COMBINED 3D ELECTROMAGNETIC AND SPIN RESPONSE SIMULATION OF MRI SYSTEMS
WHITEPAPER

Modern MRI systems are highly complex devices, and the interaction between the body and MRI coils introduces additional challenges into the design process. The body, with its complicated heterogeneous interior, causes major disturbance to the homogeneity of the magnetic fields, while energy absorbed by the body can cause harmful heating. Experimental measurement of these effects is often impossible, but simulation with 3D body models can help the engineer identify risks to the patient and suggest ways to reduce them.

MRI is a fundamental part of modern diagnostic imaging. With MRI, structures inside the body, even those made of soft tissue, can be imaged relatively quickly at a good resolution. However, the machinery needed to produce these images is complex. For reasonable image quality, the fields inside the scanner should be very homogeneous in the area of interest, which means that the magnets and coils need to be carefully designed to give the right field distribution. One way of testing the field distribution is to use prototypes. Unfortunately, MRI scanners are filled with large, precision-made components, which make building multiple prototypes difficult and expensive.

Simulation offers a much cheaper and faster way to test a design than repeatedly creating and testing new prototype coils and making incremental changes. As well as being a very flexible and fast approach for testing the properties of a design, it gives additional insight into the functional mechanisms of the system, allowing automatic optimization and tuning schemes. Alongside time and budgetary considerations, safety also has to be taken into account. Although MRI is usually thought of as safe by comparison to X-rays and CT scans, it can pose its own risks to the patient. During an MRI scan, the patient is subjected to significant RF fields, typically with frequencies on the order of hundreds of megahertz. These fields deposit energy in the body, and this causes heating. If the energy absorbed by the body exceeds the safety limits, significant damage can be done to tissues.

Heating and energy absorption are very difficult to test experimentally, except for by using very simplified phantoms – homogeneous models filled with fluid. Because phantoms are so simple, the energy absorption experienced by a phantom may have little relation to the energy absorbed by an actual person, as the complicated structures within the body reflect and focus the RF fields in hard-to-predict ways. Figure 1 shows the simulated

(a)   (b)   (c)

Figure 1: Simulated SAR values in: **(a)** a cylindrical phantom **(b)** a homogeneous head phantom and **(c)** a heterogeneous head model.

specific absorption rate (SAR) distribution for two phantoms and a heterogeneous head model. The distribution of energy absorption is very different in the full model, and the simulation reveals a critical hotspot in the back of the brain that did not appear in the phantoms.

To complicate matters further, the thermal properties of living tissue are not the same as those of a simple fluid. Cells generate some heat themselves as they metabolize, while blood flow disperses temperature hotspots and sweating promotes heat loss through the skin. A phantom cannot replicate these so-called "bioheat" effects, as demonstrated in Figure 2, but at the same time it is also impossible to measure the temperature distribution within a living patient undergoing an MRI scan to the required accuracy.

Although they are useful tools for investigating the behavior of a scanner or coil, on their own phantoms are not enough for examining the effects of MRI on the body. Simulation offers a much more reliable way of calculating the risks MRI poses to a patient, by taking into account the properties of all the different materials and structures that make up the human body.

SETTING UP: CHOOSING APPROPRIATE BIOLOGICAL MODELS

To simulate the fields inside a person accurately, we need a model that describes the incredible complexity of the human body – after all, a simulation is only as good as the model it's using. There are a range of different model types, with different advantages and disadvantages.

The simplest model is a homogeneous body; essentially a phantom in the shape of a human. Homogeneous models are widely available, and they are usually jointed and easy to pose. However, as already shown, a homogeneous phantom – even one shaped like a person – is an inaccurate representation of the body. The main advantage to using homogeneous models in MRI simulations is that they can usually be simulated relatively quickly, and the results from a simulation can be easily compared to the results from a real measurement using a phantom, as will be demonstrated later in the paper. If the measured results match the simulated results, that is a good indication that no errors

Figure 2: Simulated temperature distribution in (a) a homogeneous phantom, and (b) a heterogeneous model taking bioheat into account.

**Figure 3:** Complete set of CST Studio Suite Voxel Family models, showing the range of ages and body shapes available.

Figure 4: A voxel model of a rat for use in veterinary MRI.

were made and the simulation will deliver reliable results even for more complex models.

For a full simulation of the internal fields, we need a heterogeneous voxel model. Voxel models are built up of small cubes, with each cube representing the tissues in a

Figure 5: A comparison of the images from a low field and a high field MRI, showing a cross-section of the hippocampus.

small volume of the body. These produce accurate models even of quite small structures inside the body, including their thermal and electromagnetic properties.

Individuals can affect the RF field very differently – a change in size or shape can have a big effect on the field distribution and the heat generated in the body. This is where the range of body shapes provided by the CST Studio Suite Voxel Family[1] comes in useful. The CST Studio Suite Voxel Family comprises seven individuals, both male and female, of a variety of ages and body shapes, including a baby and a pregnant woman – Figure 3 shows the models available in the family.

The electromagnetic characteristics of a tissue can vary depending on the frequency of the applied field, so CST Studio Suite includes macros to generate the correct permittivities and loss values for any frequencies that are of interest. These macros can take into the account the loss of water from cells caused by aging, and even the difference between the tissues of the fetus and the mother. CST Studio Suite® supports other voxel models as well, such as the Visible Human model (HUGO) along with general voxel models imported from other sources, such as the rat model shown in Figure 4.

SOLVER CHOICE AND COMPLETE TECHNOLOGY

MRI is a interdisciplinary field, drawing from a number of different areas of physics. CST Studio Suite offers an integrated design environment for multiphysics simulation to allow all the different simulations to be carried out as part of one workflow, with the results from each simulation forming the basis of the next. The superconducting magnets in an MRI machine generate a large static magnetic field, which is simulated using the magnetostatic solver in CST Studio Suite; the gradient coils produce dynamic but low-frequency fields which are best simulated by the magneto-quasistatic solver and the RF coils generate high-frequency fields which can be calculated by both the time domain and the frequency domain solvers. As well as the EM effects, the patient in the machine is subject to thermal and biophysical effects, which can also be simulated in CST Studio Suite.

The choice of model and the type of problem in turn influence the choice of solver. This paper mainly focuses on


Figure 6: An MRI scan degraded by interference.

the design of the RF coils and, for high frequency problems, CST Studio Suite has two very powerful solvers at its core: the transient solver and the frequency domain solver. Both can be used to solve most MRI problems, but for some situations, the transient solver will work faster than the frequency domain solver, or vice versa.

The transient solver works on a hexahedral mesh, dividing the object up into cuboid mesh cells. The Perfect Boundary Approximation (PBA)® and Thin Sheet Technique (TST)™ methods allow the simulation to capture the shape of the object at boundaries without having to use inaccurate staircase cells or a very fine mesh size. The time taken for the solver to run scales linearly with the number of mesh cells. This makes the transient solver useful for solving problems that are large and detailed, as many MRI problems are. The transient solver can also take advantage of GPU computing, which can significantly speed up the simulation process.

The frequency domain solver on the other hand can use either a hexahedral mesh or a tetrahedral mesh. The tetrahedral mesh follows the geometry of the object, building up a simulation model out of tetrahedrons. With it, there is no need for hexahedral meshing, and the cells can even be curved to provide a better representation of round surfaces as often seen in MRI systems. This mesh is not compatible with voxel models, but it is suitable for modeling coils and homogeneous phantoms. If the model has multiple ports, the frequency domain solver can handle all of them at once, and it can calculate the fields in highly resonant structures quickly.

For multi-coil systems, the solvers can be supplemented by a circuit simulator such as CST Studio Suite, which can match, tune and link the individual field solutions of every coil without having to run a full 3D simulation at every step.

To ensure the safety of the MRI scan, it's particularly important to examine its thermal effects. The power losses

Figure 10: The fields around each of the 8 different transmitters in a head coil.

Figure 7: A model to test an 8 channel head coil, including the magnet bore, the gradient coils and the HUGO voxel model.

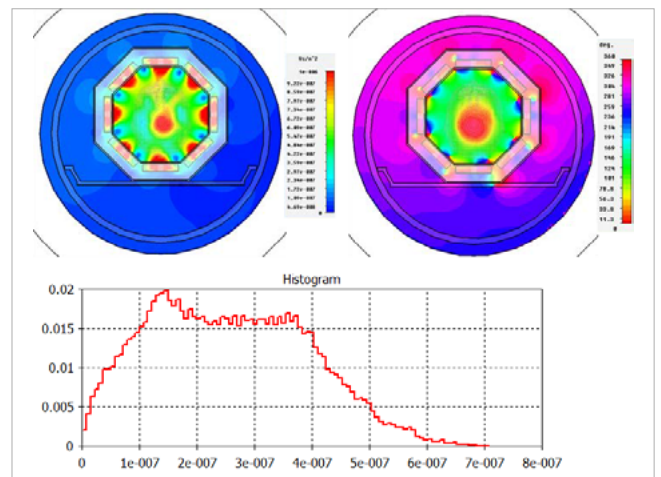**Figure 11:** Amplitude (upper left) and phase (upper right) of the B1⁺ field generated by a tuned multi-channel headcoil. Both are relatively homogeneous around the center of the brain. The lower picture shows a histogram of the B1⁺ distribution inside the brain.
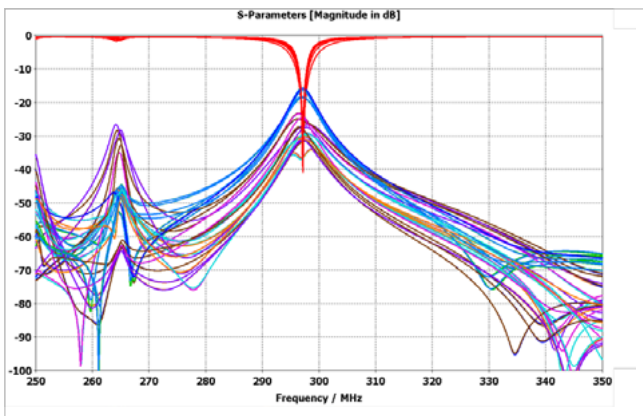

Figure 8: A multi-channel coil connected to a tuning circuit in CST E.

Figure 12: A simulated phantom in a head coil and a real phantom.

Figure 9: The S-parameters for a tuned coil, showing good matching at the desired frequency of 297.2 MHz and low coupling to other coils.

Figure 13: A comparison of the simulated (left) and measured (right) B1+ field distributions

calculated during the initial simulation can be used as a thermal source in the next simulation allowing for a direct simulation of the heating experienced by the body.

CST Studio Suite offers two thermal solvers: a thermal stationary solver and the thermal transient solver, both of which can take bioheat into account when calculating temperaturedistributions. Thermal simulations provide an additional layer of safety beyond the SAR data, by highlighting any spots, where tissues might be heated past their safe limits.

DESIGNING AND TUNING AN ULTRA-HIGH-FIELD MRI COIL

Safety and efficiency are important when designing any MRI device, but the new generation of ultra-high-field devices, with B-fields of 7 tesla, 9.4 tesla or more, introduces a new set of problems for MRI designers to consider. High-field devices produce images with much better resolutions than their low field counterparts, as shown in Figure 5, but the high fields increase the resonant frequency of the protons, and higher frequencies mean greater power dissipation – while the signal to noise ratio improves linearly over frequency, the SAR value is proportional to the square of the frequency, and this can increase the heating effect substantially.

Higher frequencies also mean shorter wavelengths – for a typical 7 T device, these will be around 13 cm, comparable to the size of structures in the body. At such short wavelengths, interference has a major effect on the quality of the images. Figure 6 shows one possible result of such interference; in this case, a long black shadow on the left-hand side of the image. These shadows can hide clinically relevant details, so it is important that these be reduced as much as possible.

To reduce interference and improve patient safety while still experiencing the benefits of high-field MRI, radiographers use multi-channel coils that allow very precise control of the RF fields. Simulation can be used to calibrate these coils, to prevent crosstalk between the separate parts of the coil and to make sure the fields produced are homogeneous in the area of interest.

To simulate and tune a multi-channel coil, a model of the coil needs to be either created in Studio Suite or imported in from another CAD package. The coil's inputs and outputs are modeled with ports, which allow them to be linked together by a circuit using CST E, as shown in Figure 8.

A typical simulation workflow would now always follow these steps:

• Simulate S-parameters and field solutions of every individual channel with a 3D field solver
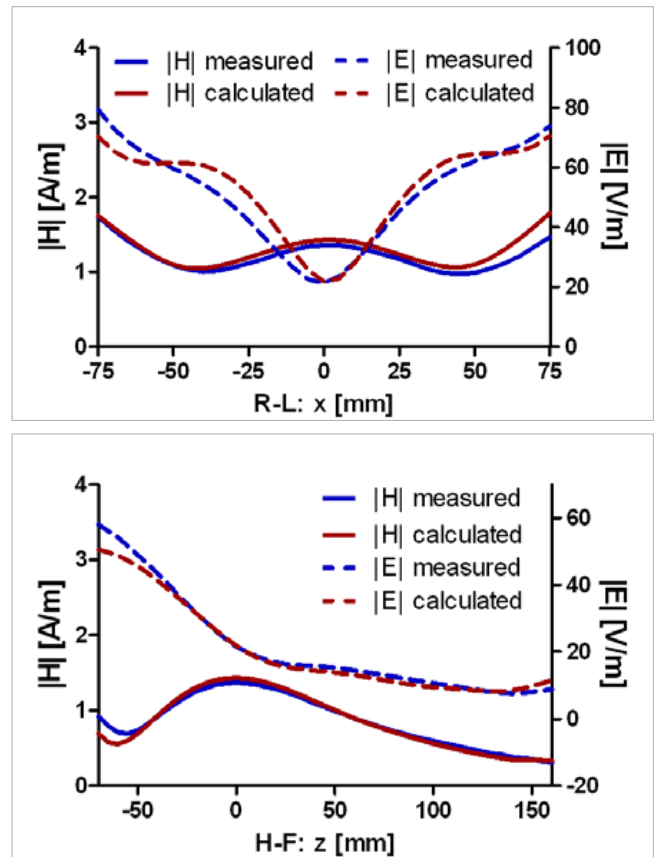

Figure 14: H and E-field distributions along the x (top) and z (bottom) axes highlighted in Figure 13.

• Tune, match and decouple the individual coils in the circuit simulator using the previously calculated S-Parameters
• Rescale and combine the individual field solution based on the tuned S-Parameters and the phases and amplitudes of the external ports
• Run postprocessing methods to extract relevant data like $B1^+$ and SAR, or carry out thermal calculations on combined fields

These steps are described in more detail for an eight channel head coil model as seen in Figure 7 below. (A photograph of the coil is shown in Figure 12).

The most computationally demanding part of the process is simulating each transmitter in turn using the transient solver. The coil used in this example, with 16 ports (two single-ended ports per channel) and 40 million mesh cells, took around 10 hours to simulate, but once run, the transient solver generates S-parameters describing how RF energy propagates through the multi-port network.

CST E can use these S-parameters in the second step to tune the circuit quickly without having to run the simulation again, automatically changing the circuit elements around each coil. This gives us one tuned and matched field per channel, as shown in Figure 9 and Figure 10. Figure 9 shows the S-Parameters all tuned to 297.2 MHz (the proton resonance frequency of a 7 tesla system), Figure 10 the 8 individual magnetic fields.

Figure 15 Top: Simulated field distribution in three patients from identical applied fields in a body coil.
Bottom: Measured field distribution, as measured by flip angles, in three subjects similar to those simulated at the top.

The fields from each transmitter can then be combined and, by optimizing the phase and amplitude of the signal at each external port of the coil, the combined field can be made homogeneous in the area of interest. Once generated, the combined field can be used for the postprocessing or thermal simulations. Figure 11 shows the combined $B1^+$ field for the head coil in use in a 2D cutplane and as a statistical evaluation of the homogeneity of the B1+ field inside the brain in form of a histogram.

Further postprocessing can be applied to the final combined results, as will be shown later.

VALIDATING SIMULATIONS

Before carrying out safety-relevant postprocessing, it is worth double-checking if the simulated results agree to measured ones, if possible.

The frequency domain solver and the transient solver use two very different methods for simulating electromagnetic fields, and running both solvers on one model is a good way of validating the simulation, even without measured data being available. If both solvers produce similar results, it is likely that the simulation settings are properly set up.

However, for further validation, simulations can also be tested against real world measurements. Homogeneous phantoms are widely available in MRI labs and easy to test, and their simple structure makes them easy to simulate too. A simulation using a model of a phantom should, if the simulation is set up correctly, produce very similar results to an experiment using the same phantom.

Figure 12 shows an experiment set up to test the simulation of the field distribution inside a head coil. A cylindrical head and shoulder phantom filled with tissue-simulating fluid was used, and a model of it was constructed in CST Studio Suite with the same electromagnetic properties. The phantom was exposed to a field from the head coil, and the $B1^+$ field distribution within the phantom was measured. The 2D cutplane results, shown in Figure 13, and the 1D results plotted in Figure 14 along the two dashed lines seen in Figure 13 (left), are very close to the simulated measurements both qualitatively and quantitatively. While the magnetic fields represent the fields needed to create the image, the E-fields also need to be controlled in order to keep the power absorbed by the body low.

Testing the coil empty will not give the full picture of how it will work in practice, however. As already mentioned, putting a patient in the coil introduces sources of interference, and this will affect the field distribution. Figure 15 shows the different field patterns in three different subjects for a simulated body scan and an actual body scan.

Three subjects similar (but not identical) in body shape to the models used in the simulation shown in Figure 15 were found, and scanned using a body coil at 7 tesla. The $B1^+$ field distribution was measured for an applied field with the same phase and amplitude relations to that used in the simulation as shown in Figure 15, and the measurements agree quite well with the simulated results. The models used, although they were not exact matches of the patients, still predicted the key features of the field, showing which fields would be homogeneous and which would not.

At present, this is most useful at the design stage, to ensure that the coil will produce usable images from a wide range of body shapes. These simulations may one day also be used to create kinds of look-up tables, so that radiographers will be able to use the results from these simulations to calibrate their coils before the patient even enters the scanner.

POSTPROCESSING: SAR AND THE MRI TOOLBOX

CST Studio Suite contains a toolbox of postprocessing templates to calculate relevant quantities needed to characterize an RF coil. One such quantity is the distribution of the B1+ and B1- fields. To make sure that the coil produces an uniform field inside the subject, we can use an optimizer to find a usable B1+ field, taking into account the electromagnetic properties of the body.

Once we have produced a homogeneous field, whether with a single-channel or multi-channel coil, and verified our result, we then have to make sure this field cannot harm the patient.

Because it is very difficult to determine the temperature distribution inside a patient, the standard measure used when quantifying MRI safety is the specific absorption rate, or SAR.

The SAR is the power absorbed by the body per mass of tissue. There are a number of ways of presenting the SAR: averaged over a volume containing some mass, over a whole organ or structure, or simply as a point-by-point value. SAR can be calculated at certain frequencies from both TD- and FD simulations as well as time averaged for broadband periodic pulse excitations. There are legal limits on the whole-body, partial-body and 10g-/1g-averaged SAR values to prevent patients from being exposed to too much RF power during the scan. The SAR, in all its forms, has to be carefully monitored during the simulation.

Because the MRI Toolbox uses postprocessing methods, it can be run after the simulation, and multiple results generated without having to start the solvers again. The SAR templates generate field distribution plots, like those produced by field monitors, as well as log files. These log files contains statistics about the input power, reflected power, power absorbed in tissue and the dielectric losses, as well as the SAR distribution across either the total volume or a subvolume, including the total absorbed power, the total SAR and the highest averaged SAR and point SAR.

These templates give a good estimate of how much power will be absorbed by the patient when the coil is working as expected during an MR examination with constant amplitude and phase weights of the channels. In the case that different waveforms are played out in the individual channels, the computational cost of the underlying algorithm may be too high. To allow for an efficient SAR analysis, "Virtual Observation Points" (VOPs) can be utilized in such cases. Instead of checking SAR in millions of mesh-cells, only about 70 points are necessary in the given example, if a SAR overestimation of 10% is accepted, or about 400 if the maximum overestimation is 5%. Hence, the VOP approach allows for online local SAR monitoring during MR scans and RF pulse optimization under local SAR constraints.

If the phase relation is fully unknown, a worst case scenario can be considered, which assumes constructive interference of the individual transmit channels in every mesh cell. In this example, this leads to an entirely new SAR hotspot forming in the left side of the patient's head, as shown in Figure 17.

More critically, the worst case SAR peak for this coil is 2.5 times higher than the SAR value in the regular circular polarized operation mode. So that we can work out what risk a malfunction or a mistake when setting up the phases poses to patients, the MRI Toolbox also includes a template for calculating the worst case SAR.

Using the multiple voxel models available in the CST Studio Suite Voxel Family, it is also possible to examine the power loss and SAR values for multiple body shapes. As seen earlier, different body shapes affect the fields in different ways, and this can have a major effect on the way that power is absorbed by the body.

The best practice for calibrating a coil is to run simulations using a number of voxel models, evaluate SAR for each one, and then, for each model, find the type of SAR that poses the most risk to the patient. For example, if the worst case SAR includes a small but severe hotspot, the point SAR or 10 g average SAR is the most critical one to reduce, while if the power loss is evenly distributed across a large area, the SAR averaged across the whole body or a part of it is more likely to be important. If the scan region includes a particularly sensitive organ, the SAR across particular tissue types can be calculated as well in order to work out how much energy that organ absorbs.

Once we know which SAR values are most important, we can then set about reducing them. Since the power absorbed scales linearly with the power of the coil, it's relatively simple to just decrease the power of the coil until all the SAR values are well within the safe region. This makes it unlikely that the coil will ever produce enough RF power to harm a patient, but also means that the image quality may be reduced or the duration needed to take the image may be increased. With simulation, it may one day be possible to set up coils at different strengths for different patients, so that every scan can produce a good quality image without exposing the patient to unnecessarily high field strengths.

In the example illustrated in Figure 18, three different models were exposed to the same field from a head coil. In the models shown in Figs. 18(a) and 18(b), the SAR distributions are quite uneven, with some noticeable hotspots. In these images, the most critical SAR value is the 10 g averaged SAR, and the maximum permissible average power is 25 W for the model in Fig. 18(a) and 33 W for the model in Fig. 18(b). In the model in Fig. 18(c), the SAR distribution is very different. There are no hotspots: the power loss throughout the imaging region is fairly constant. In this case, both the SAR averaged across the entire head and the 10 g averaged SAR are important, and the maximum permissible average power is 35 W. Based on these results, the coil's average power over the cycle should be kept below 25 W, the lowest of the permissible power values.
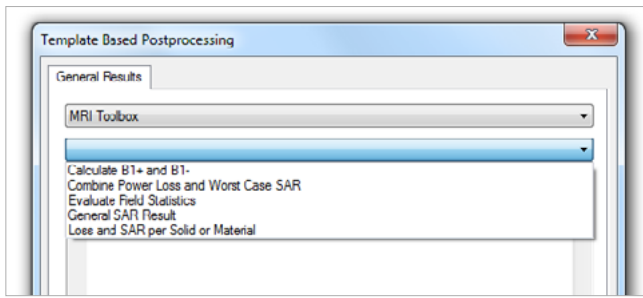
**Figure 16:** The MRI Toolbox in CST Studio Suite.

Figure 17: 10 g average SAR distributions for (a) a tuned coil and (b) a worst case tuning.
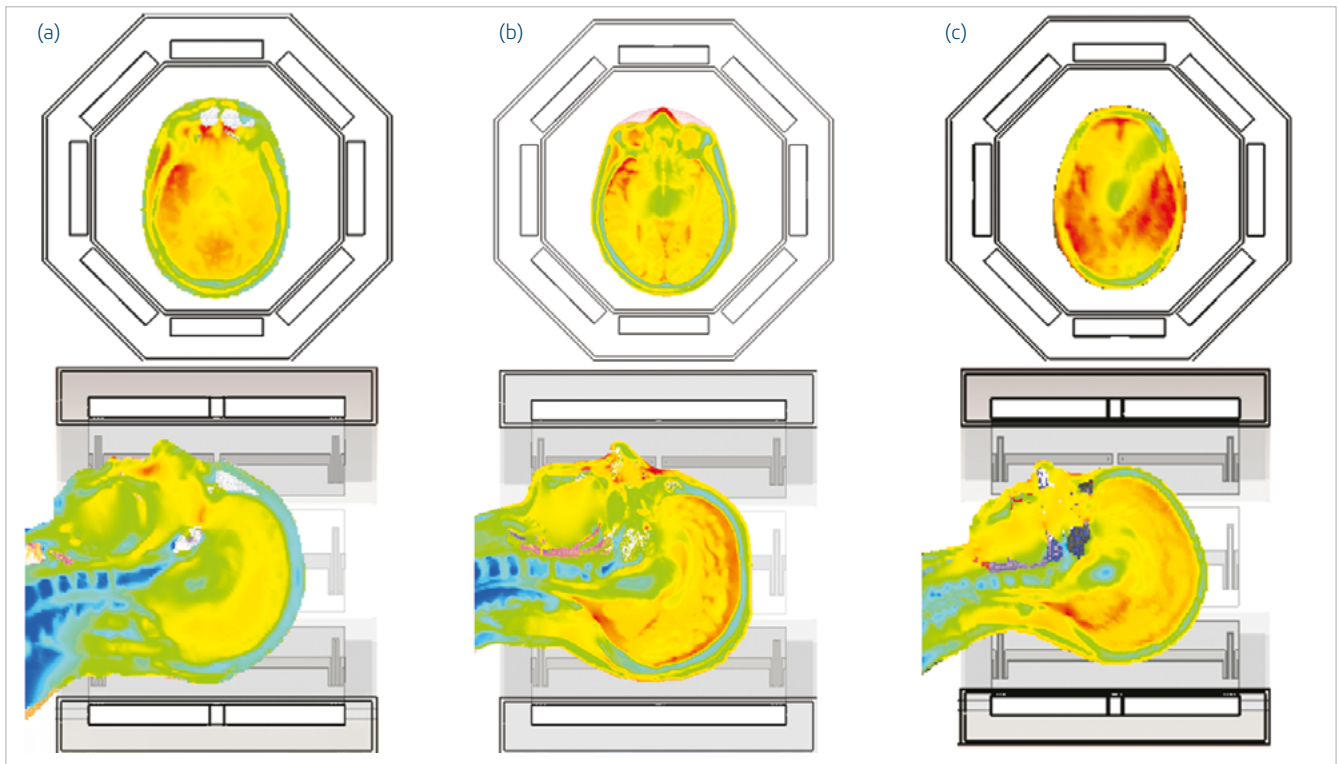
(a) (b) (c)

Figure 18: 10 g averaged SAR distributions in three different models in the same head coil, in transverse (top) and sagittal (bottom) planes.

One other form of SAR value which is sometimes used is the realized B1+ per 1 W/kg SAR. This normalizes the power of the B1+ field so that it is 1 W at the point of maximum SAR, and this is sometimes used to compare the efficiency of different coils.

TEMPERATURE AND BIOHEAT

SAR was originally introduced as a measure to estimate the heating inside the body, but it can be more appropriate to use the CST Studio Suite bioheat solver to directly calculate temperature distributions. In a thermal simulation, the EM losses act as a heating source. As well as classical parameters like thermal conductivity and capacity of materials, the bioheat solver also considers thermal properties of living tissues such as the metabolic heating from chemical processes in the body and the bloodflow, which typically acts to reduce temperatures above the basal temperature of 37°C. To avoid overheating of tissues, the bloodflow rate tends to increase strongly for temperatures above 37°C, a

process called thermo-regulation. This effect is considered alongside the temperature dependency of permittivity, electric and thermal conductivities. Finally, sweating can be modelled via a surface convection.

Generally, a tissue temperature of up to 39°C is considered to be acceptable. Practical experience shows that the temperature criterion is actually less tight than the SAR one in many cases, which means that it is often possible to apply more power to the MRI coil while still maintaining patient safety.

One further example of a situation where thermal simulation is very useful is the study of implant safety. Metal implants, illustrated in Figure 19, are widely used in surgery to support bones and joints, but they can intensify the fields around them, increasing the local SAR, while eddy currents caused by oscillating fields can cause heating. Thermal simulation, as demonstrated in Figure 19, can be used to test implant design and placement to make sure that any heating is minimal.

Figure 19: Left: A model in a head coil with three implants affixed to the side of the skull Right: Temperature distribution inside the skull with the implants. The temperatures remain well below the limit of 39°C specified in the safety guidelines.

IMAGE SIMULATION

Obviously a homogeneous field distribution is a good indicator for high quality MR images. However, if a perfect field distribution cannot be achieved, it is useful to see what impact this has on the final image. In addition, the selected imaging sequence influences both the image contrast and possible imaging artefacts. The expected quality of the resulting image can only be evaluated using a dedicated MRI simulator such as the Juelich Extensible MRI simulator (JEMRIS – www.jemris.org). JEMRIS is a free open source program for solving Bloch equations which describe the behavior of spin ensembles, for example in biological tissue.

The calculations can be based either on ideal field distributions or ones imported from CST Studio Suite, and on a scan sequence that can be described in JEMRIS

This simulates the imaging process taking into account the effects of the transmitting (B1$^+$) field distribution as well as the receiving (B1-) field distribution calculating relevant quantities such as the field of view (FOV) of the coil, its geometry factor (gfactor) for parallel image acquisition and the noise covariance between the different channels, as demonstrated in Figure 20. JEMRIS is extendible and can take advantage of parallel processing and cluster computing, making it both flexible and powerful. CST Studio Suite and JEMRIS are tightly coupled, offering a complete MRI simulation workflow. New equipment can have its EM field characteristics simulated in the CST Studio Suite, and its image properties tested in JEMRIS, purely based on a virtual computer model of the MRI system.

SUMMARY AND CONCLUSION

All the examples described in this paper demonstrate the power of simulation when applied to MRI problems. With the right choice of models and solvers, simulation can accurately represent the behavior of the magnetic fields used in MRI, and therefore can be a tremendous design help during MRI coil development. In addition, simulation is capable of predicting the effects these fields will have on living tissue. Safety is of great importance in the imaging


Figure 20: Sample outputs from JEMRIS for a multi-channel head coil, showing the voxel model used, with colors indicating different tissue types (top left), the simulated MR image (top right), the coupling between the individual coil elements described through the noise covariance matrix (bottom left) and the g-factor, 256x256, 4 fold acceleration (bottom right) across a transverse slice of the head.

process and, due to the absence of measurement options inside the living human, simulation is the only way of estimating safetycritical MRI characteristics with high accuracy, thereby helping the engineer and the radiographer reduce the risk to their patients.

REFERENCES

[1] CST Studio Suite, https://www.3ds.com/products-services/simulia/products/cst-studio-suite/

AUTHORS

Andreas Bitz, Erwin L. Hahn Institut, Essen, Germany
Jörg Felder, Forschungszentrum Jülich, Germany
Tilmann Wittig, SIMULIA Senior Technical Sales Manager, Dassault Systèmes
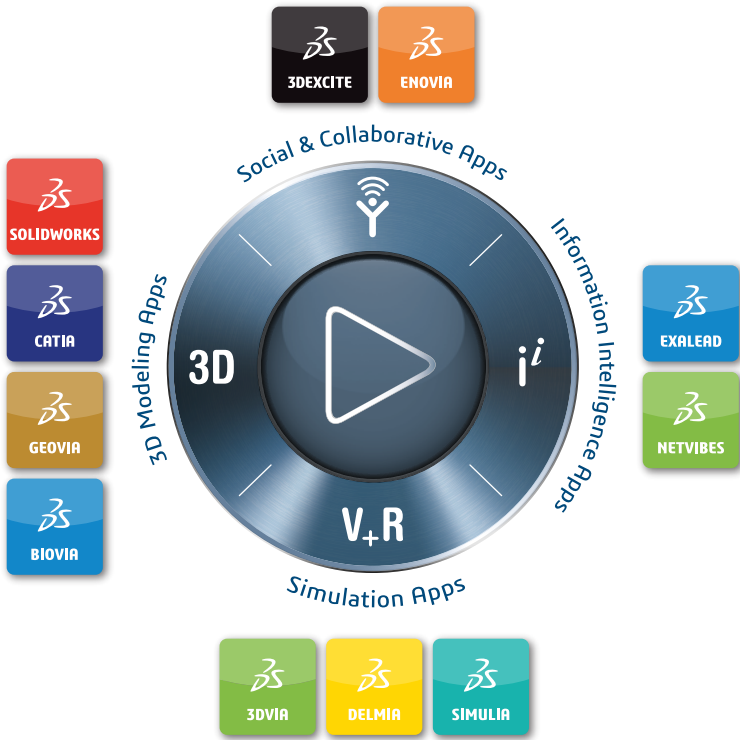
Our **3D**EXPERIENCE® platform powers our brand applications, serving 12 industries, and provides a rich portfolio of industry solution experiences.

Dassault Systèmes, the **3D**EXPERIENCE® Company, provides business and people with virtual universes to imagine sustainable innovations. Its world-leading solutions transform the way products are designed, produced, and supported. Dassault Systèmes' collaborative solutions foster social innovation, expanding possibilities for the virtual world to improve the real world. The group brings value to over 210,000 customers of all sizes in all industries in more than 140 countries. For more information, visit **www.3ds.com**.

Europe/Middle East/Africa
Dassault Systèmes
10, rue Marcel Dassault
CS 40501
78946 Vélizy-Villacoublay Cedex
France

Asia-Pacific
Dassault Systèmes K.K.
ThinkPark Tower
2-1-1 Osaki, Shinagawa-ku,
Tokyo 141-6020
Japan

Americas
Dassault Systèmes
175 Wyman Street
Waltham, Massachusetts
02451-1223
USA