Technologies for Precision Magnetic Field Mapping

Philip Keller
Metrolab Instruments, Geneva Switzerland
www.metrolab.com

Table of Contents

Abstract

Mapping a magnetic field, for example to verify a computed field model, is a common problem, yet it defies a common solution.

We will examine three technological solutions, with specific application examples:
- Three-axis Hall sensors for mapping undulators in a synchrotron light source;
- Fluxmeters with rotating coils to characterize accelerator dipole magnets; and
- NMR field mappers to shim superconducting magnets used in MRI.

For each measurement technology, we describe:
- its principles, advantages and disadvantages;
- how one typically constructs a field mapper; and
- the specific implementation for the application presented.

Comparing best practices from different application areas is intended to promote dissemination and cross-fertilization of known magnetic field mapping techniques.

1   Technology Overview

For more than 170 years, some of the world's best scientists have worked on the problem of accurately measuring a magnetic field. This has resulted in a panoply of measurement techniques, very nicely summarized in Figure 1.
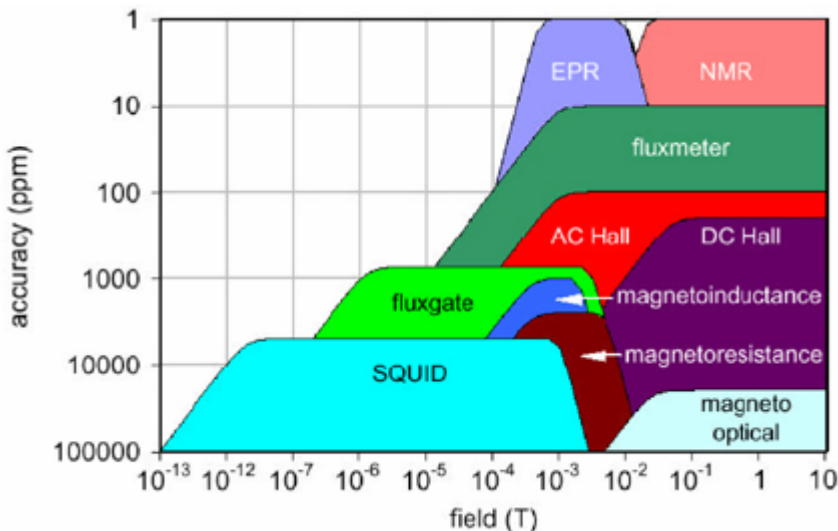
Figure 1. Overview of magnetic measurement technologies.
**From a presentation by Luca Bottura of the CERN.**[i]

This figure classifies the different techniques by their accuracy and the field range they cover. We can learn much from this diagram:

- The biotope of magnetic measurement techniques is indeed rich and diverse. It ranges from classical to high-tech, from well-known to abstruse, and from widespread to highly specialized.

- We can measure an impressive range of fields: witness the horizontal axis that spans fourteen orders of magnitude. That's like going from picoseconds to minutes. No wonder we need several different techniques; not many clocks will cover that sort of dynamic range either.

- Under the right circumstances, we can measure magnetic field strength with astounding accuracy. In fact, if we replace the term "accuracy" by "resolution," we would have to add about 50% to the vertical axis; NMR measurements are getting close to ppb (parts per billion) resolution. That's like measuring the distance from here to the moon with a resolution of under a meter – and that with a commercially available instrument.

- However, these resolutions are only achievable at high field strengths. In particular, for anything below the strength of the earth's magnetic field (about .5 mT), one is hard pressed to achieve better than .1% accuracy.

In this article, we will focus on three technologies, suitable for field strengths above about 10 mT: **Hall**, **fluxmeters** and **NMR**. In addition, rather than measuring the magnetic field strength at a single point, our goal is to generate a **magnetic field map**, showing how the magnetic field strength varies throughout a volume.

2 Mapping with Hall Devices

2.1 Hall Magnetometer Technology
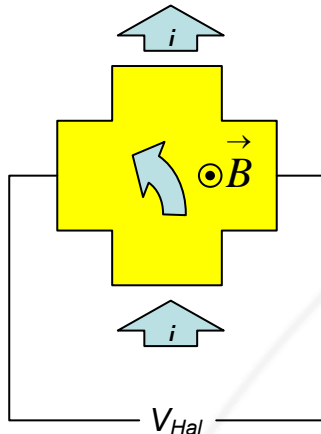
Most of us are familiar with the functioning of a Hall device. The classical Hall device is a thin plate of semiconductor material with four terminals. As shown in Figure 2, a current is made to flow through two opposing terminals. A magnetic field perpendicular to the plate make the conductors want to drift towards one side, thereby generating a potential difference, measurable as a voltage – the Hall voltage – across the other two terminals.

Figure 2. Functioning of a Hall device.

Hall devices present important benefits that have made them the most commonly used measurement technology in the range above 10 mT:

- Easy to build: Hall devices have benefited from the overall evolution of the semiconductor industry. We now have devices that are compact, low cost, low noise and highly reproducible.

- Easy to use: Hall devices are readily integrated into electronic systems.

- Fast: For most practical purposes, measurements can be considered instantaneous.

- Variable sensitivity: By varying the current injected, we can adapt the sensitivity of the Hall device to the strength of the magnetic field to be measured.

- Single vector component: The Hall device is only sensitive to the component of the magnetic field vector that is perpendicular to the plate (well, almost – see below).

However, we must keep in mind that Hall devices also have limitations:

- Offset: Because of geometric and material imperfections, the Hall voltage will almost always be non-zero in a zero field. In the latest generation of devices, spinning-current techniques, whereby the current is injected into each terminal in turn, are used to cancel most of the offset; but the problem persists at a lower level.

- Nonlinearity: The Hall voltage has only an approximately linear dependence on the applied field strength.

- Temperature sensitivity: The number of charge carriers in a semiconductor depends on the temperature, and the sensitivity of a Hall device is therefore strongly temperature dependent.

- Drift: As the semiconductor material ages, the offset, nonlinearity and temperature sensitivity all change.

- Angular error: It is practically impossible to mount a Hall device exactly level, especially since the Hall plate itself is generally mounted inside a protective package and is therefore hidden from view.

- Cross-talk between axes: A magnetic field in the plane of the sensor will also generate a small Hall voltage – this is the Planar Hall Effect. In fact, Hall elements often display even higher-order cross-talk terms between the axes, and for high-precision measurements, these effects become significant and must be corrected for.

- Inductive pickup: The voltage-sensing leads of a Hall device form a small pickup coil, so moving a Hall element in a magnetic field, or measuring an AC field, generates a voltage independent of the Hall voltage. Depending on the strength of the field, the rate of motion and the precision of the measurement, this error can be significant.

If they become noticeable (and they usually do), most of these problems can only be solved by periodically recalibrating the device in a reference field, usually a magnet controlled by NMR. A three-axis device can be used to address angular errors and axis cross-talk.

2.2 Hall Mappers

To build a Hall mapper, we need:

- Three-axis Hall sensor or sensor array: Generally one uses a commercially available three-axis Hall magnetometer, allowing all three components of the magnetic field to be measured simultaneously. If the direction of the field is precisely known, it is possible to use a single- or dual-axis device – but see the caveats in the previous section concerning angular error. Using an array of sensors will speed up the mapping process.

- Positioning jig: This is almost always a custom-designed mechanical arm that allows a systematic scan of the volume of interest. Alternatively, the jig can move the magnetic system itself past a fixed sensor. In larger jigs, it can be quite challenging to minimize unwanted motion like bounce, sag and backlash. If the jig is motorized, we have to be careful to avoid interference between the motors and the magnetic system being measured; for very strong magnetic fields, we can choose between a long drive shaft, pneumatic- or piezoelectric motors. Whether motorized or not, it is usual to measure the position of the sensor with an independent set of linear or angular encoders.

- Temperature control: There are three basic approaches to dealing with the temperature sensitivity of Hall sensors: air-condition the room; use Hall probes encased in a temperature-controlled oven; or measure the sensor temperature and correct the measurements using calibration data.

- Calibration: Calibration may be in-house or performed by the manufacturer. Unless we can place the whole mapper in the reference magnet, we must be able to remove the sensor to perform a periodic calibration. The mounting system must allow the sensor be installed reproducibly in the same position and orientation. Ideally, we use the same measurement electronics (current source, signal amplifiers, etc.) for calibration as for measurement.

The output of a Hall mapping system is straightforward:
- three components of the magnetic field vector,
- measured with low (typically 1%) to medium (typically .01%) precision, and
- sampled at points throughout the volume of interest.

Often this is not what we are really interested in; for example, our real goal might be to compute a line integral to determine the deflection of an electron beam, or the gradient to compute the force exerted on a piece of iron. In this case, it is important to keep in mind the additional loss of precision associated with these computations.

2.3 Hall Mappers – an Example

In recent years, Hall mappers have found widespread use for characterizing insertion devices for synchrotron light sources. Throughout the world, dozens of such facilities are now in operation, under construction or in the planning stages. Our example comes from the Laboratório Nacional de Luz Síncrotron (LNLS – National Laboratory of Synchrotron Light) in Brazil, one of the very first such facilities in the developing world.

Figure 3. LNLS: storage beam, booster and 12 experimental lines.

Synchrotron light is light emitted when a high-energy electron beam is forced to bend in a magnetic field. Because the electrons are travelling at relativistic speeds, the classical dipole radiation pattern collapses into a tightly collimated beam. As a result, synchrotron radiation is by far the most intense light source available to us. Its frequency can be tuned over a wide range, from the IR to X-rays, and it can be elliptically or linearly polarized. Because the light is produced by batches of electrons, synchrotron light sources generate sub-nanosecond pulses. These properties make synchrotron light interesting as a probe in a variety of disciplines. For example, the LNLS has research programs on materials science, structural molecular biology, nanostructures and micromachining.

The current generation of light sources use special magnets called undulators or wigglers to cause the electron beam to radiate. As shown in Figure 4, magnets with alternating polarity force the beam to assume a serpentine path, radiating at each turn. The undulator is designed so as to have no net effect on the electron beam in the storage ring; ideally, it makes no difference to the beam whether the undulator is inserted or not – hence the name "insertion device."
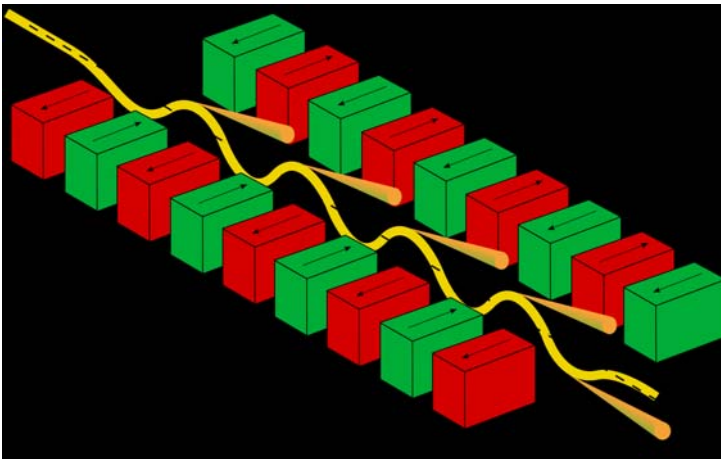
Figure 4. Principles of operation of an insertion device (from the Wikipedia).

The magnetic field of an insertion device is usually mapped with two technologies, to answer two different questions. One goal is to verify that the device will in fact not disturb the beam. The key to answering this question is the field integrated along the entire length of the device, and is best measured with a rotating coil – this we will discuss in the next section. The other goal is to verify the uniformity of the alternating field that the beam will be subjected to. To achieve the high spatial resolution required this is done with a Hall mapper, as shown in Figure 5.[ii]
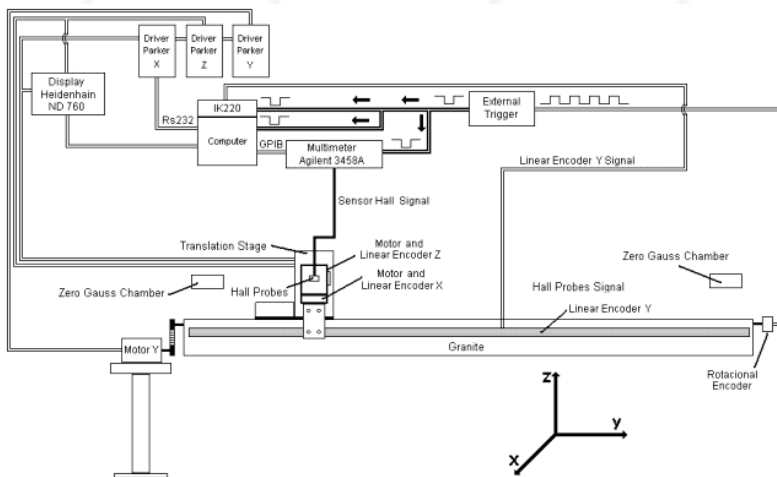

Figure 5. Diagram and photograph of LNLS Hall mapper.

LNLS decided to use three Hall probes and a multimeter rather than a commercial three-axis Hall magnetometer. The jig is mounted on granite to minimize vibrations. Two motors on the carriage allow selecting a position in X and Z; the scan is then performed in Y. The range in Y is 4.2 m, which, at the minimum step size of 0.28 mm, can yield 15,200 points per scan. The maximum scan rate is 75 mm/s, meaning that a scan takes nearly a minute.

The linear encoders used to measure X, Y and Z have an accuracy of ± 5 μm and a resolution of 0.1 μm. This may seem like overkill, but in high-gradient fields, the overall precision is dominated by the positional rather than field measurement precision – see below. An interesting feature of this mapper are the zero-Gauss chambers on each end, permitting a dynamic offset drift correction along the length of the scan.

The LNLS team performed an in-depth analysis of all imaginable error sources, including electrical, transport, temperature, offset, zero Gauss, probe angle & position, eddy current

and calibration. Thanks to this detailed characterization, they were able to achieve excellent precision, with measurement deviations of $\pm 2 \times 10^{-5}$ T in the 2 T "constant field" regions. In the transition regions, where the gradients attain 66 T/m, the standard deviation increased to $\pm 7 \times 10^{-4}$ T, which was attributed to vibration.

3 Mapping with Fluxmeters

3.1 Fluxmeter Technology

Fluxmeters are a direct application of Faraday's Law of Induction. They were the instrument of choice in the nineteenth and early twentieth century, and are still used, for example, to measure the earth's magnetic field or to measure hysteresis curves. Yes, after all this time, they are still relevant, in particular for mapping applications. In fact, they often give results that are clearly superior to those of Hall mappers.

The diagram below illustrates the basic principle:
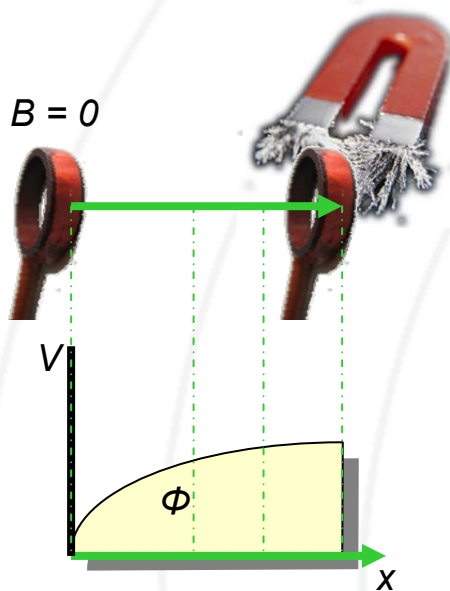
- The sensor is a coil.

- Suppose that the coil starts in a region of zero field, and is then moved towards our magnet. We record the induced voltage during this motion.

- The integral of the voltage is equal to the change in magnetic flux, $\Delta\Phi$. However, since we started in a region of zero field, $\Phi_0 = 0$ and $\Delta\Phi = \Phi$, the flux at the terminal point.

- Dividing the flux by the area of the coil gives the magnetic flux density, B.

- We can compute B anywhere along the path, without any additional measurements, simply by using the partial integral up to that point.

Figure 6. Using a fluxmeter as a mapper.

Using fluxmeters for mapping offer some outstanding benefits:

- Very flexible. People continue to find variations on the basic approach outlined above. Note that $\Delta\Phi$ should be completely independent of the path taken – this gives us quite some latitude when designing our jig! Alternatively, instead of moving the coil, we could simply rotate it by 360°; this will give us the local field strength as well as its direction in two dimensions. Or consider a "coil" consisting of a single wire with which we sweep out the aperture of a magnet; in this case, it is the area of the coil that changes rather than its position or orientation. In AC fields, we have an induced voltage without changing anything at all – how convenient. As a final example: how would you measure the flux density inside a solid chunk of iron? With a movable coil wrapped around the iron, of course.

- Precise: We have learned a thing or two since the nineteenth century, and modern electronics gives us a big leg up on Faraday. Fluxmeter mapping systems regularly achieve a precision of .01%, and in the hands of an expert, even a factor of ten better.

- Mapping "built in": As described above, a single set of measurements can yield a map with arbitrary spatial resolution, as long as we can get all the partial integrals.

- Can measure integrated field or gradient directly: As mentioned at the end of Section 2.2, we may be interested in a line integral to determine the deflection of an electron beam, or a gradient to compute the force exerted on a piece of iron. By designing a custom coil, fluxmeters allow us to measure these quantities directly, thereby greatly improving both measurement speed and precision.

A fluxmeter does place high demands on its user:

- Requires thought: A modern Hall instrument can be used "out of the box" and is accessible to any technician. This is not the case for fluxmeters.

- Requires a custom coil: Good coils can be very high-tech. They should slip into tight places, be absolutely stable both mechanically and electrically, and not tangle up their own cables. These conflicting goals have led magnetic cartographers to the very edge of what can be achieved with wires, exotic materials and machining.

- Requires good technique: Just to name a few:
    - Drift compensation: Any voltage offset, however small, will be integrated forever, and will eventually show up in the measurement. It is therefore essential to have a drift compensation strategy. For example, in the little experiment described at the beginning of this section, we should return the coil to a region of zero field; the residual $\Delta\Phi$ can then be used to compensate for the drift on all the preceding measurements.
    - Noise: If you want to achieve 10 ppm precision with a voltage input limited to ±5 V, your allowable noise floor is well under 50 µV. Good luck.
    - Coil motion: A little bit of coil vibration can turn into a big noise signal. See above.
    - Thermal expansion: It is a well-known fact that magnets are always stronger in the afternoon. Well, not really; but with a fluxmeter they appear to become stronger as the lab heats up and the area of the coil expands.
    - Thermocouple effects: At this level of precision, any connector acts as a thermocouple. As long as this voltage offset is constant, it should be corrected by the general drift compensation. However, all bets are off if it varies during the measurement because of temperature changes.

3.2 Fluxmeter Mappers

Figure 7 illustrates the elements required in a fluxmeter mapping system:
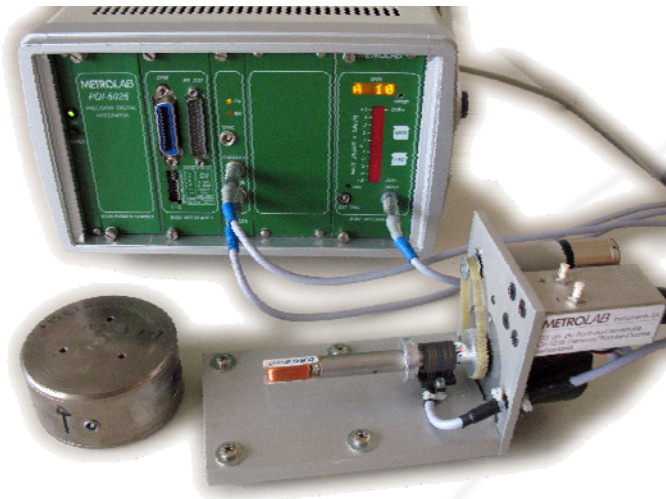
- A coil

- Positioning jig: rotational in this case.

- Integrator: Includes motor control module.

- Temperature control: None - relies on room temperature control.

- Calibration: With NMR calibration magnet shown (bottom right).

Figure 7. Fluxmeter mapping system.

Even the small jig shown in Figure 7 is not nearly as simple as it looks:

- A high-precision, geared DC motor, visible just behind the buttons, and a toothed belt assure a positive, slip-free drive.

- The belt allows considerable backlash; however, this is not important since a precision angle encoder (black cylinder mounted below the buttons) is coupled directly to the shaft and triggers the acquisitions by the integrator. A reference output on the encoder provides a reproducible starting angle.

- The buttons allow the shaft to be turned manually forwards or backwards. This is useful to verify the signal amplitude (and thus the required amplification) before launching a measurement. In normal measurement mode, the motor is driven under computer control via the motor controller built into the integrator, thus minimizing the system complexity.

- On the shaft just to the left of the toothed drive wheel, one can see a set of slip rings, used to transmit the signal from the turning coil. Without these, cable twist would limit the rotation to a most a few turns. In this case, two slip rings per signal are used to minimize the impedance and noise.

- Just to the left of the slip rings is a spring-loaded break in the shaft, allowing the shaft to be "wiggled" into small apertures. Precision-machined mating surfaces are required to ensure that the composite shaft is perfectly straight, and a pin & notch are used to reliably transmit the torque.

The jig is almost always custom-designed for the application, and, as this little example shows, requires a substantial design and construction effort.

The integrator, on the other hand, can be bought off-the-shelf. There are three basic designs:

- The brute-force approach uses an ADC to sample the coil voltage, and sums the output values to compute the integral. This approach suffers from sampling error as well as

cumulated numerical errors, but is becoming more viable as ADC sample rates and output widths improve.

- The classical approach uses an analogue integrator circuit and an ADC on the back end. This approach is limited by the frequency response, dynamic range and noise of the analogue integrator.

- The third approach feeds the input voltage into a Voltage-to-Frequency Converter (VFC); the integration is simply a counter on the VCF output. This clever hybrid of analogue and digital design was conceived over twenty years ago at the CERN, and has long since become the standard of performance for integrators.

The output from a fluxmeter mapper is more complex than that of a Hall mapper:

- The flux Φ at each sample position. Sometimes the flux is exactly what is desired; often we want flux density and need to divide by the (calibrated) area of the coil; in other cases we measured a gradient (or higher-order term), and need to apply the appropriate normalization.

- Measurements are medium to high precision, roughly between 1000 and 10 ppm.

- Unless we use a coil specifically designed to approximate a point sample (fluxball), coils are generally planar and our data is "slice-sampled" rather than point-sampled.

Subsequent computational processing depends entirely on the application. In high-energy physics, it is common to expand the angular map of an accelerator magnet in a Taylor series that directly represents the dipole, quadrupole, hexapole, octupole, … components of the field.

3.3 Fluxmeter Mappers – an Example

Fluxmeter mappers have traditionally been used in the high-energy physics community, to map accelerator magnets and thereby to be able to predict the beam deflection. The most impressive recent project of this type is undoubtedly the Large Hadron Collider (LHC), currently the most powerful accelerator in the world, due to go online in November 2007 near Geneva at the European Organization for Nuclear Research (CERN).

The LHC is a proton or heavy-ion collider ring, 27 km in circumference. The bulk of this distance is made up of 1232 superconducting dipole magnets, each 15 m long. These dipoles produce 8.3 T, with opposing fields in two adjacent apertures – one aperture for each of the two beams circulating in opposite directions. The LHC dipole magnet and the facility where they were tested is shown in Figure 8.
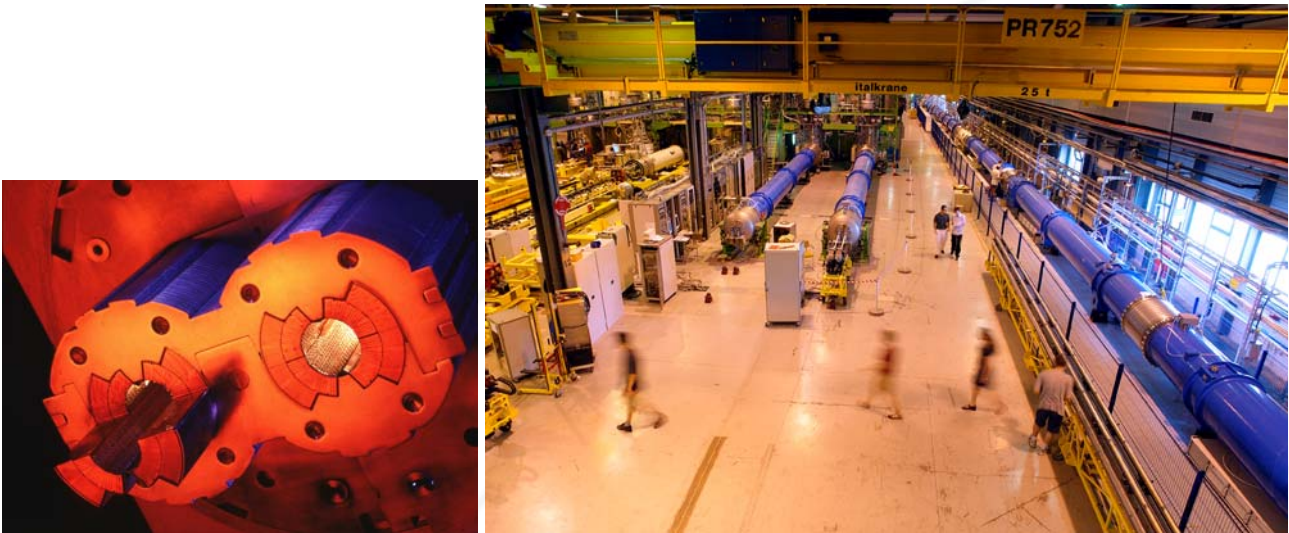
Figure 8. The LHC dipole magnets and the facility where they were tested.

Every one of the 1232 LHC dipoles was measured with a rotating-coil system of amazing sophistication.[iii] One of these systems can be seen in the centre of Figure 8, and in greater detail in Figure 9. Twin rotating shafts, 16 m long and 36 mm OD, map both apertures at the same time. In order to accommodate the ever-so-slight bend in the magnet, each shaft is composed of 13 segments, ceramic tubes coupled with titanium bellows and resting on ceramic roller bearings.

Two tangential coils are mounted on opposite sides of the tube. In fact, only one is needed, but the other one serves as a spare and to equilibrate the shaft. A coil in the centre – a "bucking" coil – is used to subtract the dominant dipole signal, allowing the full dynamic range of the tangential coils to be devoted to the low-level angular variations.

Given the number of dipoles to be tested, performance was at a premium. A complete map could be generated in typically ten seconds: three turns forward (one to accelerate, one to measure, and one to decelerate) and three turns backwards (to repeat the measurement and unwind the cables) – and that was that. The reproducibility of the dipole measurement was estimated at 4×10^{-5} T – not bad for a "nineteenth-century" method!
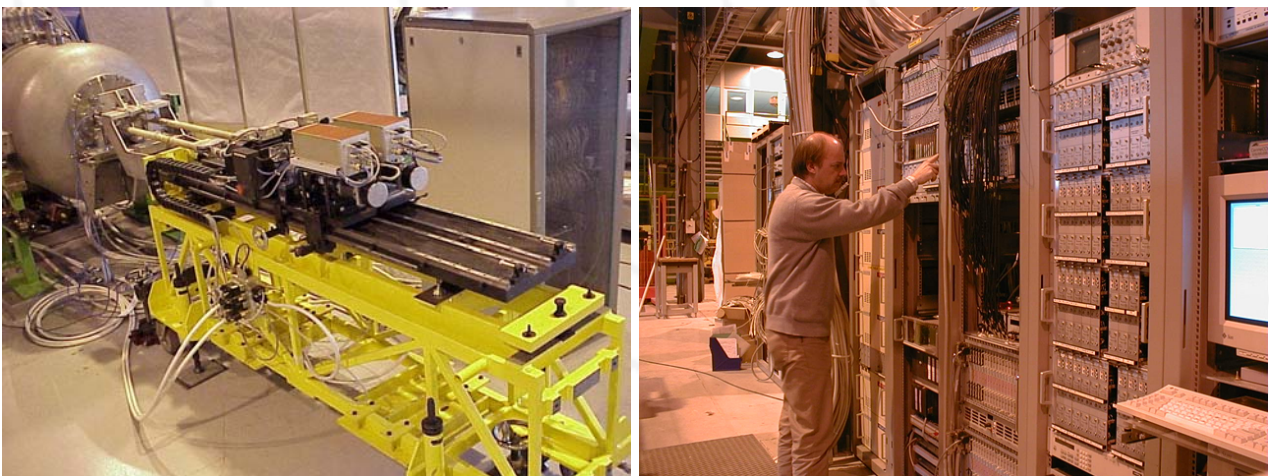

Figure 9. LHC dipole mapping system.

4 Mapping with NMR Teslameters

4.1 NMR Teslameter Technology

With a resolution approaching 1 ppb (part per billion), Nuclear Magnetic Resonance (NMR) is the gold standard of magnetic measurement. Figure 10 illustrates the technical principle:
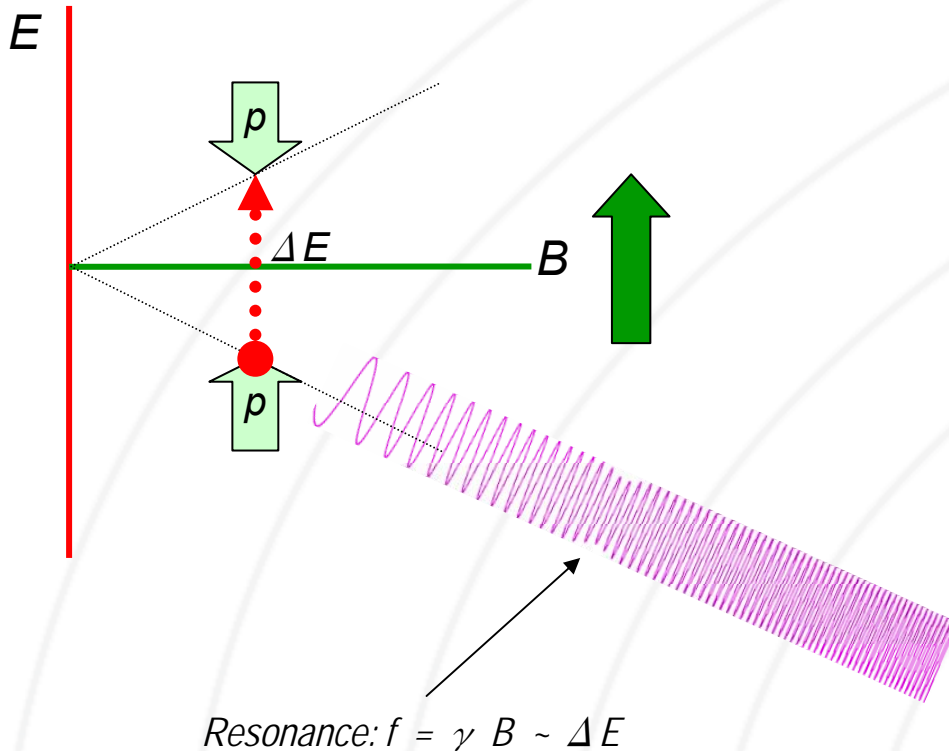

$$\text{Resonance:} \quad f = \gamma\ B \sim \Delta E$$

Figure 10. Principles of NMR magnetic field strength measurement.

A proton (hydrogen nucleus) has two possible spin states. In the absence of an external magnetic field, these states are degenerate, but as soon as we apply a field, the spin-aligned state is lower energy than the unaligned state. The energy differential ΔE between the two states depends linearly on the magnetic field strength. This exact linear dependence, very nearly independent of all other external factors such as temperature, forms the basis for NMR teslameters.

In an NMR teslameter, a coil is wrapped around a sample material rich in hydrogen nuclei – ordinary water works great. This is the "Nuclear" part of NMR. The sample is placed in the magnetic field to be measured ("Magnetic") and irradiated with an RF signal whose frequency gradually increases. When we hit the frequency that corresponds to ΔE, the sample suddenly absorbs energy ("Resonance"). For a field strength of 1 T, this resonance occurs at about 42.5 MHz; this is the gyromagnetic ratio (γ) for hydrogen nuclei. The resonance is extremely narrow, and so the frequency provides an extremely precise measure of the magnetic field strength.

NMR has some overwhelming advantages over other measurement technologies:
-    extremely precise;
-    no drift; and
-    measures total field.

That last point is perhaps not obvious. In fact, the protons align themselves with the magnetic field, and the resonance always occurs at *B/γ*, regardless of how we orient our coil. It turns out, however, that the best coupling between the RF field and the proton spins is achieved when the magnetic field is perpendicular to the coil's axis. In fact, NMR has a blind spot when the coil's axis is exactly aligned with the magnetic field.

NMR also has substantial limitations:

- Uniform field only: In a field gradient, one edge of the sample will resonate at one frequency, and the opposite edge at another. As the gradient increases, the resonance signal becomes broader and shallower, until it can no longer be detected; this limit typically lies in the range of several hundred to several thousand ppm/cm. In some situations, gradient correction coils can improve this situation.

- DC / low-AC fields only: NMR is a relatively slow measurement method, making it unsuitable for measuring rapidly varying fields. Here, we have to distinguish search mode from measurement mode: scanning through a large frequency range to find the resonance requires the field to remain stable for seconds; but once the search is completed, we can track the field and achieve measurement rates of up to about 100 measurements per second.

- Low fields require a large sample, ESR or pre-polarization: As the field and, consequently, the energy gap between the spin states tends towards zero, the populations of the two states tend to become equal, and the strength of the resonance diminishes. There are three possible remedies:
  - o Use a very large sample to increase the number of spin-flips.
  - o Pre-polarize the sample in a strong field to populate the spin-aligned state, and then remove this bias field during the measurement.
  - o Use ESR, i.e. the spin of orbital electrons instead of protons, because these have a gyromagnetic ratio (*γ*) three orders of magnitude higher than protons.

4.2 NMR Mappers

To build an NMR mapper, we need:
- NMR probe or probe array. As we will see in the example, a probe array dramatically speeds up the acquisition and simplifies the jig.
- Positioning jig.
- NMR teslameter (the electronics).

The output of an NMR mapper is straightforward:
- total field B,
- point-sampled,
  In fact, the size of the "points" are determined by the size of our NMR sample, which is typically a few millimetres in diameter.
- very high precision.
  Well below 1 ppm, in the best case approaching 1 ppb.

Computational post-processing depends very much on the application. The typical post-processing for MRI applications is described in the following section.

Figure 11. Close-up of an NMR probe array. Each probe – one of them is circled – is mounted on its own circuit card. Such probe arrays typically contains 24 or 32 probes (maximum 96).

4.3 NMR Mappers – an Example

The most widespread use of NMR mapping relates to the development, manufacturing and installation of MRI (Magnetic Resonance Imaging) magnets. These can be permanent, resistive, or superconducting magnets, with a typical field strength between 0.2 and 7 T. Technically, the preferred shape is a torroid, but for improved access to the patient, or for a less claustrophobic feel, C-shaped magnets are sometimes used. Most modern full-body systems from the major manufacturers use superconducting torroids, 1.5 or 3 T, as illustrated in Figure 12.

Figure 12. Example of modern full-body MRI scanner (Philips Achieva).

The homogeneity of the field in the imaging region must typically be guaranteed to within a few ppm. Since minor, unavoidable manufacturing variations cause inhomogeneities

generally two orders of magnitude greater than that, the homogeneity must be adjusted in a process called shimming.

The overall shimming procedure is to map the field, analyze the map, adjust the field with iron shims and/or shim-coil current adjustments, and to remap the field to verify the results. If necessary, this process is repeated until the specified homogeneity is achieved. The same process, and usually the same equipment, is used in R&D, manufacturing and installation. The measurement is extremely sensitive, and obscure effects such as the passing of an overhead crane or the return current of a nearby train can cause huge errors.

Figure 13. NMR mapping jigs for torroid- and C-shaped MRI magnets, respectively.

The probe arrays and jigs shown in Figure 11 and Figure 13 generate a map of the total field on the surface of a sphere. As long as there are no currents or magnetic material inside this sphere, and assuming the direction of the field is known, the field at every point inside the sphere can be calculated from Maxwell's equations. Generally we use an expansion in spherical harmonics to perform this computation.

It is worth noting that the jigs shown are much simpler that in the previous examples. Fundamentally, there are two reasons for this:
- Using a probe array reduces the required motion to a single dimension.
- Partly because the field is so uniform, and partly because the NMR sample points are fairly large, the requirements on the positioning precision are much less – typically 0.1 mm. Compare that to the 0.1 µm for the Hall mapper example!

5   Questions to Ask

To summarize our observations from these different examples of magnetic mapping, here is a small check-list of key questions to ask if you want (rather, need) to build a mapper:

- Use: R&D, production, field service?
  It makes a big difference if a system is to be used once in a while for R&D, or all the time in production, or packed up in a suitcase for field service.

- Measurement: field components, total field, integral, gradient?
  Do you really need Bx, By, Bz, or do you really need B, or an integral or gradient of **B**?
  By choosing the appropriate measurement technology, you may be able to speed up
  the acquisition as well improve your precision.

- Field: strength, uniformity, AC/DC, stability?
  The characteristics of your field also play an important role in the choice of
  measurement technology.

- Precision: 10% or 10 ppm?
  … as does the precision required.

- Positioning: access, range, precision, reproducibility?
  The design of your jig is as important as the choice of magnetic sensor. What sort of
  access do you have to the region to be measured? What positioning precision do you
  need? How reproducible must the positioning be?

- Environment: vacuum, cryogenic?
  The environment changes dramatically depending on whether you're in a lab, on a
  production floor or in the field. In addition, magnetic systems are often combined with
  vacuums and cryogenic temperatures. Hall effect sensors, for example, don't work at
  liquid-helium temperatures.

- Speed: cost, external error sources, human error?
  Making a fast mapper may not cost more; for example, we saw that the simplification
  of the jig may more than make up for the costs of a sensor array. In addition, the
  improved speed dramatically reduces external and human error.

[i] L. Bottura, "Field Measurement Methods." Presentation at CERN Accelerator School (CAS) on
Superconductivity, May 8-17, 2002. http://m.home.cern.ch/m/mtauser/www/archives/events/erice/MM.pdf.
[ii] G. Tosin, J.F. Citadini, R. Basílio, M. Potye, "Development of Insertion Device Magnetic
Characterization Systems at LNLS." Proceedings of the European Particle Accelerator Conference (EPAC)
2006. http://accelconf.web.cern.ch/AccelConf/e06/PAPERS/THPCH134.PDF.
[iii] J. Billan, L. Bottura, M. Buzio, G. D'Angelo, G. Deferne, O. Dunkel, P. Legrand, A. Rijllart,
A. Siemko, P. Sievers, S. Schloss, L. Walckiers, "Twin Rotating Coils for Cold Magnetic Measurements of 15
m Long LHC Dipoles." Proceedings of 16th International Conference on Magnetic Technology (MT-16).
http://doc.cern.ch/archive/electronic/cern/preprints/lhc/lhc-project-report-361.pdf.