

Switching Circuit Optimization for Matrix Gradient Coils

Stefan Kroboth¹, Kelvin J. Layton², Feng Jia¹, Sebastian Littin¹, Huijun Yu¹, Jürgen Hennig¹, and Maxim Zaitsev¹

¹Department of Radiology, Medical Physics, Medical Center University of Freiburg, Faculty of Medicine, University of Freiburg, Freiburg, Germany and ²Institute for Telecommunications Research, University of South Australia, Adelaide, Australia

Corresponding Author:

Stefan Kroboth, MSc
Killianstr 5a Freiburg, Germany 79106;
E-mail: stefan.kroboth@uniklinik-freiburg.de

Key Words: magnetic resonance imaging, matrix gradient coil, combinatorial optimization, switching circuit, nonlinear encoding

Abbreviations: Field of view (FOV), simulated annealing (SA), genetic algorithm (GA)

ABSTRACT

Matrix gradient coils with up to 84 coil elements were recently introduced for magnetic resonance imaging. Ideally, each element is driven by a dedicated amplifier, which may be technically and financially infeasible. Instead, several elements can be connected in series (called a “cluster”) and driven by a single amplifier. In previous works, a set of clusters, called a “configuration,” was sought to approximate a target field shape. Because a magnetic resonance pulse sequence requires several distinct field shapes, a mechanism to switch between configurations is needed. This can be achieved by a hypothetical switching circuit connecting all terminals of all elements with each other and with the amplifiers. For a predefined set of configurations, a switching circuit can be designed to require only a limited amount of switches. Here we introduce an algorithm to minimize the number of switches without affecting the ability of the configurations to accurately create the desired fields. The problem is modeled using graph theory and split into 2 sequential combinatorial optimization problems that are solved using simulated annealing. For the investigated cases, the results show that compared to unoptimized switching circuits, the reduction of switches in optimized circuits ranges from 8% to up to 44% (average of 31%). This substantial reduction is achieved without impeding circuit functionality. This study shows how technical effort associated with implementation and operation of a matrix gradient coil is related to different hardware setups and how to reduce this effort.

INTRODUCTION

Recently, matrix gradient coils (also termed multi-coils or multi-coil arrays) were introduced for both magnetic resonance image acquisition and B_0 shimming (1-6). Almost 2 decades earlier, a similar concept called *matrix shim coil system* was introduced for nuclear magnetic resonance (7). Such coils consist of a multitude of (up to 84) compact coil elements, which are in the most trivial configuration individually supplied by a current amplifier. The created field is a superposition of the fields produced by individual coil elements. This multiple-coil topology allows for the creation of a large variety of different field shapes by adjusting the currents in the individual coil elements. The use of multiple coils has certain advantages over that of conventional gradient coils and it opens up new possibilities to imaging and shimming. In most realizations, coil elements of a matrix gradient coil are smaller than typical dimensions of conventional gradient coils, which allows for faster switching of currents. One problem associated with conventional linear gradient coils is the requirement of highly linear fields across the entire field of view (FOV). With matrix gradient coils, encoding fields can be adapted to the actual FOV needed for the given imaging

application. This may potentially be more accurate and/or more efficient (4), particularly for small FOVs.

Matrix gradient coils are furthermore favorable for nonlinear spatial encoding schemes such as PatLoc (8-15), O-Space (16-19), or FRONSAC (20). Imaging techniques that incorporate nonlinear gradients into the magnetic resonance pulse sequence can utilize the matrix gradient coil’s ability to create a wide variety of field shapes whose nonlinearity extends way beyond that of the quadratic fields used in previous work. For magnetic resonance imaging applications that focus on local anatomy, nonlinear gradients allow for fast sequences that produce high-resolution images in a target area and low-resolution images everywhere else (13). Example images acquired with a matrix gradient coil can be found in Littin et al.’s study (2), but owing to the novelty of the concept, many topics mentioned above require further research and have not made it into clinical applications yet.

Design strategies for unconventional gradient coils are still a topic of active research. For classical linear gradients, the entire gradient system is typically split into 3 gradient channels that correspond to principal spatial directions X, Y, and Z. Each

independently driven channel lies on a separate surface within the gradient coil structure and is designed mostly independent of the other channels. Variations of the target field method (21-23) are typically used combined with modern numerical optimization techniques, allowing for a simultaneous control of the manufacturing constraints such as minimal conductor spacing, electrical characteristics such as coil inductance, resistance and minimal mutual coupling, as well as fidelity of the achieved field profiles. Less conventional gradient system concepts pursue different design strategies, typically including multiple steps. For example, a linear Z gradient coil with a dynamically movable imaging region (24) was realized by designing a multitude of current density patterns for varying FOV offsets, followed by a singular value decomposition step to compress these current densities into 3 electrical channels to be controlled by individual gradient amplifiers. In contrast to the aforementioned design strategies, when designing a matrix coil, the target field or set of fields is in general unknown. One strategy to address this challenge is to split the poorly defined and computationally intractable problem into a set of smaller tractable problems. For the matrix coil in Littin et al.'s study (2), these steps included the following:

1. definition of the number of elements to balance the field strength and flexibility against power dissipation (25),
2. optimization of the geometry and winding patterns of individual elements to achieve maximum gradient strength and sufficient shielding (6),
3. obtaining a method to drive the matrix gradient coil with substantially fewer amplifiers than coil elements by assigning several coil elements to each of the amplifiers (called a configuration) while preserving the matrix gradient coils' ability to accurately create a certain field shape for a given application (26),
4. finding a suitable strategy to switch between different configurations obtained by step 3 within an imaging pulse sequence or between the pulse sequences for different applications or imaging regions,
5. finding excitation (27), signal conditioning (28), and sampling (15) strategies that optimally use the flexibility provided by the matrix coil.


Although the latter task does not strictly belong to the coil design, it provides the initial motivation as well as the valuable feedback to the entire chain, and in particular to step 4, as the different imaging applications mentioned demand for switching between several varying field shapes. The present paper concerns itself primarily with step 4 from the above listing and builds on top of step 3 which was addressed in previous work (26).

Although the matrix gradient coil design in (2) aimed at maximum flexibility in creating a wide variety of field shapes to assess the advantages and disadvantages of this approach, future coils could be specifically designed for a certain range of applications or use cases. One example for a matrix gradient coil specifically designed for high gradient strength with acceptable switching rates, as required for high b-factor diffusion imaging, are the linear gradient coils of the Human Connectome Project

(29), which use 4 independent coil elements per linear gradient axis and are driven by a total of 12 amplifiers. Apart from imaging, matrix coils or multi-coils can be used for B_0 shimming and both imaging and B_0 shimming at the same time (30). Although shimming requires only low current and therefore relatively low-cost power supplies, encoding requires substantial currents, which necessitates powerful and therefore expensive amplifiers. For shimming, each coil element can be driven by an individual amplifier (4, 3), but for imaging applications, which is the focus of this work, it is desirable to keep the number of amplifiers as low as possible.

In an earlier work (26), we presented a method to use the matrix gradient coil with fewer amplifiers than coil elements. When using fewer amplifiers than coil elements, the coil elements are distributed over the available amplifiers such that several coil elements are assigned to a single amplifier. All coil elements assigned to a certain amplifier are connected in series and supplied with the same current from the corresponding amplifier. The coil elements assigned to a single amplifier are termed a *cluster*, and a set of N_A clusters is called a *configuration*, where the number of amplifiers is defined as N_A . In our previous work, we have shown how such a configuration can be optimized by assigning coil elements to clusters such that a desired target field shape can be approximated with high accuracy. We have shown that this combinatorial optimization problem can be solved such that even with a reduced number of amplifiers, a large range of field shapes can be created accurately. For instance, for an 84-channel coil and only 12 amplifiers, it was possible to create spherical harmonics up to third order with a median normalized least squares error of $\sim 5\%$. With ≥ 24 amplifiers, the median normalized least squares error reduces to $\approx 2.5\%$, while it increases to $\approx 10\%$ for 6 amplifiers. Note that a configuration creates the desired target field with a particular ratio of the currents supplied by the amplifiers. In principle, this restriction can be lifted by changing the ratio to create a wide range of different field shapes with a single configuration. This means that in practice, each configuration is more powerful than what it was optimized for; however, exploitation of this degree of freedom requires further investigation in future work. Details of the implementation and the performance analysis of the solver are presented in Kroboth et al. (26, 31). In Kroboth et al.'s study (26), each configuration was optimized to create a single target field; however, typical imaging applications require the application of several distinct encoding fields, where each generally corresponds to a different configuration. Different encoding fields usually require different configurations; therefore, to switch between encoding fields, a mechanism is required to switch between configurations. Essentially, this means opening and closing connections between coil elements, depending on the requirements of the desired configuration. This can be achieved via a so-called switching circuit (32), which can be cast into the epoxy resin together with the matrix gradient coil. If space limitations become a problem, the circuit can be located outside the coil. An example of such an approach, which relies on a dense switching network, was proposed recently for B_0 shimming and is termed *dynamically controlled adaptive current network* (33). In such a network—owing to a

Figure 1. A potential pulse sequence with 3 distinct target fields T1, T2, and T3, each created by a corresponding configuration consisting of N_C coil elements and driven by N_A amplifiers (where the ratio of all currents is predefined). Because configurations differ for different target fields, the sequence requires a mechanism to switch between different configurations. This can be achieved with a switching circuit. Changing the state of the switches requires $<10 \mu\text{s}$. This example shows 3 target fields; however, the number (and shape) of the target fields can be arbitrary. Note that by changing the ratio of the currents supplied to the individual clusters, other target fields can be created, which were not part of the optimization process of the configurations. Therefore, the actual capabilities of such a setup are much larger than this figure may suggest at a first glance.



particularly simple geometry of the B_0 -contributing elements, which are straight wire segments orthogonal to the B_0 direction—a 2D network of metal-oxide-semiconductor field-effect transistor (MOSFET) switches connecting each element to its neighbors and a single amplifier is sufficient to realize any discrete current path. However, the achievable surface current density in such a network is not sufficient for imaging applications requiring strong local gradients. Therefore in our research, we focus on multiturn element geometries. A switching circuit for the available 84-element matrix gradient coil, in its straightforward implementation, can be designed to connect every terminal of a coil element to all terminals of every other coil element. In addition, every terminal of each coil element has to be connected with every terminal of every amplifier. This hypothetical switching circuit would allow for arbitrary current paths through a network of coil elements. By changing the state of the individual switches, every configuration can be realized, allowing for sequences such as the one illustrated in Figure 1, where the configuration for the next encoding field can be switched at zero-current conditions. Unfortunately, this would require 18 060 switches for the case of $N_C = 84$ coil elements and $N_A = 12$ amplifiers, as the number of switches necessary to connect all coil element terminals to each other is given by $N_C \cdot (2 \cdot N_C - 1) = 14028$, and the number of switches necessary to connect all amplifier terminals to all coil element terminals is given as

$4 \cdot N_C \cdot N_A = 4032$. Not only is this technically infeasible owing to the complexity of the network and the dimensions of the individual switches, it is also potentially prohibitively expensive owing to the requirements on the switches (high currents, fast switching times). In this work we use the property that if a finite set of configurations to be switched is known *a priori*, a full switching circuit is not necessary. For instance, connections that are connected in all configurations can be hard-wired (no switch required), and connections that are not present in any configuration do not require a switch (and no hard-wiring). Only connections that are closed in some but not all configurations require a switch. The same applies to the switches that connect the amplifiers to the coil element network. The number of required switches depends on the configurations and the ordering of elements within a cluster. There are numerous switching circuits for a given set of configurations that are all capable of switching between the configurations.

We propose an optimization algorithm that finds, in the set of all possible switching circuits, one which requires a low number of switches for a given set of configurations. We exploit the fact that the ordering of coil elements within a cluster does not change its electromagnetic properties owing to the fact that the elements are connected in series and that any amplifier can supply any cluster (assuming that the amplifiers have equal specifications). In our model, the coil elements are characterized

by individual resistances and an inductance matrix and we further assume perfect switches. The configurations are designed before the optimization of the switching circuit, for instance, by the method introduced in Kroboth et al.'s study (26). *Note that the optimization procedure presented in this work does not affect the magnetic field created by the given configuration.* The coil design (2) used in this work is a cylindrical head-insert with a length of 70 cm and an inner diameter of 39 cm. It consists of 84 coil elements (6), distributed on 7 rings with 12 elements each (Figure 2A). In line with a previously published work (26), we assume that each coil element is equipped with a bridge consisting of 4 switches (Figure 2B). This allows for the current to be routed through the coil winding in both directions. Throughout this work, we treat the bridge as an integral part of each coil element. The method presented here is also applicable to other matrix gradient coil or multi-coil designs. Ideally, the problem presented in this work would be optimized jointly with the optimization of the configurations. However, as both problems are combinatorial and as the joint optimization would have multiple objectives, this was considered intractable given the currently available computational power.

METHODOLOGY

Theory

Minimizing the number of switches for a given set of configurations can be split into 2 combinatorial NP-hard (non-deterministic polynomial-time hardness) optimization routines which are executed sequentially. The first minimizes the number of switches between coil element terminals, then, based on the result, the second minimizes the number of switches which connect the amplifiers to the matrix coil. The used variables are summarized in Table 1.

Minimizing the Number of Switches Between Coil Elements. The problem of reducing the number of switches between the elements is a distant relative of the well-known traveling salesman problem (TSP) (34), where a salesman aims at visiting several cities in minimum time or minimum travelled distance. However, here we have N_F companies (configurations C_k , where $k = 1, \dots, N_F$), each employing N_A salesmen (amplifiers). Each company assigns N_C cities (coil elements) to its salesmen such that all cities are assigned and every city is assigned exactly once for each company. Given such a setup, the aim is to find a path for each salesman such that the total number of streets (connections) between cities (coil elements) and the number of necessary junctions (switches) are minimal.

A configuration defines how coil elements are grouped into clusters. Each cluster is a set of coil elements that are all supplied with the same current. When implementing a configuration, all coil elements of a cluster are therefore connected in series in an arbitrary order and supplied with current by a single amplifier. For each configuration, there is a graph that models how coil elements are connected. The ordering of elements within a cluster does not change the electromagnetic properties of the cluster (and hence the configuration); however, ordering does change the associated graph. This means that changes to the ordering do *not* affect the ability of the configuration to create a certain target field at a certain accuracy and field strength. Let C_k be the k^{th} of N_F configurations, which is optimized to create a certain target field T_k . Furthermore, let X be a variable that indicates how coil elements within the clusters are ordered for

all k . Then there is an associated adjacency matrix $A(C_k, X)$ that indicates which terminals of the coil elements are connected with each other for the configuration C_k given the ordering X . The adjacency matrix contains the terminals of the coil elements, which allows us to account for polarity. Therefore, the size of $A(C_k, X)$ is $2N_C \times 2N_C$. Furthermore, let $\Sigma(X, G)$ be the sum of a set of N_F general adjacency matrices $G(C_k, X)$:

$$\Sigma(X, G) = \sum_{k=1}^{N_F} G(C_k, X). \quad (1)$$

The nonzero entries of $\Sigma(X, G)$ indicate connections. Let $\sigma^{ij}(X, G)$ be the element of $\Sigma(X, G)$ in the i^{th} row and the j^{th} column. Every entry where $\sigma^{ij}(X, G) = N_F$ indicates a hardwired connection because this connection exists for every configuration in case of the ordering X ; therefore, there is also no switch required. Finally, the number of necessary switches is given by the following equation:

$$s(\Sigma(X, G)) = \sum_{i,j:j<i} \begin{cases} 0 & \text{if } \sigma^{ij}(X, G) = 0 \\ 0 & \text{if } \sigma^{ij}(X, G) = N_F. \\ 1 & \text{else} \end{cases} \quad (2)$$

Applying equation (2) to the adjacency matrices $A(C_k, X)$ serves as the cost function for the optimization problem as follows:

$$X^* = \underset{X}{\operatorname{argmin}} \{s(\Sigma(X, G))\}. \quad (3)$$

For optimization, the following degrees of freedom can be exploited, which are reflected in X :

- The order in which the elements are connected within a cluster can be arbitrary.
- Each element is equipped with a bridge consisting of 4 switches that allow current to be routed through the element in both directions (Figure 2B). Hence every element can be connected in both orientations provided the current direction in the coil windings is adapted accordingly by using the bridge switches.

To change X , either the ordering of elements within a cluster can be changed or the direction in which a coil element is connected. Changing X affects the individual adjacency matrices $A(C_k, X)$, which furthermore affects $s(\Sigma(X, A))$.


Figure 3 illustrates this principle using a toy example with a matrix gradient coil with $N_C = 4$ coil elements (E1, E2, E3, E4), $N_A = 2$ amplifiers, and $N_F = 3$ configurations.

Minimizing the Number of Switches from the Amplifiers to the Coil Element Network. Once a favorable ordering of coil elements is obtained, the number of necessary switches from the terminals of the amplifier to the entry and exit terminals of the individual clusters of all configurations can be minimized. Let $B(C_k, Y)$ be adjacency matrices that indicate the connections between amplifiers and clusters of configuration C_k given Y that defines which amplifier is assigned to which cluster. Using equations (1) and (2) with the adjacency matrices $B(C_k, Y)$, the global optimizer Y^* can be obtained via the following equation:

$$Y^* = \underset{Y}{\operatorname{argmin}} \{s(\Sigma(Y, B))\}. \quad (4)$$

Under the assumption that all amplifiers have equal properties and that they can adapt to the actual load, it is irrelevant which cluster is driven by which amplifier. Furthermore, the positive and negative terminals of the amplifier can be connected to either the entry- or exit-terminal of a cluster because the am-

A



B




Figure 2. Implementation of the matrix gradient coil shown at the backside of a Siemens Trio scanner (A). To conduct measurements, the coil is pushed into the isocenter of the scanner. The black connections feed current to the individual coil elements. The water cooling is supplied by the red and blue hoses. On the side of cart are connector boards with connections for every coil element. A schematic representation of a coil element E which consists of the coil windings L equipped with a bridge switch (B). This allows to route current through the element in both directions. One direction is achieved by having the switches S_1 and S_3 in on-state and S_2 and S_4 in off-state. The other direction is achieved the other way around. In addition, the bypass mode can be achieved by having control switches S_1 and S_4 in on-state and switches S_2 and S_3 in off-state, or the other way around.

Table 1. List of Variables

Variable	Explanation
E_1	Coil element 1
S_1	Switch 1
Amp1	Amplifier 1
p_1^2	Terminal 2 of coil element E_1
N_A	Number of amplifiers
N_C	Number of coil elements
N_F	Number of fields
N_D	Number of different (random) configurations
N_R	Number of repetitions (optimization algorithm runs)
T_k	Target field k
C_k	Configuration k
X, Y	Orderings
$G(C_k, X)$	General adjacency matrix for a configuration C_k and an ordering X
$A(C_k, X)$	Adjacency matrix for the connections between coil element terminals
$B(C_k, Y)$	Adjacency matrix for the connections between coil element terminals and amplifier terminals
$\Sigma(X, G)$	Sum of adjacency matrices
$\sigma^{ij}(X, G)$	Element (i, j) of sum of adjacency matrices $\Sigma(X, G)$
$s(\Sigma(X, G))$	Number of switches for $\Sigma(X, G)$
T^i	Temperature at iteration i (Simulated Annealing)

plifiers are bipolar and therefore the direction of the current can be changed. These degrees of freedom are exploited to minimize the number of switches that connect the amplifiers to the coil element network. The concept is illustrated in Figure 4, where changing how amplifiers are assigned to the network of coil elements leads to an ordering that requires 2 fewer switches.

Matrix Gradient Coil and Other Hardware

The algorithm is applied to the matrix gradient coil design introduced in Littin et al.'s study (2). The gradient coil design dictates the basic conditions for optimization such as the number of coil elements; however, in principle, the presented method can be applied to any matrix gradient coil or multi-coil design. For the given coil, eddy currents are <1% of the original field owing to the use of Litz wire, shielded element design, and a substantial distance to the cryovessel. They can therefore be ignored for most applications. Mutual coupling has been assessed in Littin et al.'s study (2). Both eddy currents and mutual coupling are ignored in the present work because they are not affected by any of the operations performed during the optimization. The amplifiers further impose restrictions, owing to the limitations regarding the load. This is considered via the maximum and minimum number of coil elements per amplifier in the optimization of the configurations as shown in Kroboth et al.'s study (26). The amplifiers used to drive the matrix gradient coil can be adapted to load changes. Therefore the amplifiers need to

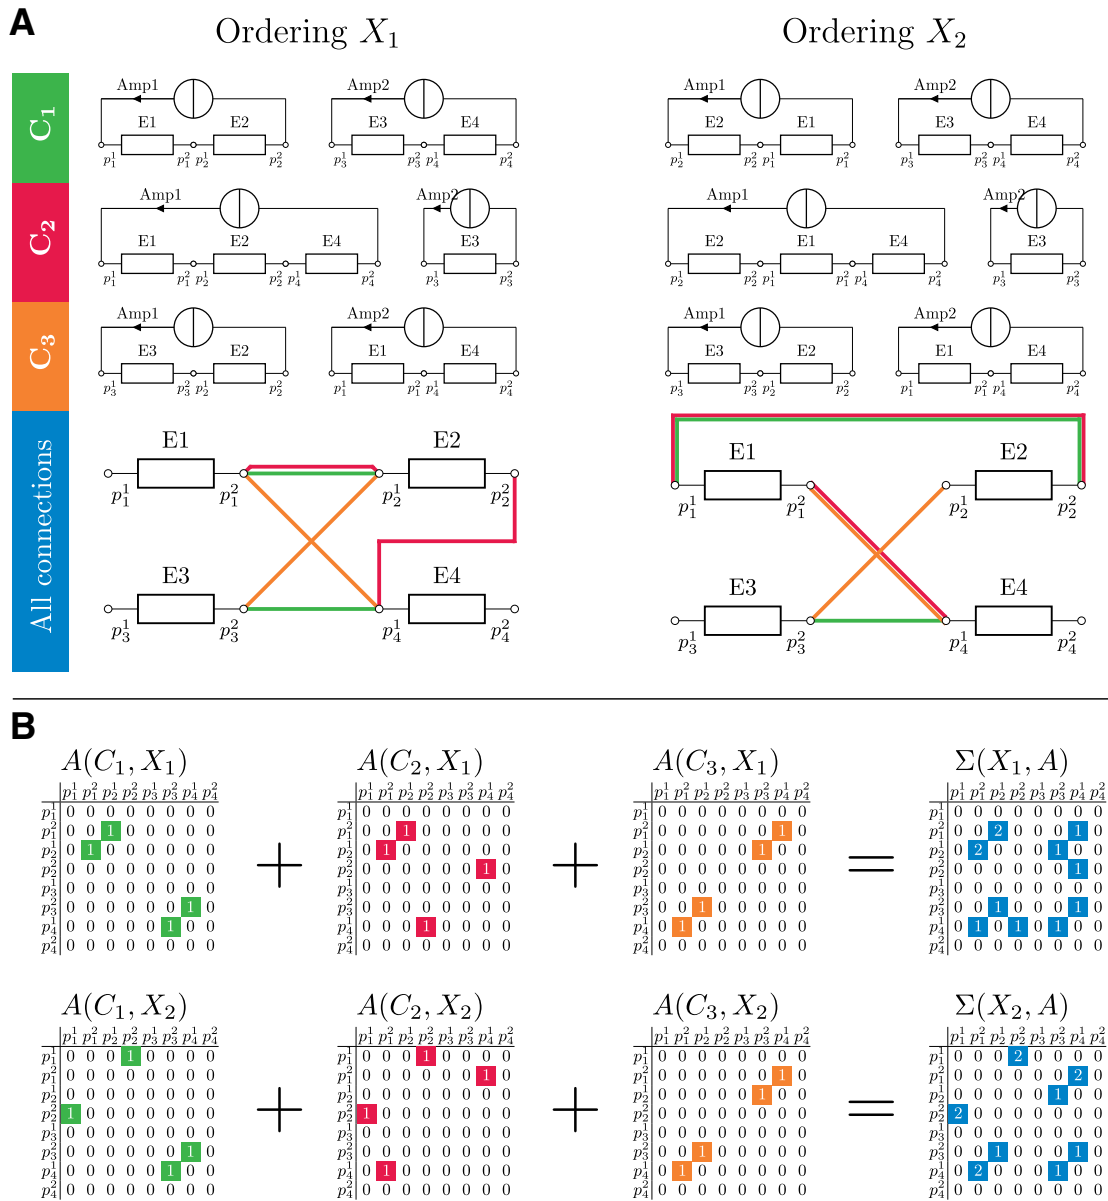


Figure 3. (A) Illustration of how changing the ordering of coil elements within a cluster can affect the number of switches for a matrix gradient coil with 4 coil elements (E1, E2, E3, E4; each box represents a circuit as shown in Figure 2B), 2 amplifiers (Amp1, Amp2), and $N_F = 3$ configurations (C_1, C_2, C_3). (B) The initial ordering X_1 and the updated ordering X_2 are shown in the left and right columns, respectively. The first 3 rows show how the coil elements are assigned to and supplied by amplifiers to create the N_F target fields. The fourth row summarizes all connections between coil elements for each setup. Terminals are labeled with the symbol p . Observe that in C_1 the ordering of E1 and E2 changes in the first cluster, whereas in C_2 the ordering of E1, E2, E4 changes to E2, E1, E4. Configuration C_3 remains unchanged. As can be seen in the last row, X_1 has only 1 connection between terminals which is used by 2 configurations, whereas X_2 has 2 shared connections. While the ordering X_1 requires 5 switches, the ordering X_2 requires only 4 switches. Corresponding adjacency matrices of the undirected graph representing 2 configurations, X_1 (top) and X_2 (bottom). Each nonzero entry represents a connection between 2 terminals of the coil elements. Applying equation (2) to each of the 2 adjacency matrices ($X_1; A$) and ($X_2; A$) gives the final number of switches, which is 5 for X_1 and 4 for X_2 . Similar strategies can be followed for the switches connecting the amplifiers to the coil element network (Figure 4).

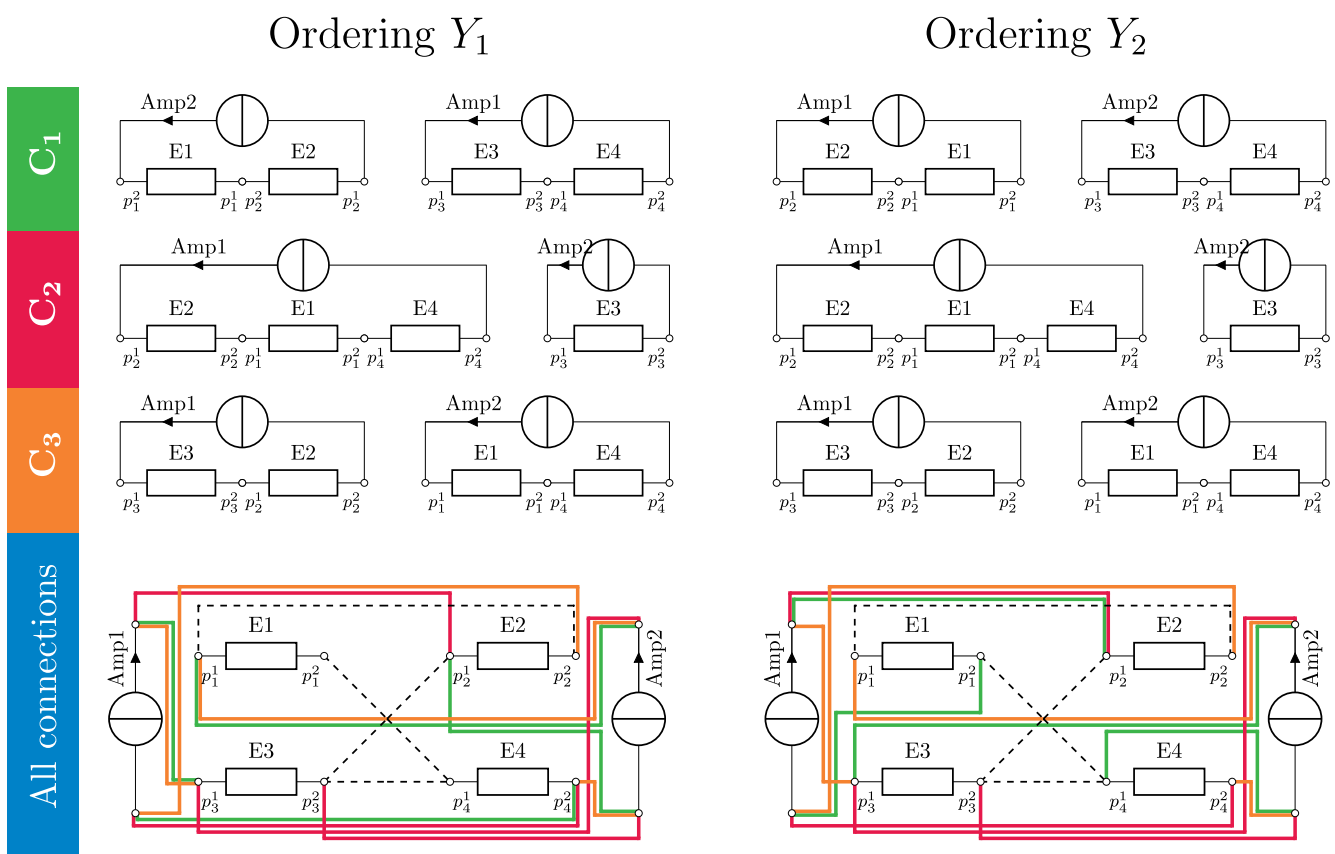


Figure 4. This figure illustrates possible moves taken in the optimization of the switches connecting the amplifiers to the coil element network. The configuration Y_1 is equivalent to the configuration X_2 of Figure 3. The difference between Y_2 and Y_1 is that in configuration C_1 , Amp1 is connected the other way around. Note that the terminals of E1 and E2 are still connected in the same way as in Y_1 . In addition, in configuration C_2 , the amplifiers are swapped such that Amp2 now powers the first cluster and Amp1 the second cluster. None of the possible moves this optimization can perform can influence the result of the previous optimization illustrated in Figure 3. The final row shows all connections for the different configurations. For this toy example, the maximum number of switches is 12. The ordering Y_1 has 3 shared connections (from terminal 1 of Amp1 to p_2^1 , from terminal 1 of Amp2 to p_3^1 , and from terminal 2 of Amp2 to p_4^2). Owing to 3 shared connections, the ordering number of switches reduces to 9. In Y_2 , there are again 3 shared connections (from terminal 1 of Amp1 to p_3^1 , from terminal 1 of Amp2 to p_3^1 , and from terminal 2 of Amp2 to p_1^1 of E1 p_4^2). The connection from Amp2 to p_4^2 is special, as this is wired for every configuration, meaning that there is no switch necessary, and instead, it can be hard-wired. In total, the number of switches for Y_2 reduces to 7.

be reprogrammed whenever the load changes during switching. Switching is performed at zero current conditions only (32).

Algorithm

The optimization problems addressed in this work [equations (3) and (4)] are combinatorial and NP-hard. Performing an exhaustive search over the parameter space is too expensive for practical use. For instance, consider the optimization of a switching circuit that is designed to switch between 3 configurations given a 84-channel coil and 12 amplifiers (where for simplicity, the coil elements are equally distributed amongst the amplifiers). This would require the evaluation of $((84/12)!)^{12} \approx 2 \cdot 10^{133}$ different orderings for the minimization of the number of switches in between coil terminals alone. This number is given by $(84/12)! = 7!$ possible orderings of coil elements within a

single cluster to the power of the number of clusters/amplifiers ($N_A = 12$) and finally to the power of the number of fields ($N_F = 3$). This calculation does not even consider the bridge switches, which would further increase the number of necessary evaluations. The additional complexity owing to the nonconvex and multimodal (many local minima) nature of the optimization problems further aggravates finding the global minimizer with conventional gradient-based methods. Therefore a probabilistic metaheuristic called simulated annealing (SA) (35) is used to find a (local) minimum. This iterative and stochastic optimization method is inspired by the annealing of metals, where the metal is heated to a certain temperature and then cooled down according to a schedule. The energy state of the resulting crystals strongly depends on the initial temperature and the cooling schedule. These ideas can be transferred to optimization prob-

Algorithm 1. Annealing Function 1

```

1: for  $j < \lfloor T^i \rfloor + \text{random\_integer}(1, 5)$  do
2:    $C_r \leftarrow$  random configuration
3:    $K_r \leftarrow$  random cluster in  $C_r$ 
4:    $e_1 \leftarrow$  random element index 1 in  $K_r$ 
5:    $e_2 \leftarrow$  random element index 2 in  $K_r$ 
6:   if  $\text{random}(0,1) > 0.5$  then
7:      $X^i \leftarrow \text{swap\_elements}(X^i, e_1, e_2)$ 
8:   else
9:      $X^i \leftarrow \text{flip\_element\_orientation}(X^i, e_1)$ 
10:  end if
11:   $j \leftarrow j + 1$ 
12: end for
13: Return  $X^i$ 

```

Algorithm 2. Annealing Function 2

```

1: for or  $j < \lfloor T^i \rfloor + 1$  do
2:    $C_r \leftarrow$  random configuration
3:    $v_1 \leftarrow$  random amplifier index 1
4:    $v_2 \leftarrow$  random amplifier index 2
5:   if  $\text{random}(0, 1) > 0.5$  then
6:      $Y^i \leftarrow \text{swap\_amplifiers}(Y^i, C_r, v_1, v_2)$ 
7:   else
8:      $Y^i \leftarrow \text{flip\_amplifier\_orientation}(Y^i, C_r, v_1)$ 
9:   end if
10:   $j \leftarrow j + 1$ 
11: end for
12: Return  $Y^i$ 

```

lems, where instead of metals, parameter vectors are “heated up” to a temperature T^0 , and then cooled down by reducing the temperature over the course of the iterations. The parameter vectors are randomly modified in each iteration by a degree proportional to the temperature. These modifications (sometimes also called *moves*) are performed by a problem-specific *annealing function*. After every iteration, the temperature is reduced according to a cooling schedule. The cooling schedule used in this work is T^0/i , where i is the current iteration number, which has worked well during our simulation experiments. This means that in early iterations with high temperature, the parameter space will be explored randomly by performing big “jumps,” while later iterations with lower temperature will focus on improvements in close proximity to the current position of the parameter vector in the parameter space. The quality of every parameter vector is calculated by the so-called *cost function*. Better solutions are always accepted; however, worse solutions may also be accepted with a probability given by $P = 1/(1 + \exp(\delta/T))$, where δ is the absolute difference of the previous and current cost function value. This is beneficial for overcoming local minima. The acceptance probability increases with temperature and small δ . When the current best cost function value has not improved for N_{SAmax} iterations, we consider it converged and stop the optimization. Note that this method is not guaranteed to find the global optimum.

Annealing Function

The annealing function defines how a parameter vector is modified in each iteration of the SA algorithm. It depends on the temperature T^i of iteration i . If the temperature is high, more modifications will be performed. In the case of minimizing the number of switches between coil elements (Algorithm 1), the annealing function randomly performs 1 of 2 operations with the same probability. It either picks 2 random elements within a randomly chosen cluster of a randomly chosen configuration and swaps them to change the ordering or reverses the polarity of a randomly chosen coil element (which can be accounted for in the actual implementation of the circuit by adapting the state

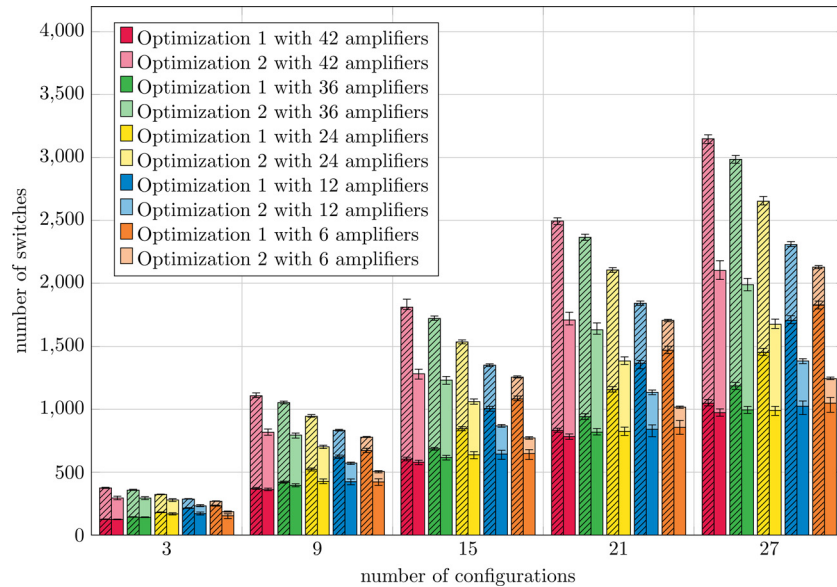
of the switches of the corresponding bridge switch). This step is repeated several times depending on T^i .

When minimizing the number of switches from the amplifiers to the coil element network, the annealing function (Algorithm 2) randomly changes how amplifiers are assigned to clusters within randomly chosen configurations. In other words, within a configuration, the positions of some of the amplifiers in the network are randomly permuted as this may affect the number of switches. The concept is illustrated in Figure 4 using a toy example, where swapping amplifiers Amp1 and Amp2 in configuration C_2 and changing the direction how Amp1 in C_1 is connected to the first cluster reduced the number of necessary switches by 2. The number of times this procedure is repeated depends again on the temperature.

Optimization Settings

To assess the performance of the optimization algorithm, randomly generated configurations were used. The configurations were generated for $N_F = \{3, 9, 15, 21, 27\}$ (number of fields) and $N_A = \{6, 12, 24, 36, 42\}$ (number of amplifiers/clusters). The range of the number of fields is not based on specific applications, but it was chosen to cover a large range with equidistant spacing. All configurations are based on a matrix gradient coil with $N_C = 84$ coil elements. For each setting, $N_D = 100$ random configurations were created. Each optimization is run for $N_R = 4$ times to account for the statistical nature of the solver. For creating the random configurations, the minimum and maximum numbers of coil elements per amplifier are set to $N_{Cmin} = 2$ and $N_{Cmax} = 15$, respectively. This large amount of randomly created configurations covers a large range of possible cases and is used here as a means to assess the performance of the algorithm. Note that these randomly created configurations are likely to not produce useful magnetic fields. In addition, the algorithm has been tested on a subset of the (realistic) configurations obtained in Kroboth et al.’s study (26) for an $N_C = 84$ channel matrix gradient coil. These configurations were designed for $N_A = \{6, 12, 24, 42\}$ (number of amplifiers/clusters) to create spherical harmonics up to full third order (15 fields). Here we optimize a switching circuit for the first 3, 9, and 15 configurations. The number of configurations were chosen

Figure 5. Comparison of the unoptimized (hatched bars) with the optimized case. The mean values of the number of switches between coil elements (before and after optimization 1) and the number of switches from the amplifiers to the coil elements (before and after optimization 2) are stacked to show the total number of switches. Different setups (3 to 27 configurations and the number of amplifiers ranging from 6 to 42) were evaluated. Per bar, 100 random configurations were created where each configuration uses all 84 coil elements. Each optimization instance was run 4 times, therefore the total number of data sets per bar is 400. The variation over all runs is indicated with error bars.



such that the results are comparable to the results obtained with the random configurations (they do not correspond to the number of basis functions contained in certain orders of spherical harmonics). For all runs, the initial temperature was set to $T^0 = 10^4$, and the maximum number of iterations where the best cost function value has not changed was set to $N_{SAmax} = 2 \cdot 10^4$ (stopping criterion). These parameters were experimentally found to obtain good results. The software has been implemented in the Matlab (The MathWorks, Natick, MA) programming language using Matlab’s Global Optimization Toolbox. Each optimization was executed on a single CPU core without parallelization.

RESULTS

Combinations of different hardware setups with a large range of numbers of configurations have been investigated. Figure 5 shows the comparison between the unoptimized and optimized cases for randomly generated configurations. Each bar includes 100 different optimization instances where each is run 4 times, leading to a total of 400 data points per bar. As it would not be fair to compare the optimized switching circuit to the full switching circuit, the optimized circuits are compared with the unoptimized switching circuit. The unoptimized switching circuit is based on the ordering given by the method that produces the configurations. In case of the random configurations, the initial, unoptimized ordering of coil elements within clusters and the assignment of amplifiers to the coil element network is random. For the realistic configurations, the optimization algorithm of Kroboth et al.’s study (26) does not consider the ordering of the coil elements within clusters and the assignment of amplifiers to the network; therefore, the ordering can also be

considered random. Note that therefore the unoptimized case is rather an average case than a worst-case scenario.

Optimization 1 – Switches Between Coil Elements

As expected, the number of required switches increases with the number of configurations that are to be switched. This is because the number of different pathways through the coil element network increases with the number of configurations. The switching circuit needs to be able to realize all these pathways, which in general, comes with the need to increase the number of necessary switches. For the investigated range of configurations, the number of switches after optimization ranges from ≈ 200 switches (3 configurations) to ≈ 1000 switches (27 configurations). The influence of the number of amplifiers (which is equal to the number of clusters, as each amplifier is connected to a cluster) is rather small, at least for the optimized case. In the unoptimized case, fewer amplifiers/clusters lead to an increase in the number of switches. The optimization however is capable of reducing this substantially, see, for instance, the case of 15 configurations and 6 amplifiers. In comparison, for a large number of amplifiers (ie, 15 configurations and 42 amplifiers), the optimization only leads to a small reduction. This suggests that the number of configurations to be switched is the dominating factor in terms of number of switches between coil elements, while the actual composition of the individual configurations plays a minor role.

Optimization 2 – Switches Between Amplifiers and Coil Elements

The results of this part also depend on the number of configurations; however, this time there is also a pronounced depen-

dependency on the number of amplifiers: more amplifiers require more switches. However, the reduction of switches from unoptimized to optimized suggests that the more amplifiers, the higher is the potential for minimization, as there are more degrees of freedom to exploit. This means that this optimization is more effective for a large number of amplifiers, but effectively, a lower number of amplifiers leads to fewer necessary switches. Over all investigated data sets, the number of switches between amplifier and coil element terminals after optimization ranges from ≈ 30 to slightly over ≈ 1130 .

Combined Results of Optimizations 1 and 2

For a given number of configurations, it is favorable to have a low number of amplifiers. However, from the results in Kroboth et al.'s study (26), we know that for a low number of amplifiers, it is in general more difficult to obtain a configuration which is capable of accurately creating a desired target field. This is because lowering the number of amplifiers essentially limits the number of different current values flowing through the coil and hence restricts the coil's capabilities to accurately reproduce target fields. Therefore, the number of amplifiers is usually dictated by the desired field accuracy in the optimization of the configurations, as shown in Kroboth et al.'s study (26). The optimization algorithms themselves perform best for a low number of amplifiers, leading to a reduction (compared to the unoptimized case) ranging from $\approx 8\%$ to $\approx 44\%$. The average reduction across all investigated data sets is $\approx 31\%$.

Realistic Configurations

The algorithm has furthermore been tested on a subset of the configurations obtained in Kroboth et al.'s study (26). These configurations were designed such they are able to create spherical harmonics up to full third order for 6, 12, 24, 42 amplifiers. The median least squares error of all 15 configurations compared with the desired target fields was $\approx 10\%$ for 6 amplifiers, $\approx 5\%$ for 12 amplifiers, and $\approx 2.5\%$ for ≥ 24 amplifiers (including all 84 channels driven individually). Figure 6 shows the results for switching circuits optimized to switch between the first 3, 9, and 15 configurations. These results are in line with the results shown in Figure 5. This affirms the assumption that in practice, the number of switches depends stronger on the hardware setup such as the number of amplifiers and the number of configurations than on the actual composition of the configuration.

DISCUSSION

Many parameters require consideration when designing a matrix gradient coil. Mainly, the coil should be suitable for the task at hand, but also the technical effort and costs associated with it cannot be neglected. The parameters that influence the coil's capabilities and technical effort are (among others) the number of coil elements, the number of amplifiers, and the complexity of the switching circuit. This work in combination with Kroboth et al.'s study (26) aims to shed light on the relationship between those parameters in terms of both technical effort and ability to create fields. In Kroboth et al.'s study (26), we showed (in case of the matrix gradient coil) that even with a reduced number of

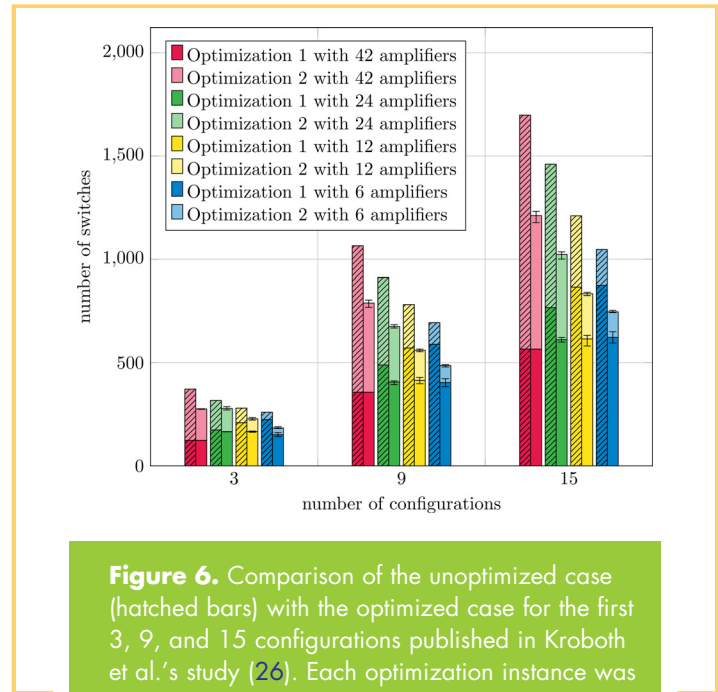


Figure 6. Comparison of the unoptimized case (hatched bars) with the optimized case for the first 3, 9, and 15 configurations published in Kroboth et al.'s study (26). Each optimization instance was run 100 times to account for the statistical nature of the solver. The variation of these runs is indicated with error bars. The general trends match Figure 5.

amplifiers, a large range of field shapes can be created accurately. Here we extend the previous work by showing how the number of switches needed for a switching circuit can be reduced. By changing the order of coil elements within clusters and by varying which amplifier supplies which cluster of a configuration with current, we were able to reduce the number of switches by up to $\approx 44\%$ without affecting the configurations' ability to create a target field at a desired accuracy. Many random configurations were used to assess the algorithms on a wide range of different situations. To address the statistical nature of the optimization algorithms, each instance was optimized 4 times. In addition, a subset of the configurations optimized in Kroboth et al.'s study (26) was used to show the performance of the algorithms on a realistic setup. The problem of minimizing the switches was split into 2 sequential optimization problems. At first, given a set of configurations, the number of switches connecting the terminals of all the coil elements with each other is minimized. Based on the results, the number of switches that connect the amplifiers to the clusters is minimized. Both problems pose challenging combinatorial optimization problems, which were solved with the use of SA. In the early stages of development, the problem was also solved with a genetic algorithm (GA); however, initial tests showed that performance was similar compared with SA. Given that GA has substantially more parameters to tune, SA was chosen. With proper tuning, we expect that GA can outperform SA in terms of speed; however, finding the perfect tuning is difficult given that it is likely problem-dependent. One major advantage of GA is its potential ability to extensively explore the parameter space owing to cross-over. Similar effects can be achieved by high

initial temperatures and/or multiple runs of SA in parallel. Unfortunately, the solution space of the second optimization may, in principle, be constrained by the solution of the first in an unfavorable way. To avoid this, both problems can be solved in a single optimization problem. However, this problem is much more complex, and the results of a pilot study were consistently worse than that of the proposed methods (data not shown).

In addition to reducing the number of switches, the results of this work also expose the relationships between the number of configurations and the number of amplifiers. Figure 5 gives insight into how the complexity of the switching circuit scales with different hardware setups. Running both optimizations requires ≈ 30 minutes to ≈ 8 hours depending on the parameters. The biggest influence on the runtime is the number of configurations. Because switching circuits are not changed or modified often, these long run-times are acceptable. The proposed method operates on configurations only and is therefore independent of the actual matrix gradient coil design, as configurations can be computed for any design.

The resulting optimized numbers of switches in the range of ≈ 200 to ≈ 1200 may appear to be high. Indeed the gradient array method (24) allows to create one localized Z gradient, which is movable along the Z direction without the need for switches (but using 3 gradient amplifiers). It is currently unclear how many gradient channels and amplifiers would be required to extend this approach of moveable FOVs to 3 spatial dimensions. Contrary to that, the *dynamically controlled adaptive current network approach* (33), which was recently introduced for shimming, relies on a single current amplifier, but it requires a substantial number of switches. As noted by the authors, for a body-sized coil with 2-cm wire spacing, 6000 switches would be required. To achieve surface current densities comparable to that of the present matrix coil (2) the number of switches would increase to 25 000. Compared with that, the matrix coil setup with a switching circuit, capable of generating a full third-order spherical harmonic set, driven by only 12 amplifiers and controlled by ≈ 900 switches, appears to be favorable for imaging applications. It is further to be noted that none of the clusters in any of the configurations contained > 15 elements connected in series, which is advantageous with regard to the voltage drop and the power dissipation in the coil.

One potential future step could be combining the optimization of configurations (26) with the optimization of the switching circuit presented in this work. However, this would be a substantially harder optimization problem, which additionally has 2 objectives (field accuracy and number of switches). Alternatively the configuration optimization and switching circuit optimization could be merged partially. For instance by favoring

configurations that offer good starting conditions for the switching circuit optimization. These conditions need to be defined, and their effectiveness and their feasibility have to be assessed in future work. In this current work, it was assumed that each coil element is equipped with a bridge (Figure 2B); however, not every element may require such a bridge. In future work, the switches of the bridges could be included into the optimization to further reduce the technical effort. Furthermore it was assumed that the amplifiers are bipolar which is common for gradient channels. With adaptations, the optimization could also be performed with unipolar amplifiers. This also remains part of further research.

The setup shown in Figure 1 does not allow for the superposition of fields created by different configurations; however, each configuration can potentially create many different field shapes by changing the currents through the individual clusters. Therefore there are many more degrees of freedom for encoding than this diagram may suggest. Theoretically, the switching circuit offers more paths through the network of coil elements than it is optimized for which can be exploited for the creation of fields. Unfortunately, finding all possible or reasonable paths is also a computationally complex problem and the obtained additional functionality most likely does not scale well with the invested effort.

Also the dynamically controlled adaptive current networks B_0 -shimming method (33) can benefit from the methodology presented here. In the original approach, desired current patterns are approximated by a dense network of small wire paths connected with MOSFETs, leading to very high flexibility for the creation of a wide range of current densities, with many of them leading to similar field shapes. The method presented here may be used to sparsify the dense switch network by eliminating the switches that are irrelevant for effective field creation and hardwiring the connections that are persistent across the set of the relevant target fields. The methodology and results presented in this work can further serve as a valuable tool for the search of such tradeoffs by combining it with the various approaches from the literature (2, 4, 24, 33).

CONCLUSION

We presented a method which reduces the effort associated with a switching circuit for matrix gradient coils by minimizing the number of switches for a given set of configurations without affecting the configurations' ability to accurately create a desired target field. The results give insights in the relationship between the number of switches and different parameters, which are important for the design of matrix gradient coils.

ACKNOWLEDGMENTS

This work was supported by the European Research Council (ERC) under grant ERC-POC 755466 "mrSANE" and by the German Research Foundation (DFG) under grant ZA 422/6-1.

Disclosures: No disclosures to report.

Conflict of Interest: The authors have no conflict of interest to declare.

REFERENCES

- Wintzheimer S, Driessle T, Ledwig M, Jakob PM, Fidler F. A 50-channel matrix gradient system: a feasibility study. In Proceedings of ISMRM; Stockholm, Sweden; 2010. p. 3937.
- Littin S, Jia F, Layton KJ, Kroboth S, Yu H, Hennig J, Zaitsev M. Development and implementation of an 84-channel matrix gradient coil. *Magn Reson Med*. 2017; 79:1181–1191.
- Juchem C, Umesh Rudrapatna S, Nixon TW, de Graaf RA. Dynamic multi-coil technique (DYNAMITE) shimming for echo-planar imaging of the human brain at 7 Tesla. *Neuroimage* 2015;105:462–472.
- Juchem C, Nixon TW, McIntyre S, Rothman DL, de Graaf RA. Magnetic field modeling with a set of individual localized coils. *J Magn Reson*. 2010;204:281–289.
- Juchem C, Green D, Graaf RA de. Multi-coil magnetic field modeling. *J Magn Reson*. 2013;236:95–104.
- Jia F, Littin S, Layton KJ, Kroboth S, Yu H, Zaitsev M. Design of a shielded coil element of a matrix gradient coil. *J Magn Reson*. 2017;281:217–228.
- Konzbul P, Sveda K, Srnka A. Design of matrix shim coils system for nuclear magnetic resonance. *IEEE Trans Magn*. 2000;36:1732–1735.
- Hennig J, Welz AM, Schultz G, Korvink J, Liu Z, Speck O, Zaitsev M. Parallel imaging in non-bijective, curvilinear magnetic field gradients: a concept study. *MAGMA*. 2008;21:5–14.
- Schultz G, Ullmann P, Lehr H, Welz AM, Hennig J, Zaitsev M. Reconstruction of MRI data encoded with arbitrarily shaped, curvilinear, nonbijective magnetic fields. *Magn Reson Med*. 2010;64:1390–1403.
- Gallichan D, Cocosco CA, Schultz G, Weber H, Welz AM, Hennig J, Zaitsev M. Practical considerations for in vivo MRI with higher dimensional spatial encoding. *MAGMA*. 2012;25:419–431.
- Gallichan D, Cocosco CA, Dewdney A, Schultz G, Welz A, Hennig J, Zaitsev M. Simultaneously driven linear and nonlinear spatial encoding fields in MRI. *Magn Reson Med*. 2011;65:702–714.
- Schultz G, Gallichan D, Weber H, Witschey WRT, Honal M, Hennig J, Zaitsev M. Image reconstruction in k-space from MR data encoded with ambiguous gradient fields. *Magn Reson Med*. 2015;73:857–864.
- Layton KJ, Gallichan D, Testud F, Cocosco CA, Welz AM, Barmet C, Pruessmann KP, Hennig J, Zaitsev M. Single shot trajectory design for region-specific imaging using linear and nonlinear magnetic encoding fields. *Magn Reson Med*. 2013; 70:684–696.
- Testud F, Gallichan D, Layton KJ, Barmet C, Welz AM, Dewdney A, Cocosco CA, Pruessmann KP, Hennig J, Zaitsev M. Single-shot imaging with higher-dimensional encoding using magnetic field monitoring and concomitant field correction. *Magn Reson Med*. 2015;73:1340–1357.
- Layton KJ, Kroboth S, Jia F, Littin S, Yu H, Zaitsev M. Trajectory optimization based on the signal-to-noise ratio for spatial encoding with nonlinear encoding fields. *Magn Reson Med*. 2015;104–17.
- Stockmann JP, Ciris PA, Galiana G, Tam L, Constable RT. O-space imaging: Highly efficient parallel imaging using second-order nonlinear fields as encoding gradients with no phase encoding. *Magn Reson Med*. 2010;64:447–456.
- Stockmann JP, Galiana G, Tam L, Juchem C, Nixon TW, Constable RT. In vivo O-Space imaging with a dedicated 12 cm Z2 insert coil on a human 3T scanner using phase map calibration. *Magn Reson Med*. 2013;69:444–455.
- Tam LK, Galiana G, Stockmann JP, Tagare H, Peters DC, Constable RT. Pseudo-random center placement O-space imaging for improved incoherence compressed sensing parallel MRI. *Magn Reson Med*. 2015;73:2212–2224.
- Wang H, Tam L, Kopanoglu E, Peters DC, Constable RT, Galiana G. Experimental O-space turbo spin echo imaging. *Magn Reson Med*. 2016;75:1654–1661.
- Dispenza NL, Littin S, Zaitsev M, Constable RT, Galiana G. Clinical potential of a new approach to MRI acceleration. *Sc Rep*. 2019;9:1912.
- Turner R. Gradient coil design: A review of methods. *Magn Reson Imaging*. 1993;11:903–920.
- Lemdiasov RA, Ludwig R. A stream function method for gradient coil design. *Concepts Magn Reson Part B Magn Reson Eng*. 26B:67–80.
- Poole M, Bowtell R. Novel gradient coils designed using a boundary element method. *Concepts Magn Reson Part B Magn Reson Eng*. 31B:162–175.
- Smith E, Freschi F, Repetto M, Crozier S. The coil array method for creating a dynamic imaging volume. *Magn Reson Med*. 78:784–793.
- Jia F, Schultz G, Testud F, Welz AM, Weber H, Littin S, Yu H, Hennig J, Zaitsev M. Performance evaluation of matrix gradient coils. *MAGMA*. 2016;29:59–73.
- Kroboth S, Layton KJ, Jia F, Littin S, Yu H, Hennig J, Zaitsev M. Optimization of coil element configurations for a matrix gradient coil. *IEEE Trans Med Imaging*. 2018;37:284–292.
- Weber H, Gallichan D, Schultz G, Cocosco CA, Littin S, Reichardt W, Welz A, Witschey W, Hennig J, Zaitsev M. Excitation and geometrically matched local encoding of curved slices. *Magn Reson Med*. 69:1317–1325.
- Witschey WR, Cocosco CA, Gallichan D, Schultz G, Weber H, Welz A, Hennig J, Zaitsev M. Localization by nonlinear phase preparation and k-space trajectory design. *Magn Reson Med*. 2012;67:1620–1632.
- Van Essen DC, Ugurbil K, Auerbach E, Barch D, Behrens TEJ, Bucholz R, Chang A, Chen L, Corbetta M, Curtiss SW, Della Penna S, Feinberg D, Glasser MF, Harel N, Heath AC, Larson-Prior L, Marcus D, Michalareas G, Moeller S, Oostenveld R, Petersen SE, Prior F, Schlaggar BL, Smith SM, Snyder AZ, Xu J, Yacoub E. The Human Connectome Project: A data acquisition perspective. 2012;62: 2222–2231.
- Umesh Rudrapatna S, Fluerebrock F, Nixon TW, Graaf RA de, Juchem C. Combined imaging and shimming with the dynamic multi-coil technique. *Magn Reson Med*. 2019;81:1424–1433.
- Kroboth S, Layton KJ, Jia F, Littin S, Yu H, Zaitsev M. Optimization of matrix gradient coil switching for a limited number of amplifiers. In: Proceedings of ISMRM. ISMRM. Toronto, 2015.
- Yu H, Huette F, Littin S, Layton KJ, Kroboth S, Jia F, Hennig J, Zaitsev M. An improved design of multi-channel switching circuit for matrix gradient coil. In: Proceedings of ISMRM. ISMRM. Toronto, 2015.
- Harris CT, Handler WB, Chronik BA. A new approach to shimming: the dynamically controlled adaptive current network. *Magn Reson Med*; 2014;71:859–869.
- Applegate DL, Bixby RE, Chvatal V, Cook WJ. The Traveling Salesman Problem: A Computational Study. Princeton Series in Applied Mathematics. 2007;608.
- Kirkpatrick S, Gelatt CD, Vecchi MP. Optimization by simulated annealing. *Science*. 1983;220:671–680.