Parallel Transmit Technology for High Field MRI

Lawrence L. Wald[1, 2, 3]; Elfar Adalsteinsson[1, 3, 4]

[1]*Athinoula A. Martinos Center for Biomedical Imaging, Department of Radiology, Massachusetts General Hospital, Boston, MA, USA*
[2]*Harvard Medical School, Boston, MA, USA*
[3]*Harvard-MIT Division of Health Sciences and Technology, Boston, MA, USA*
[4]*Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Boston, MA, USA*

Introduction

The success of parallel reconstruction methods and their impact on image encoding has sparked a great deal of interest in using the spatial distribution of transmit coils in an analogous fashion. Namely, by breaking down the transmit field into multiple regions each controlled by a separate transmit channel, spatial degrees of freedom are created that allow the spatial information in the array to be exploited in the excitation process. While a homogeneous birdcage-type body-coil driven by a single RF pulse waveform has served the MR community well, it does not possess spatial degrees of freedom, and subsequently works best for uniform excitations. Parallel excitation arrays and the potential to utilize the spatial information in an array during RF transmission offer the possibility to move beyond the uniform slice-select excitation, and to generate spatially tailored RF pulses; excitation pulses with a carefully controlled spatially varying flip angle or excitation phase that can mitigate artifacts or isolate specific anatomy.

While the concept of spatially tailored excitation pulses has been known for some time, the implementation of such pulses is largely impractical on conventional single-channel excitation systems, and only by introducing the additional spatial degrees of freedom in a transmit array do they achieve practical durations for clinical imaging. An early application of spatially tailored parallel excitation was to mitigate the non-uniform flip angle problem created by RF wavelength effects at high field (Fig. 1). These non-uniformities arise when the wavelength of the RF approaches the dimension of

1 Flip angle inhomogeneity resulting from wavelength effects in the brain at 7T (central brightening) and liver at 3T (drop-out). Spatial variations in the transmit efficiency, and therefore the flip angle, are more problematic than the receive inhomogeneities since they lead not only to image shading, but more importantly, image contrast alterations.
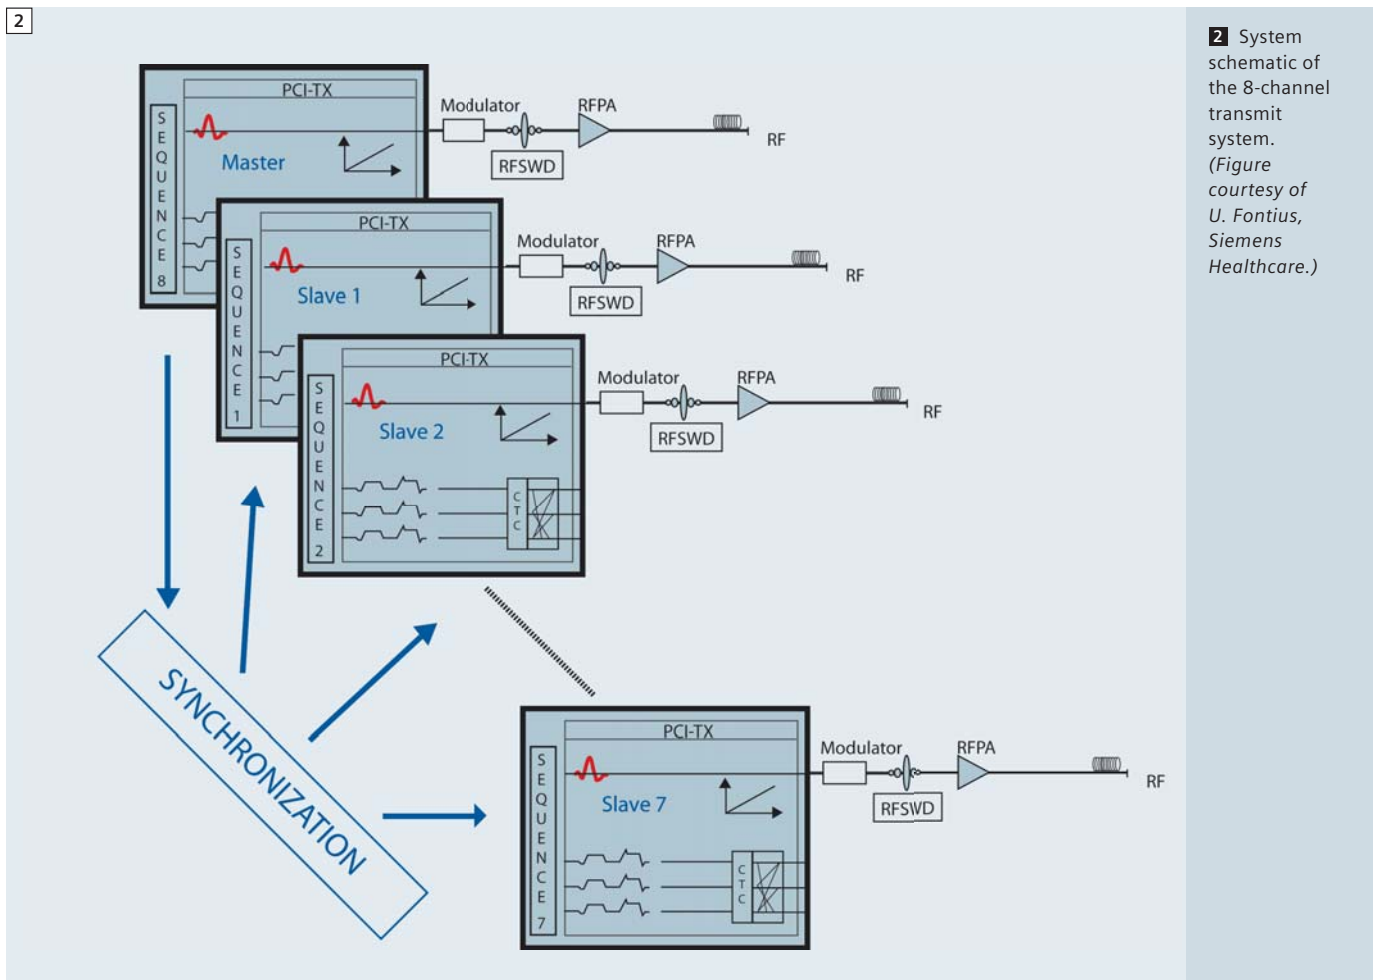
the human head or body and create destructive excitation field interference among sections of a conventional transmit coil. This is most noticeable in the head at 7T where a strong center brightening is a typical feature (perhaps more properly termed peripheral darkening) and in the body at 3T where shading is seen near large regions of non-fatty tissue in the abdomen. Unlike detection inhomogeniety that manifests primarily as image intensity shading, a non-uniform transmit B_1 field results in spatially dependent tissue contrast and therefore reduced diagnostic power, which cannot be recovered with an image normalization scheme. A spatially tailored excitation mitigates this problem by anticipating the flip angle inhomogeniety and compensating for it in the spatial profile of the excitation. Once the technology is in place for spa-

tially tailored RF excitations, the ability to generalize excitation profiles beyond the slice-select pulse offers many exciting opportunities for selective excitations of anatomically tailored volumes. An anatomy-specific excitation could potentially reduce image encoding needs (e.g. for cardiac or shoulder imaging) by reducing the effective field-of-view, it could enable more accurate CSI exams in tissues with many interfaces like in the prostate, and allow selective spin-tagging excitations (potentially allowing vessel territory perfusion imaging), or simply provide clinically useful but non-traditional excitations such as curved saturation bands for the spine or brain.

In this article we review some of the progress which has been made with a prototype 8-channel parallel transmit system integrated into a Siemens

MAGNETOM Trio, A Tim System, and a 7 Tesla MAGNETOM system. We discuss some of the recent advances in calculating parallel transmit RF pulses for spatially tailored excitation and show examples of B_1 transmit mitigation at 3T and 7T. Further, we describe some of the recent advances in methodology as well as some of the outstanding issues that must be overcome for routine application.

Experimental Setup

Flexible delivery of independent RF waveforms to each channel of the array is needed to realize the full potential of parallel transmission. Additionally, fast gradient trajectories are required during the RF pulses to modulate the spatial profile of the excitation. Since eddy current compensation is performed during the RF waveform generation using knowledge of the gradient history, each RF channel

needs to be fully integrated into the full waveform generation system of the scanner. To achieve this, a prototype 8-channel transmit system was set up in a master-slave configuration with each channel capable of running an independent pulse sequence, and importantly, independent B_0 eddy current compensation. Finally each channel utilized a separate RF power amplifier (8 kW each in the 3T case and 1 kW each in the 7T case) and fully independent SAR monitoring on each channel.

Which RF transmit array configurations capture the maximum ability to capitalize on the parallel nature of the excitation? This question is central to the optimal design of a flexible parallel excitation system and remains an open research problem. Two principles guide our design of the array configuration:

- obtaining the maximum benefit from the limited (expensive) number of excitation channels, and
- retaining the simplicity of birdcage-like excitation in one channel.

These two goals are elegantly achieved when a so-called "Butler matrix" [1] is inserted in the path from the RF amplifiers to the coil elements to drive a ring of excitation coils on a cylindrical former. In contrast to a direct drive of the coil elements by the RF amplifiers, the Butler matrix transforms the phase relationship

of the array elements so that each RF excitation channel drives not just a single RF element but influences all of them in a specific (and familiar) relationship; namely the spatially orthogonal modes of a birdcage coil. The spatial patterns of these modes and the phase relationships needed to generate them are well known from birdcage theory [2], and when achieved, have several benefits. Firstly, it allows the "master" channel of the array to operate as a uniform birdcage-like excitation coil. Although at high field the so-called uniform birdcage mode generates significant field variation (one of the original motivations for parallel TX technology), it is useful in
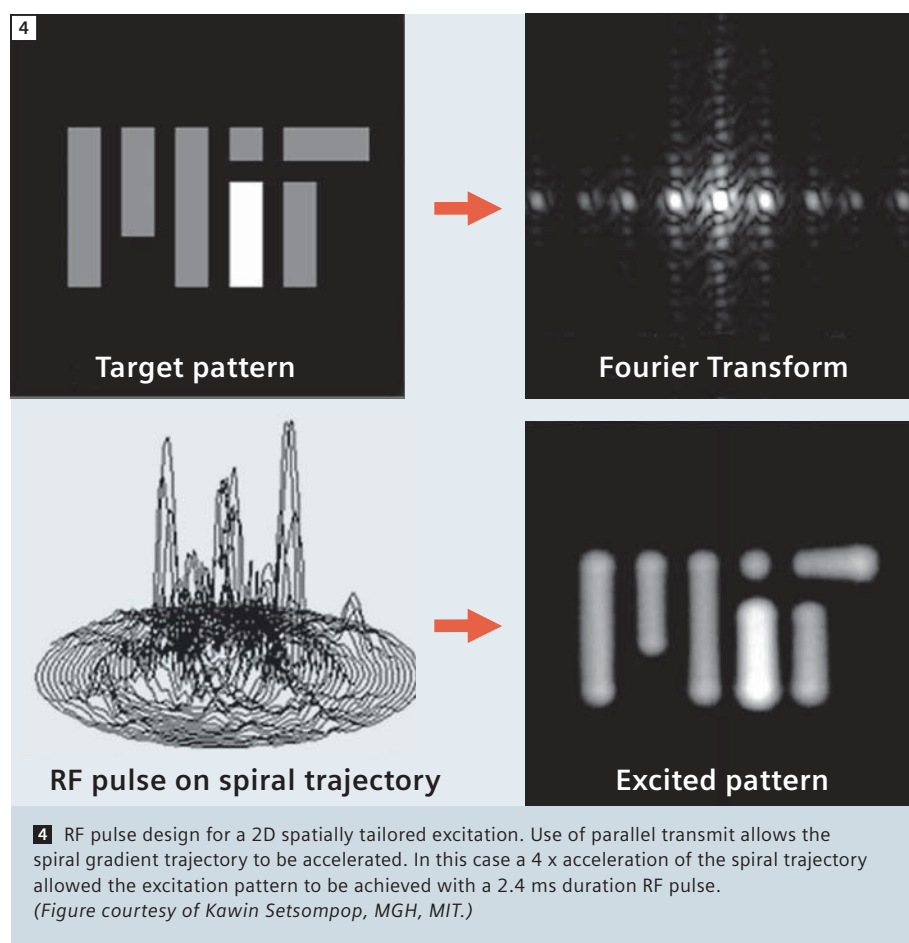

3 A 16-channel 7T strip-line array for the head (upper left) and a 16 x 8 Butler matrix (upper right). Below are the magnitude and phase B_1 maps of each of the 16-channels for both the "strip-line" basis set, and the "birdcage" basis set. While excitation ability is roughly equally divided among the strip-line modes, it is concentrated in a few valuable modes in the birdcage basis set. We can choose which 8 modes to drive based on their performance. *(Figure courtesy of Vijay Alagappan, MGH.)*

practice to have one of the channels operate from this well established and efficient starting point. The other channels then span progressively higher-order modes of the birdcage coils with their spatially specific amplitude and phase variations.

Since the Butler matrix achieves these modes through simple linear combinations, at first blush it would appear that this "basis set" of excitation patterns would be no better or worse for accelerating spatially tailored excitation than simply driving the array of elements one at a time. The superiority of the Butler matrix driven array becomes apparent when only a subset of the array modes is chosen for excitation. In practice this allows the benefit of a larger array to be captured in a system with fewer transmit channels (i.e. lower cost) by capturing a majority of the transmit efficiency and acceleration capabilities in a valuable subset of the channels (and ignoring the less valuable channels). We explored this "array compression" principle by driving a 16-channel stripline array for 7T head transmit with a 16 x 16 Butler matrix connected to the 8-channel transmit system [3], and demonstrated the theoretically predicted tradeoffs. The excitation configuration that integrates a Butler matrix in this manner allowed us to pick and chose among the modes of a 16-channel array and drive only the best subset of the 16 available modes with our 8 transmit channels. The choice of the optimum 8 birdcage modes compared to 8 strip-line elements allowed a flip-angle inhomogeneity mitigating excitation to achieve a 43% more uniform excitation and 17% lower peak pulse power in a water phantom at 7T [3].

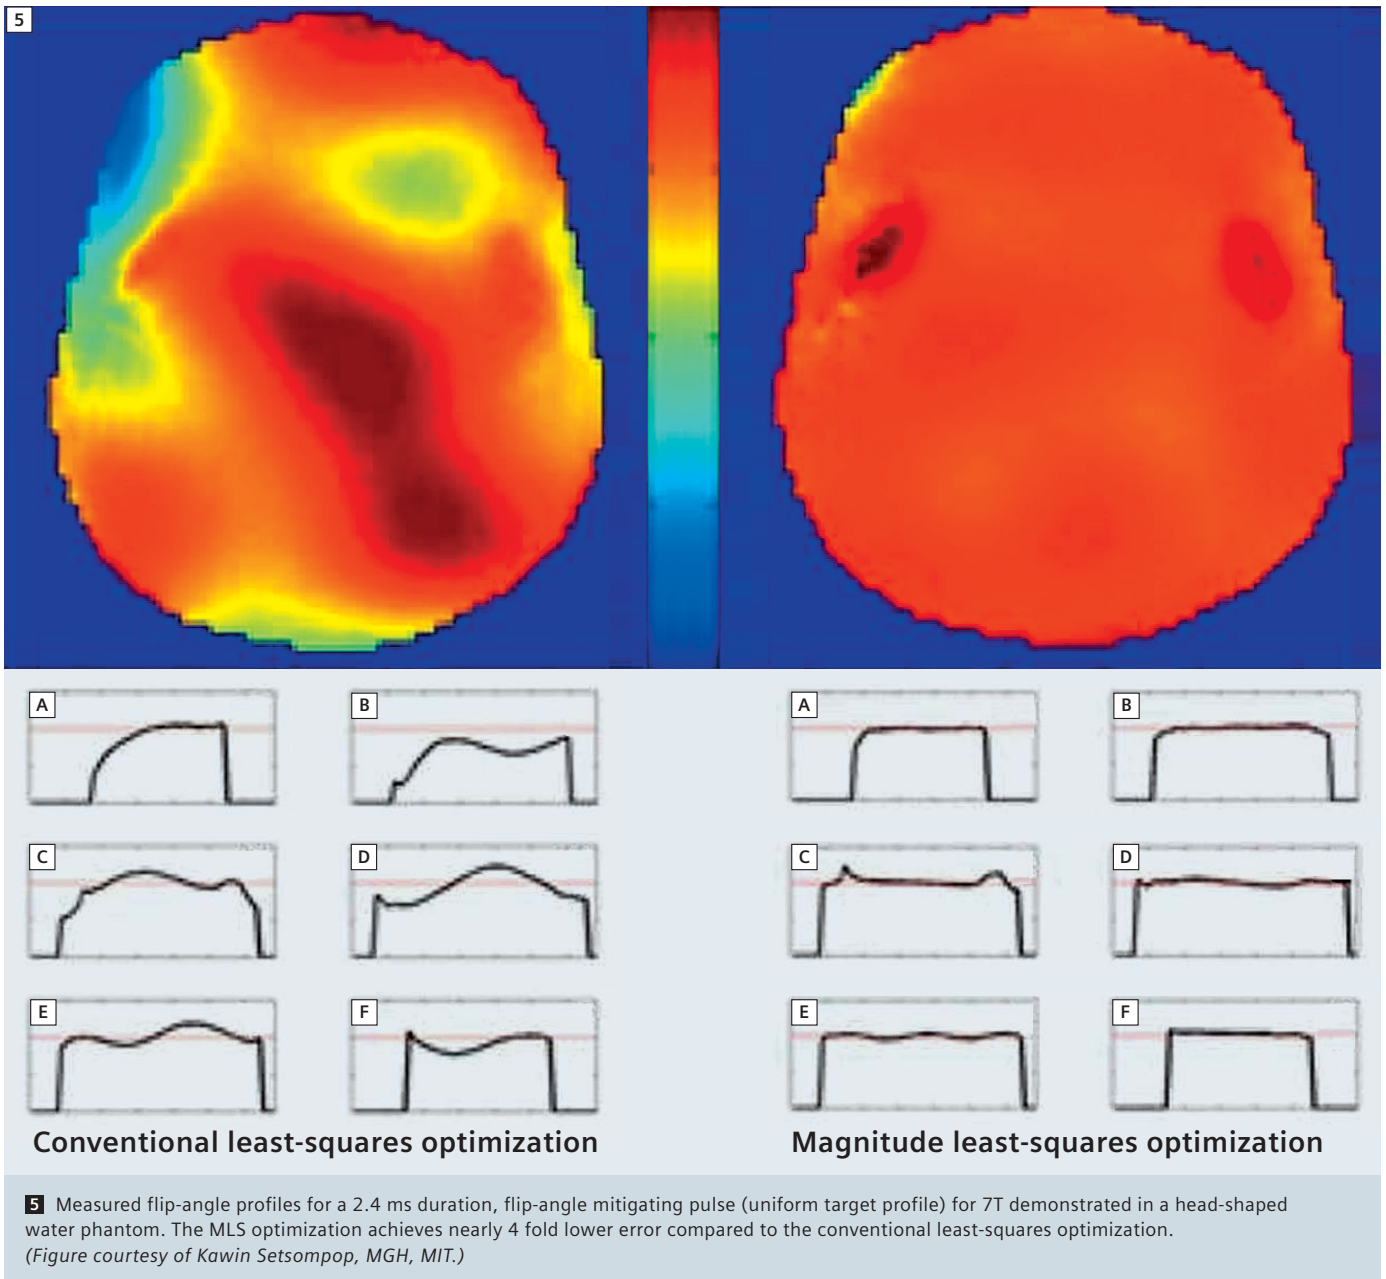## Spatially tailored RF excitation

RF excitations appropriately modulated in amplitude and phase during time-varying gradients offer the potential of spatially tailored RF phase and amplitude in the excitation [4]. Although such pulses have been demonstrated for many years, the lengthy encoding period needed duration of these pulses (as long as 50 ms) has precluded their routine use. Parallel transmission addresses

Target pattern → **Fourier Transform**

RF pulse on spiral trajectory → **Excited pattern**

4 RF pulse design for a 2D spatially tailored excitation. Use of parallel transmit allows the spiral gradient trajectory to be accelerated. In this case a 4 x acceleration of the spiral trajectory allowed the excitation pattern to be achieved with a 2.4 ms duration RF pulse. *(Figure courtesy of Kawin Setsompop, MGH, MIT.)*

this limitation by accelerating the excitation encoding gradient trajectory analogously to how parallel receive provides unaliased images with an accelerated encoding trajectory [5, 6]. A practical goal is to achieve 3D excitation pulses in less than 5 ms with a spatial profile that can mitigate the observed B_1 pattern in the head or body. This short duration is needed to be useful in common anatomical imaging sequences such as TSE, MPRAGE and FLASH.

Shaping the 2D spatial flip-angle distribution of an RF excitation requires modulated RF amplitude and phase while the gradients trace an excitation k-space trajectory, typically a spiral or echo-planar path. In practice, we first choose a target magnetization map, which is proportional to the flip angle map for small flip angles. For example, the target magnetization map might be a uniform flip-angle distribution or a selected region around

the anatomy of interest (for zoomed imaging). The calculation of the corresponding RF waveform is greatly simplified in the low flip angle case where it can be reduced to a k-space or Fourier analysis [4]. The RF excitation during such a gradient traversal is viewed as a series of short, small flip angle excitations. The phase and amplitude of these small RF pulses is altered so that the deposition of RF energy in the "excitation k-space" matrix is the Fourier transform of the desired spatial flip-angle map. In parallel transmit, the pulse duration is significantly reduced since an accelerated, under-sampled excitation k-space trajectory is used. The missing information is provided by incorporation of the spatial profiles of the multiple transmit array elements in the design process so that an unaliased excitation pattern is achieved.

Conventional least-squares optimization   Magnitude least-squares optimization

5 Measured flip-angle profiles for a 2.4 ms duration, flip-angle mitigating pulse (uniform target profile) for 7T demonstrated in a head-shaped water phantom. The MLS optimization achieves nearly 4 fold lower error compared to the conventional least-squares optimization. *(Figure courtesy of Kawin Setsompop, MGH, MIT.)*

Regularization of Specific Absorption Rate (SAR) and relaxation of phase constraints

A critical observation about the parallel transmit pulse design problem is that there are many different solutions for the RF pulses that achieve a very similar fidelity to the target excitation pattern. Knowing this, it is beneficial to choose a solution which produces a "close enough" pattern but minimizes SAR. This can be achieved by explicitly penalizing pulse amplitude when solving for the optimal pulse shapes, thus resulting in a significantly lower global SAR with little loss of excitation pattern fidelity [7]. Another important observation that yields significant payoff in the RF design is that the vast majority of MR applications ignore the phase in the final image (only magnitude images are viewed). In this case, the excitation can tolerate a slow phase roll across the FOV with no impact on the final image. We have capitalized on the relaxation of the phase restraint by developing a "Magnitude Least Squares" (MLS) algorithm for solv-

ing the parallel transmit pulse design optimization [8]. In this scheme, the algorithm attempts to achieve the target excitation pattern in magnitude, but allows slow phase variations across the FOV. The relaxed constraint allows a higher fidelity in the magnitude pattern or a lower pulse power (i.e. low global SAR). Figure 5 compares the MLS result for a slice-selective excitation with uniform target flip-angle distribution to the conventional Least Squares optimization. A two-fold improvement in target magnitude fidelity was achieved with similar SAR. Conversely, the same target fidelity could be achieved with a ~2 fold reduction in SAR.
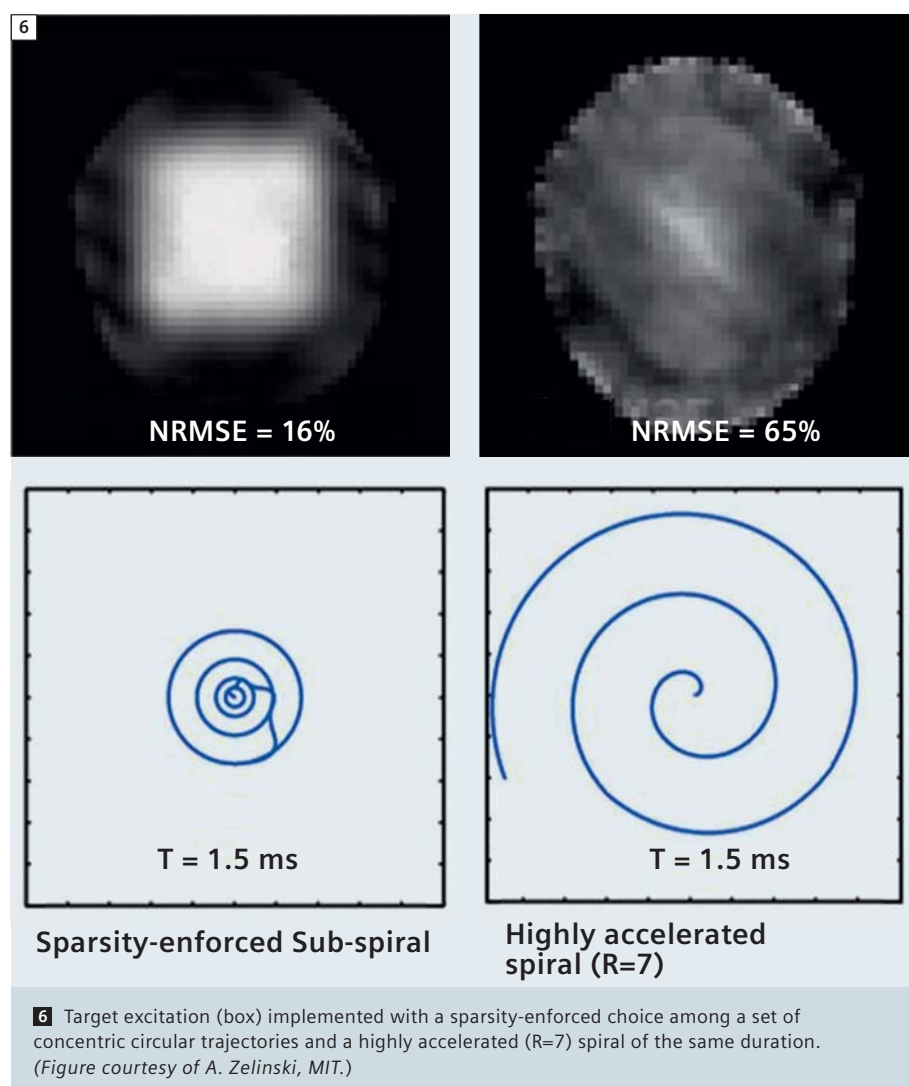
3D shaped excitations

Extending accelerated spatially tailored pulses to general 3D shapes requires covering k-space in 3 directions and is consequently very time consuming. Nonetheless, full 3D excitation is required for many applications, including adding in-plane flip-angle modulation to the traditional slice-selective excitation. We explored the capabilities for these 3D shaped excitations by using a variant of the echo-volume or "spokes" trajectory. This design class of RF excitation pulses can be viewed as multiple slice-selective RF pulses in z that are played out with different amplitude and phase modulations for each (k_x, k_y)-location, providing a conventional slice-selection in z, but with spatial modulations in the image plane, (x, y). Since the conventional axial slice select gradient can be viewed as a line segment in excitation k-space along k_z, multiple such lines look like a collection of spokes orthogonal to (k_x, k_y) when viewed in k-space. The parallel transmit problem then reduces to determining the number and (k_x, k_y) location of such spokes, as well as the calculation of the phase and amplitude of each transmit channel for each spoke to achieve the desired modulation in the x, y plane.

The excitation trajectory design problem is guided by our knowledge of the desired target excitation pattern and B_1 profiles of the transmit array elements, and further augmented by a SAR penalty term in the optimization cost-function. Thus, the k-space trajectory and RF pulse can be jointly optimized to produce a higher fidelity excitation pattern while satisfying constraints on overall SAR. When a mode-mixing strategy is employed in the transmit array, we can additionally choose which modes to connect to the transmit channels based on the excitation trajectory. Since the excitation k-space amplitudes and phases are related to the target pattern by the Fourier transform (for low flip-angle excitations), the k-space trajectory can be limited to regions with the largest magnitude Fourier coefficients. However, this does not take advantage of the "don't care" regions outside the body but within the FOV. A better strategy is to let a sparsity-enforcing algorithm choose the trajectory from among a discrete set of k-space grid points allowing an explicit trade-off between excitation fidelity and pulse length [9]. The simulation in Figure 6 demonstrates the advantage of choosing a subset of circular trajectories
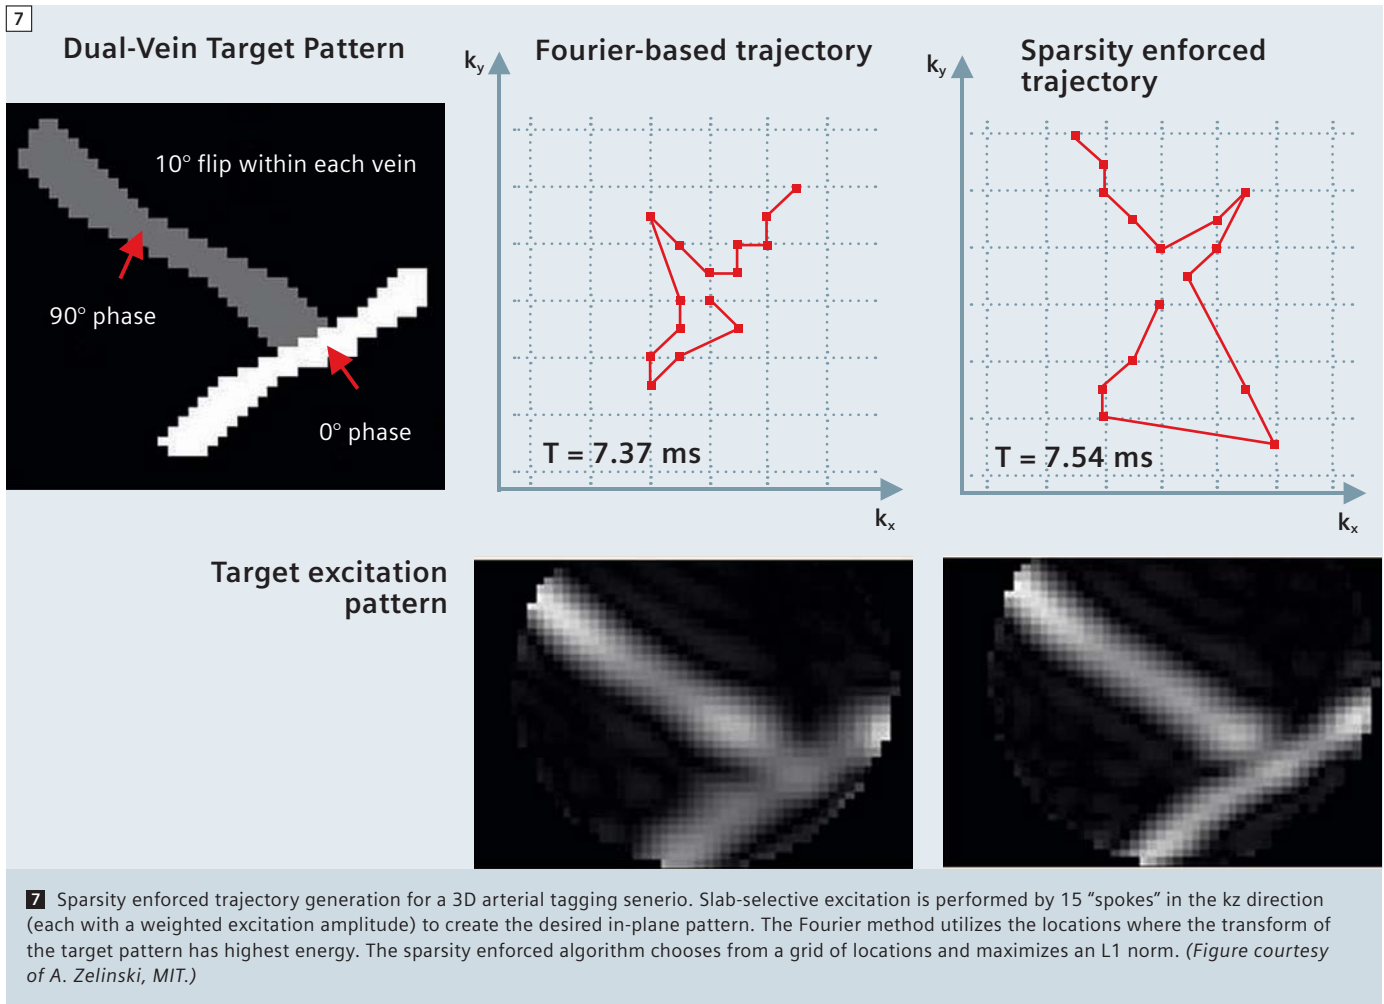

NRMSE = 16%	NRMSE = 65%
T = 1.5 ms	T = 1.5 ms
Sparsity-enforced Sub-spiral	**Highly accelerated spiral (R=7)**

6 Target excitation (box) implemented with a sparsity-enforced choice among a set of concentric circular trajectories and a highly accelerated (R=7) spiral of the same duration. *(Figure courtesy of A. Zelinski, MIT.)*

among a candidate set of circles in the k_x, k_y plane compared to a highly accelerated spiral excitation of the same overall duration. Figure 7 shows a similar optimization from among a grid of possible locations in the k_x, k_y plane of the "spokes" trajectory. The target excitation is slab-selective in z, and selectively excites the two simulated arteries in-plane, such as might be used for a vessel-selective arterial spin labeling experiment. In this case the two crossing vessels are tagged with excitations differing in RF excitation phase by 90°, demonstrating tight RF control in both magnitude and phase for a challenging 3D excitation target.

B$_1$ mitigation at 7T

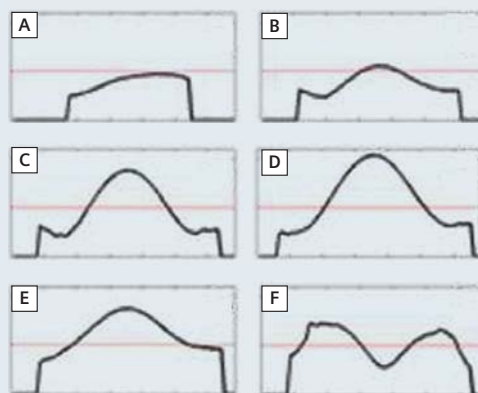Parallel excitations were performed on both head-shape water phantom and in

vivo human studies at 7T [10]. Both slice-selective B$_1$ mitigated excitation and arbitrarily shaped volume excitations were created and validated via a 16-element degenerate strip-line array coil driven with a Butler matrix utilizing the 8 most favorable birdcage modes. RF and gradient excitation waveforms were designed using the MLS optimization, and a spokes' placement optimization algorithm. With this design method, optimized parallel excitation waveforms for human B$_1$ mitigation were only ~50% longer than conventional single-channel slice-selective excitation while significantly improving flip-angle homogeneity. We compared the B$_1$ mitigation performance by parallel transmission to "RF shimming," which can be viewed as a simplified form of parallel transmit where the array elements are driven

with individual amplitudes and phase shifts but not separate pulse shapes. For the slice-selective excitation, the RF shimming can be viewed as a special case of the "spokes" trajectory where only a single spoke (at the center of (k_x, k_y)-space) is employed. Thus, RF shimming utilizes the spatial patterns of the transmit array, but not the encoding ability of the gradient trajectory. Figure 8 shows the measured B$_1$ map for the "uniform" birdcage excitation, the RF shimming and pTX with spokes trajectory (all slice selective excitations). The full pTX method clearly demonstrates superior B$_1$ mitigation performance. The phantom inhomogeniety is similar in shape to that of the head, but exhibits more severe field variations than in the human head; a 3 fold variation in flip angle across the slice. Nevertheless, the
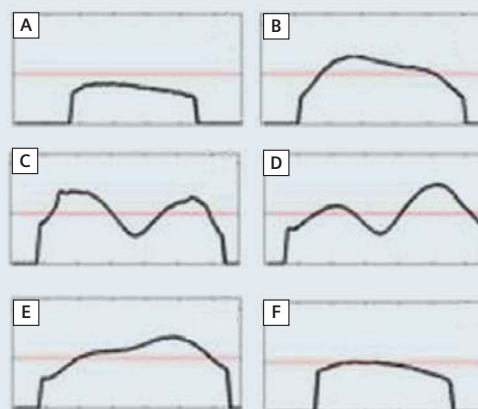

7 Sparsity enforced trajectory generation for a 3D arterial tagging senerio. Slab-selective excitation is performed by 15 "spokes" in the kz direction (each with a weighted excitation amplitude) to create the desired in-plane pattern. The Fourier method utilizes the locations where the transform of the target pattern has highest energy. The sparsity enforced algorithm chooses from a grid of locations and maximizes an L1 norm. *(Figure courtesy of A. Zelinski, MIT.)*

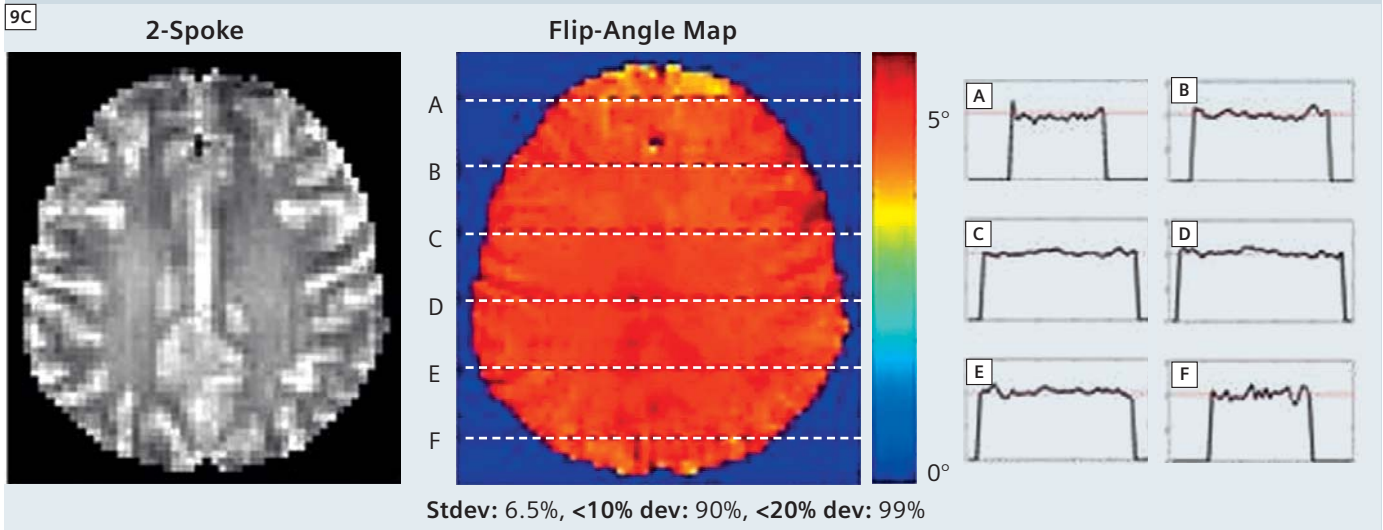## B_1 maps in a head shaped water phantom at 7T.


**8A** **Birdcage mode**



8B **RF shimming**

8C **3 spoke pTX**

8 Comparison of B_1 inhomogeneity mitigation in slice-selective excitations at 7T (head-shaped phantom) using **A:** conventional birdcage excitation (SD of B_1 across head = 42% of mean), **B:** RF shimming (identification of optimal amplitude and phase excitation settings) with 8-channel transmit array (SD = 29% of mean) and **C:** parallel TX with 3-spoke trajectory (2.4 ms duration) (SD = 5% of mean). *(Figure courtesy of K. Setsompop MGH, MIT.)*

9A Birdcage — **Flip-Angle Map**

Stdev: 14.8%, <10% dev: 45.2%, <20% dev: 78.7%

9B Shimming — **Flip-Angle Map**

Stdev: 12.2%, <10% dev: 54.1%, <20% dev: 90.3%

9C 2-Spoke — **Flip-Angle Map**

Stdev: 6.5%, <10% dev: 90%, <20% dev: 99%

9 Flip-angle inhomogeniety at 7T in the human head using 3 methods (conventional birdcage excitation, RF shimming and pTX with spokes trajectory.) Subject #5 (who displayed the most severe inhomogeniety) is shown. Gray scale images show a proton density-weighted low flip-angle image with the receive profile divided (leaving only variations due to transmit.) Color scale image is a quantitative flip angle map acquired with each of the 3 methods. *(Figure courtesy of K. Setsompop MGH, MIT.)*

10
3T flip angle maps in abdomen

CP birdcage   "True Form" drive   pTX, 2 spokes

pTX, 3 spokes   pTX, 4 spokes
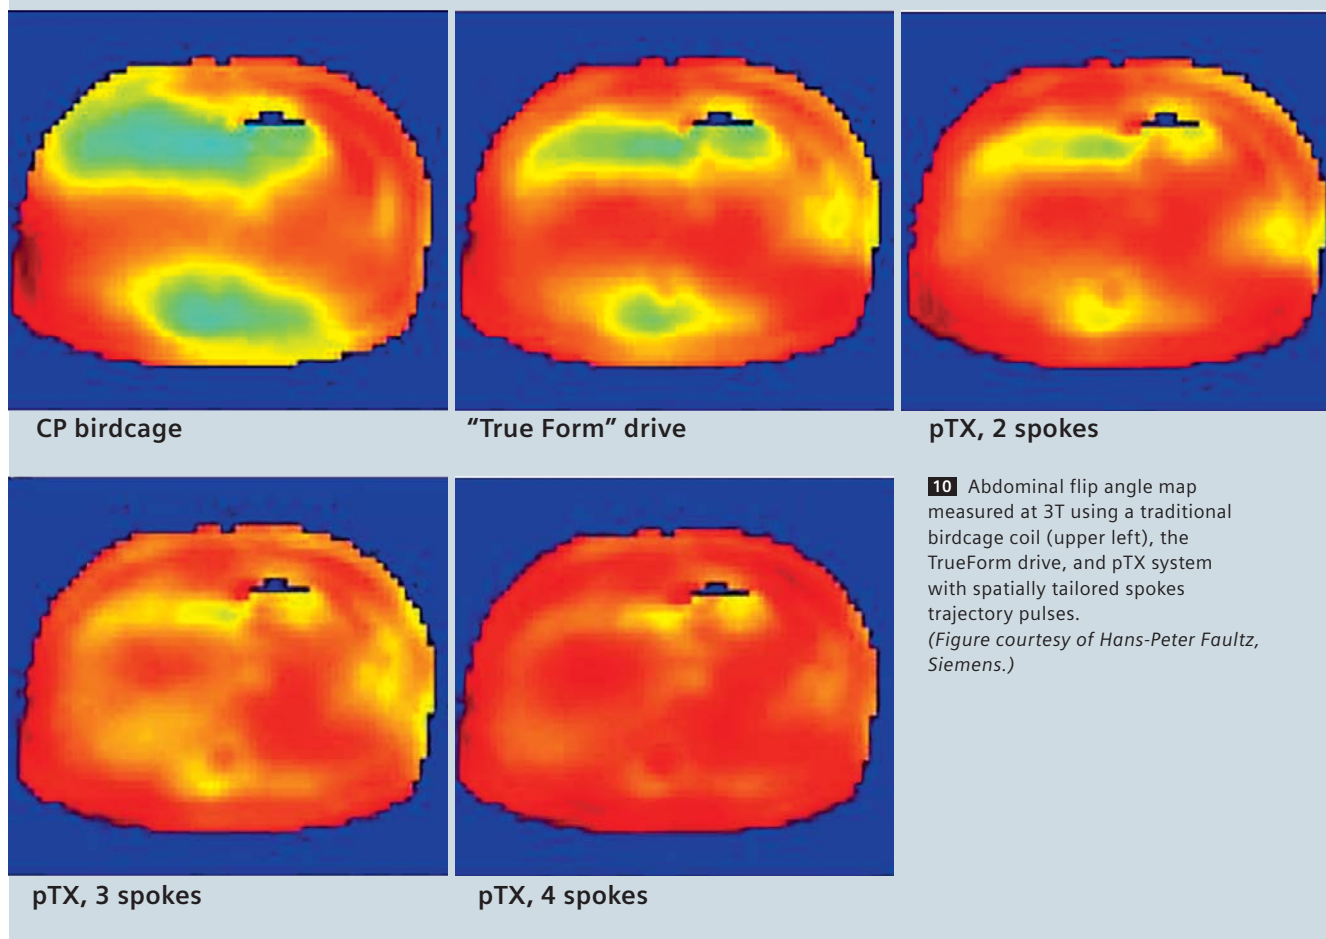
10 Abdominal flip angle map measured at 3T using a traditional birdcage coil (upper left), the TrueForm drive, and pTX system with spatially tailored spokes trajectory pulses.
(Figure courtesy of Hans-Peter Faultz, Siemens.)

3-spoke trajectory and 8-channel array is sufficient to remove the vast majority of the inhomogeneity.

Figure 9 shows B_1 maps obtained from one of 6 healthy subjects (studied with institutional approval and informed consent). In this case a 2.3 ms duration 2-spoke slice-selective trajectory was used with the 8-channel system and the MLS design method. The birdcage and RF shimming acquisition used a 1.4 ms sinc-like excitation. While some contamination from anatomy is seen in the B_1 transmit maps, the pTX method significantly reduced the B_1 inhomogeniety (standard deviation (Stdev.) across slice was 8% of the mean compared to 21% for the birdcage excitation and 14% for the RF shim).

Figure 10 shows the parallel transmit method applied to a similar flip-angle inhomogeneity problem; the abdomen at 3T. Here a similar wave-cancelation occurs in body imaging where the size of the body becomes comparable to the wavelength of the RF. The conventional circularly polarized birdcage can be significantly improved upon by optimizing the phase relationship between the drive ports of the coil to produce a more uniform and efficient elliptical polarization tailored to the body. Further gains in uniformity were realized with parallel transmit and the 3D spokes spatially tailored excitation pulses calculated based on knowledge of the B_1 field profiles of the transmit array elements.

SAR considerations

While the results in Figure 8 demonstrate the ability of the parallel transmit method to mitigate the inhomogeneous flip-angle distribution at high field, the spokes pulses used more RF energy to achieve the desired flip angle than the simple birdcage transmit (but less than the RF shim). The total pulse energy for the birdcage, RF shim, and pTX-spokes methods were 10.7 mJ, 24.8 mJ, and 21.8 mJ respectively. This suggests that the parallel methods achieve uniformity only with some degree of self-cancellation among the fields or excited magnetization. A similar effect is seen in the 2D spiral trajectories, where pulse energy significantly increases with acceleration,
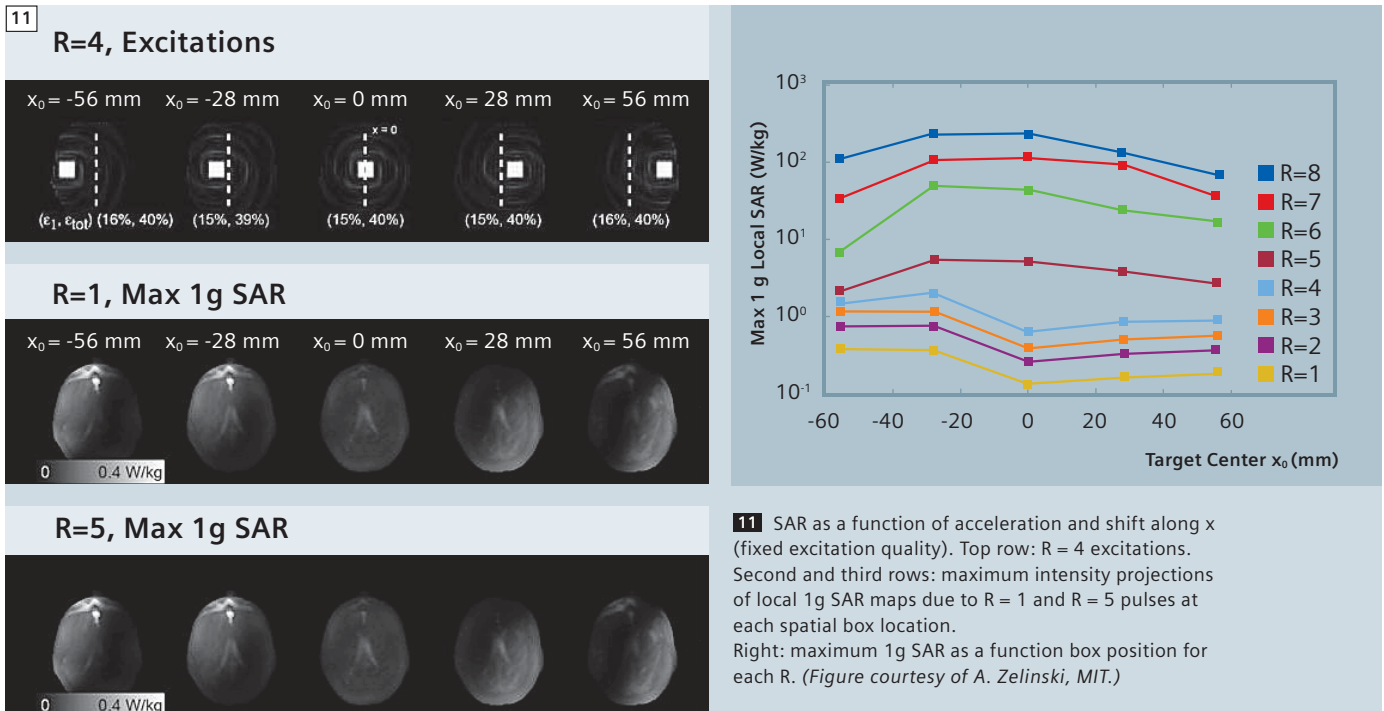
even with an explicit B_1 amplitude penalty in the pulse design cost-function and local SAR levels are difficult to predict [11]. An example is shown in Figure 11 where the local SAR is calculated for a series of box-shaped excitations placed at 5 different positions in the head (left to right). The spatially tailored 2D pulses used an 8-channel array and spiral trajectories with accelerations ranging from R=1 to R=8. The pulse design maintained a constant fidelity to the target pattern by trading off the pulse amplitude constraint and the fidelity constraint. The local SAR was calculated from the E_1 fields in a multi-tissue head model for the array and pulse. The first observation based on these results is the enormous cost in local SAR incurred by keeping the fidelity constant in the face of increasing parallel transmit acceleration (nearly 3 orders of magnitude variation in local SAR!). The second observation is that local SAR varies significantly with the position of the excitation box. For low accelerations the central box positions have the lowest SAR, while the higher accelerations, the peripheral

positions have lower SAR.
For evaluation and monitoring of SAR, the main concern for human imaging is the potential for the E_1 fields from the array elements to constructively superimpose locally, creating a local SAR hot spot. A simple estimate demonstrates how serious the "worst-case" superposition can be. If the E_1 fields from the eight elements superimpose and generate an 8 fold increase compared to a single element, then the local SAR at that location will increase 64 fold. Similarly, electric fields can destructively interfere. This means that if one channel stops transmitting due to equipment failure, the local SAR can actually increase. Therefore, in addition to monitoring the average power from each channel, a pTX system must make an estimate of local E_1 fields and how they superimpose so that the local SAR limits are not exceeded. As the pulse design becomes increasingly tailored to the individual patient, the local SAR check must also move in this direction. This will require fast local SAR calculation methods based on the field patterns calculated for the array

and the specific pulse designed for that subject. Preliminary work has exploited the ability to penalize high-amplitude RF pulses in the pulse design optimization, but significant future development is needed to explicitly include local SAR regularization in the design of the RF pulses and enable a flexible trade-off between RF excitation properties (due to B_1) and local SAR distribution (due to E_1).

Remaining challenges

In addition to the SAR estimation and monitoring problem, several other outstanding challenges must be solved before accelerated 2D and 3D spatially tailored excitations can be routinely employed. The method relies on a fast but accurate mapping of the B_1 transmit field in the subject, which is an intense and ongoing area of innovation with several promising methods being proposed in the literature. A second area of innovation is the calculation of high flip-angle spatially tailored RF pulses. Most of the work performed to-date has assumed the small flip angle approximation. While this approximation provides

11 SAR as a function of acceleration and shift along x (fixed excitation quality). Top row: R = 4 excitations. Second and third rows: maximum intensity projections of local 1g SAR maps due to R = 1 and R = 5 pulses at each spatial box location.
Right: maximum 1g SAR as a function box position for each R. *(Figure courtesy of A. Zelinski, MIT.)*

for elegant and computationally tractable RF designs with familiar tradeoffs based on well known Fourier transform properties, large flip-angle pulses are central to many clinical pulse sequences and the low flip-angle constraint needs to be addressed for general applicability of pTX. This computational problem is now just starting to be addressed.

Conclusions

Theoretical work on parallel RF transmission and recent experimental validations on 8-channel prototype systems at 3T and 7T indicate that parallel excitation has the potential to overcome critical obstacles to robust and routine human scanning at high field strength. As these developments are extended, high-field human imaging will be possible with essentially constant flip angle, and therefore no compromise in signal strength or clinical contrast, across the human head and body with RF pulse durations comparable to current slice selective pulses. While most work has been concentrated on head-sized transmitters at 7T, the methods are readily translatable to body transmit coils at 3T. Of course, intriguing research questions remain open in several areas, including optimal coil array designs that minimize element couplings and maximize spatial orthogonality of individual channels; the estimation of local SAR from a subject-specific spatially tailored RF pulse; and the development of rapid and robust RF pulse designs that extends the current low-flip angle domain to arbitrary excitation angle, such as spin echoes, saturation, and inversions pulses. However, with continued active research in these areas, progress is likely to accelerate, and logical extensions of the architecture of a current clinical scanner readily accommodates the requirements of a general parallel RF excitation system supported by fast, subject and application tailored RF pulse design software capable of extending MR excitation from the simple slice-select to the more generally tailored anatomy- or application-specific RF excitation pattern.

WIP – Works in Progress. This information about this product is preliminary. The product is under development and not commercially available in the U.S., and its further availability cannot be ensured.

References
1 Butler, J. and R. Lowe, Beamforming matrix simplifies design of electronically scanned antennas. Electron Design, 1961. 9: p. 170–173.
2 Tropp, J., Mutual Inductance in the Bird-Cage Resonator. J Magn Reson, 1997. 126(1): p. 9–17.
3 Alagappan, V., et al. Mode Compression of Transmit and Receive Arrays for Parallel Imaging at 7T. in International Society for Magnetic Resonance in Medicine. 2008. Toronto, Canada.
4 Pauly, J., D. Nishimura, and A. Macovski, A k-space analysis of small-tip angle excitation. J Magn Reson, 1989. 81: p. 43–56.
5 Katscher, U., et al., Transmit SENSE. Magn Reson Med, 2003. 49(1): p. 144–50.
6 Zhu, Y., Parallel excitation with an array of transmit coils. Magn Reson Med, 2004. 51(4): p. 775–84.
7 Zelinski, A., et al., Comparison of three algorithms for solving linearized systems of parallel excitation RF waveform design equations: Experiments on an eight-channel system at 3 Tesla. Concepts in Magnetic Resonance Part B: Magnetic Resonance Engineering, 2007. 31B: p. 176–190.
8 Setsompop, K., et al., Magnitude least squares optimization for parallel radio frequency excitation design demonstrated at 7 Tesla with eight channels. Magn Reson Med, 2008. 59(4): p. 908–15.
9 Zelinski, A.C., et al., Sparsity-enforced slice-selective MRI RF excitation pulse design. IEEE Trans Med Imaging, 2008. 27(9): p. 1213–29.
10 Setsompop, K., et al., Slice-selective RF pulses for in vivo B1+ inhomogeneity mitigation at 7 tesla using parallel RF excitation with a 16-element coil. Magn Reson Med, 2008. 60(6): p. 1422–32.
11 Zelinski, A.C., et al., Specific absorption rate studies of the parallel transmission of inner-volume excitations at 7T. J Magn Reson Imaging, 2008. 28(4): p. 1005–18.
12 Shinnar, M. and J.S. Leigh, The application of spinors to pulse synthesis and analysis. Magn Reson Med, 1989. 12(1): p. 93–8.
13 Shinnar, M., et al., The synthesis of pulse sequences yielding arbitrary magnetization vectors. Magn Reson Med, 1989. 12(1): p. 74–80.
14 Shinnar, M., L. Bolinger, and J.S. Leigh, The use of finite impulse response filters in pulse design. Magn Reson Med, 1989. 12(1): p. 81–7.
15 Shinnar, M., L. Bolinger, and J.S. Leigh, The synthesis of soft pulses with a specified frequency response. Magn Reson Med, 1989. 12(1): p. 88–92.
16 Pauly, J., et al., Parameter relations for the Shinnar-Le Roux selective excitation pulse design algorithm. IEEE Tr Medical Imaging, 1991. 10(1): p. 53–65.
17 Mao, J.M., TH, K. Scott, and E. Andrew, Selective inversion radiofrequency pulses by optimal control. J Magn Reson, 1986. 70(2): p. 310–318.
18 Conolly, S., D. Nishimura, and A. Macovski, Optimal control solutions to the magnetic resonance selective excitation problem. IEEE T Med Imaging, 1986. MI-5(2): p. 106–115.
19 Geen, H., S. Wimperis, and R. Freeman, Band-selective pulses without phase distortion. A simulated annealing approach. J Magn Reson Med, 1989. 85(3): p. 620–627.

Contact
Lawrence L. Wald
Associate Professor of Radiology
Athinoula A. Martinos Center for Biomedical Imaging
Department of Radiology
Massachusetts General Hospital
Harvard Medical School and
Harvard-MIT Division of Health Sciences and Technology
wald@nmr.mgh.harvard.edu

Elfar Adalsteinsson
Associate Professor
Department of Electrical Engineering and Computer Science
Harvard-MIT Health Sciences and Technology
Massachusetts Institute of Technology
elfar@mit.edu