High-precision measurements

G. Fernqvist
CERN, Geneva, Switzerland

Abstract
Power converters in the accelerator context are often required to deliver a performance that is unusual in an industrial context. This paper deals with some of the important aspects of high-precision output voltage and current control and in particular the current measurement transducers and their application.

1 Introduction

Particle accelerators often require a relatively high static and dynamic DC performance from the power converters that are employed to feed the magnets and other systems. An old saying is very applicable here: “If you can’t measure it, you can’t control it”. Today all accelerators are computer controlled. The main components of the measurement chain determining this performance are the output sensing transducers, the Digital–Analog Converters (DACs), the Analog–Digital Converters (ADCs) and their respective deployment, i.e., how the environment affects their performance. This paper deals with the transducers and their applications.

2 Precision

Precision must be quantified to be useful. The International Standards Organization (ISO) has defined terms and methods to express measurement uncertainty in a standardized way [1]. Suffice it to say here that ‘precision’ is a qualitative term, whilst ‘accuracy’ and ‘uncertainty’ are quantitative terms that can be used in specifications. The International Electrotechnical Commission (IEC) has also laid down terms and methods for specifying and evaluating performance of many different types of electrical equipment. It is indeed very useful to study the relevant standards before starting any project studies.

The aim is to quantify device performance or imperfections. This must not be confused with measurement errors or measurement uncertainty, but knowledge of the latter is necessary to determine the device performance. To give an example: if the intention is to measure the accuracy and stability of the output current from a power converter, the accuracy and stability of the external measuring device must first be determined and then taken into consideration when presenting the result of the measurement.

3 User requirements

As with any project it is necessary to determine as precisely as possible what is required of the power converter. In the accelerator context this information is normally provided by the users, e.g., beam optics/magnet specialists.

- Voltage or current output? Pulsed or DC operation?
- Bi-polar or uni-polar output i.e. 1-, 2-, or 4-quadrant operation?
- Type of load? (Inductive? Resistive? etc.).
- Performance parameters, static and dynamic.
- Reliability expectations (MTBF, MTTR, etc.).
- Environmental parameters (ambient temperature, humidity, mains, cooling water temperature, etc.)

Keep this parameter list updated during the project and use it as a discussion base to renegotiate external constraints and parameters that lead to high costs.

4 Equipment specifications

It is now time to translate the user requirements into equipment specifications. Here we concentrate on the high-precision aspects and take the DC precision as an example.

Identify the components that will contribute errors and uncertainties from a block diagram. In analog regulation loops, Fig. 1, the DAC, the DC Current Transformer (DCCT) and the error amplifier in the regulation electronics will contribute. CERN has for a long time employed two DCCTs in the output to facilitate detection of faulty subsystems. It is difficult in most accelerators to detect with other means that a circuit current is outside specification.

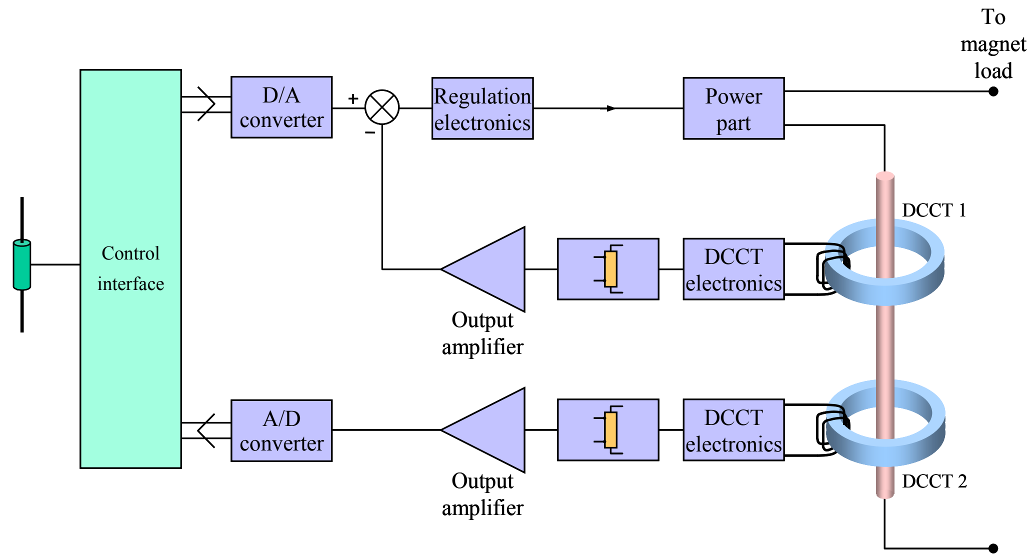

Fig. 1: Traditional analog regulation loop

With a digital regulation loop, as in the LHC example, Fig. 2, the ADC and the DCCT will contribute. The DAC will not, being in the forward path inside the loop.

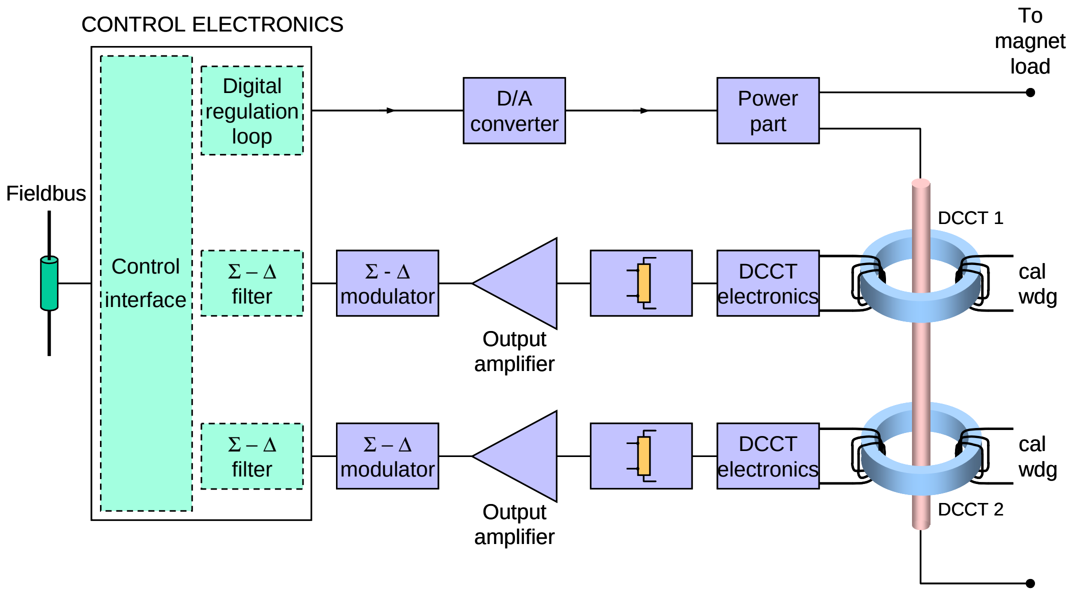

Fig. 2: Digital regulation loop for LHC

Establish an accuracy budget from preliminary specifications (Spec), example in Table 1, which will later be extended to include real obtained performance (Real) once some hardware has been built and evaluated. This will give a good overview to decide where efforts towards improvement can be profitable. Obviously the most significant contributions are attacked first.

Table 1: Accuracy budget from each device to the whole LHC machine

	Device performance				LHC machine impact					
	ppm of FS		ppm of value		1/2 hr Stability		Reproducibility 1 day		Accuracy 1 year	
Device	Spec	Real	Spec	Real	Spec	Real	Spec	Real	Spec	Real
DCCT										
Zero uncertainty (hyst etc.)	50	3			0	0	0	0	50	3
Settling after change			0	30						
Repeatability	3	3			0	0	3	3	3	3
Uncomp non-linearity	50	50			0	0	0	0	50	50
LF noise, 0.1-10 Hz	0	3	0		0	3	0	3	0	3
Stability 1/2 hr, 1-100 mHz	10	15	0		10	15	10	15	10	15
Gain drift 24 hr			10	10	0	0	10	10	0	0
Gain drift 1 year			100	100	0	0	0	0	100	100
Gain Temp Coeff			5	10	0	0	25	50	50	100
Offset drift 24 hr	10	10			0	0	10	10	0	0
Offset drift 1 year	40	40			0	0	0	0	40	40
Offset Temp Coeff	3	2			0	0	15	10	30	20
DCCT total					10	18	73	101	333	334
A/D converter, 16 bit succ. approx.										
Uncomp non-linearity	60	240			0	0	0	0	60	240
LF noise, 0.1-10 Hz	60	60	0		60	60	60	60	60	60
Stability 1/2 hr, 1-100 mHz	0		0		0	0	0	0	0	0
Gain drift 24 hr			30	30	0	0	30	30	0	0
Gain drift 1 year			100	100	0	0	0	0	100	100
Gain Temp Coeff			3	3	0	0	15	15	30	30
Offset drift 24 hr	10	10			0	0	10	10	0	0
Offset drift 1 year	50	50			0	0	0	0	50	50
Offset Temp Coeff	0.6	1			0	0	3	5	6	10
A/D total					60	60	118	120	306	490
Miscellaneous					5	5	10	10	100	100
Total					75	83	201	231	739	924
LHC committment					50	50	100	100	1000	1000
Conditions										
Temp change (K)					0	0	5	5	10	10
No special temp ctrl										

A number of specification parameters need considerations:

4.1 Stability, noise, ripple

The variation of the output quantity around its average value. Must be accompanied by a bandwidth specification. Stability is normally specified in the 1–100 mHz region. Noise is normally specified in several regions e.g. 0.1–10 Hz, 10 Hz–10 kHz, above 10 kHz. Ripple can be included in noise or specified separately, indicating the fundamental frequency. The IEC uses the all-encompassing term Periodic And Random Deviation (PARD).

4.2 Resolution

Resolution is the smallest increment that can be discerned in the output whilst changing the input. In digital systems it is often the size of the least significant bit. Generally the resolution is limited by noise and is hence also bandwidth dependent.

4.3 Accuracy

The DC accuracy of most quasi-linear systems is best described with two terms: a zero offset and a gain error. The offset is the output value for zero input. It can be expressed in volts or in ppm of the full scale value. The gain error is the difference in slope between actual and specified gain, normally expressed in ppm.

4.4 Linearity

The non-linearity is the maximum deviation of the output quantity from an ideal straight line between zero and nominal output. It is also possible to express this as the deviation from a linear regression line. However, this should be avoided, because, in practice, this complicates the set-up and evaluation and can give an overly optimistic result.

Linearity is a static quantity and great care must be exercised while measuring it, such that any settling behaviour (see below) does not invalidate the data.

4.5 Temperature behaviour

The offset and gain of the system will change over the temperature range and sometimes in a non-linear fashion. Temperature coefficients must be established. The changes originate in physical properties of resistors, capacitors, inductors, semiconductors, etc. In bad cases it may lead to instabilities and oscillations in part of the temperature range.

4.6 Settling behaviour

Description of how a measured quantity approaches its final value. One part originates with regulation characteristics (bandwidth, pole-zero configuration, etc.) and another part with any thermal changes. The latter can have unexpectedly long time-constants (hours or even days).

4.7 Repeatability and reproducibility

The dispersion obtained while returning repeatedly to the same operating point.

4.8 Long-term drift

Slow change over time (from weeks to years) due to material ageing or stress modification. The main concern is amplifiers and precision resistors. Humidity often has an influence and can cause drifts of 1–100 ppm/year.

5 Equipment

Voltage transducers do not normally present any significant problems and are not given much attention in this paper. The current transducer domain is a much more important subject and is examined in some detail. DAC and ADC technologies are described in Ref. [2].

5.1 Voltage transducers

The function of a voltage transducer is to measure the concerned voltage with a certain accuracy and bandwidth, whilst providing common mode rejection and/or isolation with respect to the control electronics. If the measured voltage is close to ground potential, the transducer may simply consist of a resistive voltage divider.

If a bandwidth of more than a few kilohertz is required, it will become necessary to complement the resistive divider with parallel capacitors to swamp any circuit stray capacitances. Bandwidths of up to the megahertz level can then be obtained with good accuracy.

If the measured voltage is not close to ground, good isolation has to be provided. Most problems can be solved with integrated circuit isolation amplifiers. They have good accuracy (~100 ppm), isolation (2 kV r.m.s.) and bandwidth (10–100 kHz).

If a higher isolation voltage is required, isolation with a magnetic circuit is a possible approach and bandwidths over 100 kHz can be obtained. Industrial units exist (LEM, Geneva), but the accuracy is generally limited to 0.1%, see Fig 3.

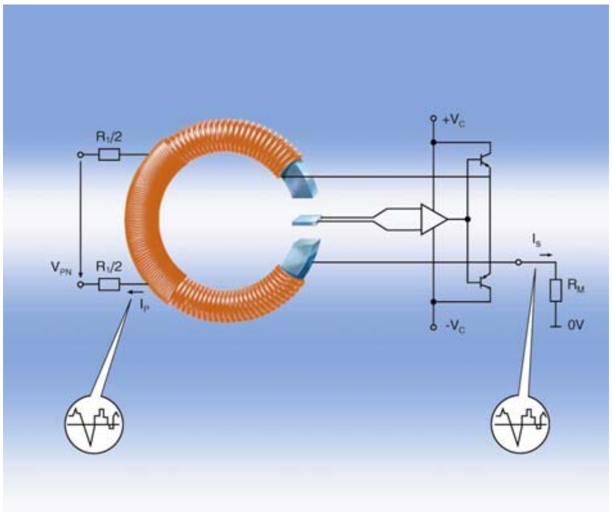

Fig. 3: LEM voltage transducer

Another option, more complex, but no longer very expensive, is an ADC floating at the measured voltage level and transmitting the measurement data over an optical fibre. At the receiving end, the information can be used directly in a digital control loop or converted back to an analog voltage with a DAC. Practically any isolation level can be reached and the only small problem may be to provide power for the electronics at the head end.

Verification of the performance of any of the above-mentioned types of voltage transducer is relatively straightforward with industrial instrumentation.

5.2 Current transducers

5.2.1 Current measuring resistors

The simplest current transducer is a resistor; it is often called a shunt. When it is deployed downstream of a current transformer, it is generally called a burden.

The electrical resistance is defined as the ratio of the voltage developed across a resistor over the current passing through it. It is a material property, a function of several parameters, not a constant. The resistance of materials is affected by temperature, but also many other factors, such as humidity, pressure, material, electric and magnetic fields, etc. It is always important to keep the voltage drop across the current measuring resistor low to minimize the power dissipation and hence the temperature rise. The minimum acceptable voltage depends on the accuracy level required. A voltage of 50–60 mV is generally used for 1% industrial panel meters, whereas 1 V would be considered the minimum for a 10^{-6} high-precision application.

The most interesting parameter when discussing dissipation is the Temperature Coefficient (TC) of the resistor. The basis for a good temperature performance is a resistive material that has an inherently low TC. Copper, aluminium, silver and gold all have a TC of about 4000 ppm/K and are not suitable. A normal carbon composition resistor has a TC of about 1000 ppm/K, a metal film resistor a TC from 25 to 100 ppm/K and very good high-precision resistors have a TC of 1–10 ppm/K. Over the years several mixtures have evolved from which resistors are made, like Manganin, Konstantan, Evanohm, Zeranin, etc.

A further refinement is to compensate artificially for the material TC over a certain temperature range. This can be achieved by mechanical, electronic or computational methods. In any case, it is VERY difficult to make a resistor with 1 ppm/K or better.

Internal heating in current sensing resistors will give rise to temperature gradients and hence resistance change. This effect is worse, the higher the power density in the resistor. If the resistor is also designed with a compensation technique to obtain a low TC, an internal gradient can become a problem which even outside cooling is unable to remove. CERN has found it very useful to measure the resistance change as a function of the dissipation (ppm/W), a ‘power coefficient’. As a result, one manufacturer (Vishay) now publishes a power coefficient for its current-sensing resistors.

To know the resistor’s total performance with temperature, self-heating due to dissipation, heating due to surrounding components, and, of course, changes in the ambient temperature have to be taken into account.

The packaging is important to the resistor performance. There is always a conflict between good electrical insulation and good thermal conductivity. No plastic package is hermetic, hence long-term drift is always better with a metal package. Examples are shown in Fig. 4. Mechanical stress on the package (even on a metal package!), which may also be temperature dependent, can also change the resistor value significantly.

The resistor leads are often made of different materials for many reasons. To ensure that the lead resistance does not adversely affect the electrical performance, a low value resistor needs separate current leads and voltage sensing leads. This configuration is called a four-terminal resistor and is generally available in industrial precision resistors for values < 10 ohms. Make a calculation of the resistance of a printed circuit board track, consider its TC of 4000 ppm/K and you will be surprised to see how quickly it becomes important when high precision is required.

Shunts have always existed for currents up to more than 20 kA, but obviously the accuracy falls off with the increasing dissipation and the devices can be complex with oil and water cooling etc. It is very difficult to obtain an accuracy better than 0.1% for high currents. Another, often serious, disadvantage is the lack of isolation between the power circuit and the measurement output.

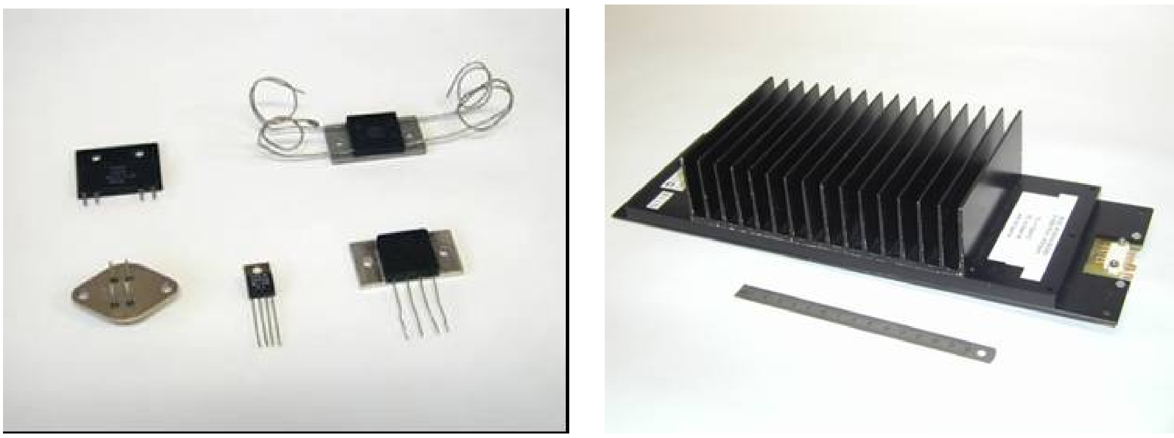

Fig. 4: Examples of burden resistors

5.2.2 Measuring transformers

The measurement transformer was introduced in AC measurements a long time ago. With good magnetic materials and design, the accuracy of the transformer can reach around 0.1%. The error is mainly due to the magnetizing energy needed for the core. Special configurations exist for standards laboratory set-ups in which the magnetizing energy is supplied from a separate source and ratio accuracies can then be brought towards the 10^{-6} level, but it is not adapted for industrial use and has limited bandwidth.

High bandwidth was always a requirement for particle-beam current transformers and Hereward (CERN) proposed an active transformer configuration around 1960. The extension of this concept to include DC was proposed by Unser (CERN) in 1969 and again published in the CERN Courier 1970 [3], where it was picked up by industry, Hazemeyer (NL). The concept was transformed into an industrial current transducer design by Hazemeyer and the first application was at DESY in Hamburg [4]. It became known as the zero-flux current transformer or DCCT and was used in large quantities in the SPS project at CERN in the 1970s.

A related cousin, called the current comparator [5], was invented in Canada at the NRC in the early 1960s by Miljanic, Kusters, and Moore for laboratory calibration purposes. It started as a DC device, but was later extended to low-frequency AC. Some of the industrial products that resulted from this development are often used at CERN today for DC calibration of the DCCTs.

5.2.3 Principle of the zero-flux current transducer or DCCT

The losses in the core of a transformer can be brought close to zero by putting it in an active feedback configuration. A sense winding detects any excursion from zero flux in the core and compensates this through driving an equal and opposite current in the secondary winding and hence balancing any flux change induced by the primary winding. The lower limit on the bandwidth can be in the region of a few hertz.

To extend this concept to DC, two methods exist, detecting zero flux with a Hall-effect device in the magnetic core or with a magnetic modulator. The Hall effect method is currently limited to an accuracy of 0.1% by the inherent temperature sensitivity (LEM and VAC) Fig. 5.

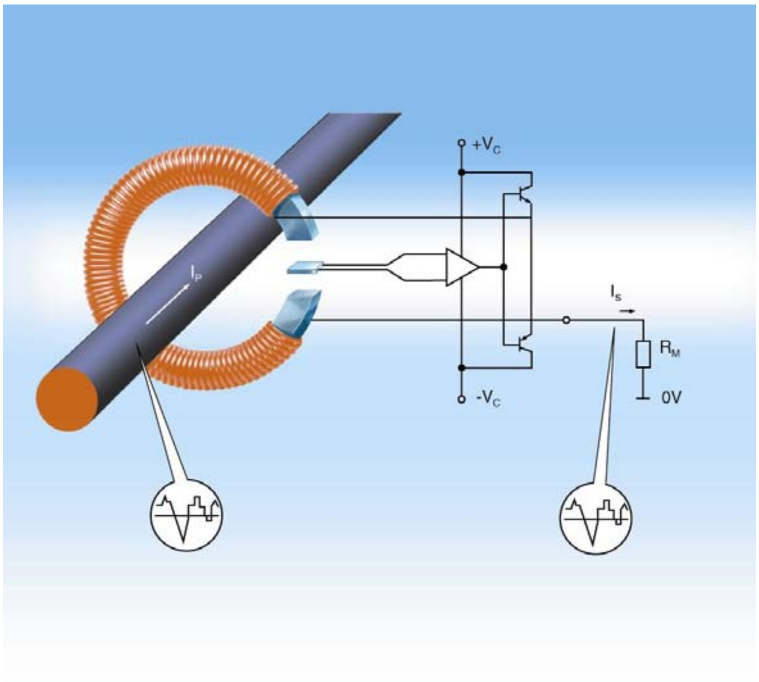

Fig. 5: LEM current transducer

The magnetic modulator is the most accurate method, able to obtain an accuracy of better than 10^{-6}. A core is excited with a voltage and a synchronous demodulator (second harmonic or peak detection) examines the magnetizing current. It can be seen from the non-linear hysteresis characteristics of the core that an asymmetry in peak current and hence the second harmonic contents will be produced as a result of any DC magnetization. Its phase will be dependent on the polarity. A synchronous detector can thus produce a signal to a feedback system providing a compensating DC current to maintain a DC flux balance. The zero-flux detector can have sensitivity down to 5–10 µAt, which provides resolution much better than 10^{-6} for medium to large DCCTs.

The AC and DC feedbacks are combined in one system, Fig. 6, and the compensation current is now a true representation of the measured primary current. The compensation current is generally passed through a burden resistor and an output amplifier brings this voltage to 10 V for nominal current.

The burden resistor will always be the Achilles heel because of the varying power dissipation. If ppm performance is required, temperature control is mandatory. In summary, it is easy to obtain a performance of 100 ppm from a DCCT, difficult to obtain better than 10 ppm and the practical limit is at a few ppm, depending on environment, recalibration periods, etc.

Other problems result from the magnetic cores being somewhat sensitive to external magnetic fields. These fields may be truly external, produced by a current return busbar in the proximity or even by the primary busbar being off-centre in the DCCT head. Local asymmetric saturation of the detector cores will introduce an offset change in the zero-flux detector. Internal magnetic shielding of the detector cores will reduce this sensitivity, but increases complexity and hence the cost of the head. It is increasingly important to preserve symmetry in any busbar/return busbar configuration with increasing primary current. In multi-kA circuits two or four return bars are necessary for high-precision transducers, clearly with good sharing between the busbars (not always a trivial problem).

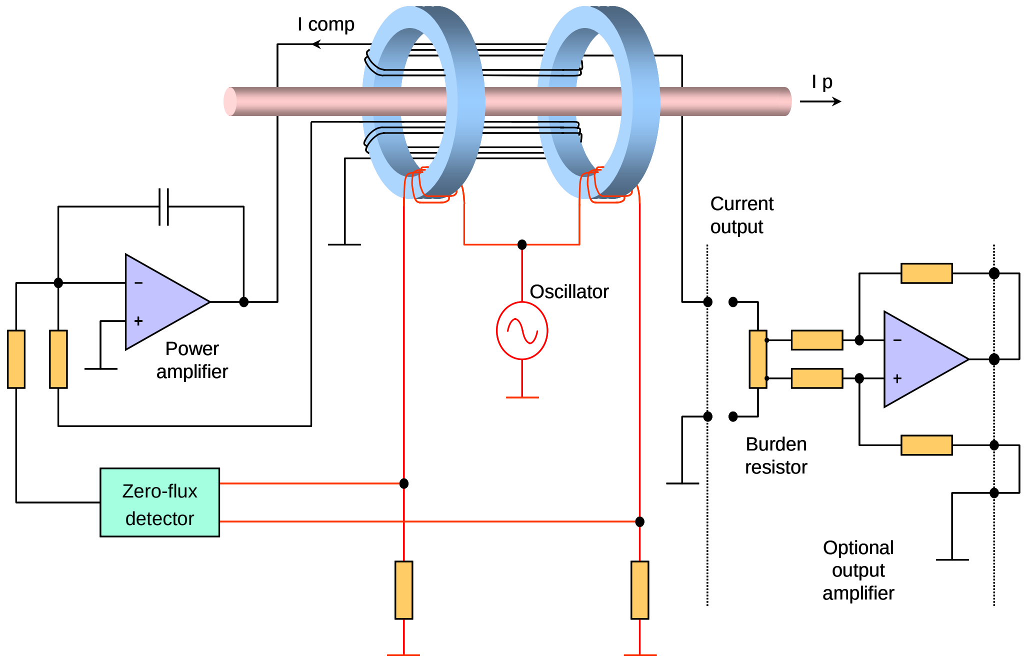

Fig. 6: Principle of a zero-flux DCCT

Different types of DCCT executions can be seen in Fig. 7. Fully integrated low-current models to the left and a separate head, high-current type to the right (Hitec, Danfysik and Ritz).

Fig. 7: DCCTs of different size and manufacture

6 Thermal management

The importance of temperature dependence of materials cannot be underestimated. In our experience, it is the single most important factor affecting the performance of high-precision electronic systems. Components are practically always specified with a temperature coefficient, so far so good, but is it constant with temperature? If it is small, probably not. What will the real operating temperature be? What is often forgotten in many designs is the internal self-heating of many components, the temperature gradients that are created, and the evacuation of this heat. Separate the dissipative components from the temperature sensitive ones wherever it is possible. Do not believe that all heating is convective, a significant portion is radiative. Put in thermal screens and use fans when necessary to control the air flow. Forced-air cooling makes a very big difference and long-life fans (50–100 000 h) are easily available.

High-precision electronics need particular precautions. Thermal emf's thrive on temperature gradients. Printed circuit designs should consider keeping critical points at equal temperature, most easily through proximity, but also through the use of thermal equalizing bridges and thermal screens to avoid air currents. Low-frequency noise can easily be generated by sensitive connections being exposed to turbulent air. A resistor can generate quite big thermal emfs with poor choice of materials between the resistance and the leads. An easy test is to connect the resistor to a microvoltmeter and then touch it with your fingers. Surprising how different resistors can be!

7 Calibration infrastructure

If a good precision level must be established and maintained, a certain infrastructure of standards and instruments must exist on-site. A Digital VoltMeter (DVM) with sufficient resolution and accuracy is a good start, but this must also be calibrated from time to time. If many DVMs are used, it is strongly recommended that voltage (most commonly 10 V) and resistance (0.1 Ω–100 kΩ in decade steps) standards are available on-site to keep calibration costs down. See Fig. 8 for examples. An air-conditioned room (23°C by international convention) is mandatory for good calibration results.

Fig. 8: Examples of voltage and resistance standards

A next step would be a power converter that can supply the full primary current for calibration of DCCTs. It can of course be implemented with several primary turns through the DCCT head to decrease the power converter current. The current should be as stable as possible and with low ripple to simplify measurement. Switch-mode converters may cause severe Electro Magnetic Compatibility (EMC) problems and should be avoided unless experience is available. At least two DCCTs should be chosen as reference DCCTs and all other DCCTs will be measured against these reference units.

The burden resistors in the reference DCCTs should be measured against standard resistors, but without special equipment this can only be done at low current levels. A complete high-precision evaluation of important burden resistors at their operating current level requires a special measurement bridge (e.g. MIL/Canada) or an accurate and stable current source. An extremely accurate current source developed at CERN is described in Ref. [6].

On-site calibration of a DCCT can be done in several ways, but is dependent on the means available. The first method would be a transportable DCCT to be connected at the power converter output terminals and verify the output current. If a precise current source is available, a second method is to inject a calibration current into the burden resistor and read the DCCT output. However, not all DCCTs provide easy access to the burden current terminals.

The third and ultimate method [6] would be to specify a calibration winding in the DCCT head in which a calibration current can be injected. The advantage is that the operation can be carried out without any disconnection of power terminals etc., and that it tests the complete function of the DCCT, including the zero-flux detector, up to nominal ampere-turns. However, it requires a special current calibrator with high impedance output and sufficient voltage capability.

8 Integration problems

An electronics designer tries to anticipate what the operating conditions will be for the future use of the equipment he is designing, but all possible combinations can not be covered.

The most common problems originate from poor understanding of grounding and shielding and poor control of ground and return currents. Remember that 10 mA going through 10 mΩ will generate a drop of 100 μV, which is equivalent to 10 ppm of 10 V. Note that each of the mentioned quantities is on its own a reasonable value in any electronics system! It is the combination in the wrong place that may cause problems. An excellent treatment of grounding and shielding is provided in Ref. [7].

Another problem comes from capacitive loading of amplifiers. If the output of an amplifier, e.g., a DCCT voltage output, is loaded with a few hundred pFs, probably nothing will happen. With increasing capacitance, the pole formed by the final stage output impedance and the load capacitance will move down in frequency and when the phase shift is big enough, the amplifier will oscillate at high frequency. It is a wise precaution to check with an oscilloscope whilst connecting a capacitor simulating output cable capacitance. The remedy is simple and can be found in many IC amplifier application notes.

9 Electromagnetic compatibility

Recently more weight has been given to interference problems and norms covering EMC levels and measurement methods are now well established by the IEC [8]. The susceptibility of the equipment that could potentially be disturbed (high-precision, low-level electronics) can now be measured in a systematic way, as can the emissions from potential sources of disturbance (switch-mode converters etc.). Several factors will influence the level of attention that must be given to these aspects: the precision desired, the switch-mode power involved and the dI/dt–dV/dt in the power circuits, i.e., the power spectrum of conducted and radiated energy. A precision in the 100–1000 ppm range requires few precautions. A 10–100 ppm precision requires many precautions and it is advisable to test integrated full-scale prototypes. A 1–10 ppm precision requires extreme precautions, thorough knowledge of EMC problems and remedies, as well as very detailed knowledge of the electronics subsystems that will be used.

How do you know if you have an EMC problem? Typical symptoms are unexpected offset voltages at zero, poor linearity over the control range, conflicting results from different measurements or instruments, and periodic or non-periodic disturbances on the signal lines.

Preliminary tests should focus on identifying the main source of disturbance and the way it is transmitted. Is it radio frequency or low frequency? Is it radiated or is it conducted? Walkie-talkies are very good noise generators to get a first idea of susceptibility. If the mains is suspect, insert a mains filter. Harmonics on the mains can easily be verified with an oscilloscope etc. RFI can also be identified with a 100–500 MHz scope and a matched divider. Spikes of mVs are not a problem, spikes of tens of mV start causing problems in analog electronics and hundreds of mV are practically impossible to live with. Check ±15 V power rail cleanliness.

Normal remedies are to check the shielding arrangements, introduce L-C filters and/or common mode chokes. If more serious problems are present, a good EMC consultant is your best bet. A more complete and general treatment of common EMC problems is given in Ref. [9].

10 Conclusions

– Establish realistic specifications from the beginning and be aware of the cost implications.

– Build conservatively with good margins.

– Watch out for ‘specmanship’ and quality control in industrial products.

– Test everything in the laboratory before installation, not in the final machine.

– Switch-mode converters increase EMC problems by at least an order of magnitude.

– Presumption is the mother of all screw-ups.

References

[1] *Guide to the Expression of Uncertainty in Measurement* (ISO, Geneva, 1995).

[2] J. Pickering, *ADCs and DACs* (ISO, Geneva).

[3] *CERN Courier*, December 1970, 380–382.

[4] J. Lisser, private communication (2005).

[5] W.J.M. Moore and P.N. Miljanic, *The Current Comparator* (IEE/Peter Peregrinus, 1988).

[6] G. Fernqvist *et al.*, A novel current calibration system up to 20 kA, *IEEE Trans. Instrum. Meas.* **52** (2003) 445–448.

[7] Henry W. Ott, *Noise Reduction Techniques in Electronic Systems*, 2nd ed. (Wiley, New York, 1988).

[8] IEC 61000, *Electromagnetic Compatibility* (IEC, Geneva).

[9] A. Charroy, *EMC* (IEC, Geneva).