential and integral operators. *J Res Natl Bur Stand* 45:255–282.
29. Golub GH, Kahan W. 1965. Calculating the singular values and pseudoinverse of a matrix. *SIAM J Numer Anal* 2:205–224.
30. Wang Z, Bovik AC, Sheikh HR, Simoncelli EP. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE Trans Image Process* 13:600–612.
31. Alagappan V, Wiggins GC, Potthast A, Setsompop K, Adalsteinsson E, Wald LL. 2006. An eight channel transmit coil for transmit sense at 3T. In: ISMRM, Seattle, Washington, USA. p 121.
32. Ibrahim TS, Abduljalil AM, Baertlein BA, Lee R, Robitaille P-M-L. 2001. Analysis of B1 field profiles and SAR values for multi-strut transverse electromagnetic RF coils in high field MRI applications. *Phys Med Biol* 46:2545–2555.
33. Collins CM, Liu W, Wang J, Gruetter R, Vaughan JT, Ugurbil K, et al. 2004. Temperature and SAR calculations for a human head within volume and surface coils at 64 and 300 MHz. *Magn Reson Med* 19:650–656.
34. Keshvari J, Keshvari R, Lang S. 2006. The effect of increase in dielectric values on specific absorption rate (SAR) in eye and head tissues following 900, 1800 and 2450 MHz radio frequency (RF) exposure. *Phys Med Biol* 51:1463–1477.
35. Katscher U, Vernickel P, Overweg J. 2005. Basics of RF power behaviour in parallel transmission. In: ISMRM, Miami, Florida. p 17.
36. Ullmann P, Wuebbeler G, Junge S, Seifert F, Ruhm W, Hennig J. 2006. SAR-analysis for transmit SENSE with a 4-channel head array at 3 T. In: ISMRM, Seattle, Washington. p 601.
37. Cunningham CH, Pauly JM, Nayak KS. 2006. Saturated double-angle method for rapid B1+ mapping. *Magn Reson Med* 55:1326–1333.