MIT Libraries | DSpace@MIT

MIT Open Access Articles

Comparison of simulated parallel transmit body arrays at 3 T using excitation uniformity, global SAR, local SAR, and power efficiency metrics

Massachusetts Institute of Technology

Comparison of simulated parallel transmit body arrays at 3 T using excitation uniformity, global SAR, local SAR and power efficiency metrics

Bastien Guérin[1], **Matthias Gebhardt**[2], **Peter Serano**[1], **Elfar Adalsteinsson**[3,4], **Michael Hamm**[2], **Josef Pfeuffer**[2], **Juergen Nistler**[2], and **Lawrence L. Wald**[1,4]

[1]Martinos Center for Biomedical Imaging, Dept. of Radiology, Massachusetts General Hospital, Charlestown USA

[2]Siemens Healthcare, Erlangen Germany

[3]Dept of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge USA

[4]Harvard-MIT Division of Health Sciences Technology, Cambridge USA

Abstract

Purpose—We compare the performance of 8 parallel transmit (pTx) body arrays with up to 32 channels and a standard birdcage design. Excitation uniformity, local SAR, global SAR and power metrics are analyzed in the torso at 3 T for RF-shimming and 2-spoke excitations.

Methods—We used a fast co-simulation strategy for field calculation in the presence of coupling between transmit channels. We designed spoke pulses using magnitude least squares (MLS) optimization with explicit constraint of SAR and power and compared the performance of the different pTx coils using the L-curve method.

Results—PTx arrays outperformed the conventional birdcage coil in all metrics except peak and average power efficiency. The presence of coupling exacerbated this power efficiency problem. At constant excitation fidelity, the pTx array with 24 channels arranged in 3 z-rows could decrease local SAR more than 4-fold (2-fold) for RF-shimming (2-spoke) compared to the birdcage coil for pulses of equal duration. Multi-row pTx coils had a marked performance advantage compared to single row designs, especially for coronal imaging.

Conclusion—PTx coils can simultaneously improve the excitation uniformity and reduce SAR compared to a birdcage coil when SAR metrics are explicitly constrained in the pulse design.

Keywords

Local SAR; global SAR; power; parallel transmit; pTx coils; excitation fidelity; spokes

Corresponding author: Bastien Guerin, Athinoula A. Martinos Center for Biomedical Imaging, 149 Thirteenth Street, Suite 2301, Charlestown, Massachusetts 02129, guerin@nmr.mgh.harvard.edu.

Introduction

The degrees of freedom (DOFs) provided by parallel transmission (pTx) coils can be used to trade-off excitation fidelity, SAR (local and global) and power. We have shown previously that the optimal operating point of a pTx array, i.e. the radio-frequency (RF) pulse achieving the optimal trade-off between the above-mentioned quantities, can be quickly computed using a constrained optimization strategy that simultaneously constrain local SAR, global SAR, peak and average power on every channel (1). In this approach, computationally efficient local SAR constraint over the whole-body is achieved using a compression of the local SAR matrices called virtual observation points (VOPs) (2). A significant advantage of the VOP compression over alternative techniques including that proposed by Sbrizzi et al. (3) is the guarantee that the compression error always results in an overestimation of local SAR. This property ensures the safety of the patient by guaranteeing that the maximum tolerated local SAR limit is never exceeded. Design of RF pulses using a constrained optimization strategy as opposed to a regularized algorithm allows generating the best possible excitations consistent with regulatory (i.e., SAR) and system (i.e., power) limits without user intervention. Regularized pulse design approaches are also able to control local SAR (4), global SAR (5) and power (6,7) however they require manual tuning of the Lagrange multipliers associated to every controlled quantity. This becomes burdensome as the number of transmit channels increases and when controlling local SAR (e.g., simultaneous constraint of local and global SAR as well as power on every channel typically requires control of a few hundreds parameters) (8). A drawback of constrained optimization is that it is slower than regularized optimization because the Lagrange multipliers are solved for in addition to the RF pulse. However, we have shown that constrained spoke pulse design with more than 11 spokes (which is much greater than what is usually required) and more than 1,300 SAR and power constraints can be made fast enough for use in clinical environment using a dedicated primal-dual algorithm (i.e., <10 seconds for a single LS spoke pulse on an Intel i7 2.80 GHz CPU and <2 minutes for an MLS spoke pulse) (1).

Intuitively, it is clear that increasing the number of transmit channels should improve the ability of such a pulse design algorithm to find advantageous trade-offs between excitation fidelity, SAR and power. However, given the cost of high-power RF amplifiers, it is important to determine the incremental performance improvement brought by additional transmit channels under realistic power budget assumptions for specific pTx coil geometries and configurations. In analogy with the optimization of receive coils, the benefit of additional transmit channels also likely depends on the imaging application, in particular the region of the body being imaged and the amount of acceleration needed (9-11). Finally, we point out that since SAR, and more specifically local SAR, is often the limiting factor in pTx imaging (1,4,12-15), it is important to evaluate not only the encoding capability of pTx coils but also their ability to decrease, or at least maintain, SAR. A good way to characterize the SAR vs. excitation fidelity trade-off of a pTx array is to plot on a single graph different pulses achieving different trade-offs between these two quantities (L-curve). L-curves associated with different arrays on a single graph allow comparing their relative performance for a wide range of operating points. It is important not to concentrate on a

single operating point for each design as L-curves may cross, indicating a reversal in the relative performance of two coil arrays for a certain threshold value of the excitation fidelity.

Harvey et al. used finite element modeling (FEM) of a 16-channels 3 T body pTx coil loaded with a simplified human torso model to assess the impact of reducing the number of transmit modes on the static RF-shimming excitation fidelity and power (16). They found that using two transmit modes (equivalent to two channels) improved the uniformity of static RF-shimming by 80% compared to conventional quadrature birdcage excitations. This benefit came at the price of a 50% increase in power consumption however. They also found that increasing the number of transmit modes beyond two did not significantly improve the quality of RF-shimming excitations but increased power consumption even further. This work was limited to static RF-shimming and did not consider other performance metrics such as global SAR and local SAR. Furthermore, the authors used an un-regularized pulse design approach, which made it impossible to explore a range of operating points of the arrays (they plot a single operating point). Kozlov and Turner proposed a co-simulation strategy allowing fast simulation of coupled pTx coils (17) and used it to compare the performance of coils with various numbers of transmit channels in both the circumferential and the z-direction (18,19). Although they did find significant differences between the fields created by each of these arrays, they did not perform pulse design and therefore were not able to assess the impact of additional transmit channels on pTx performance.

Lattanzi et al. (20,21) proposed an analytical framework for the calculation of the lowest possible global SAR consistent with a given excitation quality for static RF-shimming. To do so, they constructed a basis set of electromagnetic fields in uniform spheres and used the elements of this basis as the "transmit channels" in a pulse design algorithm. As the rank of the basis set was increased (i.e., more and more modes were added to the basis), the global SAR vs. excitation fidelity trade-off approached the ultimate trade-off achievable by any coil. PTx coils with uncoupled circular loops were also analyzed. This study showed that increasing the number of transmit channels improves the global SAR vs. excitation fidelity trade-off. However, this type of analytical approach is limited to the simulation of uniform spheres and cannot handle local SAR. Childs et al. compared the RF power vs. excitation fidelity trade-off of a coil with increasing number of transmit modes (up to 8 independent modes) in actual patients (22). They found that increasing the number of transmit channels improved the ability of the coil to produce homogeneous excitations at low RF power, which is somewhat in contradiction with the study by Harvey et al. (16) as well as our own results. They did not study the dependence of local SAR with the number of channels. Finally, Tian et al. (23) simulated stripline pTx coils with up to 48 channels arranged in up to 3 z-rows and showed that these were capable of producing more uniform static RF-shimming excitations than a birdcage coil of similar dimension at the cost of increased power consumption. Although they report local SAR and global SAR for these pulses and arrays, they did not optimize these quantities in the pulse design process and therefore could not study the optimal performance of these coils with respect to these metrics.

In this work, we used a co-simulation strategy similar to that developed by Kozlov and Turner (17) to compute the fields produced by eight pTx arrays with up to 32 independent channels arranged in up to 4 rows (z-directions). Coupling between the transmit channels

was modeled. All designs were compared to a high-pass birdcage coil of similar dimension. We performed this study at 3 T in a realistic body model. We studied the capability of the different designs to excite uniform slices in the liver (slice-selective RF-shimming and 2-spoke pulses) and to reduce global SAR and local SAR using the pulse design algorithm mentioned earlier in this section (1).

Methods

Co-simulation of coupled arrays

We simulated 8 pTx arrays shown in Fig. 1 (the notation Xr/Ycpr denotes a coil made of X rows or Y coils per row, which are in the z-direction). The coils were loaded with the Ansys body model, which contains 33 tissue classes, and were centered on the liver. All coils had a cylindrical geometry with the same overall dimensions (radius of 35.2 cm, length in the z-direction of 35.0 cm, shield radius of 37.2 cm, shield length of 158 cm). The RF shield was modeled in these simulations. Transmit elements were 5 mm wide copper loops separated by gaps of 2 cm in both the circumferential and the z-directions. The loops were broken by tuning capacitors regularly placed along their circumference (each loops of the single-row arrays were broken by 8 tuning capacitors. The loops of the multi-row arrays were broken by 4 tuning capacitors). We also simulated a high-pass birdcage coil with 16 rungs with the same radius and length. The power budget available for each pTx coil simulated is indicated on Fig. 1. Although the total power budget (i.e., the sum of the power available on all channels) was not constant across coils, these represent reasonable power limits that allow assessment of the performance of these designs under realistic conditions.

Fig. 2 shows the flowchart of the co-simulation process. This strategy is almost the same as that proposed by Kozlov et al. (17). The only difference is that we compute the electromagnetic fields and S-matrix of the loaded coil at a single frequency (the Larmor frequency), as opposed to several frequencies as performed by Kozlov et al. As explained below, in order to display tuning curves, we extrapolated the S-matrix at the Larmor frequency using an analytical formula introduced by Lemdiasov et al (24) (note that the fields and S-parameters at the Larmor frequency were known from the HFSS simulation and were not modified by this extrapolation step). In the Step 1 of Fig. 2, the entire S-matrix of the coil/shield/body model structure as well as the fields associated with each port were computed at the Larmor frequency. No lumped elements were modeled in this step; instead all capacitors and sources were replaced by ports. In Step 2, the S-matrix was estimated over a 50 MHz range around the Larmor frequency by assuming a purely inductive frequency response (24). In Step 3, the extrapolated S-matrix was loaded in the circuit simulator and lumped elements and sources were re-instated on the model. In this work we simulated lossless lumped elements with an infinite quality factor (however copper was modeled with a high but finite conductivity, therefore losses in the conductive elements of the coils were properly modeled). Tuning/matching/decoupling capacitors were optimized in this step using the optimization routines of the circuit simulator as described in the next section. Separation of the field/S-matrix calculation from the optimization of lumped elements values allows fast simulation of a large number of tuning/matching/decoupling conditions. In Step 4, the current and voltage across the ports of the coil (including lumped elements

ports) were computed in the circuit simulator and were used to scale the fields computed in Step 1 (17). The B1+ maps and electric fields created by each transmit element were then computed by exciting the voltage sources in the circuit simulator one at a time (note that for sources with a 50 Ω internal impedance, there is a constant relationship between voltage and power as shown in Supplemental Figure S5. Therefore we use the terms "power source" and "voltage source" interchangeably). Local SAR matrices were computed at every location of the body model by taking the outer product of the vector of electric fields created by all channels with itself and scaling the result by the ratio of the conductivity and the density divided by two as described in Refs. (2-5,20). Compression of the local SAR matrices and computation of excitation pulses were performed using in house programs written in C++ (see next section).

Any field solver can be used in Step 1 and any circuit simulator can be used in Step 3 of this co-simulation process. In this work, we use the High Frequency Structural Simulator (HFSS, Ansys, Canonsburg PA) as the field solver and Advanced Design System (ADS, Agilent, Santa Clara CA) as the circuit simulator. HFSS is a finite element modeling (FEM) solver that allows computation of the entire S-matrix as well as all fields associated with all ports in a single simulation run (in contrast, finite difference time domain, or FDTD, requires N simulations to estimate the S-matrix of a coil with N ports). This leads to a significant reduction in computation time when simulating RF coils with many ports. Another advantage of FEM solvers over FDTD is that they allow simulating a great variety of spatial scales without significant computational penalty. This is crucial in the simulation of RF coils used in MRI since these are commensurate with the human torso but are made of a very thin strips of copper of thickness <10 um (in FDTD, computation time is inversely related to the smallest discretization length. This essentially prohibits simulation of thin copper strips, which must therefore be approximated by elements of thickness >1mm). FEM solvers like HFSS require much more memory than FDTD solvers however. For example, simulation of the 1r/4cpr design, which contained 40 ports, took 7 hours using 32 processors (Intel Xeon X7560 2.27 GHz) and required 90 GB of memory. Simulation of the 3r/8cpr design, which contains 208 ports, took 46 hours and required 181 GB of memory (see Supplemental Figure S1 for more details on the memory requirement and convergence of the HFSS simulations).

Optimization of the tuning/matching/decoupling capacitors (Fig. 2, Step 3) was performed after the field/S-matrix computation step (Fig. 2, Step 1) in the circuit simulator. Coils were tuned to 123.2 MHz and matched to better than -30 dB by optimization of the tuning and matching capacitors using the gradient routine of ADS. This process took only a few seconds. Adjustment of the decoupling capacitors was more difficult however. We found that a good initial guess for the value of these capacitors is essential to ensure that gradient-based optimization routines converge to an acceptable decoupling of the entire coil (we attempted decoupling any pair of coils to better than -15 dB). This is due to the fact that the decoupling problem is non-convex and gradient-based optimization algorithms only guarantee convergence to a local minimum. We computed the initial guess of decoupling capacitors by decoupling only two loops of the array. This was performed by first isolating a loop of the model and tuning and matching it to 123.2 MHz and -30 dB (loops can be electrically removed from the model by setting their tuning capacitors to very small values which prevents current from flowing continuously on the copper strip). A nearest-neighbor

loop was then added to the model, causing a split in the S11 resonance of the first loop. The two peaks of the split S11 resonance were finally merged into a single peak by manual adjustment of the decoupling capacitors between the two loops and the value of the decoupling capacitors were refined further using the gradient optimization routine. Each array was also studied in an ideally decoupled configuration, which is only possible in simulation by electrically isolating each loop of the model. The cost function used for tuning, matching and decoupling with the gradient routine of ADS was the sum of the square of the difference between the target and achieved S-parameters values.

SAR and power constrained pulse design

We designed slice-selective RF-shimming (1 spoke) and 2-spoke pulses for imaging of the liver area in the transverse and coronal plane (all imaging planes passed though the isocenter). The target excitation was a flat distribution with a flip-angle of 40 degrees. The arms of the body model were excluded from the optimization mask, which followed the outer contour of the body, but were included in the SAR calculation. The objective function minimized by the pulse design algorithm was a magnitude least squares term (MLS) that did not penalize the phase of the magnetization profile (25). When designing 2-spoke pulses, we placed the first spoke at the origin of excitation k-space and the second spoke at (k_x=3, k_y=-3, k_z=0) in units of [1/m] (the torso is approximately a cylinder of diameter 30 cm, thus full field-of-view Fourier imaging requires sampling at a spatial frequency of $1/0.3 \approx 3$ [1/m]). For both RF-shimming and 2-spoke pulses, the basic pulse waveform was a 3-lobe (4 zero-crossings) Hamming apodized sinc profile of 0.8 ms duration. We designed the same pulses with the birdcage coil as well as an additional sinc pulse of 1.6 ms duration to allow an equal pulse-length comparison with the 2-spokes excitations. Slice thickness was 5 mm for all pulses.

We computed pTx spokes pulses with explicit constraint for local SAR, global SAR, peak and average power on all channels using a fast primal-dual algorithm (1). Local SAR was constrained in the whole-body in a computationally efficient manner using the virtual observation points (VOP) compression algorithm of local SAR matrices (2). The VOP algorithm is applied after 10 g averaging the local SAR matrices; therefore the VOPs also reflect the 10 g averaging process. Ten grams averaging was performed by growing concentric cubes of tissues around the voxel of interest until a region of mass 10 g was achieved. Voxels not belonging to the model (i.e., air) were excluded from the averaging region (2). We define the following power quantities:

$$
\begin{aligned}
a) & \quad P_c(t_i) = \frac{|V_c(t_i)|^2}{8Z_0} \\
b) & \quad P_c^{peak} = \max_i P_c(t_i) \\
c) & \quad P_c^{av.} = \frac{1}{N} \sum_{i=1}^{N} P_c(t_i) \\
d) & \quad P^{tot} = \sum_{c=1}^{C} P_c^{av.}
\end{aligned}
\quad [1]
$$

where c is a channel index, i is a time index and N is the total number of time points in the RF pulse (including the imaging gradients and data acquisition module, therefore this term

accounts for the duty-cycle of the RF pulse). $V_c(t)$ is the source voltage time course for channel c. The quantities a), b), c) and d) are the instantaneous power, peak power on channel c, average power on channel c and total forward power, respectively. Electrically, the voltage source and the load form a voltage divider with R1=50 Ω (internal resistance of the source) and R2 equal to the input impedance of the load. When the load is matched to the source, R1=R2=50 Ω, so that the voltage across the load is half of the source nominal voltage (this, plus the fact that we use amplitude and not RMS voltage quantities, explains the factor of 8 between voltage and power in Eq. [1] – for more details see the Supplemental Figure S5). We used the same power metrics (defined in Eq. [1]) to characterize the power consumption of the capacitively and ideally decoupled arrays.

To compare the performance of different pTx designs, we designed MLS spoke pulses subject to fixed power constraints (these are indicated on Fig. 1) but by varying the local SAR constraint. The global SAR constraint was fixed to the Federal Drug Administration (FDA) tolerated whole-body value of 4 W/kg (26). The set of pulses associated with different local SAR limits define an L-curve showing the optimal trade-off between excitation fidelity and local SAR. In this work, excitation fidelity was defined as the root mean square error between the achieved and target flip-angle maps, and was expressed as a percentage of the target flip-angle (this quantity is computed within the optimization mask that does not include the arms). We report the excitation fidelity, local SAR (computed using the original local SAR matrices averaged over 10 g, not the compressed representation via VOPs), global SAR, peak and average power of every pulse. SAR and average power values were computed assuming a 10% duty-cycle (this choice of the duty-cycle corresponds to relatively fast pulses. Slow pulses with long TRs can have duty-cycles well below 1%, in which case SAR is rarely a limitation). The MLS pulse design problem is non-convex, which can lead to L-curves that do not vary smoothly ("jumps" in SAR and/or excitation fidelity performance can occur because of convergence of the design algorithm to different local minima). In order to minimize this effect that could potentially make it difficult to compare the L-curves of different coils, L-curves were computed by initializing the pulse design algorithm with the pulse computed previously.

Results

Extrapolation of the admittance matrix

We assessed the accuracy of the inductive extrapolation of the admittance matrix by comparing the extrapolated S-parameters to the actual S-parameters computed with HFSS in a 50 MHz window centered at the Larmor frequency (123.2 MHz) for the 8 channels pTx coil 1r/8cpr. The average difference between the extrapolated and the actual reflection coefficients (S11) was 0.014 dB. The average difference between the extrapolated and the actual transmission coefficients (S21) was 0.185 dB. When plotted on the same graphs, the reference and extrapolated transmission and reflection S-parameters were essentially undistinguishable.

Importance of decoupling

Matching was better than -30 dB for all coils (the real part of the input impedances were 49.6±2.1 Ω in average for all ports and coils and the imaginary parts were 0.3±2.0 Ω). Fig. 1 shows that the designs 1r/4cpr, 1r/8cpr, 1r/16cpr and 2r/4cpr could be well decoupled by adjustment of the decoupling capacitors placed between nearest neighbors using ADS. Our decoupling strategy failed however for the other designs (these showed residual coupling of -5 dB or worst). This indicates two possible problems. First, there is no guarantee that our decoupling procedure, despite our efforts to initialize it properly, converged to the global optimum. Second, it is not clear that this decoupling problem (decoupling the entire array by decoupling the nearest-neighbors) admits a solution for coils with more than 16 coils per row and/or more than 2 rows. This is because as the number of coils increases, the number of pair of coils to decouple increases quadratically whereas the number of decoupling degrees of freedom (DOFs) increases linearly. Ideally, for perfect decoupling of the arrays, a capacitor should be placed between any two coils. This and more sophisticated types of decoupling strategies are being explored (27-30) but in this work we used the more common approach consisting in placing decoupling capacitors only between nearest neighbors. Thus the designs 1r/4cpr, 1r/8cpr, 1r/16cpr and 2r/4cpr were studied in both their capacitively ideally decoupled configurations and the designs 1r/24cpr, 2r/8cpr, 3r/8cpr and 4r/8cpr designs were studied only as ideally decoupled.

Also shown in Fig. 1 is the cumulative sum of the singular values (SVs) of the B1+ maps created by each coil in a volume of interest containing the torso but excluding the arms. The cumulative sum of the SVs was computed by stacking the B1+ maps of each coil in a matrix of size N×C (N is the number of pixels, C is the number of channels), plotting the cumulative sum of the SVs of this matrix and normalizing the result setting the first SV of each coil to 1 (31). If the B1+ maps created by the loops of a pTx array were non-overlapping, the cumulative sum of the normalized SVs would be equal to the number of transmit channels. Instead, there is significant overlap between the excitations created by the different channels and the sum of the relative SVs is smaller than the channel count. For the designs 1r/16cpr and 1r/24cpr, the cumulative distribution of SVs plateaus, indicating that using more than ~12 channels in a single row did not significantly increase the encoding capability of the array. Instead, it seemed beneficial to distribute additional channels beyond the first 12 in the z-direction (i.e., multi-row arrays). This was confirmed by the pulse design trade-offs shown in Figs. 5-8. Another interesting remark is that residual coupling between transmit elements almost always decreased the rank of the B1+ basis set created the coil (this is true for all designs except 2r/4cpr). This is likely due to the fact that coupling tends to make the B1+ profiles "more similar".

Fig. 3 shows SAR, power and excitation fidelity trade-offs for slice-selective RF-shimming with the designs 1r/8cpr and 2r/8cpr, both capacitively decoupled and ideally decoupled. The design 1r/8cpr could be well decoupled by adjustment of the decoupling capacitors whereas the array 2r/8cpr could not (see Fig. 1). The presence of relatively large residual coupling in the 2r/8cpr array increased the power consumption compared to its ideally decoupled counterpart. Because the power available on each channel was limited (to 8 kW peak and 800 W average), this essentially restricted the regions of the local SAR/global

SAR/excitation fidelity space accessible to the coupled array. In contrast, good decoupling of the 1r/8cpr design lead to a performance close to ideal both in term of SAR (global and local) and power. Supplemental Figure S6 shows similar L-curves for 2-spoke pulses.

Performance comparison of pTx coils

Fig. 4 shows transmit efficiencies and local SAR to global SAR ratios for the birdcage coil and the pTx coils driven in their uniform birdcage mode. The transmit efficiency is the average B1+ created by each array in a volume of interest corresponding to the torso of the body model but excluding the arms. The birdcage coil had the greatest transmit efficiency and the smallest local to global SAR ratio. These results show that birdcage coils outperform pTx coils in exciting birdcage modes and pTx coils should be driven using optimized RF pulses that make full use of their additional degrees-of-freedom. It is interesting to note that small residual coupling did not necessarily result in a smaller transmit efficiency of the uniform birdcage mode (small residual coupling reduced the transmit efficiency for 1r/8cpr and 2r/4cpr but not for 1r/4cpr and 1r/16cpr). Indeed, complex spatial interferences between the magnetic fields created by the channels of a pTx coil make it difficult to predict the net resulting strength of the uniform birdcage mode B1+ (coupling affects the shape of the fields associated with every channels, which may or may not result in constructive interferences boosting the total B1+). However, strong coupling between the transmit channels always resulted in power loss and therefore decreased the transmit efficiency significantly (e.g., 2r/8cpr).

Fig. 4 also shows that small residual coupling could have a significant influence on the local to global SAR ratio. Supplemental Figure S7 shows that the maximum SAR hotspot of the birdcage modes of the coils occurred at the same location whether the arrays were ideally or capacitively decoupled configurations, except for 2r/4cpr. For this array, small residual coupling moved the SAR hotspot from the armpit (ideal decoupling) to the elbow (capacitive decoupling). This new location of the hotspot (elbow) corresponds to the junction between the two rows of this coil, where even relatively small residual coupling can cause the distant coils to create non-negligible electric fields.

Figs. 5 shows trade-offs between local SAR, global SAR, peak and average power and excitation fidelity for slice-selective RF-shimming in the transverse plane. PTx coils with increasing number of transmit channels could create high quality excitations at lower local and global SAR than the birdcage coil (Fig. 9a). Increasing the number of transmit channels also allowed creation of more uniform excitations at a given SAR (Fig. 10a). For the same pulse length and same excitation fidelity, the designs 1r/4cpr, 1r/8cpr, 1r/16cpr, 1r/24cpr, 2r/4cpr, 2r/8cpr, 3r/8cpr and 4r/8cpr could reduce local SAR by 30%, 70%, 74%, 75%, 62%, 75%, 81% and 82%, respectively, compared to the birdcage coil. The 3r/8cpr design was able to reduce local SAR from 14.1 W/kg to less than 2.5 W/kg, a dramatic reduction that could potentially allow a 5.6-fold reduction in TR or a similar increase in the number of slices at constant acquisition time if local SAR is the only limit to reduction of the sequence duration. This improvement in the SAR vs. excitation fidelity performance came at the price of a large increase in power consumption however. At constant excitation fidelity, the total average power consumed by the designs 1r/4cpr, 1r/8cpr, 1r/16cpr, 1r/24cpr, 2r/4cpr, 2r/

8cpr, 3r/8cpr and 4r/8cpr, in this order, were 2.4, 3.9, 9.9, 10.5, 4.5, 4.6, 5.9 and 10.3 times greater than that of the birdcage coil. The fact that more power was needed to reduce local SAR and global SAR can seem counterintuitive. Part of this confusion stems from our definition of power. The total forward power defined in Eq. [1] does not reflect the power actually absorbed by the load because this metric is computed without consideration of the phase relationship between the waveforms played on different channels. Instead, this quantity represents an upper bound of the power actually transmitted to the load. It is interesting to note that the 1r/24cpr design did not perform better than 1r/16cpr. This is partly due the fact that the additional channels of the 1r/24cpr design did not significantly increase the pTx DOFs as shown in Fig. 1. Also, it is evident from its peak power vs. excitation fidelity L-curve that this design was strongly power limited. This is partially due to the fact that the smaller loops of the 1r/24cpr design have a lower B1+ penetration in the body than the loops of the 1r/16cpr design and therefore need to be driven with more power (this is reflected in the B1+ transmit profiles shown in Supplemental Figure S3). The white bar graphs on the right side of the flip-angle maps of Fig. 9 indicate that local SAR and peak power were the limiting constraints (note that this result is specific to the target flip-angle used and the duty-cycle and the specific shape of the slice-selective sinc pulse simulated).

Figs. 6 and 10 show that using two spokes allowed the pulse design algorithm to improve excitation uniformity and lower SAR compared with RF-shimming. Note that the duty-cycle of all pulses is constant in this study (10%) so that the 2-spoke pulse length was twice the duration of the RF-shimming pulse. If the RF-shimming pulse length was increased 2-fold to match that of the 2-spoke excitation, this would lower SAR 4-fold and would achieve a lower SAR than that of the spoke excitation for most operating points. However, the 2-spoke pulses would still produce more uniform excitations.

When comparing the arrays driven by 2-spoke pulses, we calculated three different pulses for the birdcage coil: (i) RF-shimming with a pulse duration equal to that of a single spoke sub-pulse; (ii) RF-shimming with a pulse duration equal to the duration of the 2-spoke pulse; (iii) 2-spoke pulses as designed on the pTx arrays. All pTx coils except 1r/4cpr beat the birdcage coil pulses both in term of SAR and excitation fidelity for both RF-shimming and 2-spoke excitations. The best excitations in terms of homogeneity vs. SAR trade-off were obtained using 2-spoke MLS pulses played on pTx arrays with multiple rows of at least 8 channels per row.

Figs. 7 and 8 show similar SAR/power/excitation fidelity trade-offs for imaging in the coronal plane. The results are similar: Increasing the number of transmit channels improved the ability of the pulse design algorithm to find favorable SAR vs. excitation fidelity trade-offs. This came at the price of increased power consumption however. Unlike in transverse imaging, there was a clearer distinction in performance between single row arrays and those with multiple rows stacked in the z direction. Supplemental Figures S8 and S9 show flip-angle maps (at constant local SAR) and SAR maximum intensity projection maps (at constant excitation fidelity) corresponding to the RF-shimming and 2-spoke pulses shown in Figs. 7 and 8.

Discussion

We have analyzed the performance of pTx coils with up to 32 channels arranged in up to 4 rows in the z-directions using excitation fidelity, local SAR, global SAR, peak and average power as performance metrics. Our simulations show that pTx coils with increasing number of channels can create more uniform excitations than a birdcage coil of similar dimensions at much lower local SAR and global SAR. The "price to pay" for such performance improvement is an increase in power consumption by the pTx coils. These results are in agreement with the work of Lattanzi et al (20,21) and extend their conclusions to non-uniform body models and more realistic coil models operating under limited power budget. In a separate work, they also showed improving RF shimming in both axial and coronal planes with increased number of channels in multi-row geometries (32).

To assess the performance of different transmit coils, we used a co-simulation strategy allowing simulation of coupled arrays and fast optimization of tuning, matching and decoupling capacitors (17). Pulses were designed using a constrained optimization algorithm that we recently developed which allows design of magnitude least squares (MLS) pulses with the best possible flip-angle distribution consistent with the main regulatory (i.e., local and global SAR) and system limits (i.e., peak and average power on every channel). By varying the local SAR constraint limit and keeping all other constraints constant, we were able to compare the performance of the different pTx coils and the birdcage coil for a range of operating points in term of local SAR, global SAR, peak and average power and excitation fidelity while assuming a realistic power budget. A limitation of this study is that the locations of the spokes were fixed for all coils. This is potentially a confounding effect as our specific choice of the spoke locations may be more or less optimal for different coils. Ideally, spoke locations (when using two spokes and more) should be optimized for every coil in order to compare their true optimal performance. Optimization of the spoke locations is still a subject of active research however and in this work we chose to use a simpler design strategy with fixed spoke locations.

We attempted decoupling the coils simulated by adjustment of capacitors placed between nearest neighbors (capacitive decoupling). We found that optimization of these decoupling capacitors, which is non-convex, was challenging and that a good solution could only be obtained using a gradient-based optimizer algorithm by proper initialization of the decoupling capacitor values. We computed this initial guess by decoupling only two loops of the models. Even using this strategy, we were not able to decouple the designs 1r/24cpr, 2r/8cpr, 3r/8cpr and 4r/8cpr. A better strategy, which we did not evaluate, may consist in decoupling nearest-neighbors within the same row using capacitors and in decoupling adjacent rows by overlap as studied by Gilbert et al. (33). There are also indications that inductive decoupling may be more efficient that capacitive decoupling in decoupling multi-row arrays (19,34). Additionally, we found that the close proximity of the shield to the RF coil (2.2 cm gap in this work) significantly increased coupling between the transmit channels as the shield essentially acts as additional loops strongly coupled to the pTx channels. Other, more advanced decoupling strategies are being investigated that show promises in solving this problem (29,30), but we did not investigate those in this work. In order to evaluate pTx coils with multiple rows and many channels even when those could

not be well decoupled by adjustment of nearest neighbor capacitors, we studied these arrays in their ideally decoupled configurations.

We found that the main impact of the presence of residual coupling between transmit channels was an important increase in the power consumption of the arrays. With limited input power, the SAR vs. excitation fidelity performance of coupled arrays worsened because they reached their peak and average power limits more quickly than if they were well decoupled. Analysis of the singular values of the B1+ maps of coupled and ideally decoupled pTx coils also revealed that coupling between the transmit channels decreased the encoding capability of the array. These two phenomenon are likely related as the loss of pTx DOFs due to coupling caused the spoke pulse design problem, which is already highly underdetermined, to be less well conditioned. This, in turn, likely caused the pulse design algorithm to seek high-power pulses that decreased the excitation error and SAR only marginally. In actual coils, the power transmitted from one channel to another due to coupling can be absorbed via circulators that protect the RF power amplifiers from power "going upstream" (this power is lost for excitation). These observations are in agreement with simulation results by Kozlov and al. showing that SAR and the power consumption of pTx arrays depends significantly on their tuning, matching and decoupling conditions (35). Note that substantial residual coupling at the level seen in the 2r/8cpr design had significant impact on SAR (Figs. 3 and 4) but that small residual coupling (design 1r/8cpr, Fig. 3) did not worsen the SAR vs. excitation error tradeoff significantly.

Another conclusion of this work is that birdcage coils largely outperform pTx coils when these are driven in their uniform birdcage mode. This suggests that if pTx arrays are employed, their excitation DOFs should be exploited using optimized pulses. We point out that this result could be specific to the loop coils investigated in this work and may not generalize to all pTx geometries. For example, degenerate birdcage pTx coils (36) driven in their uniform birdcage mode may be expected to more closely resemble birdcage coils. This paper was limited to the evaluation of loop coil arrays. Because of gaps present between transmit loops in the designs simulated in this work, which are not present in a birdcage coil, the transmit efficiencies of the pTx coils driven in their birdcage mode differed from the efficiency of the birdcage coil. Other designs, such as TEM and striplines coils may perform differently and may be preferable for high field imaging. We also found that the local to global SAR ratio, which is widely used to characterize the SAR distribution of birdcage coils, is an unstable metric for pTx coils driven in their birdcage mode. This metric could vary greatly under small perturbations of the coupling and loading conditions. This is especially true for multi-row arrays where local SAR hotspots are more easily formed at the junction between two rows (Fig. 4 and Supplemental Figure S7). Again, this indicates that pTx coils should be driven using optimized pulses with explicit constraints for SAR.

An important conclusion of this work is that increasing the number of transmit channels improved the ability of the pulse design algorithm to create uniform magnitude excitations at low local and global SAR. Note that local SAR and global SAR must be constrained in the pulse design process for these results to hold. If unconstrained, it is in fact likely that the local and global SAR produced by pTx coils become greater than those created by a birdcage coil. This result is in agreement with the work of Lattanzi et al. (20,21) but is

somewhat at odds with the study of Harvey et al. (16), which found that increasing the number of transmit modes beyond two increased power consumption significantly but not the excitation uniformity. This is likely due to the fact that they used an un-regularized pulse design strategy in which power was allowed to increase rapidly and that they did not control or report SAR.

The improved SAR vs. excitation fidelity of pTx coils compared to the birdcage coil came at the price of an increase in power consumption. As noted in the Results sections, the fact that more power is needed to reduce SAR can seem counterintuitive and stems from our definition of power (1). The total forward power defined in Eq. [1] does not reflect the power actually absorbed by the load. This metric is computed without consideration of the phase between the waveforms played on the different channels and therefore represents an upper bound of the power actually absorbed by the load. Additionally, coil losses due to the finite conductivity of copper are included in the total forward power but not in the power absorbed by the load. Depending on the drive configuration, electric fields interfere destructively in the conductive sample (resulting in low SAR) while still requiring significant current in the conductive coil elements, which results in significant coil losses. Nonetheless, we retain our simple definition of total forward power as our main power metric because it is an important quantity used to characterize the power limits of RF amplifiers.

In this study, we evaluated the impact of coil placement not only around the patient circumference but also in the z-direction. Examination of the singular values of the B1+ maps created by coils with single and multiple z-rows indicate that there is little advantage in using more than 12 coils per row of the design studied. Instead, placing additional channels in the z-direction seemed to maximize the encoding capability of the array. Analysis of the trade-offs between SAR, power and excitation fidelity confirmed this trend, especially for non-axial slices: Multi-row designs with 8 coils per row performed consistently better than single-row designs with the same total number of channels. This performance improvement was most pronounced for coronal imaging. It is likely that off-center axial slices would also benefit from the additional pTx DOFs in the z-direction provided by coils with multiple rows. This conclusion is in agreement with a recent study by Tian et al (23) who observed in simulations that pTx coils with multiple z-rows produce more uniform RF-shimming excitations that birdcage coils and single-row pTx coils, especially for non-axial slices. Our work extends this result to multiple spoke pulses and to a wide range of operating points of the arrays corresponding to different global and local SAR levels. In other words, multi-row arrays do not only improve the excitation fidelity, they also have a better ability to reduce both global and local SAR than single-row coils. Lattanzi et al also studied the performance of multi-row pTx coils for body imaging (uniform cylinder) but did not compare their performance to that of single-row designs (32). Nonetheless, they also found that increasing the number of transmit channels improved the ability of the pulse design algorithm to produce uniform RF-shimming excitations.

Although the coils simulated in this work were only evaluated as transmit arrays, they can also be used for reception of the MR signal. Lattanzi et al. (37) have studied the SNR of coils with a similar geometry than those modeled in this work (i.e., cylindrical geometry,

several staggered loops in the z-direction) and have shown that these allow approximating the ideal current pattern corresponding to the "ultimate SNR" quite well as the number of elements increases. This indicates that these RF structures, which are general and allow imaging of any region of the body, may be efficient for reception as well as for transmission.

Finally, we point out that it is essential to evaluate the performance of pTx coils under realistic power budget constraints. This point is illustrated by the better performance of the 2r/8cpr design over the 3r/8cpr and 4r/8cpr designs for MLS RF-shimming in the coronal direction (Fig. 7) despite the smaller number of channels of the former coil. Analysis of the value of the Lagrange multipliers indicates that, although the three designs reached their respective maximum power constraints, the 3r/8cpr and 4r/8cpr designs were "more power constrained" than 2r/8cpr. Indeed, the maximum Lagrange multipliers associated to the peak power constraints were 0.72, 6.17 and 6.17 for the 2r/8cpr, 3r/8cpr and 4r/8cpr designs, respectively. We remind the reader that the unit of a Lagrange multiplier λ is such that the product λf, where f is the constraint function (here peak power expressed in Watts), has the same unit as the cost function. Therefore a Lagrange multiplier indicates the "cost" of enforcing a constraint (strictly speaking, the Lagrange multiplier is, at the optimal solution, the derivative of the cost function with respect to the constraint). Thus, for this particular imaging situation, raising the peak power limit improve the excitation quality for all designs but this improvement would be 8.5 times more pronounced for the 3r/8cpr and 4r/8cpr designs than for the 2r/8cpr design. This observation underscores the need for designing pTx pulses while constraining simultaneously SAR and power. This could not be done in the studies of Harvey et al. (16), Lattanzi et al. (20,21), Child et al. (22) and Tian et al. (23) because these authors used regularized pulse design strategies that made it difficult to control several quantities at once.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

References

1. Guérin B, Gebhardt M, Cauley S, Adalsteinsson E, Wald LL. Local specific absorption rate (SAR), global SAR, transmitter power, and excitation accuracy trade-offs in low flip-angle parallel transmit pulse design. Magnetic Resonance in Medicine. 201310.1002/mrm.24800

2. Eichfelder G, Gebhardt M. Local specific absorption rate control for parallel transmission by virtual observation points. Magnetic Resonance in Medicine. 2011; 66(5):1468–1476. [PubMed: 21604294]

3. Sbrizzi A, Hoogduin H, Lagendijk JJ, Luijten P, Sleijpen GLG, van den Berg CAT. Fast design of local N-gram-specific absorption rate–optimized radiofrequency pulses for parallel transmit systems. Magnetic Resonance in Medicine. 2012; 67(3):824–834. [PubMed: 22127650]

4. Lee J, Gebhardt M, Wald LL, Adalsteinsson E. Local SAR in parallel transmission pulse design. Magnetic Resonance in Medicine. 2012; 67(6):1566–1578. [PubMed: 22083594]

5. Zhu Y. Parallel excitation with an array of transmit coils. Magnetic Resonance in Medicine. 2004; 51(4):775–784. [PubMed: 15065251]

6. Grissom W, Yip C, Zhang Z, Stenger VA, Fessler JA, Noll DC. Spatial domain method for the design of RF pulses in multicoil parallel excitation. Magnetic Resonance in Medicine. 2006; 56(3): 620–629. [PubMed: 16894579]

7. Setsompop K, Wald LL, Alagappan V, Gagoski B, Hebrank F, Fontius U, Schmitt F, Adalsteinsson E. Parallel RF transmission with eight channels at 3 Tesla. Magnetic Resonance in Medicine. 2006; 56(5):1163–1171. [PubMed: 17036289]

8. Boyd, SP.; Vandenberghe, L. Convex optimization. Cambridge Univ Pr; 2004.

9. Wiggins GC, Polimeni JR, Potthast A, Schmitt M, Alagappan V, Wald LL. 96-Channel receive-only head coil for 3 Tesla: Design optimization and evaluation. Magnetic Resonance in Medicine. 2009; 62(3):754–762. [PubMed: 19623621]

10. Keil B, Alagappan V, Mareyam A, McNab JA, Fujimoto K, Tountcheva V, Triantafyllou C, Dilks DD, Kanwisher N, Lin W. Size-optimized 32-channel brain arrays for 3 T pediatric imaging. Magnetic Resonance in Medicine. 2011; 66(6):1777–1787. [PubMed: 21656548]

11. Keil B, Blau JN, Biber S, Hoecht P, Tountcheva V, Setsompop K, Triantafyllou C, Wald LL. A 64-channel 3T array coil for accelerated brain MRI. Magnetic Resonance in Medicine. 2012; 70(1):248–258. [PubMed: 22851312]

12. Graesslin, I.; Steiding, C.; Annighoefer, B.; Weller, J.; Biederer, S.; Brunner, D.; Homann, H.; Schweser, F.; Katscher, U.; Pruessmann, K. Local SAR constrained hotspot reduction by temporal averaging. Proceedings of the International Society for Magnetic Resonance in Medicine; Stockholm, Sweden. 2010. p. 4932

13. Zelinski AC, Angelone LM, Goyal VK, Bonmassar G, Adalsteinsson E, Wald LL. Specific absorption rate studies of the parallel transmission of inner-volume excitations at 7T. Journal of magnetic resonance imaging. 2008; 28(4):1005–1018. [PubMed: 18821601]

14. Seifert F, Wübbeler G, Junge S, Ittermann B, Rinneberg H. Patient safety concept for multichannel transmit coils. Journal of Magnetic Resonance Imaging. 2007; 26(5):1315–1321. [PubMed: 17969165]

15. Graesslin, I.; Glaesel, D.; Biederer, S.; Vernickel, P.; Katscher, U.; Schweser, F.; Annighoefer, B.; Dingemans, H.; Mens, G.; Yperen, G. Comprehensive RF safety concept for parallel transmission systems. Proceedings of the ISMRM; Toronto, Canada. 2008. p. 74

16. Harvey PRPC, Simons J. B1 shimming performance versus chanel/mode count. Proceedings of the International Society for Magnetic Resonance in Medicine. 2010; 18:1486.

17. Kozlov M, Turner R. Fast MRI coil analysis based on 3-D electromagnetic and RF circuit co-simulation. Journal of Magnetic Resonance. 2009; 200(1):147–152. [PubMed: 19570700]

18. Kozlov M, T R. Analysis of RF transmit performance for a multi-row multi-channel MRI loop array at 300 and 400 MHz. Proceedings of the Asia-Pacific Microwave Conference. 2011:1190–1193.

19. Kozlov, M.; Turner, R. Optimization of a near-field array. IEEE; 2012. p. 815-817.

20. Lattanzi R, Sodickson DK, Grant AK, Zhu Y. Electrodynamic constraints on homogeneity and radiofrequency power deposition in multiple coil excitations. Magnetic Resonance in Medicine. 2009; 61(2):315–334. [PubMed: 19165885]

21. Lattanzi R, Sodickson DK. Ideal current patterns yielding optimal signal-to-noise ratio and specific absorption rate in magnetic resonance imaging: Computational methods and physical insights. Magnetic Resonance in Medicine. 2012; 68(1):286–304. [PubMed: 22127735]

22. Childs AS, Malik SJ, O'Regan DP, Hajnal JV. Impact of number of channels on RF shimming at 3T. Magnetic Resonance Materials in Physics, Biology and Medicine. 2013:1–10.

23. Tian, J.; DelaBarre, L.; Strupp, J.; Zhang, J.; Pfeuffer, J.; Hamm, M.; Nistler, J.; Ugurbil, K.; Vaughan, J. Searching for the Optimal Body Coil Design for 3T MRI. Salt Lake City: 2013. p. 2746

24. Lemdiasov RA, Obi AA, Ludwig R. A numerical postprocessing procedure for analyzing radio frequency MRI coils. Concepts in Magnetic Resonance Part A. 2011; 38(4):133–147.

25. Setsompop K, Wald L, Alagappan V, Gagoski B, Adalsteinsson E. Magnitude least squares optimization for parallel radio frequency excitation design demonstrated at 7 Tesla with eight channels. Magnetic Resonance in Medicine. 2008; 59(4):908–915. [PubMed: 18383281]

26. Zaremba L. Guidance for industry and FDA staff: Criteria for significant risk investigations of magnetic resonance diagnostic devices. US Department of Health and Human Services, Food and Drug Administration. 2003:14.

27. Hollingsworth, N.; Moody, K.; Nielsen, JF.; Noll, D.; McDougall, M.; Wright, S. Tuning ultra-low output impedance amplifiers for optimal power and decoupling in parallel transmit MRI. IEEE; 2013. p. 946-949.

28. Kirilina E, Kühne A, Lindel T, Hoffmann W, Rhein K, Riemer T, Seifert F. Current CONtrolled Transmit And Receive Coil Elements (C2ONTAR) for Parallel Acquisition and Parallel Excitation Techniques at High-Field MRI. Applied magnetic resonance. 2011; 41(2-4):507–523. [PubMed: 22162914]

29. Jevtic, J. Ladder networks for capacitive decoupling in phased-array coils. ISMRM; Glasgow, Scotland: 2001. p. 17

30. Mahmood Z, Guerin B, Adalsteinsson E, Wald L, Daniel L. An Automated Framework to Decouple pTx Arrays with Many Channels. 2013:2722.

31. Breuer, F.; Blaimer, M.; Mueller, M.; Heidemann, R.; Griswold, M.; Jakob, P. The use of principal component analysis (PCA) for estimation of the maximum reduction factor in 2D parallel imaging. ISMRM; 2005. p. 2668

32. Lattanzi, R.; Zhu, Y.; Sodickson, D. Electrodynamic analysis of SAR and transmit homogeneity for RF shimming on a dielectric cylinder. ISMRM; 2009. p. 4510

33. Gilbert KM, Curtis AT, Gati JS, Klassen LM, Menon RS. A radiofrequency coil to facilitate B 1+ shimming and parallel imaging acceleration in three dimensions at 7 T. NMR in Biomedicine. 2011; 24(7):815–823. [PubMed: 21834005]

34. Shajan G, Kozlov M, Hoffmann J, Turner R, Scheffler K, Pohmann R. A 16-channel dual-row transmit array in combination with a 31-element receive array for human brain imaging at 9.4 T. Magnetic Resonance in Medicine. 201310.1002/mrm.24726

35. Kozlov, M.; Turner, R. Effects of Tuning Condition, Head Size Position on the SAR of a MRI Dual-row Transmit Array at 400 MHz. Taipei: PIERS; p. 422-426.

36. Alagappan V, Nistler J, Adalsteinsson E, Setsompop K, Fontius U, Zelinski A, Vester M, Wiggins GC, Hebrank F, Renz W. Degenerate mode band pass birdcage coil for accelerated parallel excitation. Magnetic Resonance in Medicine. 2007; 57(6):1148–1158. [PubMed: 17534905]

37. Lattanzi, R.; Grant, A.; Sodickson, D. Approaching ultimate SNR and ideal current patterns with finite surface coil arrays on a dielectric cylinder. ISMRM; 2008. p. 1074
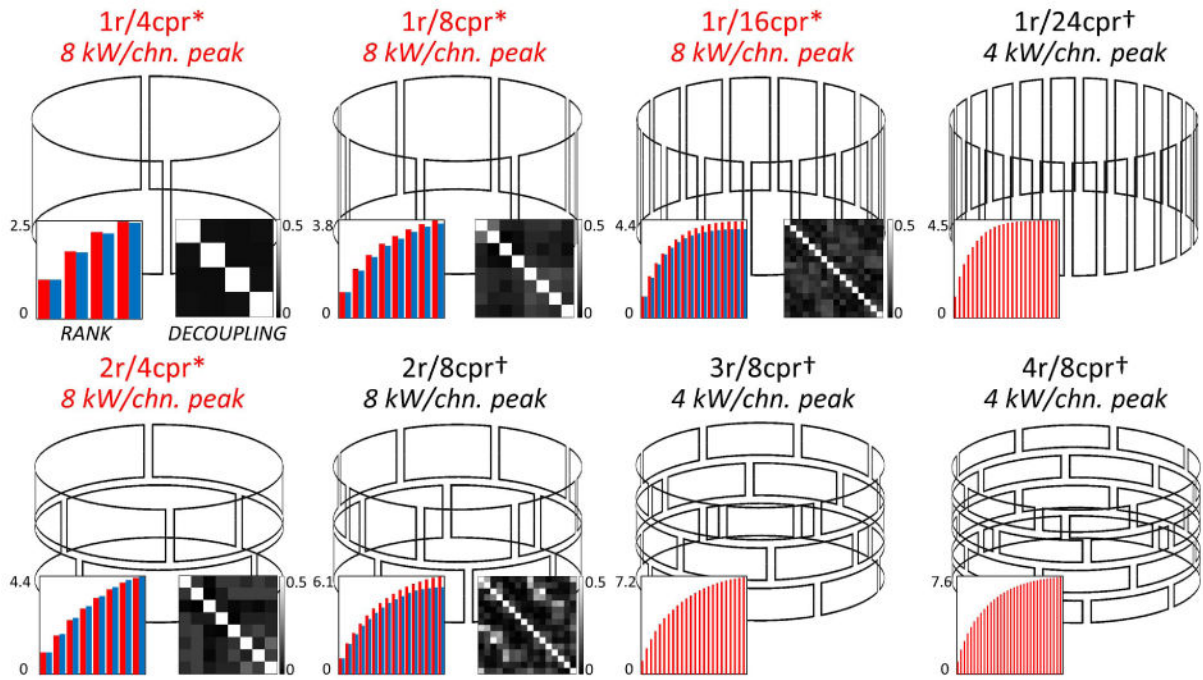
*Coil was fully decoupled by adjustment of capacitors placed between nearest-neighbors
†Coil could not be fully decoupled by adjustment of decoupling capacitors and was ideally decoupled in simulation

Fig. 1.
PTx coils simulated. The notation Xr/Ycpr indicates a coil made of X rows of Y coils per row (total number of channels is X×Y). The maximum peak power available on each channel is indicated next to every coil. The maximum average power available on each channel is 10 times smaller than the maximum peak power. The designs 1r/4cpr, 1r/8cpr, 1r/ 16cpr and 2r/4cpr could be decoupled by adjustment of decoupling capacitors placed between nearest neighbors (capacitive decoupling). For these coils, the coupling matrix is shown (0 denotes perfect decoupling, 1 denotes perfect coupling). The diagonal elements are equal to 1. Other designs could not be well decoupled by our nearest neighbors decoupling capacitor adjustment algorithm. These designs (with no coupling matrix shown), were simulated as perfectly decoupled (ideal decoupling). Note that the designs which could be decoupled were also evaluated as ideally decoupled. The other subplot shown is the cumulative sum of the singular values for the B1+ maps of each coil in the torso (excluding the arms). Red bars correspond to the B1+ maps of the coils ideally decoupled in simulation. Blue bars correspond to the arrays decoupled using capacitors. In general, the presence of residual coupling decreased the rank of the B1+ basis set created by the coils. Using more than 16 channels per row did not increase significantly the rank of the B1+ basis set. Additional channels beyond the first 16 should be instead distributed axially (z-direction).
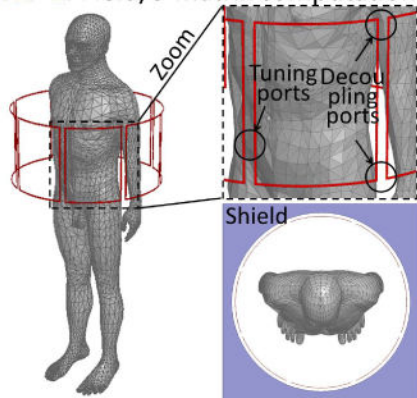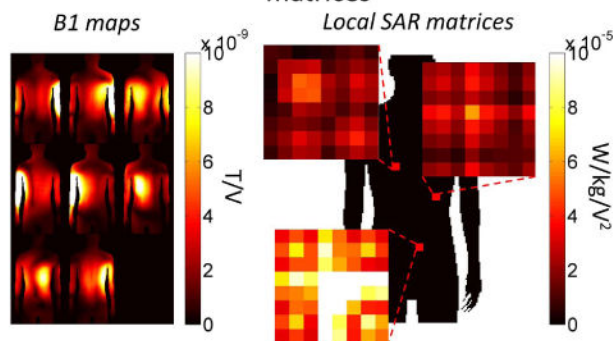
Fig. 2.
Flow chart of the co-simulation of coupled arrays. In step 2, the admittance matrix **Y** is extrapolated around the Larmor frequency using the assumption that the capacitor-less structure is purely inductive. This allows generating S-parameters in a range of frequencies while only having to simulate the coil at the Larmor frequency (thus reducing computation time).
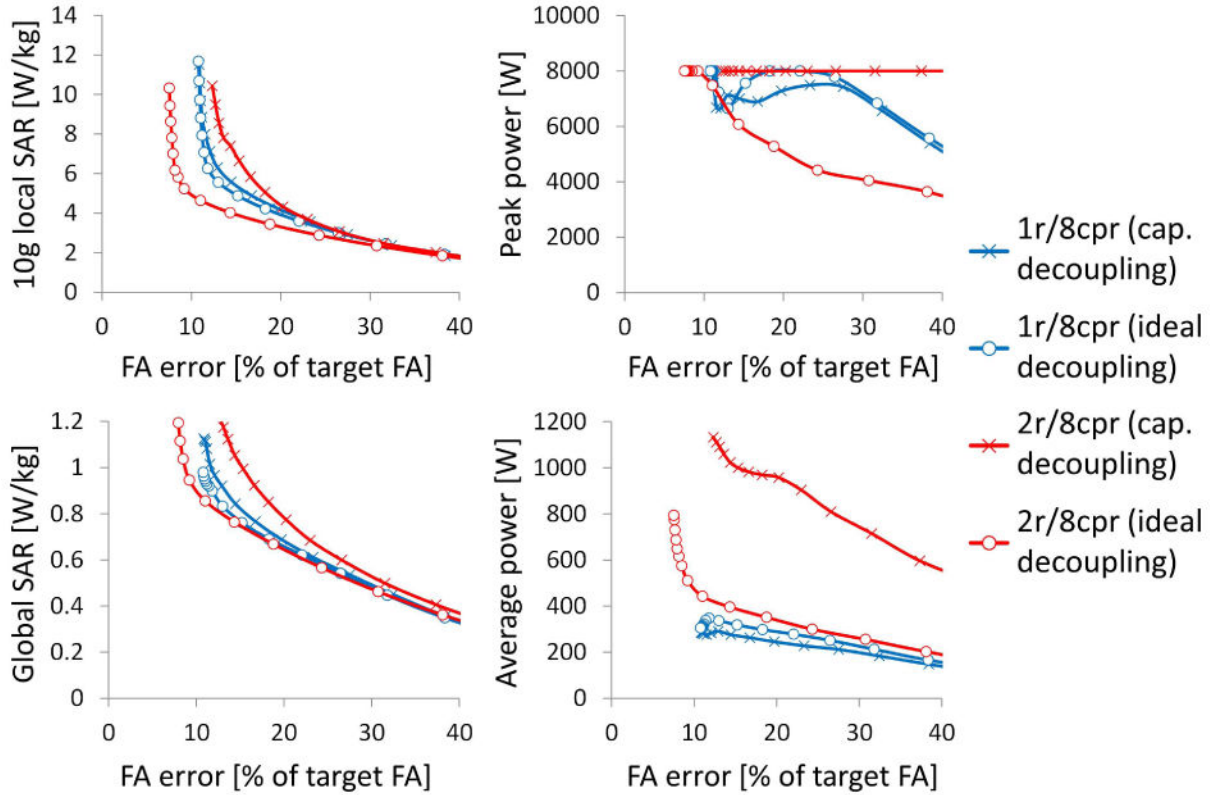
Fig. 3.
MLS RF-shimming L-curves for the 1r/8cpr (capacitive and ideal decoupling) and 2r/8cpr (capacitive and ideal decoupling) designs in a transverse slice at isocenter. The local SAR limit was varied. Peak power was constrained to 8 kW/channel and average power to 800 W/channel for both coils. Global SAR was constrained to be below 4 W/kg. The 1r/8cpr design could be well decoupled using decoupling capacitors placed between nearest neighbors, therefore there is little difference between the power consumption and the local and global SAR created by this coil and its ideally decoupled counterpart. In contrast, the 2r/8cpr design was not well decoupled using capacitors placed between nearest neighbors, likely because the staggered transmit channels in different rows are close to each other but are not connected by a decoupling capacitor and can therefore not be efficiently decoupled. The coupled coil consumed more power than its ideally decoupled counterpart, which limited its ability to create quality excitations at low local and global SAR.
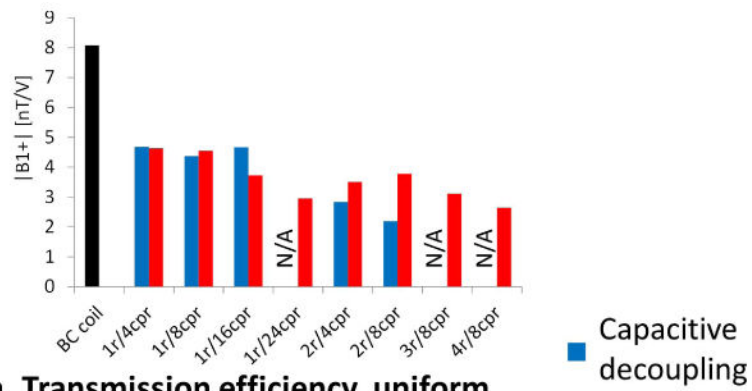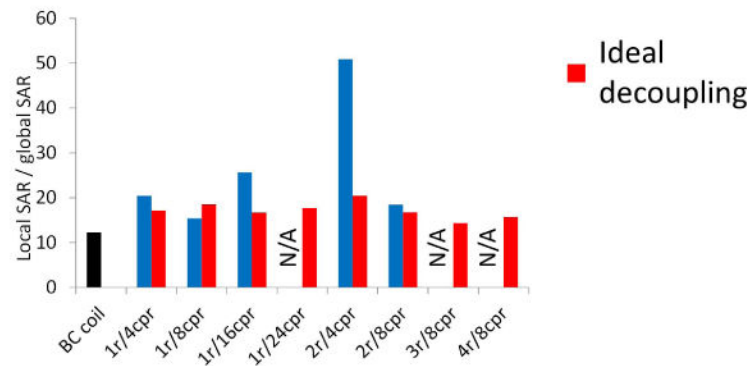
Fig. 4.
a: Transmit efficiencies in nT/V of the pTx coils driven in their uniform birdcage mode.
Also shown is the transmit efficiency of the birdcage coil. **b**: Local to global SAR ratios of
the birdcage coil and the pTx coils driven in their uniform birdcage modes.
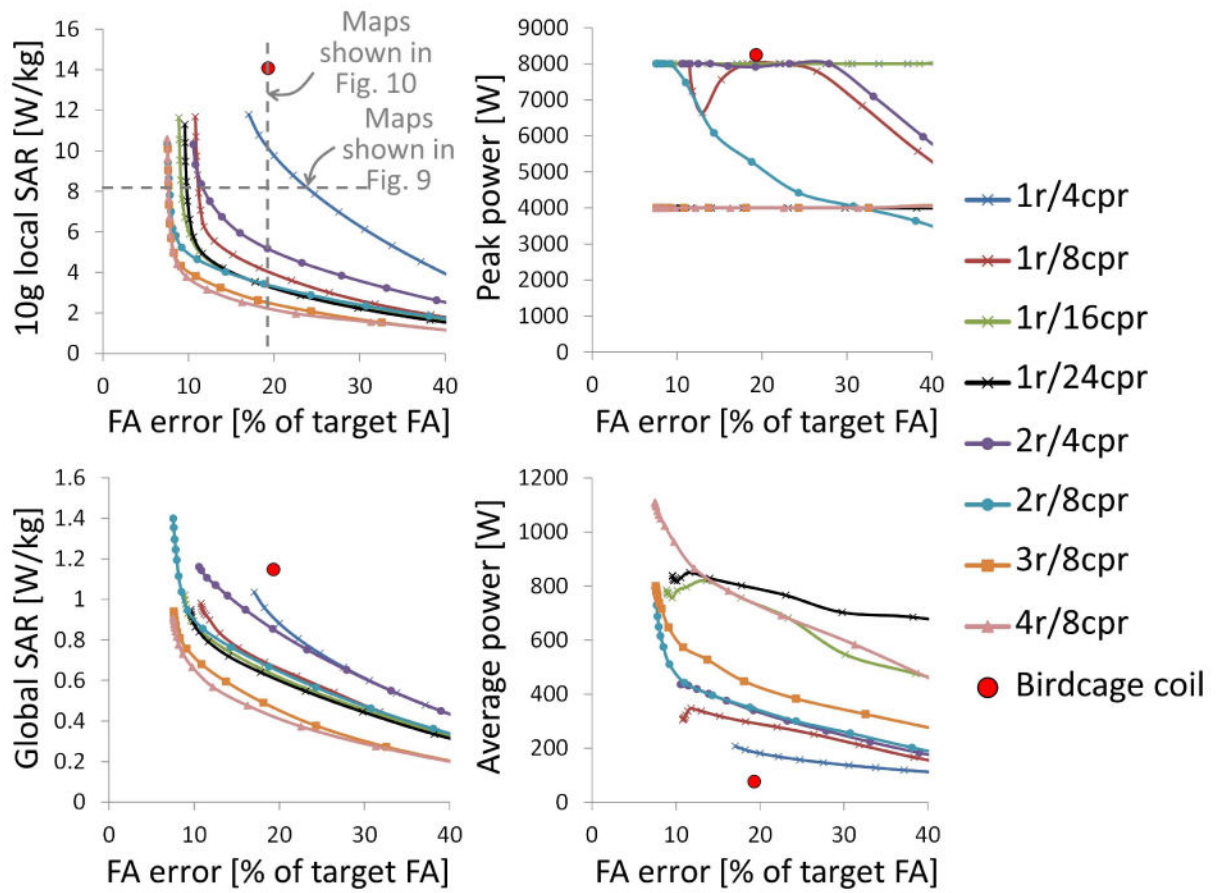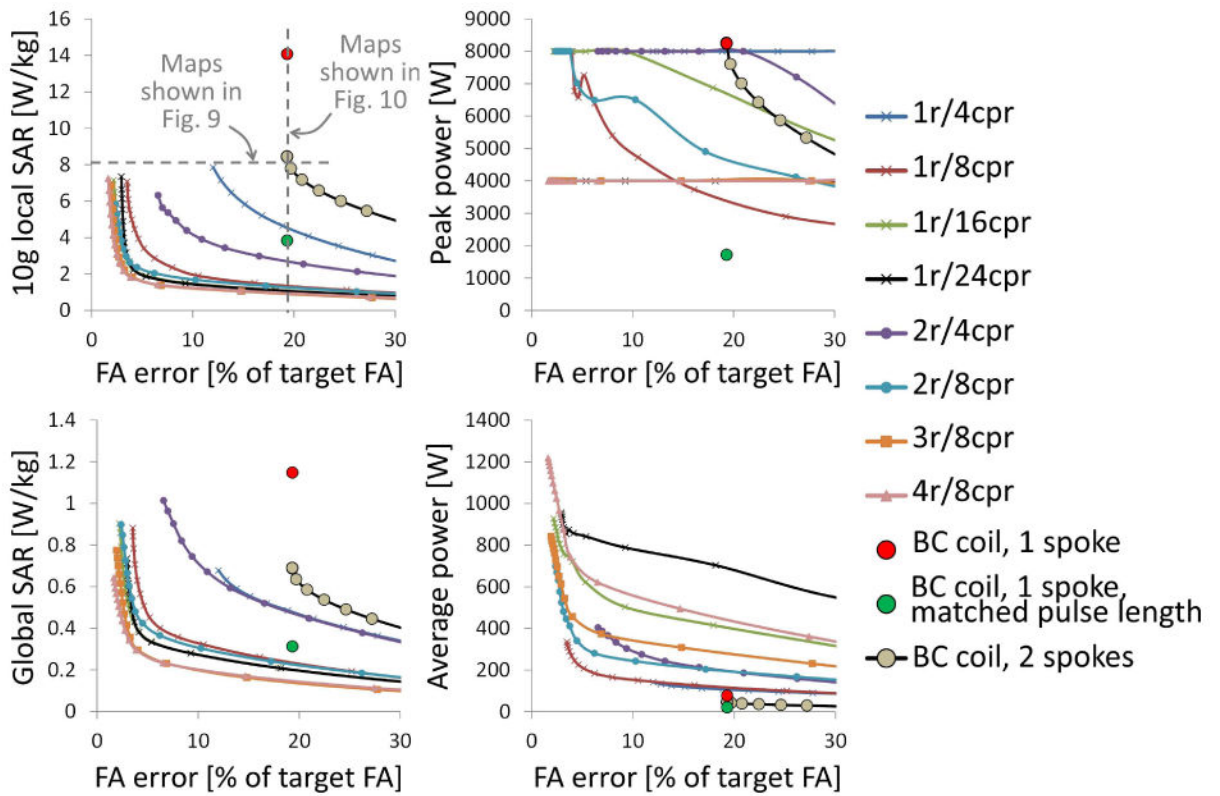
Fig. 5.
MLS RF-shimming L-curves for all pTx coils in a transverse slice at isocenter (all coils in the ideally decoupled condition). The local SAR limit was varied. The average power and peak power limits were set to the values indicated in Fig. 1 (different coils have different input power limits). The global SAR limit was set to 4 W/kg for all coils. These L-curves show that pTx coils could create more uniform excitations at lower local and global SAR than a birdcage coil of the same dimension. This improvement in the SAR vs excitation fidelity trade-off required an increase in input power however. Increasing the number of transmit channels improved the SAR vs excitation fidelity trade-off but also increased power consumption.

Fig. 6.
MLS 2-spoke L-curves for all pTx coils in a transverse slice at isocenter (all coils in the ideally decoupled condition). The local SAR limit was varied. The average power and peak power limits were set to the values indicated in Fig. 1. The global SAR limit was set to 4 W/kg for all coils. The "BC coil, 1 spoke" dot corresponds to driving the birdcage coil with an RF pulse of the same duration a single spoke sub-pulses (short pulse). The "BC coil, 1 spoke, matched pulse length" corresponds to driving the birdcage coil with an RF pulse with duration equal to the total duration of the 2-spoke pulse played on the pTx arrays (longer pulse). Finally, the "BC coil, 2 spokes" L-curve was obtained by designing 2-spoke pulses for the birdcage coil and varying the local SAR limit exactly as performed on the pTx arrays.
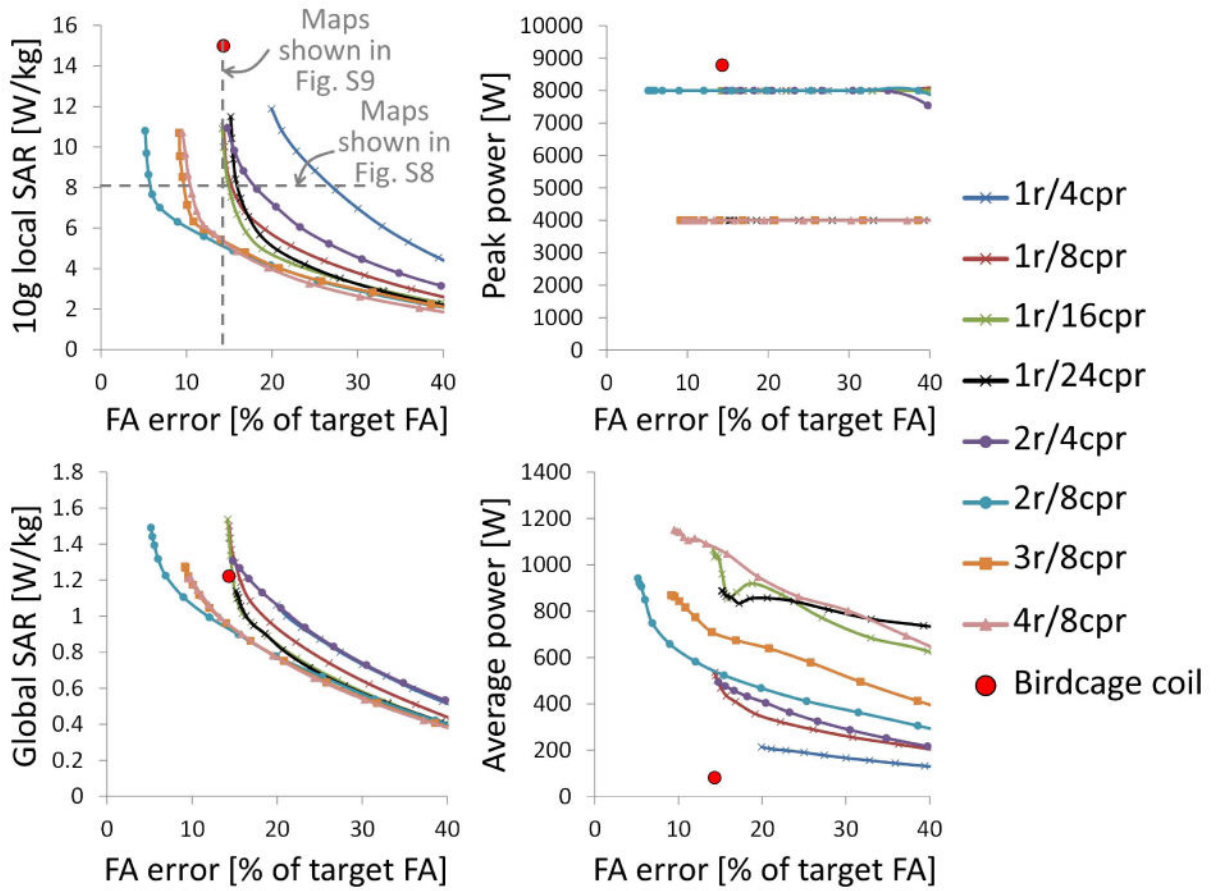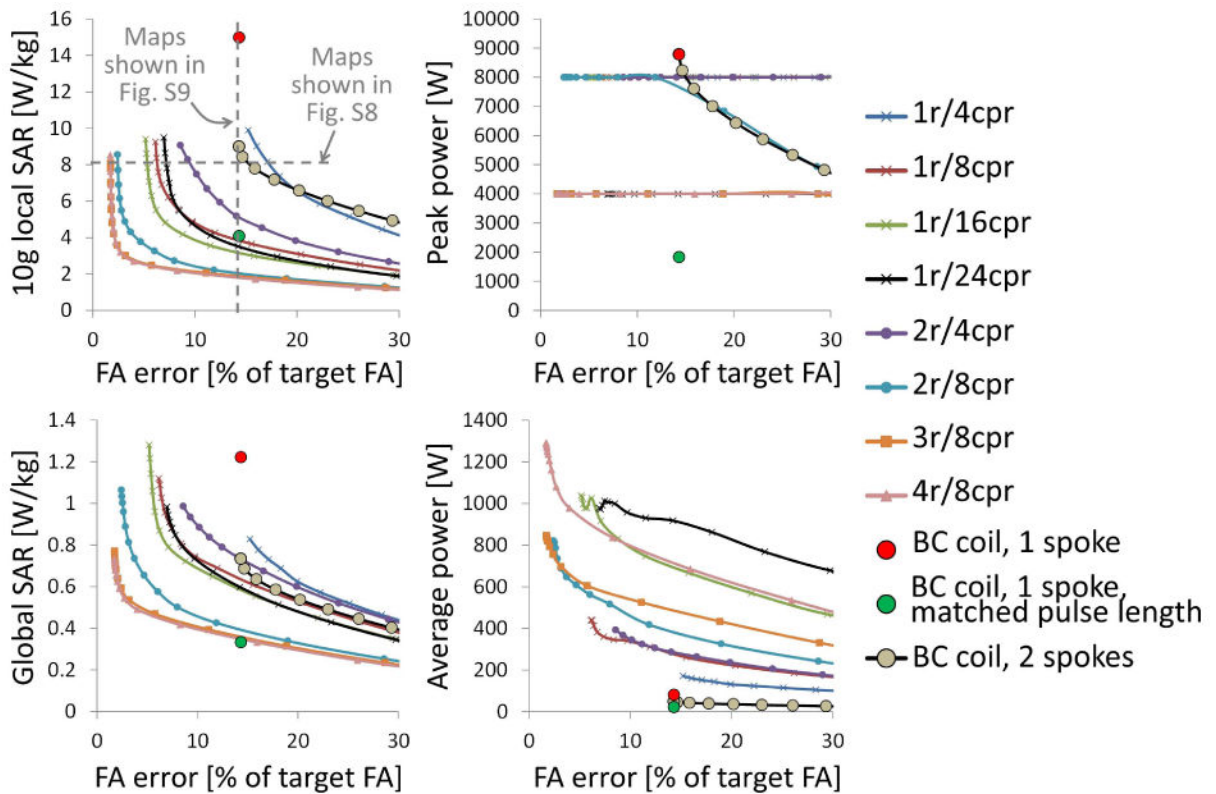
Fig. 7.
MLS RF-shimming L-curves for all pTx coils in a coronal slice at isocenter (all coils in the ideally decoupled condition). The local SAR limit was varied. The average power and peak power limits were set to the values indicated in Fig. 1 (different coils have different input power limits). The global SAR limit was set to 4 W/kg for all coils. Like for transverse imaging, the additional degrees-of-freedom of pTx coils could be used to create more uniform excitations than a birdcage coil as well as to reduce both local SAR and global SAR. This improvement in the SAR vs excitation fidelity trade-off came at the price of an increase in power consumption. Increasing the number of transmit channels improved the SAR vs excitation fidelity trade-off but also increased power consumption. For imaging in the coronal plane, a clear performance gain is visible for pTx arrays with multiple coils in the z-direction.

Fig. 8.
MLS 2-spoke L-curves for all pTx coils in a coronal slice at isocenter (all coils in the ideally decoupled condition). The local SAR limit was varied. The average power and peak power limits were set to the values indicated in Fig. 1. The global SAR limit was set to 4 W/kg for all coils. The "BC coil, 1 spoke" dot corresponds to driving the birdcage coil with an RF pulse of the same duration than a single spoke sub-pulse (short pulse). The "BC coil, 1 spoke, matched pulse length" corresponds to driving the birdcage coil with an RF pulse with duration equal to the total duration of the 2-spoke pulse played on the pTx arrays (longer pulse). Finally, the "BC coil, 2 spokes" L-curve was obtained by designing 2-spoke pulses on the birdcage coil and varying the local SAR limit exactly as performed on the pTx arrays. PTx coils were able to create almost perfect excitations (see Supplemental Figure S8) at significantly lower local SAR than the birdcage coil (Supplemental Figure S9). This improvement in the SAR vs excitation fidelity trade-off came at the price of increased power consumption. Increasing the number of transmit channels improved the SAR vs excitation fidelity trade-off but also increased power consumption. Like for RF-shimming in the coronal direction, important performance gains in the SAR vs excitation fidelity trade-off were made by distributing transmit channels in the z-direction in addition to around the circumference of the patient's torso.
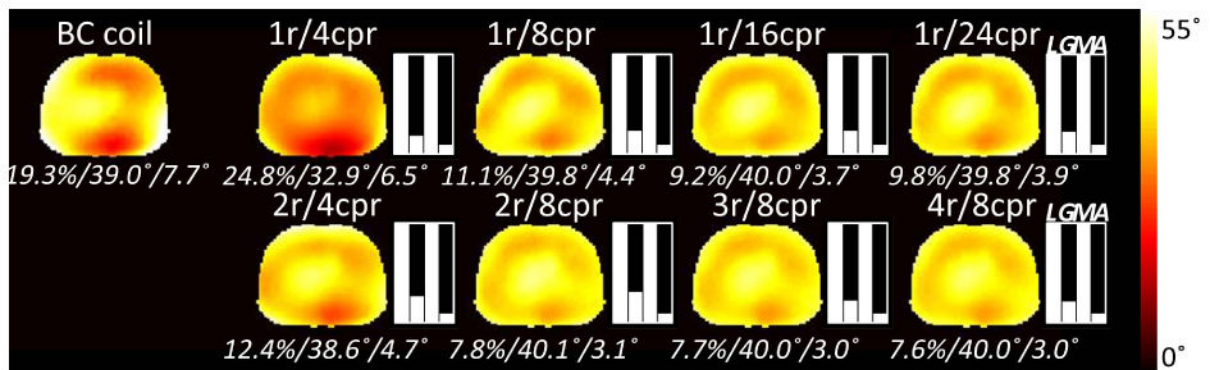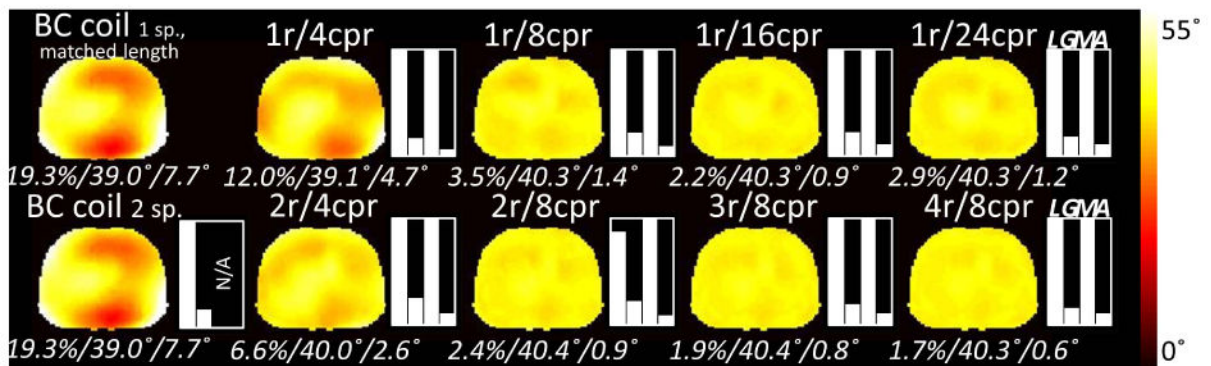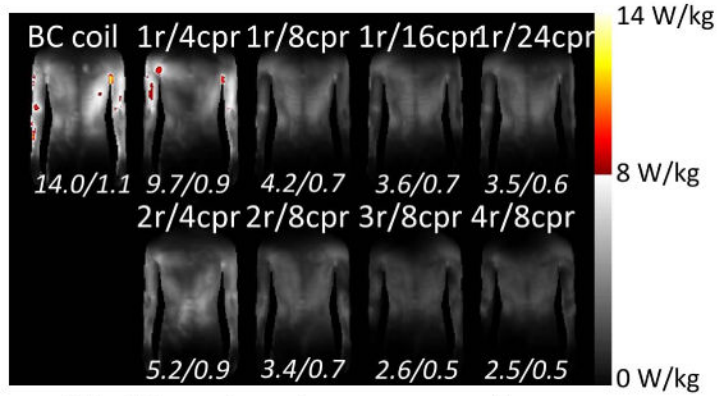
Fig. 9.
MLS flip-angle maps for imaging of a transverse slice at isocenter passing through the liver (all coils in the ideally decoupled condition). **a**: RF-shimming results for all pTx arrays. **b**: 2-spoke results. These maps are single data points from Fig. 5 and 6 and correspond to a local SAR limit of 8 W/kg (FDA limit) – except for the birdcage coil which has a local SAR of 14.1 W/kg for RF-shimming. The white bar graphs on the right of each flip-angle map indicate how far the local SAR ("L"), global SAR ("G"), peak power ("M") and average power ("A") constraints are from their respective limits. The three numbers shown below each flip-angle map are the RMSE (in percent of the target flip-angle, here 40°), mean flip-angle and standard deviation of the flip-angle map. For RF-shimming, 1r/4cpr was the only design to be so strongly power limited that it could not achieve the target flip-angle while respecting the 8 W/kg local SAR limit. Other designs were also peak power limited (and local SAR limited) but achieved the desired mean flip-angle.

Fig. 10.
Maximum intensity projection of the SAR maps created by the different pTx designs for excitation of a transverse slice at isocenter passing through the liver (all coils in the ideally decoupled condition). **a**: RF-shimming pulses. **b**: 2-spoke pulses. These maps are detailed data for pulses shown in Fig. 5 and 6 and correspond to a constant excitation error of 19%. The two numbers below each SAR map are the local SAR and global SAR expressed in W/kg. The color maps are split in two (grey below 8 W/kg and color above 8 W/kg) to better visualize the SAR hotspots that are above the FDA local SAR limit of 8 W/kg. These results illustrate the important reduction of local SAR and global SAR (at constant excitation error) that can be achieved when increasing the channel count from 1 (birdcage coil) to 32 (4r/ 8cpr). Note that such SAR reductions were achieved by designing pulses that were explicitly controlled for local and global SAR.