

Durga: A heuristically-optimized data collection strategy for volumetric magnetic resonance imaging

CHRISTOPHER KUMAR ANAND*, ANDREW THOMAS CURTIS and RAKSHIT KUMAR

Department of Computing and Software
McMaster University, Hamilton, Canada

(3 October 2006)

A heuristic design method for rapid volumetric magnetic resonance imaging data acquisition trajectories is presented, using a series of second order cone optimization subproblems. Other researchers have considered non-raster data collection trajectories and under-sampled data patterns. This work demonstrates that much higher rates of under-sampling are possible with an asymmetric set of trajectories, with very little loss in resolution, but the addition of noise-like artifacts. The proposed data collection trajectory, Durga, further minimizes collection time by incorporating short, un-refocussed excitation pulses, resulting in above 98 percent collection efficiency for balanced steady state free precession imaging. The optimization subproblems are novel, in that they incorporate all requirements, including data collection (coverage), physicality (device limits) and signal generation (zeroth and higher moment properties) in a single convex problem, which allows the resulting trajectories to exhibit a higher collection efficiency than any existing trajectory design.

Keywords: Nonuniform Fourier transform; Magnetic resonance imaging; Volumetric imaging; SOCP; k-space trajectory optimization

1 Introduction

Reducing imaging time in MRI is driven by the desire to

- (i) reduce patient discomfort by reducing time in the magnet and shortening or eliminating breath-hold times,
- (ii) capture dynamic processes, such as the beating of the heart and peristaltic motion in the abdomen,
- (iii) reduce cost for the health-care system by increasing throughput, and decreasing delays caused by motion artifacted images.

*Corresponding author: Email: anandc@mcmaster.ca

Rapid imaging has been an active research topic for the last twenty years, beginning with (Mansfield *et al.* 1976). Research in this area can be roughly divided into the discovery and exploitation of

- (i) efficient data collection strategies: echo planar (Mansfield 1977) and spiral imaging trajectories, (Irarrazabal and Nishimura 1995), (Barger *et al.* 2002)
- (ii) efficient signal generation techniques: multiple echoes per excitation (Henning *et al.* 1986), steady-state imaging (Oppelt *et al.* 1986)
- (iii) relationships between signals from geometrically different antennae: SMASH (Sodickson and Manning 1997), SENSE (Pruessmann *et al.* 1999), collectively called parallel imaging.

This paper is a contribution to the search for efficient collection strategies for volumetric imaging, specifically (i) a reformulation of the accepted design criteria, (ii) a series of convex optimization problems incorporating all of the criteria as proxy constraints, and (iii) a specific collection strategy for steady-state imaging designed according to these criteria.

Parallel imaging with non-regular sampling is very expensive computationally, and still the subject of active research, (Rosenfeld 2002). One motivation for the present research was to provide non-regular data collection strategies which would result in better-conditioned inverse problems and faster iterative methods.

In section 2 the basic image reconstruction problem in magnetic resonance imaging is described, in order to establish notation and motivate the fitness criteria presented in the next section. Design constraints derived from these informal criteria in section 3 form the basis for the convex second order cone optimization problems. The implementation details are briefly described in section 5, and interpretation of the results, including a comparison of the results against previously-published work is given in section 6. In the final two sections, some open questions and possible new approaches are reviewed, and the results are summarized.

1.1 *Related work*

In her MSc thesis (Ren 2005), Ren presents a method of optimizing planar data collection by incorporating velocity insensitivity constraints into Teardrop (Anand *et al.* 2001) designs. Velocity insensitive data collection results in (approximately) the same signal, independent of the tissue velocity. Nayak *et al.* incorporate such constraints into a rewriter following spiral data collection (Nayak *et al.* 2005). Teardrop is a natural evolution of spiral data collection, which is an evolution of EPI, the first fast collection method. Spiral strategies, named for the shape of the sampling trajectory, work because, in the

plane, rotations about the origin are a group with one generator, which allows such trajectories to be efficiently packed. Others have tried to generalize such strategies to three dimensions, but there can be no equally nice analogue to the two dimensional case because here the group of rotations does not have a single generator.

Stacked spirals (Irrarrazabal and Nishimura 1995), (Thedens *et al.* 1999), cones (Gurney *et al.* 2006), and similar strategies are designed to maintain a maximum separation between local pieces of the trajectory, based on the dictates of the Nyquist sampling criteria. Computationally, lack of a nice symmetry would result in a much larger optimization problem with quadratic growth in constraints and exponential growth in solution time.

In this paper, a strategy is presented for volumetric imaging, which, instead of lamenting this lack of simplifying symmetry, looks instead to the ‘blessing’ of dimensionality: random curves in \mathbb{R}^3 have near zero probability of intersection. Instead of struggling with optimization problems with unmanageable numbers of constraints, they are pared down to a minimum, relying on the fact that the trajectories will not intersect on average unless constrained to do so. The resulting randomness means that Nyquist’s criteria does not apply, which allows significant under-sampling to be achieved. This is not the first paper to propose using randomness to combat aliasing, see for example (Nayak and Nishimura 1998).

2 Magnetic resonance imaging

In magnetic resonance imaging, one measures radio-frequency magnetic fields created by the resonances of one or more nuclei in the object, usually hydrogen (mostly water) in people. Measurable resonance occurs because the object is placed in a large homogeneous field, and excited by the momentary application of oscillating transverse fields. For a very readable and complete account of how signals are created and the complications which arise, see (Haacke *et al.* 1999). MR signals are collected using devices akin to radio antennae, commonly called *coils*. The measurements we make with these coils are not localized, but contain contributions from every nuclei in the object. Ignoring nonuniformity in the coil, and signal propagation through the sample, the signal is the sum of the magnetic fields produced by each nucleus. By working in a rotating frame of reference close to the resonant frequency, both relative frequencies and phases are encoded into complex valued signals.

Geometric encoding is achieved by inducing transient linear variations (referred to as gradients) in strength on the homogeneous field. Linear variations in field produce linear variations in resonant frequency, which over time create linear phase variations. If $\rho : \mathbb{R}^3 \rightarrow \mathbb{C}$ is the original transverse magnetic field,

the new field will be $\exp(i\langle x, k \rangle)\rho(x)$ where $x \in \mathbb{R}^3$ and $k \in \mathbb{R}^{3*}$, the element of the dual space corresponding to the accumulated phase. It follows that the measured signal

$$s(t) = \int_{\mathbb{R}^3} e^{i\langle x, k(t) \rangle} \rho(x) dx,$$

is a sampling of the Fourier Transform of the object's original magnetization. For any given trajectory $k(t)$, there is a linear transformation $\text{Map}(\mathbb{R}^3, \mathbb{C}) \rightarrow \text{Map}(\mathbb{R}, \mathbb{C})$, and if it is invertible, reconstruction of the original magnetization from the measurements is possible.

Early MR image reconstruction was constrained by the cost of computation, and focussed on making data better fit existing hardware and software Fourier transforms. Data collection was forced to be regular, and sampled on rectangular grids (first in two and later in three dimensions). Even the first image reconstructions based on non-trivial inverse problems, *e.g.* phase conjugate symmetry (Margosian *et al.* 1986), assumed regular rectangular data sampling, as did the first parallel imaging schemes. Regular sampling is an approximation of an object in an infinite-dimensional function space by a vector in a finite-dimensional vector space. All approximations introduce errors, and if sufficient breadth of sampling in k -space is not made, the reconstructed image, being the sum of only low-spatial-frequency Fourier basis functions, will lack fine detail. (How much detail is required depends on the application, but there can never be too much.) If the samples are too widely separated, the smallest common period of all the basis functions used to reconstruct the image will be too small, resulting in aliasing: the appearance of signal where it does not belong.

For non-regular sampling, both problems still occur, but the inversely-linear relationship between defects in sampling, and defects in the image does not hold. In practice, hardware sampling rates have been higher than the sampling rates dictated by object size, resulting in under-sampling gaps only appearing between trajectories. For (interleaved) spiral imaging, under-sampling as a result of excessive pitch of the spiral results in spiral artifacts in the image beginning at a radius which grows inversely with the gap between successive spiral arcs. These artifacts are also aliased signal, but their source is not recognizable, unlike the aliasing observed with regular rectangular sampling.

This simple picture is inexact. Design tolerances, manufacturing defects, and patient-dependent effects combine to produce

- (i) distortions of the designed sampling trajectory $k(t)$,
- (ii) deviations in the resonance frequency,
- (iii) signal variation over time.

In fact, imaging can be considered a four-, five- or six-dimensional problem, if one takes into account time of acquisition, velocity of the tissue (blood flows), and chemical resonance offset (between water, fat and other organic molecules). Practical implementations account for these factors during image reconstruction. A common approach to calibration and correction is to measure some data several times (most commonly, the lowest spatial frequencies, called the centre of k -space). Tissue velocity can be measured in this way, but also compensated for by arranging for the first and higher moments of the waveform $k(t)$ to be zero (Nishimura *et al.* 1986, Haacke 1987). Both of these approaches are incorporated.

3 Fitness

Fitness of sampling trajectories is universally judged first with reference to the *point spread function*, ρ_{psf} , which is the image reconstructed from the data generated by a single delta function in image space. Other measures incorporating properties of tissues, the excitation pulse, *etc.* are beyond the scope of this paper, and it is safe to say that problems in the psf will manifest themselves in any other measure of fitness. Ideally, the psf would itself be a delta function, but this is impossible given finitely many samples. The two most important features are the height and width of the central peak, and relative height and structure outside the peak. Since the final image will be, to a first approximation, the convolution of the true image with the psf, broadness of the central peak manifests as blurriness in the image, and nonzero values outside the peak as unwanted signal. Structure in the psf outside the peak will be repeated in the image, and since the human visual system is attuned to structure this is distracting in itself. Structure is associated with larger extreme values outside the peak, which could obscure fainter features in the image. Resolution can be defined to be the width of the central peak when it reaches half its peak height (called FWHM, full-width at half max), although this is open to interpretation for non-monotone and asymmetrical central peaks.

Since using the psf directly to define an objective or constraints would be very expensive, the following heuristics are used: the union of the loci of the trajectories should visit all parts of k -space within a ball $\{k : |k| < 1/\text{resolution}\}$, and should not display visible symmetries, because these will manifest as symmetries in the spurious signal, which is usually observed in structured aliasing which contains larger extreme values. Rather than trying to minimize the maximum gap between trajectories, which corresponds to the usual interpretation of the Nyquist Theorem (Ren 2005), the present method seeks to create an irregular distribution of sampling voids while avoiding very large gaps.

Specifically for the application to multi-coil reconstructions, the method

seeks psfs in which the phase of the unwanted signal is uniformly distributed in a pseudo-random manner, as this will result in the maximum cancellation of aliasing noise from different receivers. The authors believe that the resulting trajectory design will lead to faster convergence rates for iterative SENSE reconstructions, but an analysis of this is beyond the scope of this paper. Even for single-coil reconstructions, such psfs will cause aliasing noise cancellation for noise associated with large-scale features, although it will not help for small, bright features, such as blood vessels enhanced by contrast agents.

The initial design using this methodology targets steady-state imaging, so the total time for each trajectory is short. For comparison with the most efficient published trajectory, a repetition time $T_R = 5.6\text{ms}$ was chosen.

To be able to calibrate and correct for machine and patient-dependent effects, most of the trajectories are constrained to pass through $k = 0$, the point corresponding to no linear phase modulation across the imaging volume.

4 Convex subproblems

To reconstruct a volume, one or more trajectories through k -space are sampled, the data is resampled onto a rectangular, regular lattice, which is then transformed using a Fast Fourier Transform. For each individual trajectory, several parametrized convex, second order cone optimization subproblems were developed. The subproblems are solved multiple times with different parameters, and in different orders. The composition strategy for subproblems is entirely heuristic, based on the above analysis of properties of trajectories which should be expected to lead to fit psfs. It is based on refinement on two levels: the heuristics have themselves evolved based on numerous tests, and the choice of parameters and order for the subproblems is tuned to complement properties of previously designed trajectories.

4.1 Variables

Changes in linear phase variation are induced by linear field variations, which result from currents in electromagnets (so called gradient coils) designed for this purpose. The currents are in turn driven by the gradient amplifiers, which are controlled by digital electronics. So the fundamental physical variables are the piecewise-constant control functions, which are parametrized by lists of heights and widths. The gradient strengths correspond to the differences in position in k -space. Knowing the positions in k -space is therefore equivalent to knowing the gradients, and more convenient when formulating the constraints.

For a single subproblem, m , the variables are $k_{i,m}$, and the values of $k_{i,m'}$, $m' \neq m$ are either fixed or unavailable. To handle soft (penalized constraints)

$\delta_t > 0$	for the size of the time-step,
M	number of trajectories, and
N_m	number of time steps of trajectory $m \in \{0, \dots, M\}$,
G_{\max}	peak gradient strength (see (1)),
S_{\max}	peak slew rate (see (2)),
R	transform size (see (3)),
κ	half excitation duration (see (5)),
λ_0	penalty scale for zero targets (see (6)),
λ_{goal}	penalty scale for boundary sphere targets (see (6)),
λ_{nulling}	penalty scale for the first moment (see (4)),

$k_{i,m} \in \mathbb{R}^3$	$\forall m \in \{1, \dots, M\}, \forall i \in \{1, \dots, N_m\}$ - discrete k -space positions
$\tau > 0$	penalty variable

Table 1. List of all parameters and variables.

a variable is introduced for their violation τ . For reference, all parameters and variables are collected in Table 1.

4.2 Universal Constraints

Both soft and hard constraints are used. Hard constraints must be satisfied by feasible points, whereas the violation of soft constraints is penalized, but still allowed. All of the hard constraints are common to all subproblems. Constraints for readout trajectories are created to be rotatable, as these trajectories can be used to acquire data for volumes in any orientation. Less strict constraints could be formulated if the relative orientation of the rectangular imaging volume and the principal axes of the gradient coils is fixed, and the constraints on the individual gradient amplifiers are independent. In this paper only the freely rotatable case is examined.

Peak constraints: Gradient amplifiers have peak current limits which restrict the maximum absolute value of the gradient waveform amplitude. These limits can be expressed as inequality range constraints on each of the $n + 1$ points in the discrete waveform sequence as

$$\|k_{i+1,m} - k_{i,m}\|_2 \leq G_{\max}, \quad i \in \{1, \dots, N_m - 1\}, \quad (1)$$

where G_{\max} is the maximum allowable gradient amplitude.

Slew constraints: Gradient amplifiers also have limits on slew rate or rate of change of amplitude. This can be approximated as an inequality constraint on the first-order differences between adjacent discrete points as

$$\|k_{i+2,m} - 2k_{i+1,m} + k_{i,m}\|_2 \leq S_{\max}, \quad i \in \{1, \dots, N_m - 2\}, \quad (2)$$

Transform size constraints: To insure that all data is usable without increasing the computational complexity of the reconstruction algorithm, trajectories are constrained to visit only the part of k -space that will be used in a Fast Fourier transform. For rotatable trajectories, this requires that the trajectories be contained in a ball of radius equal to the resolution of the reconstructed volume. So if reconstructed voxels will be $1\text{mm} \times 1\text{mm} \times 1\text{mm}$, trajectories must remain inside the ball $\{k : |k| < 1\text{mm}^{-1}\}$:

$$\|k_{j,m}\|_2 \leq R \quad j \in \{1, \dots, N_m\}, \quad (3)$$

where the radius, R , is the resolution in m^{-1} . In practice, a slightly smaller ball (target) is desirable, so that even corrected trajectories will be inside the larger ball. This is easily done simply by adjusting R .

First moment nulling: To make the readout gradient motion-insensitive, the first moment is required to be zero,

$$\left\| \sum_{i=1}^{N_m-1} i(k_{i+1,m} - k_{i,m}) \right\| < \lambda_{\text{nulling}} \tau \quad (4)$$

Involving all the variables, this is a global constraint. This has two consequences: (i) it cannot be effectively optimized by *local* or *greedy* algorithms which consider one point at a time, and (ii) it introduces dense columns/rows into the concrete optimization problem, and warrants extra consideration if using a solver which can take advantage of sparsity. Motion insensitivity means that the phase of the magnetization measured before and after the readout will be identical for tissues which are moving at a constant velocity. Constraints to null higher moments can easily be introduced, which will prevent pulsatile flow from modifying the magnetization. Sensitivity to motion effects all MR experiments, and is in fact used to quantify flow in phase contrast angiography, but it is generally unwelcome otherwise.

Such motion compensation is necessary in the type of imaging targeted by this technique. In the balanced steady-state imaging, magnetization is not dephased (destroyed) from one readout to the next, but is modified with new RF pulses, allowing errors to build up over time, resulting in more pronounced motion artifacts.

Endpoints: Trajectories must begin and end somewhere. Balanced steady-state imaging must begin and end at the centre of k -space. Although data collection need not begin and end there, anything else is less efficient. Conventional pulse sequences also contain magnetic field gradients for the purpose of defining a slice profile. The details of this are beyond the scope of this paper, and interested readers should consult a basic textbook such as (Haacke *et al.* 1999). The endpoints of the trajectories are constrained by

$$\begin{aligned} k_{1,m} &= (0, 0, -\kappa) \\ k_{2,m} &= (0, 0, \kappa) \\ k_{N_m-2,m} &= (0, 0, -\kappa) \\ k_{N_m-1,m} &= (0, 0, \kappa) \end{aligned} \quad \forall m \in \{1, \dots, M\}, \quad (5)$$

where $\kappa > 0$ is a parameter determined by the length of the excitation. For the present application, $\kappa = G_{\max}/2$ is used, as it simplifies testing of the algorithm.

Remark: The endpoints are chosen in conjunction with the design of an excitation radio-frequency pulse. This is another optimization problem outside the scope of this paper. The choice of endpoints is consistent with a so-called hard pulse, which excites the entire volume.

4.3 Second order cone subproblem, I

Additional constraints are added depending on the type of trajectory to be designed. For the first type, constraints to insure coverage of all parts of k -space, and repeated traversal of $k = 0$. The idea is that by forcing the trajectory to visit points distributed on the boundary, and visit the origin in between, other parts of k -space will be covered too.

We define a distribution of points on the boundary sphere of radius R . Together, these points cover the sphere without large gaps. Each of these points $c_{j,m}$ will be the goal of a designated point, $i_{j,m}$, on a designated trajectory, m . The number of goals per trajectory depends on the duration of the trajectories. All results in this paper refer to five goals per trajectory, with three goals on the boundary sphere, and two (the second and fifth) at the origin. Goals are encoded into constraints

$$\|k_{i_{m,j}} - c_{j,m}\|_2 < \begin{cases} \lambda_0\tau & c_{j,m} = 0 \\ \lambda_{\text{goal}}\tau & \text{otherwise} \end{cases} \quad \forall (j, m). \quad (6)$$

Different penalty parameters are used for the points pulled towards the bound-




Figure 1. Three individual trajectories, oriented so positive z is vertical. The two left trajectories are stage I trajectories, with the crossing of two points at $k = 0$ clearly visible. The right trajectory was optimized with relaxed $k = 0$ penalties, and does not go through $k = 0$, and has no other crossings.

ary sphere and the points pulled to the origin, as trajectories must get very close to the origin if they are to be useful for calibration, whereas points on the boundary sphere are sparse, and there is little to be gained by reaching them exactly. The best method of choosing goals is an open question. In this paper, a simple pseudo-random distribution of vertices from a triangulation of the boundary is used, but the rest of the algorithm is independent of this choice.

4.4 *Second order cone subproblem, II*

It was observed that solutions of these problems were of variable quality: some trajectories passed close to their goals, and some did not (dependent mainly on the path-length between the pseudo-random target points). Since the heuristic objective is to pass through $k = 0$ often, not on *every* trajectory, a second stage is introduced, in which we chose some of the trajectories with the worst (highest) objective value, and re-optimized them with a lower value for λ_0 . This has the effect of requiring these trajectories to pass near $k = 0$ but not through it, as visible in the right-most trajectory in figure 1.

4.5 *Second order cone subproblem, III*

Although elimination of symmetries was a goal in the design of the trajectories, the choice of initial and final points along the z axis creates a preferred direction, and as such, lower k -space coverage is observed along the $z = 0$




Figure 2. Two views of the extra ten trajectories which thread through points of low density in a thickened annulus. On the left, one sees that the trajectories move initially to different sectors but follow the annulus around in the same direction. On the right, from the top view, one sees that the density is concentrated in an annulus, and that the trajectories are quite round from this point of view.

plane, which is apparent in the point-spread-function. To correct this, an additional set of trajectories are added, designed to fill holes in the sampling pattern near the $z = 0$ plane.

To find the holes, the density of the stage I and II trajectories was computed. To calculate the density, the control points $k_{i,m}$ are interpolated to a large number of equally-spaced samples. A delta function was associated with each sample in k -space, and summed over all samples. One then convolves the sum of these delta functions with a positive, continuous function, and resamples the resulting continuous function on a regular, rectangular lattice. In the experiments, 32^3 lattices are used to save time by reducing the memory footprint of the resampling operation. This is equivalent to resampling the data corresponding to the image with a single delta function at the origin in image space (the point where the gradient coils have no effect). The resampling kernel of choice is defined in (Anand *et al.* 2004), because it is convenient to reuse the same code used to reconstruct images. To find ‘holes’ near $z = 0$, and to put them into suitable trajectories, this method restricts its attention to points in the thickened annulus:

$$E = \left\{ 6 \leq \sqrt{x^2 + y^2} \leq 12, |z| \leq 5 \right\}.$$

Dividing these points into six sectors, E_1, \dots, E_6 ,

$$E_i = \{(x, y, z) \in E, -\pi + 2\pi(i-1)/6 < \text{atan2}(y, x) - \pi + 2\pi i/6\},$$

it tries to construct the longest trajectory (m) which circumnavigates the annulus visiting the point with the lowest density in each sector using the




Figure 3. Three views of complete sets of trajectories. On the left, the red trajectories are designed in stage I, and the green are the redesigned relaxed trajectories. In the middle, the extra annular trajectories are added in cyan. On the right the full set of trajectories is shown, one colour per trajectory.

following greedy algorithm:

- (i) Start with the optimization problem given by the universal constraints. (Note that it is really a feasibility, not an optimization problem at this stage.)
- (ii) Pick the point in $e \in E_j$, with the lowest density, and the next control point $k_{i,m}$.
- (iii) Augment the optimization problem by adding the constraint

$$\|k_{i,m} - e\|_2 < \tau.$$

If the solution has a value of τ below a fixed threshold, keep this constraint in the problem and pick another pair in step 2, with j and i increased by one. If not, and $i = N_m - 3$, then remove this constraint from the problem, and use the solution to the smaller problem as the trajectory. Otherwise, return to step 2 with the same e but i increased by one.

To eliminate uneven sampling, the process starts in a different sector E_i for each trajectory, repeating after all sectors have played this role. The different shape of these trajectories is visible in figure 2.

In figure 3, all of the different types of trajectories are assembled and displayed, with different colouring to show that the stage I and II trajectories in red and green are similar, without any observable geometric relationship. The extra trajectories in turquoise, however, are visibly concentrated near a plane. In the final view of figure 3, one observes that the bounding sphere constraint is tight at some points on some trajectories.

5 Implementation

In principle, the hard work is in formulating an approximation to the trajectory design problem, which consists of convex, second order cone optimization (sub)problems (SOCP/SOCO): a class of optimization problem to which efficient interior point methods are known to apply (Lobo *et al.* 1998, Boyd and Vandenberghe 2004). Interpreting the numerical results of the subproblems as trajectories is difficult. Initial attempts at formulating convex subproblems were hampered by the lack of a good visualization of the solutions, see (Ren 2005). Subsequently, undergraduate students developed a highly interactive visualizer (Matsumoto and Castura 2005) for trajectories, using OpenGL and Cocoa Widgets. The visualizer was then modified by the third author to include a user interface for modifying model parameters and to integrate two solvers. Initially, the solver IPOPT (Wächter 2002) was used because it accepts quite general non-linear optimization problems allowing flexibility in the model. It was unable to solve simple problems designed to test its performance on the types of constraints that were planned. (Note that IPOPT is a very powerful solver that has been used successfully for many other problems.)

Primal-dual solvers specialized for SOCO constraints were examined. The only open-source solver easily called from C is `socp.c` (Lobo *et al.* 1998), written to demonstrate applications of SOCP. Although the authors warn that it is not being maintained and offers limited performance relative to other solvers, it was found that it solved the test problems with hundreds of variables. Its main limitations are: the C api requires that the user provide a feasible starting point, errors in input data are not detected and reported to the user in a useful way, dependent variables must be eliminated, and it does not use sparsity, so it will not perform as well on large sparse problems as it could.

5.1 SOCP

A second order cone programming problem, as implemented in `socp.c`, is defined as

$$\begin{aligned} \min f^T x \\ \text{s.t. } \|A_i x + b_i\| \leq c_i^T x + d_i, i = 1, \dots, L. \end{aligned} \quad (7)$$

(Equivalent, more abstract definitions are more common.) The optimization variable is the vector $x \in \mathbf{R}^m$. The problem data are $f \in \mathbf{R}^m$, and, for $i = 1, 2, \dots, L$; $A_i \in \mathbf{R}^{N_i \times m}$, $b_i \in \mathbf{R}^{N_i}$ and $d_i \in \mathbf{R}$. The norm appearing in the

constraints is the Euclidean norm i.e. $\|v\| = (v^T v)^{1/2}$. The constraint,

$$\|A_i x + b_i\| \leq c_i^T x + d_i$$

is called a *second order cone constraint of dimension N_i* . Such constraints bound a cone with elliptical cross section, hence the name, but are often used to bound the interior of a sphere, which frequently occurs as a norm constraint on a vector. All of the universal constraints have this form. The penalized soft constraints have the more general form. The SOCP (7) is a convex programming problem since the objective is a convex function and the constraints define a convex set, so each of the subproblems has a unique connected minimal set. This guarantees a unique minimum. In practice it is always helpful to start a search at a feasible point (a trajectory satisfying the constraints). For the solver used, this is a required measure, and a method adapted to our problem is given in the appendix.

A complete stage I subproblem is given below, represented in standard form.

For a given design subproblem for trajectory m , the variables are $\tau \in \mathbb{R}$, $k_{i,m} \in \mathbb{R}^3$ for $i \in \{3, \dots, N_m - 2\}$, where the initial and final positions $k_{i,m}$ are constant and thus excluded. The problem is

$$\begin{aligned}
 & \min \tau \\
 & \text{s.t. } \|k_{i_m,j,m} - 0\| \leq \lambda_0 \tau, & \forall j \in \{2, 5\} \\
 & \|k_{i_m,j,m} - c_j^m\| \leq \lambda_{\text{goal}} \tau, & \forall j \in \{1, 3, 4\} \\
 & \left\| \sum_{i=1}^{n-1} i(k_{i+1,m} - k_{i,m}) \right\| \leq \lambda_{\text{nulling}} \tau, & (8) \\
 & \|k_{i+2,m} - 2k_{i+1,m} + k_{i,m}\| \leq S_{\text{max}} & \forall i \in \{1, \dots, N_m - 2\} \\
 & \|k_{i+1,m} - k_{i,m}\| \leq G_{\text{max}} & i \in \{1, \dots, N_m - 1\} \\
 & \|k_{i,m}\| \leq R & i \in \{1, \dots, N_m\}.
 \end{aligned}$$

The only component of the objective is the penalty term, so if the variables are ordered

$$x = (\tau, x_3, y_3, z_3, x_4, y_4, \dots, z_{N_m-2}), \quad (9)$$

where $k_{i,m} = (x_i, y_i, z_i)$, then $f = (1, 0, 0, \dots)$. The cones in these problems are in one-to-one correspondence with the constraints, all of which have dimension 3. The cone corresponding to moment nulling is the only dense constraint. The rows of the other blocks have between one and three nonzeros. Since `socp.c`

is not sparsity-aware, the given block organization does not exploit sparse structure. Care must be taken for the peak and slew constraints abutting the end points, which have to be treated separately, since they contain $k_{i,j}$ values which are constant, and hence don't contribute to the A -block, but rather to the b -block, which is otherwise mostly zero.

At every state in development, it was extremely helpful to have an integrated visualization tool. The implementation details of the tool are not interesting, but the features may be:

- (i) trajectories drawn as cylindrical tubes,
- (ii) a user interface for trajectory coloration and selection,
- (iii) coloured control points on trajectories,
- (iv) spheres for goals, and the outer sphere,
- (v) transparency,
- (vi) cut-planes,
- (vii) partial trajectory display,
- (viii) density display (based on selected trajectories),
- (ix) psf display (based on selected trajectories).

6 Performance

Three surrogate measures of design performance were used prior to validation of our design: duty cycle, point spread function (psf), and frame rate. As concerns the solver, it has only been observed that it runs faster with large tolerances and it does not scale well to larger problem sizes. Scaling would improve by using a solver which leveraged sparsity, because computation is bound by solving the linear system for a Newton step direction.

All results are reported with the following optimization parameters, which were chosen to facilitate comparison with the work of other authors:

# basic trajectories	54
# relaxed trajectories	10
# threaded trajectories	10
time between control points	0.1ms
duration of gradient	5.5ms
peak gradient	40 mT/m
max slew per control point	150 T/m/s
target maximum resolution	1mm ⁻¹

.1ms is used as a discrete time step, and is a convenient unit for exposition and analysis.

6.1 Optimizer Performance

The quality of the solution depends on the coverage of k -space. Coverage in k -space will not be significantly altered by small changes in the individual trajectories. Thus the tolerance in solving the SOCP can be set quite high while heavily weighting the moment-nulling constraint violation, and somewhat heavily weighting the zero-crossing constraint. As a result, if the trajectories are anywhere near the outer goals, the first moment will be within measurement error of zero. Increasing both the relative and absolute duality-gap tolerances used for termination from 10^{-6} and 10^{-4} (suggested defaults) to $1/10$ was found to reduce the execution time for each subproblem by a factor of two, from four seconds to two seconds per subproblem, on average. (Times refer to measurements on a 2.5GHz PowerMac G5.) Total optimization time for the first two stages, including set-up and display was under a minute. Total optimization time for the density threading stage was half an hour. Neither time is reasonable for on-line optimization in a clinical setting while a patient is in the MRI machine, but both are reasonable for one-time optimization to set up a protocol.

After 10 trajectories were recalculated with relaxed penalties on the $k = 0$ goals, the maximum resolution in k -space sampled by each trajectory averaged 0.940mm^{-1} , with the maximum being 1.003mm^{-1} and the minimum 0.824mm^{-1} . The worst-case distance between the selected point and it's goal was between 0.042mm^{-1} and 0.534mm^{-1} , with an average of 0.263mm^{-1} . Changes in the choice of goals for a single trajectory and flexibility in the point pulled to each goal could improve these numbers, but of many experimental adjustments, the ones which made a noticeable difference to the psf were the relaxations and the additional density threading trajectories.

6.2 Time efficiency comparison

The simplest efficiency comparison is sampling duty-cycle, calculated by dividing the total time data is collected by the repetition time for the pulse sequence. The same k -space trajectory will result in different duty-cycles when used with different pulse sequences, although a better trajectory in one usage is usually better in most usages. Durga is designed to be used with volumetric, fully-balanced pulses, specifically balanced steady-state free precession (bSSFP).

In figure 4, pulse sequence diagrams are given for an efficient waveform design for planar bSSFP imaging, see (Anand *et al.* 2001), and for the Durga waveform design for volumetric imaging. The first line shows the envelop of the transmitted radio-frequency pulse used to excite the magnetization. (In a steady state experiment, many repetitions are required to set up the steady state.) No data can be collected during this time on a conventional imager,




Figure 4. Pulse sequence diagrams for an optimized Teardrop planar bSSFP readout on the left, and for a single Durga trajectory on the right. The left diagram is a screen capture from the console of a Picker Medical Systems Infineon scanner, and the right is a plot generated with gnuplot. The trace order is the same in both cases (rf, slice, phase and read gradients, and data acquisition state).

although in the more controlled environment of an NMR spectrometer. If the rf pulse occurs when all the gradients are zero (no current), then the pulse excites magnetization uniformly across the imaging volume. In Teardrop and Durga, the rf pulse corresponds to the period where the slice gradient (the second trace) is nonzero. Such combinations are called selective excitations, because they excite the volume differently at different slice positions, with the profile of the excitation well-approximated by the Fourier transform of the rf pulse envelope.

In the planar case, time is required to ramp up and down the slice gradient, because during data collection, no phase variation in the slice direction (transverse to the slab being measured) is desired. For the pulse used, phase variation in the slice direction is created by the main lobe, so negative lobes must surround the excitation lobe. After 1.47ms, excitation is complete, and the Teardrop readout waveform begins, lasting 3.43ms. Data collection is indicated by the step function in the bottom trace. The trajectory is the integral of the read and phase gradients, which correspond to x and y in our notation.

The displayed Teardrop sequence does not have zero first moment, so flowing tissue will not reconstruct properly. In her master’s thesis, (Ren 2005), Ting Ting Ren developed an optimization model and sequential SOCP implementation capable of incorporating this and other global constraints, and found that doing so resulted in a 3 percent decrease in resolution.

On the right, the Durga pulse sequence design is quite different. Being designed for volumetric imaging, it is necessary to dephase in the slice direction at some point to be able to resolve the z direction. To save time, the sequence utilizes dephasing (linear phase variation, $\exp(iz)$, of the magnetization) which

	Teardrop	HNC	HNC-imp	Durga
readout (ms)	3.43	2.40	2.40	5.50
excitation (ms)	1.47	1.20	1.20	0.10
rewinder (ms)	-	1.40	1.40	-
TR (ms)	4.93	5.00	5.90	5.60
duty-cycle (fraction)	0.78	0.48	0.41	0.98

Table 2. Duty-cycle calculations for three balanced k -space sampling patterns, designed for bSSFP imaging. The HNC pattern uses extra time for a rewinder to balance the first moment, the other strategies incorporate this as a constraint on the readout trajectory. The repeat time (TR) used for Teardrop is the actual time, including dead time required by the scan controller to switch modes. For HNC, both the ideal and implemented time with overhead are given, in separate columns. For Durga, no measurements have been made, so no dead time is added. Durga’s efficiency comes from what is not there: time to rewind the first moment, and time to rewind the slice excitation.

would be caused by any finite-time pulse, and starts sampling at $k = (0, 0, c_0)$, finishing sampling at $k = (0, 0, -c_0)$, thereby removing the need for the negative slice gradient lobes for rewinding. In figure 1, this behaviour is visible as a constant gap between the endpoints of the trajectory. In the pulse sequence diagram, it is visible as a short period with zero x - and y -gradients, and maximum value for the z -gradient.

Some authors refer to experiments as ‘volumetric’ when the excitation pulse is not accompanied by any gradient activity. Here it is used in the more general sense of exciting a large volume uniformly, which is possible (albeit more difficult) if gradients are also active. An energy-limited pulse has been designed which excites a volume 50cm across with uniform magnitude and linearly-varying phase.

Here, Durga and Teardrop are compared to one of the most efficient alternatives in the literature, the Hargreaves-Nishimura-Conolly (HNC) design (Hargreaves *et al.* 2004), spiral pulse with extra rewinder to zero the first two moments. In Table 2, the times for data collection, excitation, and the rewinder (where necessary) are calculated. Unfortunately, on real imagers there are always some switching delays when going from rf transmit (excitation) to rf receive (data collection), when loading waveform data, etc. For Teardrop, this makes a small difference in efficiency, but for HNC it is larger, so the ideal and the implemented repetition times are split into two columns for comparison. Although Teardrop and HNC are designed for planar imaging, they can both be adapted to volume imaging with very little extra cost (Hargreaves *et al.* 2004). Teardrop gains in efficiency over HNC by incorporating first-moment nulling as a constraint on the readout trajectory, and avoiding the time lost for the rewinder. Durga has the same advantage, and by being designed only for (large) volumetric imaging, it also saves time taken by slice rewinders, resulting in the best overall efficiency.

Duty-cycle is not the only factor influencing efficiency, if this is taken to be

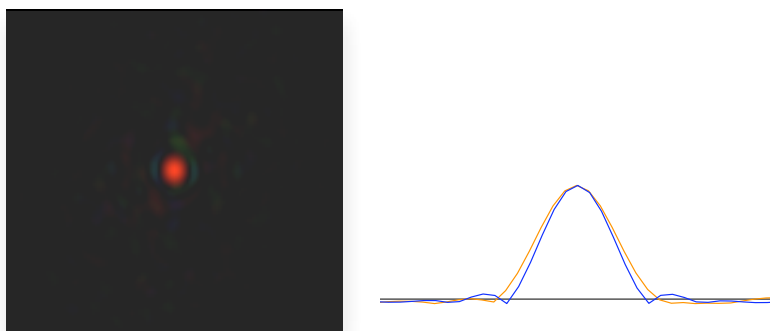


Figure 5. Close-up of the psf restricted to $x = 0$. In the image, hue is phase and brightness is magnitude. Red is positive real, and turquoise is negative real. Brightness has been lightened to make the background visible. On the right, cross sections through the centre are plotted, with the y axis in red and the z axis in blue.

information gained per unit time. The sampling pattern in k -space, and the signal to noise ratio of the sampled data are also important. The effectiveness of the sampling pattern is largely captured by the quality of the psf, and Durga performs well in this respect, as is demonstrated below. Signal to noise ratios are in turn dependent on other factors, most importantly on the type of pulse sequence (the sequencing of excitations with different energies and phases, gradient waveforms, and data collection) and the volume of tissue excited. Most noise sources are not dependent on the amount of tissue excited, so the signal to noise ratio (in the measurements) is approximately proportional to the volume excited, so Durga will benefit from being volumetric. The signal to noise ratio in the reconstructed image will also depend inversely on the volume of the voxels.

6.3 Point Spread Function

The point spread function (psf) is the image which would be constructed if a delta function were being scanned. If the image reconstruction process is linear, then the measured image will be the convolution of the psf with the true image. In practice, the approximation is good, and its failings are well understood, if not always easy to correct. Figure 5 shows the $x = 0$ slice of the psf obtained using the 54 trajectories after stage II, and plots of the cross sections along the axes. The central peak in the psf is not symmetric, with the y (vertical) axis being longer than the z (horizontal) axis. This would produce more blurring in the y direction (and, similarly, the x direction which is not shown) than in the z direction. The contour at half the peak height has a radius of 1mm. When enlarged, the first aliasing ring has a height at most 8




Figure 6. Comparison of 54 trajectories on the left and 54 + 10 density-filling trajectories on the right. Part of an $x - y$ cross section of a numerical phantom containing cubes with different phases is displayed, with cross-sectional plots in both the x and y directions. In the images, brightness corresponds to magnitude, and hue corresponds to phase. In the cross-sectional graphs, magnitude is black, and real and imaginary parts are blue and red, respectively. The effect of the extra trajectories on the resolution in the $x - y$ plane is visible in two ways: in the image, the squares on the right are clearly more square, and in the graphs, the definition and separation of the two objects is clear, especially when looking at the magnitudes, and looking at the inflection point on the real-part of the x cross section between the objects (the x sections are the adjacent plots between the images).

percent of the height of the central peak. Other aliasing, which will appear as noise, decays quickly with some points of height 2 percent.

In figure 6, we show an $x-y$ cross section of a numerical simulation of two cubes. The cubes are meant to have sharp edges, and the simulation shows that the edges are the psf predicted. On the right, the same reconstruction is done including data from the 10 extra trajectories, and one can easily see that the resolution in this plane is increased by their addition, although with an increase in apparent noise.

The Durga acquisition scheme aims to trade scan time for apparent resulting noise. Regular trajectories, whether sampled on lattices, spirals or other nested shapes have lower limits on acquisition time imposed by the Nyquist sampling theorem. If this doesn't hold for Durga, then one would expect that as the density of sampling (*i.e.* number of trajectories) decreases, the apparent noise will increase, but nothing regular like aliasing ghosts or spiral artifacts appear. This is what we observe in comparing 1 trajectory to 54 trajectories in figure 7.

Although adding extra trajectories can reduce asymmetry in the central peak, there is still asymmetry which may contribute to distortion of reconstructed images, see figure 8, but overall, the psf promises sharp images (figure 9).




Figure 7. A comparison of psfs corresponding to the first trajectory and all 54 trajectories designed in the first phase. The corresponding (sets of) trajectories are shown on the left.




Figure 8. Three planar cross sections of the psf: $x-y$, $x-z$, and $y-z$, enlarged by zero padding (using a larger Fourier transform than necessary).

Not visible in in figure 9, due to Fourier transform size limitations, there is a region of structured noise in the psf at a radius of 30cm. The largest magnitude values observed after examining several different narrow transforms (*i.e.* 2048×64) was half a percent of the peak value. The structured ringing is not created by the irregular trajectories but rather from the regular sampling along the trajectory: sampling once per $2\mu\text{s}$ (500kHz), Nyquist aliasing should occur at

$$\frac{1}{40\text{mT/m} \cdot 2 \cdot 10^{-6} \text{ s}} = 0.298\text{m}, \tag{10}$$

i.e. at a radius of thirty centimetres. Since the apparent noise is small, this




Figure 9. Grayscale surface plot of psf for 64 trajectories.

would be hard to distinguish from other sources of noise in imaging, but it does point to the need for megahertz sampling in modern MR imagers.

6.4 *Frame Rate*

In applications to dynamic imaging (beating heart, flowing blood), the most important factor is frame rate, the rate at which data for an entire plane or volume can be collected. The authors of the missile-guided trajectory (MG) compared six methods of volumetric imaging, using machine constraints of 10mT/m and 30mT/m/ms. Comparison of the published numbers with the trajectories presented here is accomplished by a five-fold time-dilation.

Table 3 repeats the better reported results from (Mir *et al.* 2004) adding the equivalent numbers for Durga. Durga is five times faster than the nearest competitor, and a comparison of the psfs indicate Durga would display similar amounts of blurring, but larger amounts of apparent noise, as a result of the lower sampling rate. It is unclear how MG would perform under similar under-sampling conditions, since it is irregular as Durga, but is designed with very different objectives, and doesn't include velocity insensitivity as a constraint. So for spoiled scans (in which the signal is zeroed and regenerated from scratch for each repetition), MG may perform as well as Durga and we encourage the authors to test it under extreme under-sampling.

There are other undersampled acquisitions, most notably undersampled radial trajectories using Projection Reconstruction and different types of data-sharing. Du *et al.* (Du *et al.* 2004) report a frame rate of 0.21 frames per second for (PR)Hyper-TRICKS, which is roughly comparable to $1/(64 \times 0.0056s) = 2.79$ frames per second for Durga (see table 2). After scaling for different gradient performance, this is in the range reported for the fully-sampled trajectories in table 3. This comparison must be put into context: (PR)Hyper-TRICKS is being used for sequence types to which Durga is not applicable, at least as currently designed, and calibration and artifact reduction for projection

Trajectory	Total number of shots	Scan time (s)
Cylindrical EPI	537	27.0
Stack of Spirals	215	11.0
Missile Guidance	357	18.0
Durga	64	1.8

Table 3. Scan times for different rapid volumetric trajectories. See (Mir *et al.* 2004).

reconstruction is much simpler and more effective than general non-raster reconstruction, which must be used with Durga.

7 Future Work

Development of this method continues the following list of priorities:

- (i) find more efficient ways of threading through points of lower sampling density
- (ii) compare performance of this model with an LP model for the problem of designing trajectories which are not rotatable
- (iii) incorporate sparsity into the solver, or benchmark on another solver
- (iv) add constraints to the current design problems to further improve the psf
- (v) try multiple assignments of goals to trajectories, searching for a set of assignments which produces a better psf
- (vi) reduce and compensate for gradient waveform distortion by incorporating it into the model

While it has been stated that a goal for this method is the ability to trade off noise for reduced scan time, the current work concentrates on the trade going one way: towards a minimum scan time producing an acceptable image. The performance of this approach remains untested when longer scan times are possible. Given that existing methods work well in this case, and longer scan times also mean longer design times, this end of the spectrum is less compelling.

Finally, for those who want both shorter scan time *and* high image quality, future work will incorporate iterative SENSE (Pruessmann *et al.* 2001) image reconstruction techniques, which seek to take advantage of differences in the geometric encoding of different receiving antennae.

8 Summary

A set of small second order cone problems has been formulated, which taken together design a very efficient set of k -space trajectories for sampling volu-

metric MRI data. Very high frame rates – above what the Nyquist sampling theorem indicates – are possible because the sampling is irregular. The amount of blurring, as shown by analysis of the point spread function, is comparable to trajectories an order of magnitude slower. Some of this efficiency comes from the integration of both first-moment nulling and slice rewinding into the readout trajectories, allowing 98 percent sampling efficiency, as measured by duty cycle. Some of this efficiency comes at the price of artifacts which appear to be background noise. There is some structure to the apparent background noise, but given the reduction in sample time, the defects are minor.

Acknowledgements

The authors thank Mark Haacke, Paul Margosian, Michael Noseworthy, Michael Thompson, Dee Wu, Ian Young and Yuri Zinchenko, for research suggestions and comments on the manuscript. Additionally, the authors thank NSERC, CFI and OIT for research support.

Appendix A: Pseudo-random goal generation

To find well-distributed points on the boundary sphere:

- (i) Start with a triangulation of one face of the hexahedron formed by joining two triangular pyramids at the bases, with an edge in the plane $y = 0$. Without loss of generality, assume the bases on the plane $z = 0$ and the apex of the pyramid at $(0, 0, R)$, scaling the pyramid if necessary, so that the bases have radius R . Here one triangle is used, and different initial triangulations would result in different numbers of goals on the boundary sphere.
- (ii) Use mid-point subdivision to generate more and more dense triangulations of this face, until the desired number of points is reached, and project all of the points radially onto the boundary sphere. By rotating this set through $\pm 2\pi/3$ radians about the z axis, and reflecting in the $x - y$ plane, a reasonably uniform covering of the boundary sphere is generated. To achieve ‘pseudo-random’ trajectories, the goals need to be ‘mixed up’. For instance, the method used by the displayed trajectories:
- (iii) Eliminate points on $z = 0$ and $y = 0$, since these points overlap with other faces.
- (iv) Sort the points lexicographically by z and then y co-ordinates.
- (v) Create triplets of points, discarding extras. Triplets correspond to trajectories.
- (vi) From the m th triplet, (S_1, S_2, S_3) , form goals

- $c_{1,m} = S_1,$
- $c_{2,m} = 0,$
- $c_{3,m}$ = rotation of S_2 by $2\pi/3$ about z -axis and reflection in $z = 0,$
- $c_{4,m}$ = rotation of S_2 by $-2\pi/3$ about z -axis,
- $c_{5,m} = 0.$
- (vii) Triplicate each of these sets of goals by rotating each set by $2\pi/3$ and $-2\pi/3,$ respectively.
- (viii) Duplicate all of the above trajectories by reflecting the trajectories through $z = 0$ and swap the first and third goals in the resulting sets, that is the first two nonzero goals.

There are now $6n_1$ trajectories, where n_1 is the number of triples formed from points in the triangulation of one equilateral triangular face of the hexahedron. For example, in numerical tests a single triangle is subdivided 3 times to produce 45 points, 28 without the two edges, which result in 9 triplets, and 54 sets of goals. Two examples of such trajectories are show in figure 1.

Note that if the swapping of two of the goals for the inverted sets was omitted, the problems would be pairwise symmetric, and the resulting trajectories would intersect pairwise in the plane $z = 0.$ This reduces coverage of k -space, and introduces symmetry which is observable as a reduction in quality of the psf, similar to the type of reduction observable for other reasons in figures 5 and 6.

Appendix B: Finding feasible starting points

Since the C interface to SOCP requires a feasible starting point, the following algorithm is necessary. Note that this depends on the particular form of the optimization problem accepted by the solver.

Being a primal-dual solver, the actual problem solved by SOCP includes the dual of our optimization problem: The dual of (7) is given by

$$\begin{aligned}
 \max \quad & - \sum_{i=1}^L (b_i^T z_i + d_i w_i) \\
 \text{s.t.} \quad & \sum_{i=1}^L (A_i^T z_i + c_i w_i) = f \\
 & \|z_i\| \leq w_i, \quad i = 1, \dots, L.
 \end{aligned} \tag{11}$$

The dual optimization variables are $z_i \in \mathbf{R}^{N_i-1}, w_i \in \mathbf{R}.$ Form the vector of $z = (z_i, i = 1, \dots, L).$ The form of the dual problem is important, because

in combination with knowledge of the sparsity of every constraint other than moment nulling, simple, efficient means of computing dual-feasible points can be found.

To construct the initial (primal) feasible problem, the shape constraints are modified from hard to soft:

$$\begin{aligned} \|k_{i+2,m} - 2k_{i+1,m} + k_{i,m}\| &\leq S_{\max}/2 + \tau \\ \|k_{i+1,m} - k_{i,m}\| &\leq G_{\max}/2 + \tau \end{aligned} \quad (12)$$

which requires only a few changes in elements of matrices defining the problem. For the ranges of parameters of interest, the solution always satisfies $\tau < \min\{G_{\max}/2, S_{\max}/2\}$, ensuring a feasible point for the original constraints. The relaxed problem has a simple feasible starting point given by taking all k values to be zero. Based on the (very sparse) block structure of our problem, it is simple to find feasible initial values of the dual variables z and w by solving small linear systems for z , and using a suitably large vector w_i in the null-space of c_i .

Even if an infeasible solver were used, starting with good primal-dual feasible points saves time.

The C solver interface is actually a bit simpler. Block matrices are formed from A_i and c_i ; from b_i and d_i ; and from z_i and w_i . Each one of these contains multiple blocks, one block per cone.

REFERENCES

- Anand C., Thompson M., Wu D. and Cull T. Teardrop, a novel trajectory for TrueFISP. In *ISMRM April 2001 Conference Proceedings, Glasgow, 2001* p. 1804.
- Anand C.K., Terlaky T. and Wang B. Rapid, embeddable design method for spiral magnetic resonance image reconstruction resampling kernels. *Optimization and Engineering*, 2004. **5**(4), 485–502.
- Barger A.V., Block W.F., Toropov Y., Grist T.M. and Mistretta C.A. Time-resolved contrast-enhanced imaging with isotropic resolution and broad coverage using an undersampled 3D projection trajectory. *Magnetic Resonance in Medicine*, 2002. **48**(2), 297–305.
- Boyd S. and Vandenberghe L. *Convex Optimization* (Cambridge University Press), 2004.
- Du J., Thornton F.J., Fain S.B., Korosec F.R., Browning F., Grist T.M. and Mistretta C.A. Artifact reduction in undersampled projection reconstruction MRI of the peripheral vessels using selective excitation. *Magnetic Resonance in Medicine*, 2004. **51**(5), 1071–1076.

- Gurney P.T., Hargreaves B.A. and Nishimura D.G. Design and analysis of a practical 3D cones trajectory. *Magnetic Resonance in Medicine*, 2006. **55**(3), 575–582.
- Haacke E.M. Improving mr image quality in the presence of motion by using rephasing of gradients. *American Journal of Roentgenology.*, 1987. **148**, 1251–1258.
- Haacke E.M., Brown R.W., Thomson M. and Venkatesan R. *Magnetic Resonance Imaging. Physical Principles and Sequence Design* (Wiley-Liss (John Wiley & Sons), New York), 1999.
- Hargreaves B.A., Nishimura D.G. and Conolly S.M. Time-optimal multidimensional gradient waveform design for rapid imaging. *Magnetic Resonance in Medicine*, 2004. **51**(1), 81–92.
- Hennig J., Nauwerth A. and Friedburg H. Rare imaging: a fast imaging method for clinical mr. *Magnetic Resonance in Medicine*, 1986. **3**, 823–833.
- Irrazabal P. and Nishimura D. Fast three-dimensional magnetic resonance imaging. *Magnetic Resonance in Medicine*, 1995. p. 33.
- Lobo M.S., Vandenberghe L., Boyd S. and Lebret H. Applications of second-order cone programming. *Linear Algebra Appl.*, 1998. **284**(1-3), 193–228. ILAS Symposium on Fast Algorithms for Control, Signals and Image Processing (Winnipeg, MB, 1997).
- Mansfield P. Multi-planar image formation using nmr spin echoes. *Journal Physics C*, 1977. **10**(3), L55–L58.
- Mansfield P., Maudsley A.A. and Baines T. Fast scan proton density imaging by nmr. *J.Phys.E:Scient.Instrum.*, 1976. **9**, 271.
- Margosian P., Schmitt F. and Purdy D.E. Faster mr imaging: Imaging with half the data. *Health Care Instr.*, 1986. **1**, 194.
- Matsumoto G. and Castura J. 3d visualiser for trajectories. BSc Thesis Project, 2005.
- Mir R., Guesalaga A., Spiniak J., Guarini M. and Irrazaval P. Fast three-dimensional k-space trajectory design using missile guidance ideas. *Magnetic Resonance in Medicine*, 2004. **52**(2), 329–336.
- Nayak K.S., Hargreaves B.A., Hu B.S., Nishimura D.G., Pauly J.M. and Meyer C.H. Spiral balanced steady-state free precession cardiac imaging. *Magnetic Resonance in Medicine*, 2005. **53**(6), 1468–1473.
- Nayak K.S. and Nishimura D.G. Randomized trajectories for reduced aliasing artifact. In *ISMRM 1998 Conference Proceedings*, 1998 p. 670.
- Nishimura D.G., Macovski A. and Pauly J. Magnetic resonance angiography. *IEEE Transactions on Medical Imaging*, 1986. **5**(3), 140–151.
- Oppelt A., Graumann R., Barfuss H., Fischer H., Hartl W. and Shajor W. Fisp: a new fast mri sequence. *Electromedica*, 1986. **54**, 15–18.
- Pruessmann K.P., Weiger M., Börnert P. and Boesiger P. Advances in sensitivity encoding with arbitrary k-space trajectories. *Magnetic Resonance*

- in Medicine*, 2001. **46**, 638–651.
- Pruessmann K.P., Weiger M., Scheidegger M.B. and Boesiger P. Sense: Sensitivity encoding for fast mri. *Magnetic Resonance in Medicine*, 1999. **42**, 952–962.
- Ren T. *An Optimal Design Method for MRI Teardrop Gradient Waveforms*. Master’s thesis, McMaster University, 2005.
- Rosenfeld D. New approach to gridding using regularization and estimation theory. *Magnetic Resonance in Medicine*, 2002. **48**, 193–202.
- Sodickson D.K. and Manning W.J. Simultaneous acquisition of spatial harmonics (smash): fast imaging with radiofrequency coil arrays. *Magnetic Resonance in Medicine*, 1997. **38**(4), 591–603.
- Thekens D.R., Irarrazaval P., Sachs T.S., Meyer C.H. and Nishimura D.G. Fast magnetic resonance coronary angiography with a three-dimensional stack of spirals trajectory. *Magnetic Resonance in Medicine*, 1999. **41**(6), 1170–1179.
- Wächter A. *An Interior Point Algorithm for Large-Scale Nonlinear Optimization with Applications in Process Engineering*. Ph.D. thesis, Carnegie Mellon University, 2002.