

Development of high-performance gradient systems

by Doug Kelley, PhD, Global 7.0T MR Product Manager, GE Healthcare

MR imaging is unique among imaging techniques in that the fundamental spatial resolution limit is primarily controlled by one part of the system, the gradients, rather than arising from characteristics of how the object in question is illuminated. The gradient subsystem, consisting of a set of coils and amplifiers under computer control, has been a focus of technological improvement over the history of the development of MR, enabling new applications from basic science to clinical medicine at each stage of the development.

This article will discuss what MR users ask the gradient system to do, and how the technology continues to improve our ability to meet those needs. We will also discuss some of the key specifications for gradient coils, and how these specifications interact and affect an MR system's ability to deliver the images we want.

What does the gradient system do?

Before considering how a gradient coil works, we should think about what we want it to do as part of the imaging process. Any MR pulse sequence can be broken down into four basic activities — magnetization, nutation, precession and relaxation. In the magnetization phase, the spins (usually in water) become magnetized along the direction of the static field provided by the magnet (which





Figure 1. A first-generation shielded gradient coil.



	<i>SuperG</i>	<i>UltraG</i> [†]	<i>MAGNUS</i> ^{††}
Gradient amplitude, mT/m:	80	113	200
Slew rate, T/m/s:	200	260	500
Current (amp)/voltage:	1034/2324	1034/2324	620/1650
Relative coil gain, mT/m/A:	1.0	1.45	4.21
FOV, cm:	55 x 55 x 50	55 x 55 x 50	26 (DSV)

Figure 2. GE's advanced gradient coil portfolio in 2022.

happens on its own within a few seconds, about five times the T1 relaxation time) after the patient goes into the magnet bore. The gradients play no role in this phase. We call this “longitudinal magnetization,” and this magnetization doesn’t do anything until and unless we move into the other phases. Hydrogen nuclei in water have a single spin that can basically be aligned with the field from the main magnet or opposed to it. The static magnetic field creates an energy difference between those two states that defines what frequency of RF the spins will respond to (the magnetic resonance or Larmor frequency).

Nutation is where we apply radio frequency (RF) pulses to convert longitudinal magnetization into transverse magnetization, tipping some of the spins away from the z-axis. Gradients play a key role here in making some of the RF pulses spatially selective, causing them to affect only those spins in a particular location in the patient and leaving the other locations alone. We can accomplish this feat by spreading the energy of the RF pulses over a short range of frequencies, and using the gradients to change the static field so that only the spins in one region (usually

a slice or a slab) will feel the effect of the nutation pulse.

Once the nutation phase has created some transverse magnetization, this magnetization will precess, causing current to flow in the RF receiver coils. As in the nutation phase, the frequency of the precession will depend on the local static field, primarily from the magnet but changed slightly by the gradients we apply either to create some spatially dependent phase differences (following a short gradient pulse) or spatially dependent frequency (if we leave the gradients on while collecting the data). One of the most important early discoveries of MR was the recognition that all these processes can be used together to encode the details of the image.¹ By changing the phase of the signal across a particular region, the net magnetization across the region is reduced by interference, which can be undone by reversing the sense of the gradient (forming a gradient echo) or by applying another nutation pulse (producing a spin echo or a stimulated echo, depending on the amplitude of the pulse).

During the relaxation phase, two things need to happen: the existing transverse magnetization needs to be eliminated

so we can repeat the measurement with some different phase encodings, and the equilibrium magnetization needs to be at least partially restored (albeit within the limits of the desired T1 weighting). Gradients cannot help with the second part, but they can certainly help with the first part, introducing additional phase variation to ensure that signal from the transverse magnetization from one acquisition does not interfere subsequent acquisitions. A failure to meet this condition can produce some very interesting artifacts.

In summary, we want the gradients to produce a spatial variation in the magnetic field along the direction of the main field (usually called the z-direction) in a controlled fashion, a particular amplitude for a particular time. We want this modification to be the same everywhere in the field of view, and the same every time we perform the experiment.

Gradient strength and slew rate

The gradient strength is how much spatial variation we get across the field of view for a given input current, typically given in milliTesla per meter (mT/m). The slew rate is how fast we can change

[†] Cleared for clinical use in SIGNA™ 7.0T in US; investigational use in all other regions.

^{††} Technology in development that represents ongoing research and development efforts. These technologies are not products and may never become products. Not for sale. Not cleared or approved by the US FDA or any other global regulator for commercial availability.

this spatial variation, typically given in Tesla per meter per second (T/m/s). The image reconstruction will be simpler if the gradient is constant at its target amplitude (because that means the signal phase is a linear function of the position), so when we need different gradient amplitudes (say, when we want to apply a phase encoding along one axis and then readout along another), we want to switch between them as quickly as possible so we can get back to collecting signal.

The gradient strength and slew rate together are limited by two things: how much power we can deliver from the gradient amplifiers to the gradient coils and the coil's inside diameter. Based on dimensional analysis,² the product of the maximum gradient strength and maximum slew rate is proportional to the maximum power delivered to the coil divided by the fifth power of the coil diameter:

$$G_{\max} \times SR_{\max} \propto \frac{(I_{\max} \times V_{\max})}{d^5}$$

For a given coil design, we can get stronger and faster gradients by putting in more power, but if we really want more gradient performance, we will need to reduce the coil diameter (d). Wider gradient bores have a number of benefits and are essential for some applications, but they impose a steep price on gradient performance, usually requiring more powerful amplifiers.

Three of our recent gradient coils illustrate this design constraint quite well (Figure 2). The SuperG gradient coil, used in our SIGNA™ Premier system, delivers outstanding performance for a 70 cm patient bore, providing a maximum gradient strength of 80 mT/m and a maximum slew rate of 200 T/m/s. Using similar coil fabrication technology, the UltraG coil, available in our SIGNA™ 7.0T system with a 60 cm patient bore, delivers a maximum of 113 mT/m and a maximum slew rate of 260 T/m/s with the same gradient amplifiers. Our MAGNUS™ gradient coil, developed for neuroscience research (and not available for clinical use), has a 42 cm patient bore and delivers 200 mT/m at a slew rate of

500 T/m/s with our previous generation of gradient amplifiers. We anticipate a much higher performance with our current generation of gradient amplifiers. Coil geometry is a dominant factor in determining gradient coil performance.

Shielding

The magnetic field from a gradient coil extends not just inside the coil where the patient goes but outside the coil as well. When we try to change the current in the coil, the field that extends into the magnet will produce eddy currents that also affect the spins in the same way as the gradient fields we want, altering the timing of the gradient fields we want to apply. Early MR systems required extensive tuning and optimization to minimize the effects of these eddy currents, and some more elaborate techniques like echo planar imaging were, at best, impractical on most clinical MR systems as a result.

In 1986,³ GE introduced a shielded gradient coil where a second set of

gradient windings was added outside the primary windings to cancel the field at least partially outside the coil. The shield windings certainly have an effect inside the coil as well, slightly reducing the efficiency of the primary windings, but by reducing the field that extends into the magnet, the performance of the gradient coil becomes much more predictable and repeatable.

Modern design techniques allow the field to be modeled in three dimensions over the target field of view (FOV) and the surrounding space, and details of the magnet construction (location, size, and conductivity of magnet structures like the warm patient bore, the cold radiation shields and other metallic structures within the magnet cryostat) can be included in the optimization to minimize the energy deposited by the gradient coil into the magnet. These magnet-gradient interactions (MGI) can become quite significant in some cases, particularly as the performance of the gradient coils improves, and have been




Figure 3. 1.5 mm isotropic resolution track density images from $b=2000 \text{ s/mm}^2$ acquired with UltraG at 7.0T showing major fiber groups throughout the cerebrum. Decreased B_0 and B_1 homogeneity in the temporal lobes causes the loss of tracks in this region.




Figure 4. 1.5 mm isotropic resolution track density images from $b=4000 \text{ s/mm}^2$ acquired with MAGNUS at 3.0T showing major fiber groups throughout the cerebrum including the inferior temporal lobes.

anecdotally blamed for some recent magnet quenches. With careful design, though, MGI can be minimized to ensure safe operation of the gradient coil.

Linearity and peripheral nerve stimulation

While the gradient coil is designed to produce a particular field variation at the isocenter of the coil, as one looks further out radially or axially, the field will begin to vary from the ideal linear form at the center. How much the field varies from this linear ideal is to a large extent under the coil designer's control, but there are some constraints.

The biggest constraint is related to the fact that the gradient coil produces an electric field as well as a magnetic field, and the two cannot easily be separated. The electric field is highest where the magnetic field varies the most spatially.

For imaging purposes, we want a magnetic field whose z component varies linearly along one axis for each coil, so the field will be zero at the isocenter and increasing positively or negatively as one moves away from the center in the corresponding direction (ideally remaining constant if one moves in a different direction). At some point, the field will begin to fall off from this ideal linear variation, because the gradient coil is of finite length. Because the spatial variation of the magnetic field will be greatest at this point, there will be a high electric field at this location as well, time varying because the magnetic field is time varying as well.

This time varying electric field can cause peripheral nerve stimulation (PNS), which typically manifests as a twitch felt by the subject. If this sensation is painful or severe it represents a significant safety hazard, but at a low level it can still be unpleasant and unexpected. The likelihood of PNS is increased when the coil linearity holds over a larger maximum FOV, determined by the design of the coil and is reduced by smaller FOV or for head-only gradient coils where much of the field variation can be shifted away from the subject.⁴

Thermal performance

As discussed above, increasing the current to the gradient coil increases the gradient strength, but this approach has a limit. Since the gradient coil is made of a normal conductor (usually copper but occasionally aluminum if lower cost is required), the coil will heat up due to its resistance as the current flows. This heat is generally removed by active cooling with water and the heat must be conducted from the gradient conductors to the cooling water. GE's latest generation of gradient coils use hollow copper tubing for the conductors through which the cooling water circulates, allowing the heat to be extracted where it is generated.

An additional consideration is the rate at which heat can be extracted, which is generally slower than the rate at which it is generated. Other components of the system are often in thermal contact with the gradient coil and as the gradient coil heats up, so will these structures. Of particular concern are the passive magnetic shims, whose magnetization degrades with temperature, causing center frequency drift and shim degradation over time, and the RF transmitter coil, which can detune due to variation in the capacitance with temperature. Careful control of the temperature is essential to maintaining the stability of system performance.

Vibration, acoustic performance and stability

When current flows through a conductor across a magnetic field, the conductor feels a force (called the Lorentz force) perpendicular to the field and the direction of current flow. In a gradient coil, this force causes a small displacement of the conductors, much like the wires around the diaphragm of a loudspeaker. Quite a lot can be done to reduce the resulting noise,^{5,6} but not without some loss in system flexibility.

The vibration can have secondary effects as well. The conducted vibration can be somewhat unpleasant for the patient. In addition, the motion of the gradient coil can produce a slight variation in the

static magnetic field that can produce baseline distortions in ¹H spectroscopy.⁷ Suppressing the vibration is largely dependent upon coil construction; improving the stiffness of the epoxy used in the gradient coil can have a marked effect on the vibration pattern of the coil.

What does it all mean?

Not every imaging protocol needs the strongest and fastest gradient coil. Achieving that performance necessarily involves tradeoffs, including bore size, FOV and cost. Understanding those tradeoffs makes it easier to find the right combination for a particular system.

From GE's perspective, figuring out how to deliver that performance where it is needed has taught us how to improve the performance of our gradient coils across our entire portfolio. Even after all the years and effort that have gone into developing gradient technology, there is still a lot of opportunity to improve the technology and open new opportunities for basic science and clinical application. **S**

References

1. Edelstein WA, Hutchison JM, Johnson G, Redpath T. Spin warp NMR imaging and applications to human whole-body imaging. *Phys Med Biol.* 1980 Jul;25(4):751-6.
2. Lee SK, Bernstein MA. Systematic Dimensional Analysis of the Scaling Relationship for Gradient and Shim Coil Design Parameters. *Mag Res Med.* 2022;88(4):1901-1911.
3. Roemer PB, Edelstein WA, Hickey JS. Self-shielded Gradients. *Proc SMRM* 1986, 1067.
4. Lee SK, Mathieu JB, Graziani D, et al. Peripheral nerve stimulation characteristics of an asymmetric head-only gradient coil compatible with a high-channel-count receiver array. *Magn Reson Med.* 2016 Dec;76(6):1939-1950.
5. Edelstein WA, Hedeem RA, Mallozzi RP, El-Hamamsy SA, Ackermann RA, Havens TJ. Making MRI quieter. *Magn Reson Imaging.* 2002 Feb;20(2):155-63.
6. Ljungberg E, Damestani NL, Wood TC, et al. Silent zero TE MR neuroimaging: Current state-of-the-art and future directions. *Prog Nucl Magn Reson Spectrosc.* 2021 Apr;123:73-93.
7. Hurd RE, Gurr D, Sailasuta N. Proton spectroscopy without water suppression: the oversampled J-resolved experiment. *Magn Reson Med.* 1998 Sep;40(3):343-7.