Specific Absorption Rate (SAR) in Parallel Transmission (pTx)

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

[1]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, Boston, MA, USA
[2]Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Boston, MA, USA

Introduction

Parallel transmission (pTx) uses multiple excitation coils driven by independent RF pulse waveforms to subdivide the transmit field into multiple spatial regions each controlled by a separate transmit channel. Increasing the number of spatially distinct transmit elements and temporally distinct RF pulse waveforms compared to conventional single-channel RF systems creates spatial degrees of freedom that allow the spatial information in the array to be exploited in the excitation process. Previous pTx work has concentrated on the potential to utilize this additional flexibility in spatial information to move beyond the uniform slice-select excitation to generate spatially tailored RF pulses; excitation pulses with a carefully controlled spatially varying flip angle pattern or excitation phase that can mitigate artifacts or isolate specific anatomy. Operating in analogy to parallel reception, parallel transmission offers the possibility to move beyond the uniform slice-select excitation and create a more anatomy-specific excitation which could mitigate drop-out regions from inhomogeneous excitation fields at 3T or 7T, reduce image encoding needs (e.g. for cardiac or shoulder imaging) by reducing the needed field-of-view, 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.

While the development of these novel RF pulse designs and applications continue to be an area of intense development, the clinical utility of a given RF excitation pulse is characterized by more than just its spatial fidelity. The Specific Absorption Rate (SAR) of an excitation pulse is often the critical limiting factor when applied to a clinical imaging sequence. The need to stay below safe SAR limits often requires unfavorable tradeoffs in acquisition parameters such as increased TR or reduced flip angle. This is especially problematic at 3T, where the power needed for a given flip angle increases as much as 4 fold compared to 1.5T applications, and thus the SAR ceiling is acutely felt.

Fortunately, while the use of parallel transmission creates some SAR problems of its own (such as the need to carefully monitor for local SAR hotspots in a specific RF pulse design), the spatial degrees of freedom that give rise to the wealth of spatial design potential also offers novel opportunities for SAR manipulation. While attempts to minimize SAR during the design of parallel transmit pulses has been ongoing since the first pTx papers, the full potential of our ability to manipulate these degrees of freedom as well as the engineering solutions needed to realize SAR reduction and management via pTx are just appearing.

In this article we review some of the progress which has been made toward strategies for reducing SAR with pTx in a prototype Siemens MAGNETOM Trio, a Tim System, and a MAGNETOM 7T. We review the pTx SAR problem and discuss some of the recent advances in calculating parallel transmit RF pulses for spatially tailored excitation with optimization methods which penalize global and local SAR. Further, we describe some recent advances in real-time SAR monitoring as well as outstanding issues that must be overcome for routine application.

SAR in spatially tailored excitations

To achieve an arbitrary spatial flip angle distribution after excitation, we modulate the RF waveform in amplitude and phase during a time-varying gradient waveform. These "spatially tailored" RF excitations have been pursued for more than two decades [1]. The design procedure is illustrated in figure 1, where the Fourier transform of the desired pattern is sampled on a spiral excitation k-space trajectory to produce the amplitude and phase modulations needed during the gradient waveform.

Although spatially tailored excitation pulses are easily demonstrated, there are serious engineering challenges to their routine and practical use. First and foremost is the lengthy and largely impractical encoding period needed (as long as 50 ms). Parallel excitation directly addresses this problem by allowing the

1

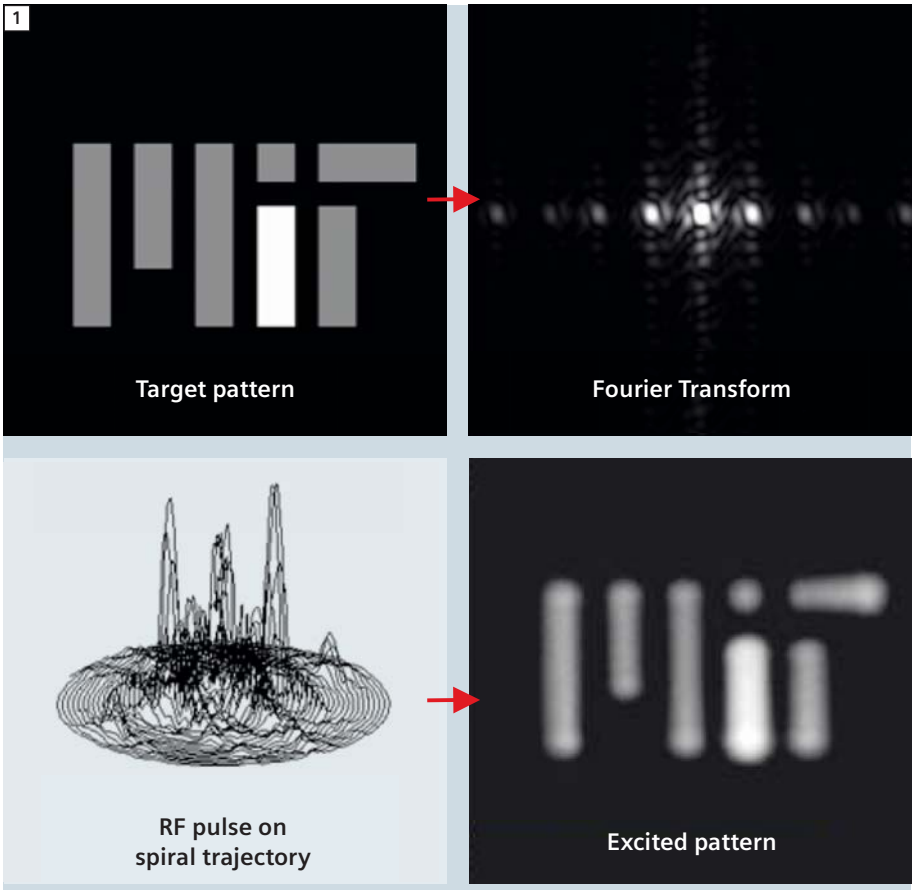

1 The design of a spatially tailored excitation starts with a choice of a spatial target pattern (in this case the MIT logo), and a gradient trajectory such as the spiral trajectory shown. The Fourier transform of the target pattern is used to modulate the amplitude and phase of the RF excitation waveform to achieve the desired distribution in excitation k-space. The RF waveform is played out during the gradient trajectory resulting in the desired excited magnetization pattern.

2

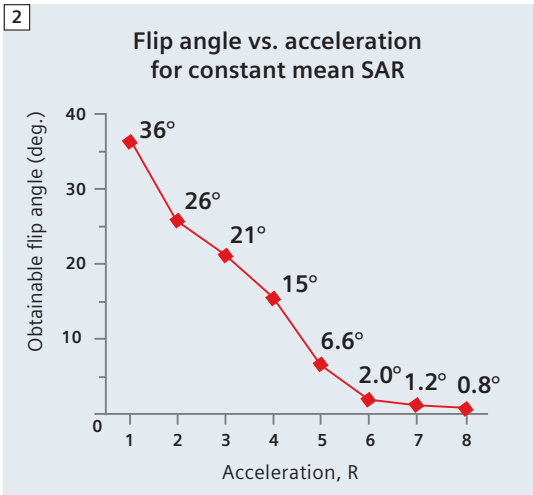

2 Obtainable flip angle for a given maximum global SAR as a function of acceleration of a 2D spatially tailored excitation using an 8-channel array coil. The obtainable flip angle is roughly linear in acceleration (and thus pulse duration) up to about R=5, showing the standard tradeoff between amplitude (and thus SAR) and pulse length (acceleration). After R=5, additional SAR penalties cause a less linear trade-off, showing the SAR penalties analogous to g-factor penalties in parallel reception.

"excitation k-space" trajectory to be undersampled, analogous to the reductions in encoding trajectory steps in parallel reception, but with some interesting contrasting properties. For instance, in accelerated parallel reception we are often limited by the image Signal-to-Noise Ratio (SNR) losses associated with the procedure. In parallel transmission, the RF waveforms take the position mathematically analogous to the final image in parallel reception and start with extraordinary high SNR since they are generated with robust digital-to-analog technology. Therefore, even with a modest 8-channel transmit array, we can reduce the duration of the pulse up to 6 fold without suffering from noise enhancement of the calculated RF pulse waveforms [2, 3]. Note that this level of acceleration would be difficult to achieve in parallel reception.

Instead, it is the SAR of the excitation that constitutes the limiting penalty in accelerating spatially tailored excitations with parallel transmission, with a secondary limitation arising from artifacts in the transmitted pattern due to imperfect knowledge of the array coil excitation profiles. The dramatic increase in SAR with acceleration can be understood intuitively. First, as the pulse length is reduced, the amplitude must be increased accordingly to achieve a given pulse area (flip angle). Increases in pulse amplitudes are expensive in SAR since the SAR increases with the square of the applied RF voltage. As a square pulse is shortened, the pulse amplitude must be increased linearly to maintain a given flip angle (pulse area). The instantaneous SAR thus increases inversely with the square of the duration, and the SAR associated with the use of this pulse within a fixed TR pulse sequence also increases linearly. Figure 2 demonstrates this effect for an eight-channel loop array at 7T as a nearly linear decrease in obtainable flip angle for a given global SAR value as a spatially tailored 2D excitation is accelerated by factor R via undersampling of its excitation k-space trajectory.

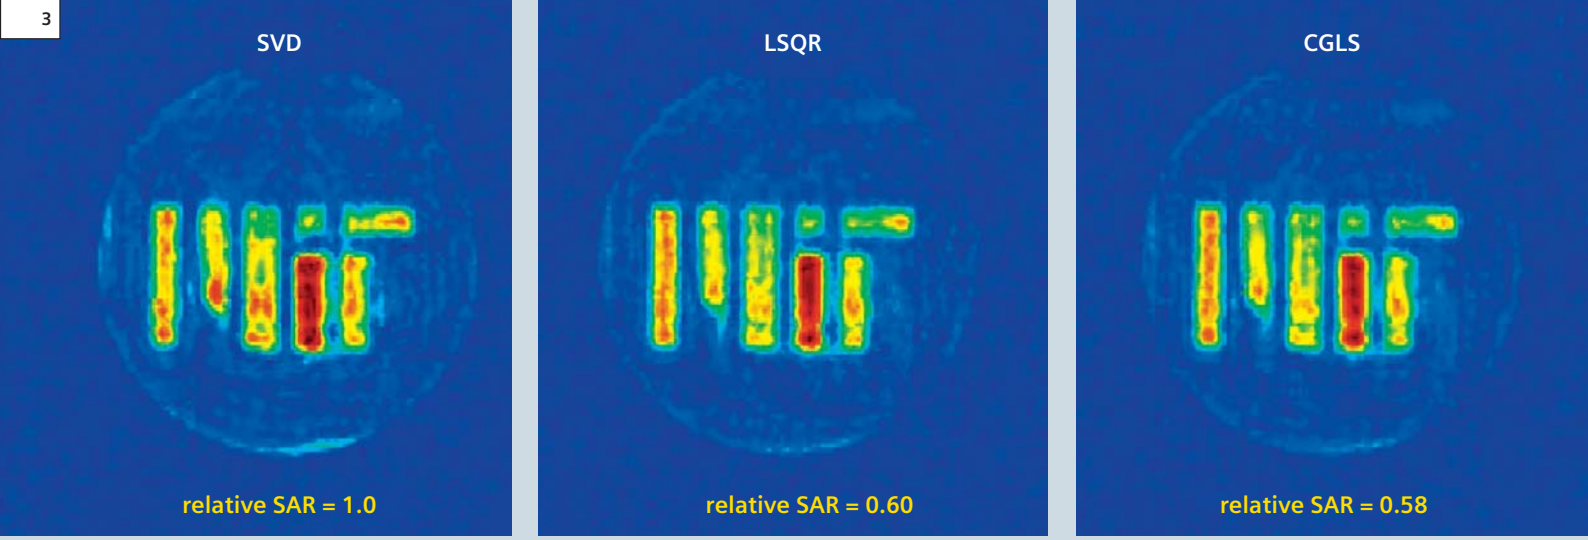

3 Accelerated spiral trajectory pulses with R=4 acceleration with approximately equal RMSE compared to the target pattern ("MIT" logo), but with the least-squares optimization performed with either the SVD, LSQR or CGLS method. In this case the global SAR was considerably lower for the LSQR or CGLS method compared to SVD. During the original work, we did not analyze the local SAR pattern, but now hypothesize that the local SAR pattern will be even more different among the three solutions.

Second, the impressive spatial definition of pTx excitation patterns obtained by these pulses come at a cost in SAR. The spatial pattern is achieved by temporally modulating the RF pulse in the presence of time-varying gradient waveforms, causing phase shifts in the excited magnetization that superimpose over time to achieve the desired pattern. Unfortunately, temporally modulating a pulse in amplitude or phase tends to create instances of higher peak power (as well temporal nulls in the pulse voltage) as the RF pulse deposits the requisite energy along the traversed excitation trajectory. Again the fact that instantaneous SAR is proportional to the square of the pulse waveform enhances the problem. Due to the nonlinear relationship between pulse amplitude and instantaneous SAR, the nulls do not counterbalance the increased peak values when the SAR is integrated over the pulse duration. This leads to increased SAR whenever a pulse is temporally modulated.

When the spatially tailored excitation strategy is generalized to multiple RF excitation channels, the excitation pattern is not simply the Fourier transform of the excitation k-space, but is determined from a linear model of the undersampled Fourier encoding and the spatial profiles of the transmit coils. This over-determined linear model is inverted to yield the needed RF waveforms. The optimization process used to invert the over-determined linear problem also gives rise to procedures for managing SAR. The simplest cost function for the optimization includes only a cost for the difference between actual excitation pattern and the original target. It is also straight forward to include a cost for the average pulse amplitude, thus acknowledging the cost of increased global SAR in the optimization. An obvious choice would be to also penalize local SAR by including the peak local SAR value in the cost function. Methods for rendering this latter approach computationally tractable are discussed below.

Aside from direct regularization procedures for global and local SAR such as the cost function modifications mentioned above, the pulse optimization offers multiple other possibilities for controlling local and global SAR. The next sections illustrate some of the approaches being pursued.

One-pulse design problem; many solutions

The first insight into controlling local and global SAR comes from the observation that the linear optimization problem generates many possible pulses with spatially similar transverse magnetization patterns, but which differ in other properties. Figure 3 illustrates one aspect of this finding by showing magnetization patterns obtained from 2D spiral-based pTx excitation patterns accelerated by R=4 for an 8 channel loop array at 7T. The optimization problem was identically posed, attempting to spatially match a target pattern (flip angle distribution in the shape of the MIT logo) with an additional term in the cost function penalizing the root-mean-squared (RMS) waveform average (proportional to global SAR). While it was expected that the three numerical approaches might obtain different degrees of accuracy, or have different computation burdens, the methods, in fact, produced very similar spatial solutions but with strikingly different global SAR. The SAR reduction amounted to 40% for the Least Squares QR decomposition (LSQR) and Conjugate Gradient Least Squares (CGLS) methods compared to the Singular Value Decomposition (SVD) method even though all of these methods attempted to penalize global SAR in the same way [4]. The fact that the solution space for the flip angle has many similar solutions but with a wide range of different power characteristics suggests that we can pick and chose from nearly equal spatial fidelities in order to improve other aspects of the excitation.

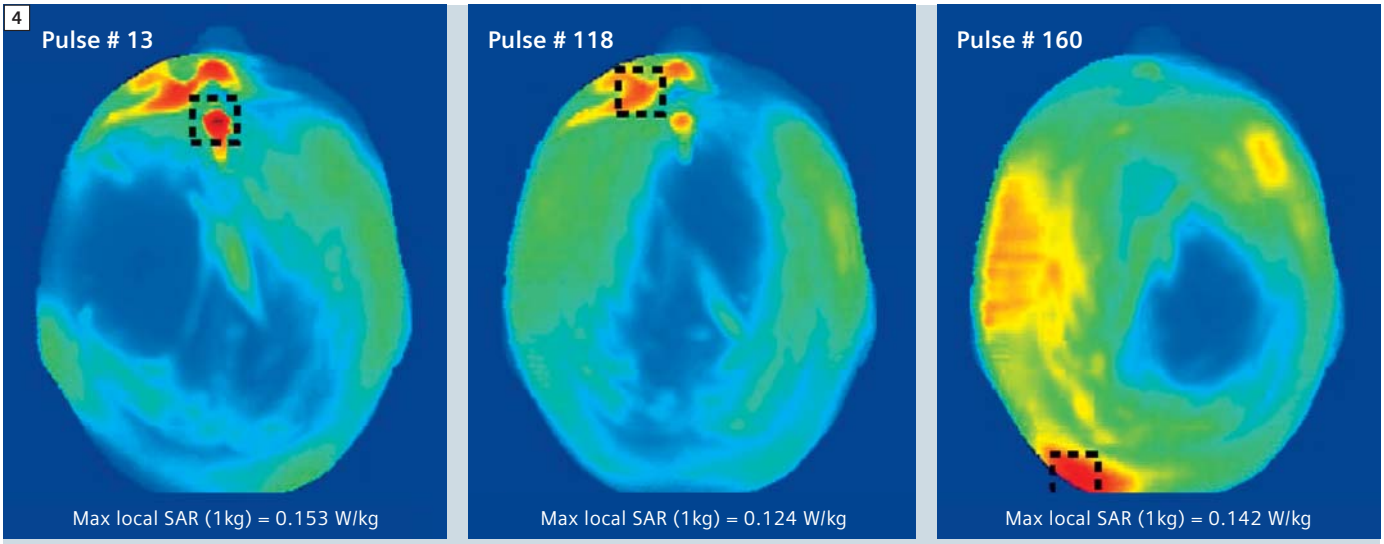

4 Maximum intensity projection of local SAR maps for 3 pulses from a set of 162 designed to produce a uniform flip angle distribution in an axial slice through the head at 7T. Maps are 1g SAR averages obtained from a tissue head model and FDTD simulation of the electric fields of the pulse. Although the pulses achieved nearly identical excitation patterns, the local SAR pattern is very different and would not coherently average with time if the acquisition could cycle through such pulses. *(Figure courtesy of Adam Zelinski.)*

SAR on the run...

A related strategy that expands on this core idea is to generate a family of solutions having similar spatial patterns but differing in the pattern of local SAR. When this set of pulses is used, for example, each pulse for a different imaging slice or even for each phase encode step, the local SAR hotspot will jump around spatially and not coherently time-average into a single long-term hotspot, thus spatially distributing the local SAR load [5-8]. Since the safety concern is local temperature increases which build up over minutes, this spatial-temporal averaging forms a meaningful solution and a viable approach to SAR management in pTx. Even if only two hotspot locations are possible, alternating between the two can lead to a 50% reduction in local SAR. An additional way to potentially force yet different local SAR patterns is to employ slightly different gradient trajectories for each pulse [8]. For example, a slice selective spokes excitation could be generated from any one of hundreds of trajectories. Combing through these solutions produces dozens of candidates within a given spatial fidelity and with different local SAR patterns. An even more sophisticated formalism computes the average SAR pattern from the collection of pulses and optimizes the individual pulses based on this time-averaged pattern [8]. Figure 4 shows the local SAR map for 3 pulses from a set of 162 calculated to mitigate B1+ inhomogeneities in the brain at 7T. The average spatial similarity of the SAR maps for a given pair of pulses was 87%, as computed from the spatial correlation of their SAR maps.

No place to hide...

In addition to forcing the local SAR hotspot to jump around in space, we pursue the regularization of local SAR in the pTx pulse design. Namely, we seek to penalize a potential pulse design during optimization by including its local SAR value into the cost function. The most straightforward way of doing this is to perform an exhaustive search of each pulse's local SAR during the iterative pulse design process. Unfortunately, computing and searching for the local SAR maximum based on the E field maps of an electromagnetic simulation is very computationally demanding, even on a relatively course grid. So while it is desirable to incorporate an exhaustive search over all spatial locations into pTx pulse design, it is currently not computationally feasible. A compromise has been proposed where

the pulse is calculated without local SAR in the cost function, note the location of the local SAR hotspot and then redo the design with a penalty for local SAR at that one location [9]. This strategy, however, is likely to simply force the local SAR peak to pop up at a different spatial location. What is needed is a comprehensive regularization which gives the local SAR no place to hide.

Towards this end, Gebhardt and colleagues at Siemens Healthcare and at the University of Erlangen-Nuernberg have noted that the E field data for a coil array is "compressible" in a sense analogous to how a compressed image can provide a nearly faithful image representation at a fraction of the disk space of the original bitmap. Using a fraction of the hundreds of thousands of phase and amplitude representations that would be needed to represent all of the possible excitation combinations of an 8-channel array, the local SAR can be reasonably predicted by a subset of 36 "virtual observation points" [10]. While this subset alone does not promise a fully accurate depiction of local SAR and would not be a replacement for actually checking the local SAR maximum of a pulse, it offers a computational approximation with dramatically reduced computational burden that can be used to penalize local SAR during pulse design while covering nearly all "hiding places." After pulse design, the accurate SAR limits of the pulse would be traditionally calculated.

Figure 6 shows the application of the virtual observation compression to regularization of the local SAR [11]. The graph shows the tradeoff between local SAR (y axis) and excitation error compared to the target pattern (x axis) as the relative weight of these two penalties is varied via the regularization parameter. Penalizing the local SAR in this way resulted in a 38% decrease in local SAR for this 8-channel array, spoke design slice-selective pulse for B1+ mitigation in the head at 7T.

5

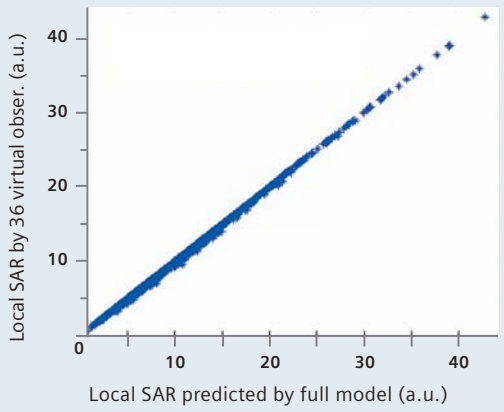

5 Comparison of the local SAR from 1000 randomly chosen transmit combinations evaluated with a full local SAR evaluation (x axis) vs evaluation through 36 virtual observation points (y axis). While the agreement between the two methods is not perfect, the vastly simpler 36 virtual observations reproduce the local SAR with good accuracy. *(Figure courtesy of M. Gebhardt, Siemens Healthcare.)*

6

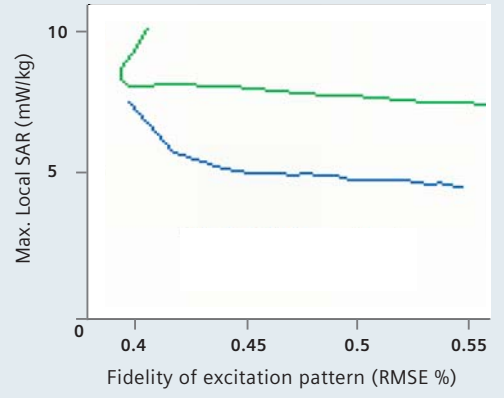

6 Trade-off between excitation error (x axis) and maximum local SAR as the regularization parameter λ is varied. Here the local SAR of the pulse is calculated from a set of 40 virtual observation points. A 38% decrease in local SAR is possible with this method.

Exploiting the choice of excitation coil mode

Since the above discussion centers on exploiting the spatial degrees of freedom in the pulse design process, it is worth considering hardware adaptations which can potentially further extend the number of spatial degrees of freedom. One obvious possibility is to increase the number of transmit channels. Some early exploration of this pathway has suggested significant SAR reduction might be possible with more highly parallel transmit arrays [12]. In this work Lattanzi and colleagues simulated increasing numbers (1, 8, 12, and 20) of circular transmit elements tiling a 30 cm diameter sphere at 7T. They compared the SAR values needed for producing a uniform excitation using RF shimming or a 2D pTx spatially tailored excitation pulse. As seen in practice [2], the pTx pulse was able to achieve a high degree of flip angle uniformity in human head at 7T even with

7

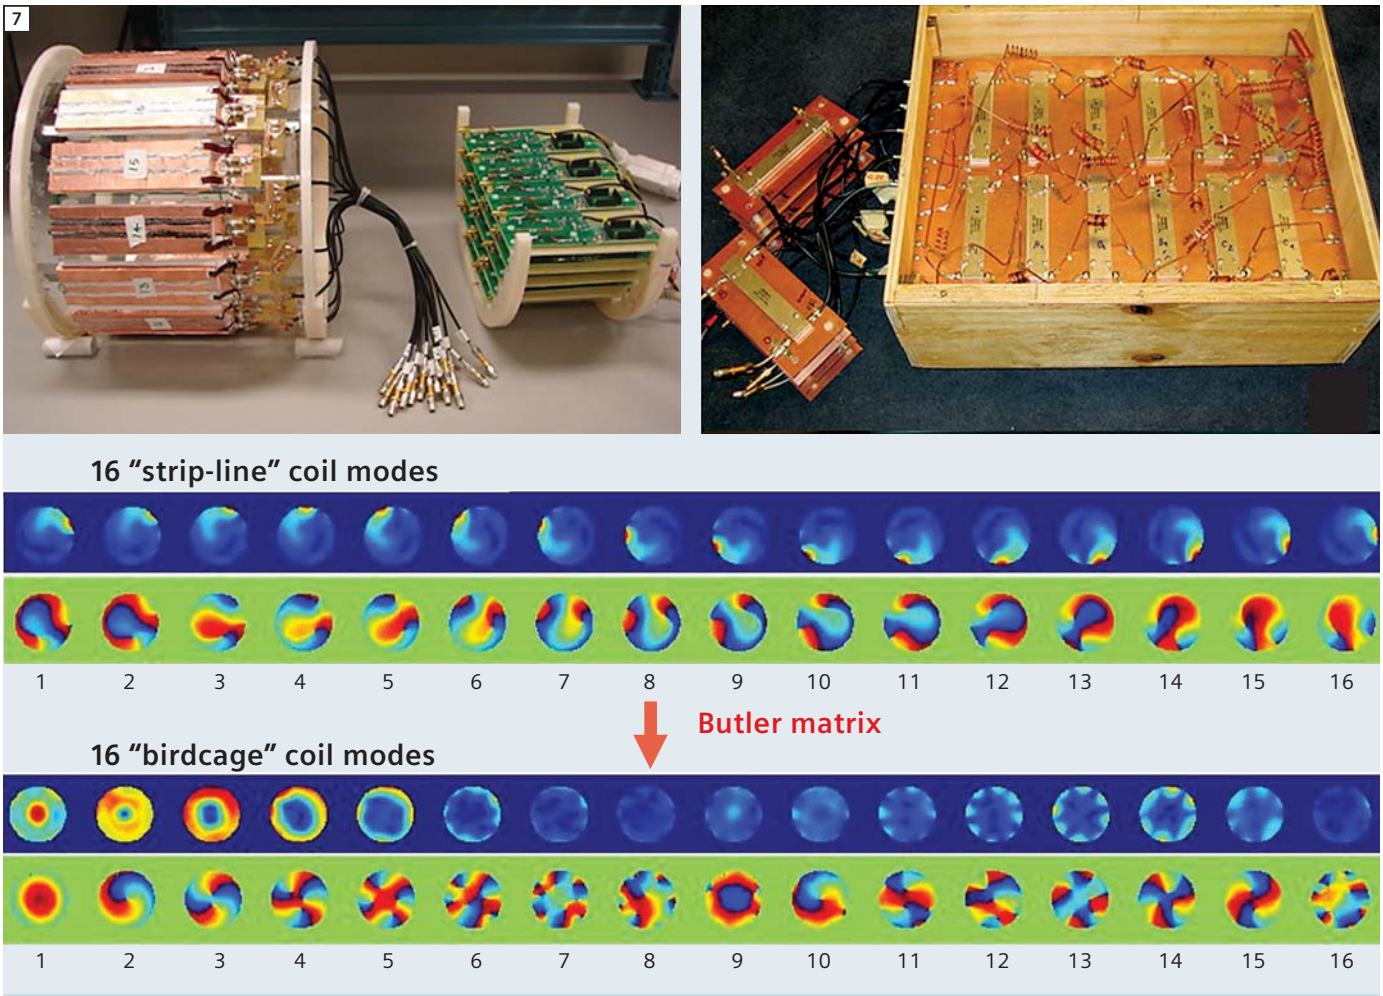

7 A 16-channel 7T strip-line array for the head (upper left) and a 16 x 8 Butler matrix (upper right). Below are the magnitude and phase B1 maps of each of the 16-channels for both the “strip-line” basis set, and the “birdcage” basis set. While excitation ability is roughly equally divided among the 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.)*

the 8 coil array, while the RF shimming approach had difficulty even with 20 array elements. Although the additional elements were not needed from a spatial excitation fidelity point of view, they were helpful in reducing global SAR. Global SAR was seen to reduce from 7.9, to 5.5, to 3.3 as the element count increased from 8, to 12, to 20 array elements. Even more interesting, Lattanzi and colleagues calculated the “ultimate intrinsic SAR” consistent with array fields derived from conductor distributions outside the body and satisfying Maxwell’s equations. Constructed in analogy to previous work calculating the ultimate SNR for array distributions, this represents the minimum global SAR possible with a theoretically perfectly optimized array (of however many elements needed). The ultimate intrinsic SAR was an additional 3 fold lower than that obtained with the 20-channel array, suggesting that if additional transmit channels can be engineered in an appropriate and cost effective way, there is potential to utilize element counts well above 20 for reduced SAR in pTx.

By constructing an array with more coil elements than channels to drive them, we can potentially provide the benefit of a higher element count transmit array without actually having to increase the number of RF transmit channels. This is a relatively cost-effective approach since the transmit channels with their power amplifiers and power monitoring equipment are generally more expensive than the array elements themselves. While

not increasing the number of transmit elements, this strategy still increases the available degrees of freedom by providing options in how we connect transmit channels to coils. It turns out that these are relatively uninteresting choices for conventional arrays, with transmit elements equally spaced around a cylinder, but it becomes more interesting when the spatial patterns of each array element are not just rotated version of their neighbors. If we think of the spatial distributions of the individual array elements as vectors, then the maximum degrees of freedom derived from this strategy occurs when these vectors are orthogonal. Thus we seek to transform our 'boring" array of elements around a cylinder into orthogonal modes by finding a suitable linear combination. For loop arrays distributed around a cylinder, an effective linear combination to achieve this goal is the discrete Fourier transform and the resulting basis set is known as the "modes" of the birdcage coil. Similar optimizations in phased array radar have even provided us with a simple hardware realization of this transformation; the Butler matrix [13], which is a generalization of a quadrature hybrid. Figure 7 shows the transformation of the spatial patterns of a 16-element strip-line array into the 16 birdcage modes via the Butler matrix. 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 strip-line array directly. The superiority of the Butler matrix driven array becomes apparent when only a subset of the array modes is chosen for excitation. The truncation captures 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 strip-line array for 7T head transmit with a 16 x16 Butler matrix connected to the 8-channel transmit system, and demonstrated the theoretically predicted tradeoffs. The excitation configuration that integrates a Butler matrix in this manner allowed us to pick and choose 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 [14, 15].
Exploiting these degrees of freedom afforded by choosing among the unequal birdcage modes even further, Zelinski demonstrated that both excitation pattern fidelity and SAR can be improved by choosing the mode set based on the desired excitation pattern [16].

Verifying and monitoring local SAR

As demonstrated in the examples above, the pTx pulse design process has unprecedented ability to impact both global and local SAR. This suggests that global and local SAR need to be carefully checked for each new RF pulse design to insure safety and regulatory compliance. In the most ambitious applications, the pulse design is envisioned to be tailored to the needs of the patient on the table, e.g. providing a volume selective excitation including only the heart, or a curved saturation band following the spine. In this case, in addition to calculating the pulse during the examination, the expected global and local SAR must be checked against regulatory limits. Since the scanner can only monitor total power absorbed by the coil/body, calculations using the pulse shape and modeled E and B fields expected in the body are the only current method available to ensure that local SAR limits are not exceeded before global SAR limits. The scanner can measure the power to each channel, but not the local superposition of E fields in the conductive tissue which gives rise to SAR. This global and local SAR simulation must be performed for every pTx pulse created since a unique spatial-temporal combination of the transmit array elements are used. This is in contrast to conventional single-channel systems where SAR need only be checked once and scaled appropriately based on the transmit power and pulse duration. Therefore, the strategy for ensuring global and local SAR compliance in a parallel transmit system is a generalization of the single transmit case. A full workflow for pTx imaging requires integrated transmit field mapping followed by RF pulse design. The SAR distribution must then be calculated for the pulse based on pre-calculated B_1 and E_1 fields in the appropriate body and coil model. Once the local maximum SAR is found, the local-to-global SAR ratio is examined and the global power limit derated so that the scanner will shut down if the real-time measure of global power to each channel indicates that either the global or local SAR limit exceeded.
An important additional safety check is to monitor in real-time what is actually being transmited by each channel using either field sensing probes or directional couplers on each transmit channel [17]. Then if a single transmit channel fails, the scan can be stopped. This is important since a given transmit channel can create an E_1 field which opposes that of the other channels. In this case, transmit power to that channel reduces local SAR. Alternatively, disruption of a single channel can cause an unpredicted increase in local SAR.
Figure 8 shows the three sections of the workflow tools designed for preliminary pTx studies [18]. The top tool generates the B_1 fields from the data provided by the mapping sequence; a step which ultimately will be performed by the image reconstruction system in an integrated prescan step. The middle panel shows the RF pulse design tool, which at this stage is intended for development use. The

8

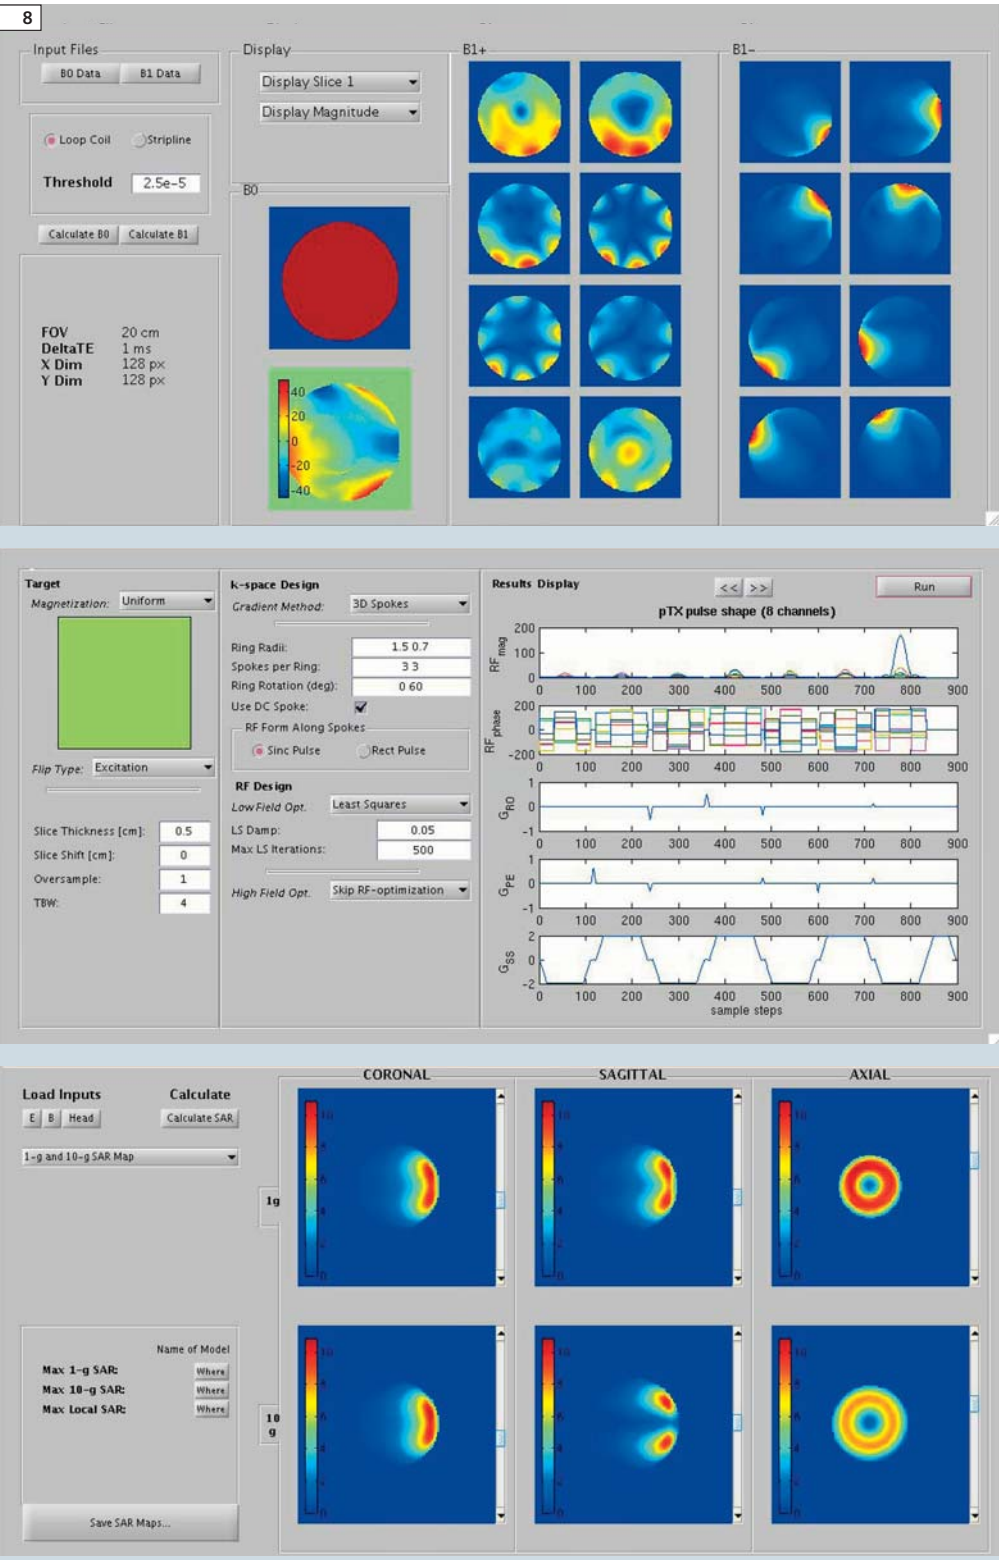

8 Workflow tools for calculating the transmit field maps (top tool), the pTx pulse based on a desired excitation pattern (middle tool) and the local and global SAR based on the RF pulse and a previously generated electromagnetic simulation of the transmit array and body model.

bottom panel shows the SAR maps generated based on the pTx pulse created in the pulse design tool using a previously calculated electromagnetic simulation of the array coil and body model. The SAR-check tool calculates the global power deratings needed in each channel to ensure that both global and local SAR are below safety limits. This information is incorporated into the header of the RF pulse file which is read by the pulse sequence. After checking that the pulse was SAR-checked and that the coil model used matches the coil on the scanner, the pulse sequence transfers the pulse waveforms to the image reconstruction computer where they are compared to the forward and reflected power from the directional couplers on each transmit channel. A disagreement between the measured and expected waveforms on any channel generates a fault which halts the scan.

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 and enable new applications based on spatially selective excitations. Furthermore, the flexible exploitation of the degrees of freedom in the pulse design problem provides the potential to significantly decrease SAR. While most work has been

concentrated on head-sized transmitters, the methods are rapidly translating to body transmit coils at 3T where clinical need is the highest and new generations of scanners offer a system architecture which is pTx-ready. Of course, intriguing research questions remain open in several areas such as the development of rapid and robust RF pulse designs that extend the current low-flip angle domain to arbitrary excitation angle, such as spin echoes, saturation, and inversions pulses. However, with continued development in these areas, progress is likely to accelerate, ultimately supporting fast, subject- and application-tailored RF pulse designs to extend MR excitation from the simple slice-select to the more generally tailored anatomy- or application-specific RF excitation pattern.

Acknowledgements

The authors would like to acknowledge the many researchers, students, and postdocs at Siemens, MGH and MIT whose work is summarized here. We especially acknowledge Kawin Setsompop, Vijay Alagappan, and Adam Zelinski whose thesis work was reviewed here. We also thank Ulrich Fontius and Andreas Potthast for their work setting up the 8-channel 3T and 7T systems and Franz Hebrank and Franz Schmitt for their leadership role in the collaboration and Michael Hamm, Josef Pfeuffer, Axel vom Endt, Mattias Gebhardt and Hans-Peter Fautz for their on-going support.
We acknowledge grant support from the NIH (P41RR14075, R01EB007942, and R01EB006847) and a research agreement and research support from Siemens Healthcare and the Siemens-MIT Alliance.

Contact
Lawrence L. Wald
Associate Professor of Radiology
Athinoula A. Martinos Center for Biomedical Imaging
Department of Radiology
Massachusetts General Hospital
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 Division of Health Sciences and Technology
Massachusetts Institute of Technology
elfar@mit.edu

References

1 Pauly, J., D. Nishimura, and A. Macovski, *A k-space analysis of small-tip angle excitation.* J Magn Reson, 1989. 81: p. 43–56.
2 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.
3 Setsompop, K., et al., *Parallel RF transmission with eight channels at 3 Tesla.* Magn Reson Med, 2006. 56(5): p. 1163–71.
4 Zelinski, A.C., 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(3): p. 176–190.
5 Adalsteinsson, A., et al., *Method For Reducing Maximum Local Specific Absorption Rate In Magnetic Resonance Imaging,* United States of America Serial No. 12/580076, Filed October 15, 2009.
6 Bornert, P., J. Weller, and I. Graesslin. *SAR Reduction in Parallel Transmission by k-Space Dependent RF Pulse Selection. in Proceedings of the ISMRM.* 2009. Honolulu Hawaii: p. 2600.
7 Graesslin, I., et al. *SAR Hotspot Reduction by Temporal Averaging in Parallel Transmission. in Proceedings of the ISMRM.* 2009: p. 176.
8 Zeliniski, A.C., *Improvements in Magnetic Resonance Imaging Excitation Pulse Design, in Electrical Engineering and Computer Science.* PhD Thesis, 2008, Massachusetts Institute of Technology: Cambridge MA. http://hdl.handle.net/1721.1/45862
9 Graesslin, I. *A Minimum SAR RF Pulse Design Approach for Parallel Tx with Local Hot Spot Suppression and Exact Fidelity Constraint. in Proceedings of the ISMRM.* 2008. Toronto Canada: p. 621.
10 Gebhardt, M., et al. *Evaluation of maximum local SAR for parallel transmission (pTx) pulses based on pre-calculated field data using a selected subset of "Virtual Observation Points". in ISMRM.* 2010.
11 Lee, J., et al. *Parallel Transmit RF Design with Local SAR Constraints. in Proceedings of the ISMRM.* 2010. Stockholm Sweden.
12 Lattanzi, R., et al., *Electrodynamic constraints on homogeneity and radiofrequency power deposition in multiple coil excitations.* Magn Reson Med, 2009. 61(2): p. 315–34.
13 Butler, J. and R. Lowe, *Beamforming matrix simplifies design of electronically scanned antennas.* Electron. Design, 1961. 9: p. 170–173.
14 Alagappan, V., et al. *A Simplified 16-channel Butler Matrix for Parallel Excitation with the Birdcage Modes at 7T. in International Society for Magnetic Resonance in Medicine.* 2008. Toronto, Canada: p. 144.
15 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: p. 619.
16 Zelinski, A., et al. *Sparsity-Enforced Coil Array Mode Compression for Parallel Transmission. in International Society for Magnetic Resonance in Medicine.* 2008. Toronto, Canada: p. 1302.
17 Gagoski, B.A., et al. *Real time RF monitoring in a 7T parallel transmit system. in Proceedings of the ISMRM.* 2010. Stockholm Sweden.
18 Makhoul, K., et al. *SAR Monitoring and Pulse Design Workflow in Parallel Transmission at 7 Tesla. in Proceedings of the ISMRM.* 2010. Stockholm Sweden.