Efficient Implementation of Hardware Optimized Gradients for Rapid Imaging Sequences

J. A. Derbyshire[1], E. R. McVeigh[1]

[1]Laboratory of Cardiac Energetics, Division of Intramural Research, NHLBI, National Institutes of Health, Bethesda, MD, United States

Introduction: Applications such as cardiac and interventional MR provide increasing demand for high frame rate real-time imaging sequences. FISP is appealing for rapid imaging because the steady-state MR signal is practically independent of TR. Hence, increased frame-rates can be obtained without SNR penalty by maintaining the acquisition window constant while minimizing the sequence *dead-times* (i.e. parts of the sequence used to traverse k-space rather than playing RF pulses or acquiring data) [1]. *H*ardware *O*ptimized *T*rapezoidal (HOT) gradient pulses minimize TRs for a given physical gradient performance specification (i.e. maximum amplitude and slew rate) by optimizing the gradient shapes so that at least one of the physical gradient axes is operating at the specified limits during the dead-times.

In HOT pulse sequences the shapes of the required gradient pulse waveforms generally change from one phase-encoding (and unencoding) step to the next for each gradient axis. Therefore, around 1000 gradient waveform shapes (i.e. 6× the number of phase encode steps) are required to image a particular scan plane orientation with a particular FOV, slice thickness etc. Scan orientation, or other changes, require the redesign and reloading of all these gradients. In an early implementation [1], all these waveforms were pre-loaded into waveform memory, but the sequence was limited due to the waveform space available. In a more recent implementation [2], the individual gradient waveforms were dynamically loaded on a TR-by-TR basis while the sequence was running, but required unconventional programming techniques such as direct sequencer hardware access resulting in a platform-dependent sequence that was difficult to maintain and extend.

Here, we demonstrate that the number of required waveform shapes can be reduced to just 12 (i.e. two per phase encoding and unencoding step, per axis) for scanners that are able to play a linear combination of two pulse shapes simultaneously on each particular gradient axis.

Theory: HOT gradient pulses are designed by considering the required starting (e.g. slice) and ending (e.g. readout) gradient amplitudes and the required time integrated pulse area (i.e. k-space offsets, $\boldsymbol{\delta k}$) for the slice, read and phase encoding axes simultaneously as 3D vectors. These vector design requirements are rotated from logical (r,p,s) gradient space into physical (X,Y,Z) gradient space where solutions can be obtained that meet the physical gradient amplitude and slew-rate specifications. For spin-warp sequences, the minimum dead-time can be determined by considering the worst case of the phase encodes extremes for each of the three physical gradient axes. A set of 1D physical gradient waveform solutions that fit in the required dead-time are then found for each of the phase encoding steps.

Define a piecewise linear gradients (PLG) to be a 1D continuous gradient waveform comprised of several segments with a constant slew rate in each segment, and where the corner points at which the slew rate changes referred to as vertices. Typically, the set of 1D solutions for a particular HOT pulse comprise PLGs with two or three segments depending on whether the PLG reaches the maximum allowed amplitude. One reason that amplitude scaling cannot be used to implement phase encoding steps is that generally all these PLGs have constant, non-zero initial and final amplitudes (e.g. the X-component of the slice select and readout gradients). It is also not generally possible to implement amplitude scaling even if the set of PLGs for a HOT pulse are rotated back into the logical coordinate system because the PLG vertices occur at different time-points for different phase-encode steps.

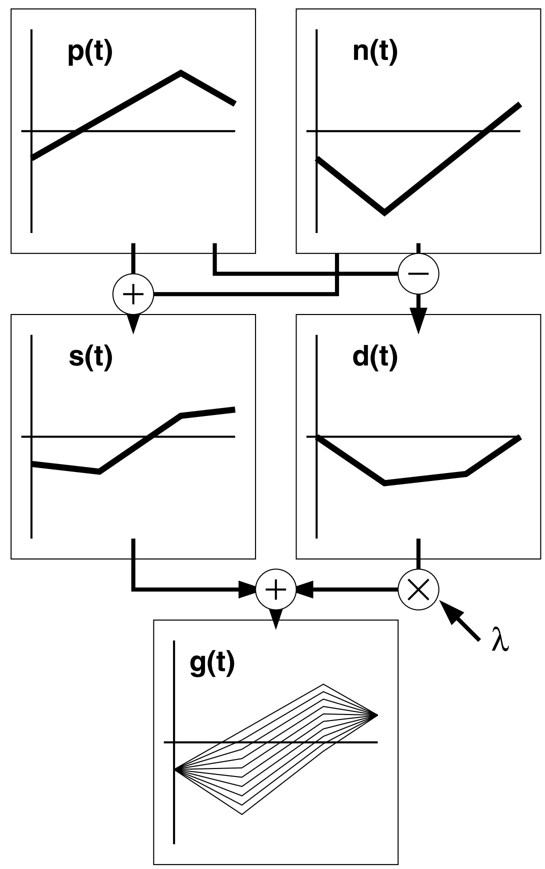

To solve this problem so that amplitude scaling can be employed for a particular HOT pulse, define p(t) and n(t) to be the PLGs corresponding to the maximally positive and negative phase encoding steps respectively. Since the linear combination of two PLGs is another PLG (albeit with more vertices), we now define:

$$s(t) = [\ p(t) + n(t)\ ] / 2 \quad \text{and} \quad d(t) = [\ p(t) - n(t)\ ] / 2 \quad \text{etc.}$$

Observe that s(t) will begin and end at the same constant value as all the PLG steps including p(t) and n(t). In fact, it can be regarded as the portion common to all the phase encoding steps. Conversely, d(t) begins and ends with zero amplitude, and represents the portion of the set of PLG solutions that varies with phase encoding step. In general, we can obtain an arbitrary PLG pulse for the r^{th} ($-N/2 \le r \le N/2$) phase encoding step of the HOT pulse by forming a linear combination:

$$g(t) = s(t) + \lambda d(t) \quad \text{where: } \lambda = 2r/N, \text{ such that } -1 \le \lambda \le 1$$

Note also that it can be shown rigorously that the slew-rate and amplitude of g(t) cannot exceed the original gradient specifications. Thus we have obtained a HOT pulse solution comprising a pair of PLGs per HOT pulse that allows simple gradient waveform amplitude scaling to be employed on scanners with pulse sequencers that permit linear combinations of two PLG type pulses on each axis simultaneously.

Implementation: A HOT pulse based FISP sequence was implemented on a Siemens Magnetom Espree 1.5T imaging system equipped with 33 mT/m, SR 100T/m/s gradients written in IDEA version VB12A. In particular, two features of the scanner were critical to the design: (i) the ability to generate aribitrary, user-desgined gradient waveforms and (ii) the ability to simultaneously play a linear combination of two such gradient pulses on each of the gradient axes.

Discussion and Conclusions: An especially efficient method for the implemention hardware optimized gradient waveforms has been demonstrated. The method requires that the scanner be able to simultaneously play a two piecewise linear gradient pulses in linear combination on each of the three gradient waveform generators. Wider implementation of hardware optimized gradient sequence methods should provide increased SNR efficiency in rapid imaging and improved frame rates in interventional real-time imaging.

Open bore systems are advantageous to interoperative MR in terms of patient accessibility, but these systems pay the price of a substantial reduction in gradient performance resulting in longer TRs and reduced frame rates. HOT pulse based sequences can mitigate the loss in gradient performance by improving efficiency.

Typically, vendor-supplied interactive real-time packages adopt a conservative approach, derating the gradient performance by √3 to allow the scan-plane to be interactively re-oriented without having to rebuild the gradient waveforms. The methods described here provide a simple way to increase performance.

References: [1] Atalar et al. MRM:32:773:1994 [2] Derbyshire et al. Proc ISMRM 2002, abstract 2359.