# McMaster University

## Advanced Optimization Laboratory


**Title:**

Energy-Constrained Pulse Design
for MRI and NMR

**Author:**

Christopher Kumar Anand and Andrew Thomas Curtis

**AdvOL-Report No. 2007/12**

July, 2007
Hamilton, Ontario, Canada

# ENERGY-CONSTRAINED PULSE DESIGN FOR MRI AND NMR

CHRISTOPHER KUMAR ANAND AND ANDREW THOMAS CURTIS
DEPARTMENT OF COMPUTING AND SOFTWARE
MCMASTER UNIVERSITY, HAMILTON, CANADA
EMAIL: ANANDC@MCMASTER.CA

ABSTRACT. Optimizing radio frequency (*rf*) pulses is of interest in both Magnetic Resonance Imaging and Nuclear Magnetic Resonance Spectroscopy. Recent research has focussed on reducing *rf* energy for imaging, and improving excitation profile uniformity for spectroscopy. This paper summarizes the different approaches used to date, and presents an approach using continuous optimization that utilizes second derivative information of the objective function, with two examples of novel, sub-millisecond pulse designs: an energy-constrained pulse for steady state imaging, and a linear-phase, broadband 90 degree pulse for NMR.

The benefits of using second derivatives of the objective function are examined, and a method is presented for efficient computation.

All optimization is done using the open-source optimization package IPOPT. Computationally efficient integration of the Bloch equations and derivative calculations are performed using code symbolically generated by Maple.

*Keywords:*radio-frequency pulse; magnetic resonance imaging; optimal control; IPOPT; SAR; specific absorption rate; slice selection; broadband pulse

## 1. INTRODUCTION

Both Magnetic Resonance Imaging (MRI) and Nuclear Magnetic Resonance (NMR) Spectroscopy measure chemical properties and geometric structure by exciting nuclear spins into higher-energy states and measuring the resulting magnetic signal.

Both techniques use magnetic fields oscillating in resonance with nuclei for excitation. The nuclear spins can be modelled using quantum mechanics, with a good classical approximation, known as the Bloch equations. At technically feasible field strengths, the oscillations are in the radio-frequency range, and the excitation waveforms are called *RF* pulses.

The classical approximation is sufficient for understanding MRI, and sufficient for designing common excitation pulses in NMR. As MRI goes to higher fields, the energy required by a given pulse increases, to the point where patient comfort and safety concerns limit the energy and hence throughput of the examination. In NMR, the span of the resonant frequencies of different nuclei increases with field strength. This is desirable because it makes it easier to separate signals from different nuclei.

It is thus useful to implement a framework that can take into account goals for excitation profile, bandwidth, and total and peak pulse energy. The examples given in this paper are targeted towards ultra-short (temporally) pulses, in the sub millisecond regime.

*Date*: 24 July 2007.

Several different alternatives are currently available for designing *rf* pulses, as sketched below. Since the problem is non-convex, constrained, and involves solving a family of systems of ordinary differential equations—the Bloch equations—it is not surprising that of the many approaches, none has been completely successful, with each method addressing different design criteria.

## 2. Related Work / Alternative Design Methods

Inverse scattering is a method of transforming algebraic data (e.g. polynomials) into solutions of differential equations. With forward and inverse transforms it is possible to quickly design pulses with discretized responses. Shinnar and Le Roux introduced this method to MR [Pauly *et al.* 1991], and it is by-far the most common method of designing pulses in MRI. This approach has been generalized by making the inverse scattering explicit, see [Magland 2005] for a recent example.

These methods are very fast, but cannot effectively address constraints on *rf* pulse amplitude, or situations in which the exact excitation profile is not important, only some properties of it. The first is a result of not treating the discretized pulse envelope as the design variable.

In both imaging and spectroscopy, hardware limits on total and instantaneous *rf* power are imposed in order to prevent meltdown of amplifiers and probes, and to limit energy deposition in the object being scanned. The VERSE pulse design concept [Conolly *et al.* 1986a] modifies pulses designed using other means (principally SLR) to reduce power by scaling/dilating portions of the *rf* and accompanying gradient waveform to produce the same excitation using less energy or less time. This approach has been used successfully to produce very short pulses for steady state imaging [Hargreaves *et al.* 2004].

The first use of basic optimization (in the low-energy limit) was made in 1973 [Tomlinson and Hill 1973]. Numerical and analytical work continued in this direction even prior to a relatively complete formulation of the optimization problem in MRI in 1986 by Conolly [Conolly *et al.* 1986b].

For some problems, quite a bit of physical insight, combined with other inverse-scattering or compositional methods are required to find a useful starting point for a local optimizer. Most approaches involve pattern search or gradient-based methods [Hargreaves *et al.* 2004, Skinner *et al.* 2004], and approximate numerical integration of the systems of Bloch equations. Not surprisingly, significant numerical/computational issues arise in the form of unacceptable execution times and convergence bottlenecks.

Previous formulations of this problem for the optimal control package SOCS(see [Betts 2001]), resulted in interesting waveform designs when starting with a zero initial waveform [Anand *et al.* 2006], but this general approach very quickly reached memory and computation limits.

One approach to reducing numerical cost is to use a mathematical description of the solution space adapted to the problem. This increases the fidelity and resolution of the solution while keeping the cost constrained. Chebyshev polynomials have been used to approximate *rf* waveforms, leading to better convergence [Ulloa *et al.* 2004], using SLR pulses as starting points.

By far the most successful approach in the design of constrained pulses involves setting up a series of constraints for the approximate integration of the Bloch equations at time points corresponding to a discretization of the control waveform, and then explicitly integrating both the magnetization (solution of the Bloch equations) and the Lagrange multipliers [Skinner *et al.* 2004, Skinner *et al.* 2003]. The approach of Skinner *et al.* and the approach we will present are more efficient than the aforementioned, as they exploit the structure of the problem. At this time, it is not clear how the two approaches could be combined.

## 3. PROBLEM

The variables $b_i \in \mathbb{R}$ determine piecewise constant functions $b(t)$. In MRI, these are the common design variables. It is also possible to co-design gradient waveforms, see [Conolly *et al.* 1986b, Anand *et al.* 2006], but for very short pulses, there is no benefit, so we omit it here. In NMR, pulses are almost always complex valued. This introduces additional freedom, but also increases the number of (real) design variables, and introduces quadratic constraints.

To define the objective we need to define intermediate variables for the final magnetization which would result from the current *rf* pulse $b(t)$, $M(s,t) \in \mathbb{R}^3$, for discrete values of $s$ determined by the particular problem and $t$ by the choice of discretization for the control waveform. The objective is then to minimize

$$F(b) = \sum_{s \in S} \left| M(s,T) - M_{\text{target}}(T) \right|^2, \tag{1}$$

at the end of the pulse (time $T$), where the sum is taken over a finite set of frequency offsets (corresponding to slice positions in the imaging case). There are no constraints on $S$, such as equal density or equal coverage. In different applications, we are free to choose the most appropriate set $S$.

The magnetization Bloch equations:

$$\frac{dM}{dt} = M \times B + \frac{1}{T_1}(M_0 - M_z) - \frac{1}{T_2}(M - M_z). \tag{2}$$

where

$$B = \begin{pmatrix} b(t) \\ 0 \\ sG(t) \end{pmatrix}, \tag{3}$$

and $G(t)$ is a constant in the NMR case or a pre-determined piecewise-constant gradient profile in the MRI case, which for implementation issues is usually also a constant, and will be a constant in all examples in this paper. In the complex case, $B = (\Re b(t), \Im b(t), sG(t))$.

The constants $T_1$ and $T_2$ are relaxation times, and for all the pulses in this paper are large enough that they can be ignored.

Path Constraints. The controls are constrained by

$$|b(t)| < M_{\text{point}}, \ \forall t \in U. \tag{4}$$

which correspond to safe operating ranges for the amplifiers.

NMR and MRI experiments are composed of multiple pulses. To simplify the exposition of this paper, we only consider the design of so-called excitation pulses,

that is, pulses designed to operate on an equilibrium initial magnetization

$$M(s,0) = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}, \forall s \in S. \tag{5}$$

3.1. **Problem.** To summarize we want to solve

$$\min \quad \sum_{s \in S} \left| M(s,T) - M_{\text{target}}(T) \right|^2, \tag{6}$$

$$\text{subject to} \quad \int_{t \in [0,T]} b(t)^2 dt < M_{\text{total}}, \tag{7}$$

$$|b(t)| < M_{\text{point}}, \tag{8}$$

$$M \text{ solves Bloch equations (2).} \tag{9}$$

## 4. Exact Integration, Gradients

The key to efficiently solving this problem is that given piecewise-constant controls, the Bloch equations can be analytically integrated. Moreover, they can be integrated for a basis of initial conditions at the start of a piecewise constant segment of $b(t)$ to another basis at the end of the segment. Let the segments be indexed by $i = 0, 1, 2, ...$, and let the linear transformation given by integrating the basis of initial conditions over interval $i$ be $E_i \in GL(R^3)$.

Using this notation, the final magnetization is

$$M(T) = E_{T-1} E_{T-2} \cdots E_i \cdots E_2 E_1 E_0 M(0), \tag{10}$$

where we abuse notation and use $T$ for the number of intervals, and the final time. The only transformation which depends on on $b(t)$ and $G(t)$ in interval $i$ is $E_i$, so the first and second derivatives are

$$\frac{\partial M(T)}{\partial b_i} = E_{T-1} E_{T-2} \cdots \frac{\partial E_i}{\partial b_i} \cdots E_2 E_1 E_0 M(0), \tag{11}$$

$$\frac{\partial^2 M(T)}{\partial b_i \, \partial b_j} = \begin{cases} E_{T-1} E_{T-2} \cdots \frac{\partial E_i}{\partial b_i} \cdots E_s \cdots \frac{\partial E_i}{\partial b_j} \cdots E_0 M(0) & i < j \\ E_{T-1} E_{T-2} \cdots \frac{\partial^2 E_i}{\partial b_i^2} \cdots E_2 E_1 E_0 M(0) & i = j, \end{cases} \tag{12}$$

where we can assume, without loss of generality, that $i \leq j$. So calculating one partial derivatives of $M(T)$ with respect to a control using (11) requires $N_T$ matrix products, just as the calculation of $M(T)$. But many of the products are repeated.

By calculating

$$U_i = E_{T-1} E_{T-2} \cdots E_{i+1} E_i \quad \in GL(R^3), \tag{13}$$

$$V_i = U(t)^{-1}, \tag{14}$$

$$W_i = E_{i-1} \cdots E_2 E_1 E_0 M(0) \quad \in \mathbb{R}^3, \quad \text{and} \tag{15}$$

once for all $t$, starting from the left for $U$ and $V$ and the right for $W$, we can do this using $N_T$ matrix-vector products and $2N_T$ matrix-matrix products. Calculate all

the partial derivatives in the gradient as

$$\frac{\partial M(T)}{\partial b_i} = U(t+1)\frac{\partial E_i}{\partial b_i}W_i, \tag{16}$$

$$\frac{\partial^2 M(T)}{\partial b_i^2} = U(t+1)\frac{\partial^2 E_i}{\partial b_i^2}W_i, \qquad \text{and} \tag{17}$$

$$\frac{\partial^2 M(T)}{\partial b_i\,\partial b_j} = U(t+1)\frac{\partial E_i}{\partial b_i}V_i U_{j+1}\frac{\partial E_i}{\partial b_j}W_i. \tag{18}$$

(With similar expressions for the derivatives with respect to the gradient waveform, if it is not held constant.) This can be performed efficiently by abutting two loops. The first loop goes backward in time, calculating $E_i$, $E_i^{-1}$, the partial derivatives of $E_i$, $U_i$, and $V_i$. The second goes forward in time, calculating $M_i$ and $W_i$, and calculating and storing the products of $W_i$ and the two partial derivatives.

The calculation of $E_t$, its inverse and partials involves many repeated calculations (including expensive sine and cosine evaluations). Automatic differentiation[1] and symbolic code generation can both reduce the effort required to produce such code. Symbolic code generation can significantly reduce the redundant computation via common sub expression identification at the symbolic level, see [Anand *et al.* 2005)] for an application to the analysis of time series with exponential components (*e.g.*, free induction decays in NMR). Symbolic code generation allows better control over access to memory, it is also much simpler than writing the code in the first place, if you have access to either Maple or Mathematica. It works very well for systems based on the Bloch equations; for more complicated systems of equations, *e.g.*, coupled spin systems, it is not always able to integrate over finite time steps [Anand *et al.* 2006], but for short pulses, the larger systems can be decomposed into Bloch equation blocks.

## 5. METHODS

Maple PDEtools and codegen packages were used for the analytic solution of the objective function, which was then converted to C code for fast execution. The code is available by contacting the first author, and quite readable to people familiar with Maple.

The peak rf-power limit for the MRI pulse designs was determined empirically in all cases by designing a pulse for a specified flip angle and trying to image a fixed, vendor-supplied phantom. Since the scan console normally adjusts *rf* pulses and gradients just before imaging (including the effect of calibrating the *rf* power to flip angle ratio for the object/antenna combination being imaged) reported scan times and flip angles were checked before and after scanning.

To verify the pulse designs, gradient echo volumes were collected on a GE 3.0T short-bore Excite-II scanner (GE Healthcare, Milwaukee, WI, USA), and compared with images acquired using manufacturer-supplied 1ms excitation pulses.

The optimization was performed using the open-source nonlinear constrained optimization software, Interior Point OPTimizer (IPOPT), [Wächter 2002].

For the examples in this paper, time steps for discretization of $b(t)$ were taken to be $2\mu$s for MRI, and $0.5\mu$s for NMR. Optimization time ranged from seconds to

[1]For a list of tools for using automatic differentiation and a list of introductory and research publications, see http://www.autodiff.org.

minutes on a 2.16GHz Core2Duo MacBook Pro, with roughly half the time spent in the Maple-generated derivative and function evaluation, and half in the optimizer. For test purposes, we used IPOPT to optimize different pulse design problems, with and without the second derivative information. Without second derivatives, the design times could reach several hours for similar levels of convergence.

## 6. Integration into Optimizer

Before integrating the derivative calculations into an IPOPT call-back function, the Maple-generated code was tested for correctness and performance. Maple integration code was slightly faster than hand-optimized C code, and calculating the gradient at the same time was less than four times the cost of the integration alone over a wide range of problem sizes. Calculating the Hessian scaled quadratically with problem size, as expected. Constructing an IPOPT call-back function involves: choosing an order for the variables, setting up index arrays, and using the chain rule to calculate two derivatives given the derivatives of $M(T)$, the final magnetization on each controlled offset frequency $s \in S$.

## 7. Results

MRI: Limited Peak and Total Energy. Two design problems are reported:

The first problem is to design a very short (100$\mu$s) broad-band excitation pulse for imaging, alongside a maximum slice gradient. Short pulses are of importance in imaging for many reasons, in this case for increasing duty cycles and reducing energy deposition. Only the target in-slice magnetization (slice thickness 32cm) is required to be uniform, with total energy constrained to be $1/2$ of what a comparable Gaussian pulse would require. It is unusual to specify a large gradient for use with a wide slice excitation, but this is useful in connection with high-efficiency $k$-space sampling [Anand *et al.* 2007]. A flip angle of 9 degrees was chosen to demonstrate constrained optimization, without triggering dilation of the *rf* pulse by the scan console. To prevent numerical problems, a constant nonzero ($10^{-6}$ of maximum) numerical value was used as an initial $b(t)$.

Optimization produced a sinc-like pulse bounded in the centre by the peak power constraint, and bounded overall by the total instantaneous power. For comparison purposes, the pulse in Fig. 1 is compared with the narrowest square-wave pulse with a 9 degree flip at the centre of the slice, and peak power under the limit. The shows the pulse waveforms, and the right the slice profile, including slice positions beyond the designed 16cm radius. Since the profile is symmetric, only one side is shown.

In Fig. 2 volume cross sections of the acquired phantom data were generated, showing the slice direction along the vertical axis. The uniformity is similar, with the optimized pulse exhibiting aliasing on the top and bottom of the phantom outside the reconstructed field of view. This shows that it is a broad-band pulse.

NMR: With and Without Second Derivatives. In the second example, a linear-phase broadband $\pi/2$ excitation pulse for NMR is designed, using machine parameters (resonance offset $\pm$20kHz, peak *rf* power of 17.5kHz, and *rf* power scaling errors of $\pm$5%) from [Skinner *et al.* 2004] but only 125$\mu$s length pulses.

Two $\pm$20kHz linear-phase broadband pulses were designed using IPOPT with identical constraints, one using first derivatives and pseudo-Hessians estimated

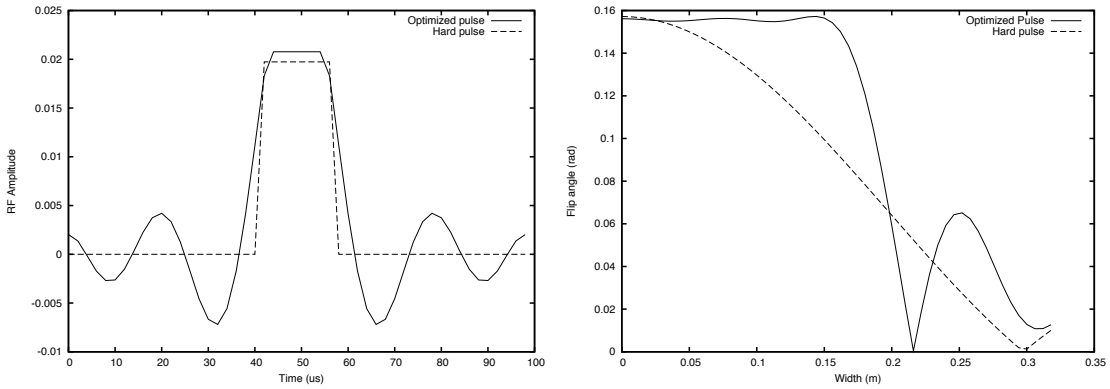


FIGURE 1. *Left:* Pulse profiles for optimized and square pulse. *Right:* Simulated slice profiles.

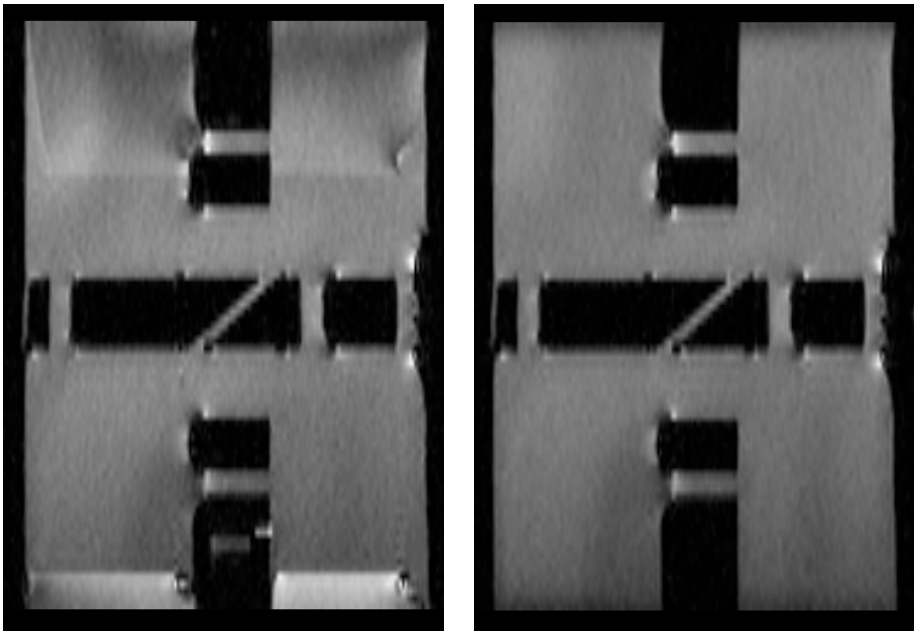

FIGURE 2. *Left:* Multi-planar reconstruction of the optimized pulse. *Right:* Multi-planar reconstruction of a standard 1ms pulse.

by IPOPT, and one with the full Hessian of second derivatives calculated as described above. Fig. 3 shows the objective function (in this case measuring essentially the least-squares deviation from the desired linear-phase excitation profile, since minimizing energy was given a very low weight), and the resulting pulse. The current software supports only real-valued pulses (no complex phase), so only the real-part is shown. The reduction of the objective shows that after 20 iterations, the optimization with second derivatives has found a high-quality pulse. The solid line ends at 250 iterations when numerical criteria used by IPOPT detect convergence. The dotted line falls much more slowly to a good value, frequently increases, indicating a failure to find improvements in the pulse which do not violate constraints, and ends after 400 iterations when IPOPT reports constraint violations which cannot be remedied. Both pulses are symmetric, which can be expected because the desired excitation profile is conjugate-symmetric, but

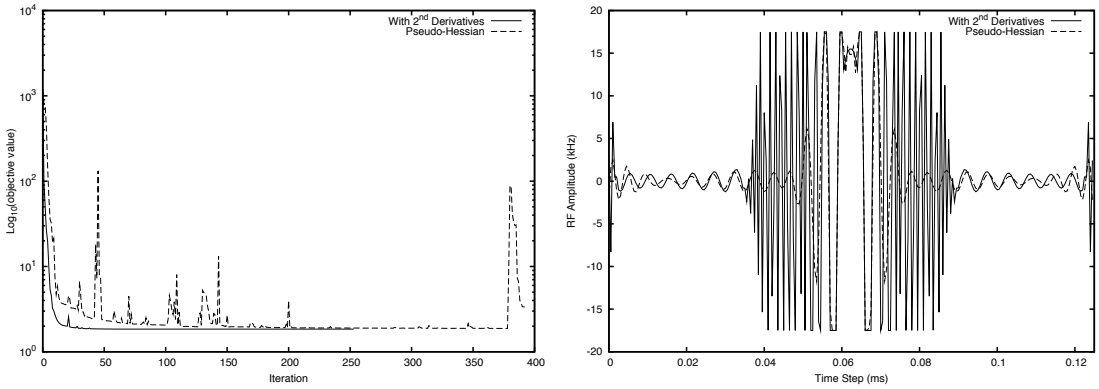


FIGURE 3. *Left:* Objective as a function of iteration, showing much faster convergence when second derivatives are used. *Right:* Pulse profiles. The pulse generated with full Hessian information reaches a higher energy (energy is unconstrained) but spends much more time at peak. Oscillations are within amplifier slew limits.

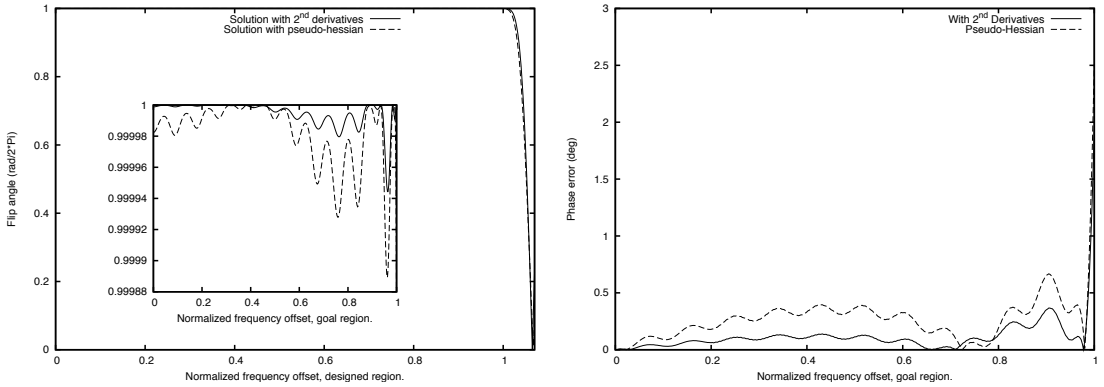


FIGURE 4. *Left:* Inset: Excitation profile over frequency range (0 to 20kHz). Full: Same profile at full scale showing falloff outside of design range. *Right:* Deviation from linearity of resulting excitation phase across the design range. In all cases we see the full Hessian perform better.

the full-Hessian design does use much more of the available peak energy. As a result, the excitation response in Fig. 4 is much closer to the ideal. The left hand plot shows the magnitude of the $x - y$ component with a scaled version inset. The right hand plot shows very low deviation from the desired linear phase.

Both relaxing the overall energy limit and reducing the width of the pulse sampling from $2\mu$s to $0.5\mu$s combine to allow designs much closer to the desired profile even at 90 degrees versus the 9 degree low-energy pulse designed for MRI.

## 8. DISCUSSION

In addition to two novel pulse designs which solve current problems in NMR and MRI, we have shown that using symbolic code generation leads to a compact optimization problem whose solution is easily within reach of the open-source

solver IPOPT. Having exact second derivatives significantly improves performance, and will allow this approach to scale to much larger problem sizes. One draw-back of using IPOPT is that sparse matrix arithmetic is always used, even though our Hessians are dense.

In addition to designing and testing variations of these pulses, we are exploring more involved pulse designs which are far more sensitive to good starting points (inversion pulses, pulse trains, *etc.*), as well as complex-valued pulse designs.

## ACKNOWLEDGMENTS

The authors thank NSERC, CFI and OIT for research support, Alex Bain, Mike Noseworthy and Mark Haacke for helpful discussions about limitations of current clinical and NMR pulse designs, and Paul Polak for pulse programming and data acquisition.

## REFERENCES

[Anand *et al.* 2006] Anand C.K., Bain A.D. and Nie Z. The simulation and optimization of NMR experiments using a Liouville space method. In I.S. Kotsireas, editor, *Maple Conference Proceedings*, 2006 pp. 203–216.

[Anand *et al.* 2005)] Anand C.K., Carette J., Curtis A.T. and Miller D. Cog-pets: Code generation for parameter estimation in time series. In I.S. Kotsireas, editor, *Maple Conference*, 2005 pp. 198–212.

[Anand *et al.* 2007] Anand C.K., Curtis A.T. and Kumar R. Durga: A heuristically- optimized data collection strategy for volumetric magnetic resonance imaging, 25 pages, to be published in *Engineering Optimization* submitted October 2006, accepted March 2007.

[Anand *et al.* 2006] Anand C.K., Stoyan S. and Terlaky T. Energy-Minimizing Slice-Select Pulses for MRI , 10 pages, accepted March 2007, to appear in H.G. Bock, E. Kostina, H.X. Phu, R. Ran- nacher (eds), *Modelling, Simulation and Optimization of Complex Processes. Proceedings of the International Conference on High Performance Scientific Computing*, March 6-10, 2006, Hanoi, Vietnam. Springer 2008.

[Betts 2001] Betts J.T. *Practical Methods for Optimal Control Using Nonlinear Programming* (Society for Industrial and Applied Mathematics, Philadelphia, PA, USA), 2001.

[Conolly *et al.* 1986a] Conolly S.M., Glover G., Nishimura D. and Macovski A. Variable-rate selective excitation. *Journal of Magnetic Resonance*, 1986a. **78**, 440–458.

[Conolly *et al.* 1986b] Conolly S.M., Nishimura D. and Macovski A. Optimal control solutions to the magnetic resonance selective excitation problem. *IEEE Trans. on Medical Imaging*, 1986b. **3**, 106–115.

[Hargreaves *et al.* 2004] Hargreaves B.A., Cunningham C.H., Nishimura D.G. and Conolly S.M. Time optimal verse excitation for 3d balanced ssfp imaging. In *ISMRM 2004*, 2004 p. 260.

[Magland 2005] Magland J. and Epstein C. Practical pulse synthesis via the discrete inverse scattering transform. *J. Magn. Reson.*, 2005. pp. 63–78.

[Pauly *et al.* 1991] Pauly J.M., LeRoux P., Nishimura D.G. and Macovski A. Parameter relations for the Shinnar-Le Roux selective excitation pulse design algorithm. *IEEE Trans Med Imaging*, 1991. **10**, 53–65.

[Skinner *et al.* 2003] Skinner T.E., Reiss T.O., Luy B., Khaneja N. and Glaser S.J. Application of optimal control theory to the design of broadband excitation pulses for high-resolution NMR. *J Magn Reson*, 2003. **163**(1), 8–15.

[Skinner *et al.* 2004] Skinner T.E., Reiss T.O., Luy B., Khaneja N. and Glaser S.J. Reducing the duration of broadband excitation pulses using optimal control with limited *rf* amplitude. *J Magn Reson*, 2004. **167**(1), 68–74.

[Tomlinson and Hill 1973] Tomlinson B.L. and Hill H.D.W. Fourier synthesized excitation of nuclear magnetic resonance with application to homonuclear decoupling and solvent line suppression. *J. Chem. Phys.*, 1973. **59**, 1775–1784.

[Ulloa *et al.* 2004] Ulloa J., Guarini M., Guesalaga A. and Irarrazaval P. Chebyshev series for designing rf pulses employing an optimal control approach. *IEEE Trans. Med. Imag.*, 2004. pp. 1445–1452.

[Wächter 2002] Wächter A. *An Interior Point Algorithm for Large-Scale Nonlinear Optimization with Applications in Process Engineering*. Ph.D. thesis, Carnegie Mellon University, 2002.