Milli-Kelvin controller stability tests

Simon Tulloch

QUCAM Technical Note

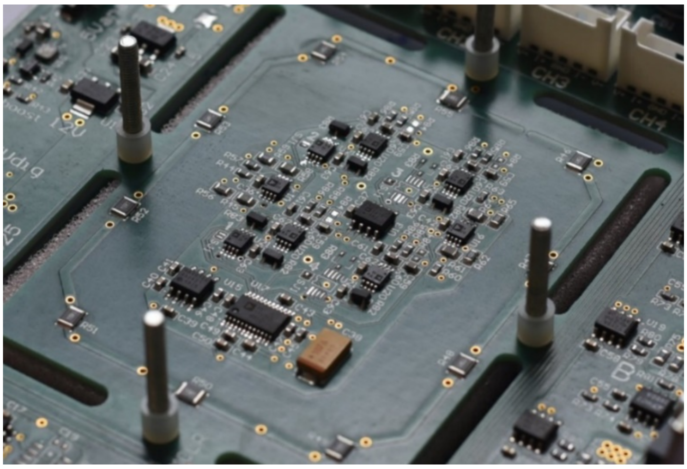

Jerez, 3 Feb 2021

1. Introduction

The Milli-Kelvin Controller will have two principal error sources. The first is offset voltage drifts in the preamplifiers due to low frequency 1/f noise and thermal coefficients. The second is offset drifts in the precision current sources used to energise the temperature sensors. These drifts cannot be characterised by simply looking at the servo stability data. Any sensor forming part of a control loop will have any drifts servoed-out and drift errors will be unobservable. Instead it is necessary to use some form of external temperature reference entirely decoupled from the control loop.

The stability of a temperature servo system should ideally be measured using an absolute temperature reference. One popular technique is to use a stirred ice-slush bath. This supposedly gives errors below +/-2mK. Unfortunately we could not reproduce this result. An alternative method was developed where the diode temperature sensor was replaced with a precision low-temperature coefficient resistor.

The resistor was chosen to give the same voltage signal as the diode sensor at approximately 274K. Since the Milli-Kelvin Controller is easily able to see the temperature coefficient of even 10 ppm precision resistors, it was necessary to mount the reference resistor in an oven stabilised to a few tens of mK over the course of the measurements. Any drifts in the indicated temperature given by the Milli-Kelvin Controller could then be interpreted as internal errors due either to offset voltage errors in the sensor preamplifiers or errors in the sensor energisation current.

Although this reference resistor produces the same voltage signal as the actual sensor diode, it differs from the diode in one important way. It will have a different sensitivity to errors in the energisation current produced by the controller. To determine how significant this difference was we had to measure the *dynamic resistance* of the diode sensor. This is explained in Section 3. It turned out that the dynamic resistance was an order of magnitude less than the reference resistor. The sensitivity of the actual temperature sensor diode to energisation current errors was therefore very much less than the reference resistor we used. We are unsure of the relevant contributions of offset voltage drift and energisation current errors. If energisation current errors are at all significant then it means our stability result must be interpreted as worst case. An actual diode sensor should have a much higher stability. In this report we are actually placing an upper bound on the system stability.

2. Diode temperature sensitivity

The temperature voltage relationship of the forward biased 1N4148 sensor diode was measured by cross calibration with a Lakeshore Pt100 sensor. The result is shown in Fig 1. This data is incorporated in a look-up table in the Milli-Kelvin controller ROM. The sensitivity is approximately -2.5mV/K and is surprisingly linear across a large temperature span.

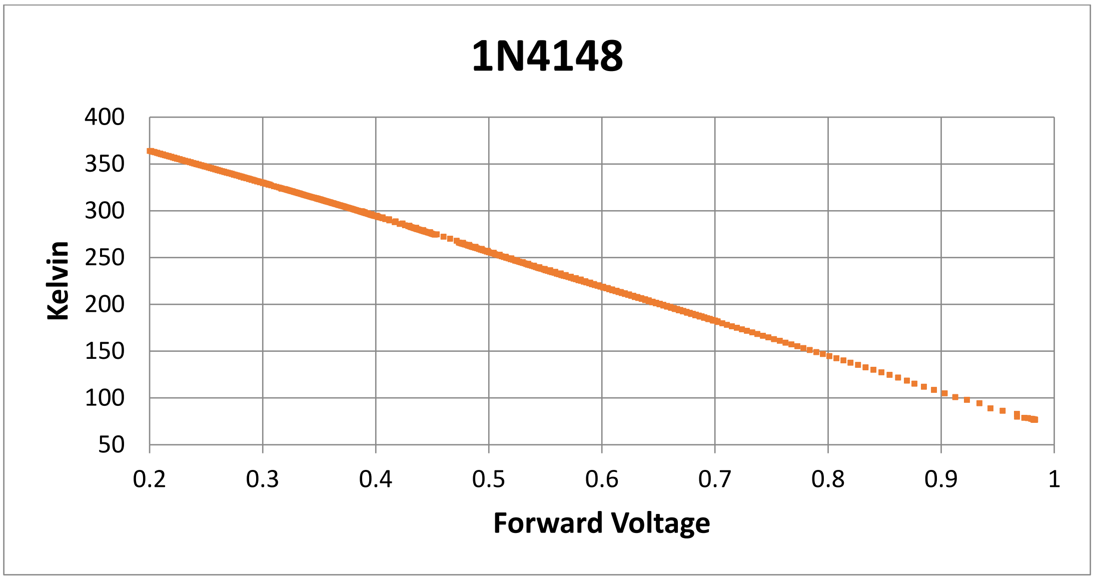

Figure 1. Temperature calibration of 1N4148 diode with 10μA energisation current.

3. Dynamic resistance of the diode

The textbook diode equation is as follows:

$$I = I_0 \exp\left(\frac{q.V}{n.k.T}\right) - I_0 \qquad \text{Eq 1)}$$

where:

$$I = \text{Forward current}$$
$$I_0 = \text{Reverse saturation current}$$
$$q = \text{Electronic charge}$$
$$V = \text{Forward voltage}$$
$$n = \text{Diode ideality factor}$$
$$k = \text{Boltzmann constant}$$
$$T = \text{temperature}$$

The Ideality factor and Reverse saturation current can be determined by measuring the forward voltages at a number of forward currents.

Eq. 1) can be rewritten as:

$$\ln I = \ln I_0 + \ln\left[exp\left(\frac{q.V}{n.k.T}\right) - 1\right] \qquad \text{Eq 2),}$$

which approximates to:

$$\ln I \approx \ln I_0 + \frac{q.V}{n.k.T} \qquad \text{Eq 3).}$$

By then plotting $\ln I$ versus V we can obtain I_0 from the y-axis intercept and n from the gradient. This is shown in Figure 2.

This gives the following result:

$$n = 1.94$$
$$I_0 = 5.2 \times 10^{-9}\text{A}$$

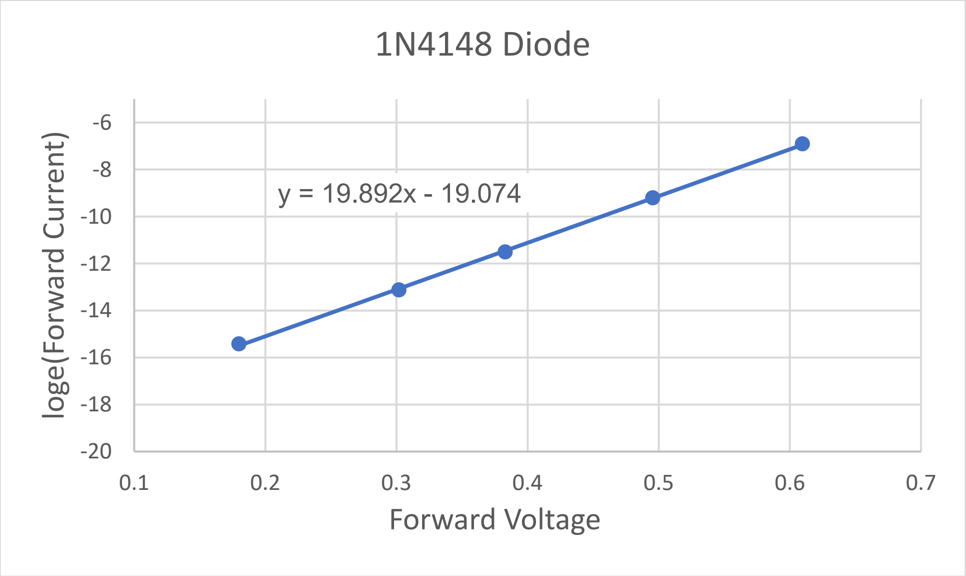

Figure 2. 1N4148 characteristics at 300Kelvin

We can now calculate the dynamic resistance of the diode as a function of the forward current. Specifically we need the dynamic resistance at 10μA, the standard operating current used for diode sensors with the Milli-Kelvin controller. Taking the derivative of Eq. 1 we get:

$$\frac{dI}{dV} = \frac{q}{n.k.T}.I_0.\exp\left(\frac{q.V}{n.k.T}\right) \qquad \text{Eq 4),}$$

combining this with Eq. 1 then gives the dynamic conductance of the diode as follows:

$$\frac{dI}{dV} = \frac{q}{n.k.T}.\left(I + I_0\right) \qquad \text{Eq 5).}$$

Taking the reciprocal of this result we obtain the dynamic resistance:

$$\frac{dV}{dI} = \frac{n.k.T}{q.\left(I + I_0\right)} \approx 5\text{k}\Omega \qquad \text{Eq 6).}$$

The value of the dynamic resistance determines how sensitive the diode is to errors in its energisation current.

4. Stability measurement technique

If we are simply trying to measure the temperature stability of the sensor preamplifiers it makes no sense to connect them to actual temperature sensors. These sensors are specifically designed to have high temperature coefficients. Instead we need to use a "dummy" sensor that gives the same voltage signal as a real sensor but has the lowest possible temperature coefficient. Any drifts we see in the preamplifier response can then be interpreted as actual electronic drifts.

We chose to use precision resistors as dummy sensors. The highest resistor stability that is commonly available is 10ppm (parts per million). A change of 1mK at a temperature of 300K corresponds to about 3ppm so it proved necessary to mount these dummy sensors in their own stabilisation oven. The arrangement is shown in Fig 3. Fig 4 shows a close-up photo of the resistor oven.

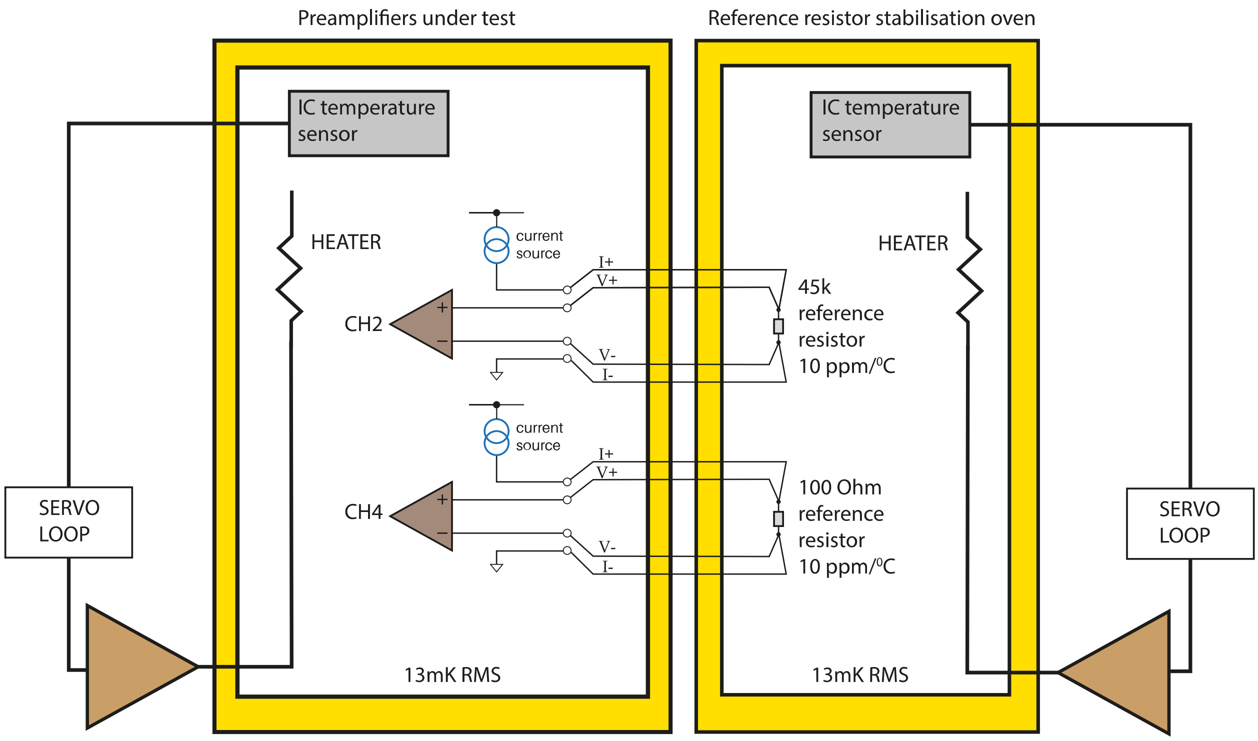

Figure 3. Measurement technique showing the two ovens, one for the reference resistors and one for the preamplifiers under test.

The choice of dummy sensor to replace the Pt100 sensor was easy, we simply chose a 100R resistor. This will clearly then have the same susceptibility to offset voltage drifts and energisation current drifts as the real Pt100 sensor. For the diode sensor channel the choice was more difficult. If we wish to replace a diode sensor with a resistor that gives the same voltage signal at 300K then we are forced to use a value of approximately 45kΩ. Note that this is almost an order of magnitude greater than the dynamic resistance of the diode. Using such a large resistor means that we will no longer have the same sensitivity to energisation current errors as for a real sensor. Our result for the diode sensor channel cannot then be interpreted as a true value for stability, instead it is a worst case result or rather, places an upper bound on its true stability.

The ease with which a Pt100 sensor can be simulated using a resistor is a significant reason for using them in high stability systems. Any such system whose performance cannot be regularly and easily demonstrated will be of limited use. The decision is not clear, however, since diode sensors are highly desirable given their greater sensitivity.

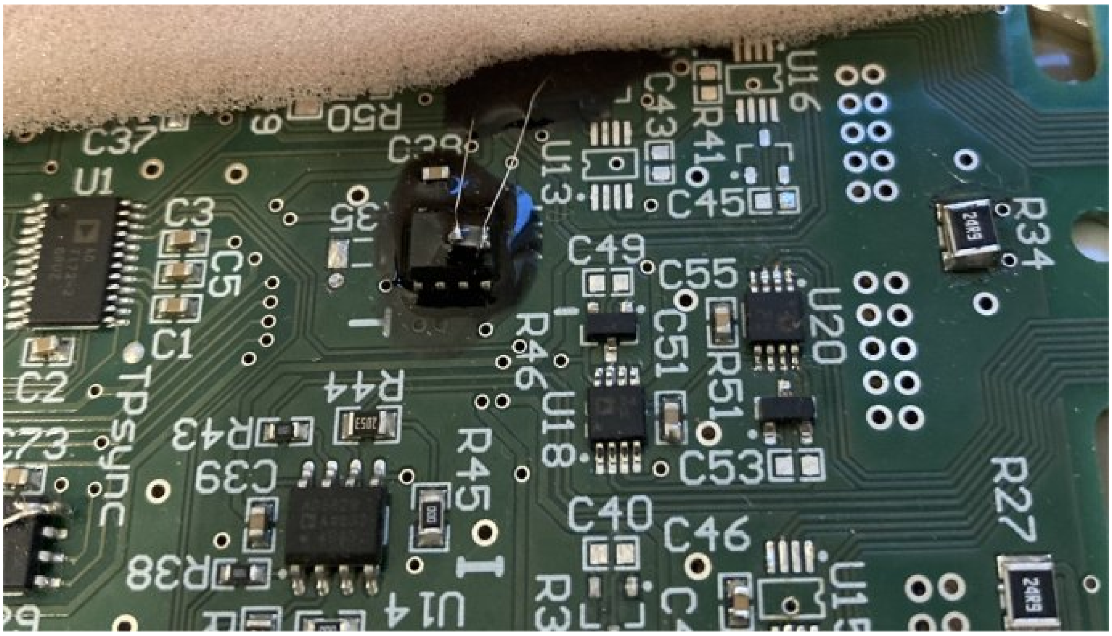

Figure 4 Reference resistor mounted in a separate temperature controlled oven. The SMD resistor is attached using black epoxy to the oven's temperature servo sensor. Oven insulation has been removed for a clearer view. A second resistor was later added to allow simultaneous measurements on two channels.

5. Stability result over 100 hours

The oven stabilisation of the reference resistors was successful with an RMS variation of only 13mK over the course of the 100 hour run. We can thus be confident that any variations in indicated temperature with the Milli-Kelvin Controller were entirely due to internal electronic drifts in its circuitry and not changes in the reference resistor.

Figures 5 and 6 show the stability data for channels 2 and 4 over the course of the 100 hour run. This raw data was later subjected to a high pass digital filter with a time constant of 24 hours. The results are shown below.

Table 1.

Sensor	Drift over 100 hours	Drift over 24 hours
CH2 (diode)	0.66mK RMS	0.52mK RMS
CH4 (Pt100)	0.97mK	0.86mK

The ambient temperature was also recorded over the course of the run. This is shown in Figure 7. There appear to be weak correlations suggesting that the preamplifier stabilisation oven wasn't working as well as expected. The correlations were very approximately on the order of +1mK/K for the diode channel and -0.5mK/K for the Pt100 channel. If the Millikelvin Controller is used in less-well regulated environments the stability will suffer.

The Pt100 and diode channels give very similar performance (sub-mK over 100 hours). This is despite the very much larger sensitivity of the diode sensor. This tends to suggest that offset voltage drifts in the preamplifiers are not a significant noise source. Instead it points to the current source drift as being the major factor. Remember the true susceptibility of a diode sensor to current source drifts will be an order of magnitude lower than what we are able to measure using the reference resistor.

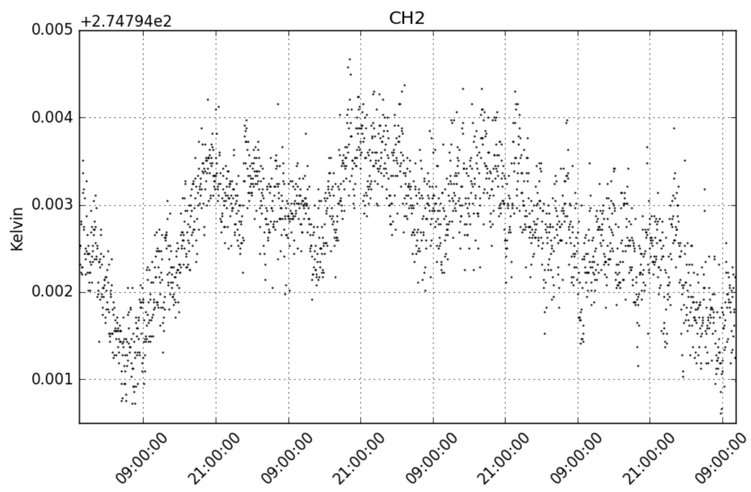

Figure 5. Raw channel 2 (diode) temperature data. The RMS deviation of this data is 0.66mK over 100 hours.

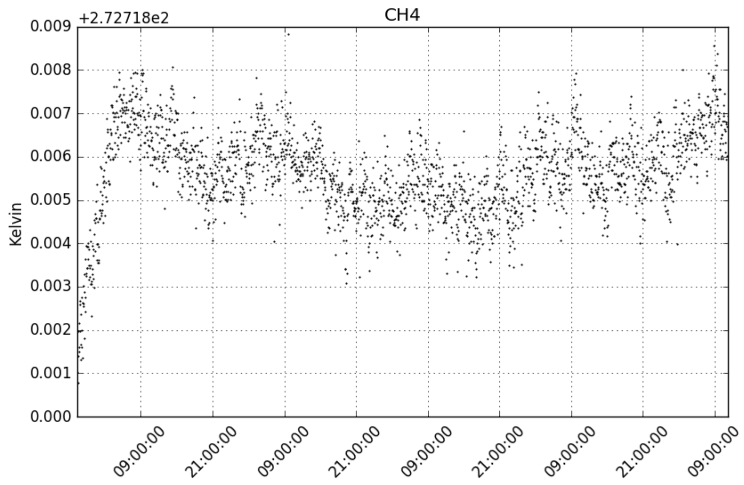

Figure 6. Raw channel 4 (Pt100) temperature data. The RMS deviation of this data is 0.97mK over 100 hours.

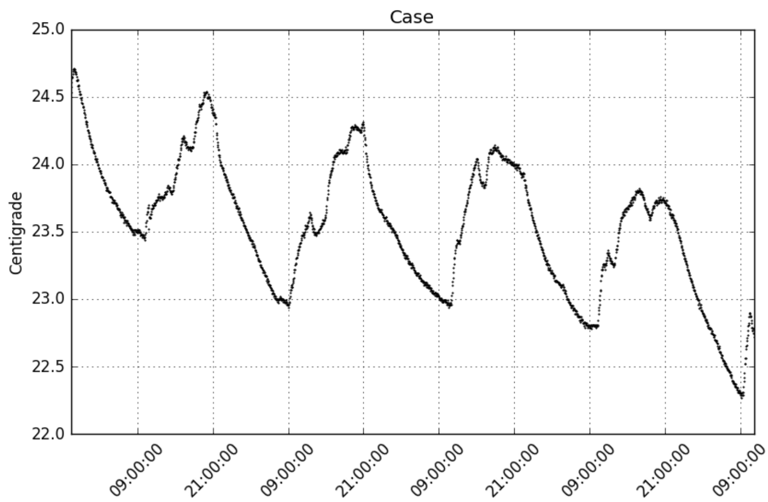

Figure7. Ambient temperature over course of runs shown in Figs 5,6.

6. Preamplifier temperature coefficients

The weak correlations between the ambient temperature and the long-term stability data shown in Figures 5 and 6 cast some doubt as to the efficacy of the preamplifier stabilisation oven. The actual oven temperature data was obtained using a temperature sensor IC mounted at the centre of the preamplifier PCB (see component U13 on Fig 11). This sensors output remained highly stable over the course of the run so we surmise that the oven contained significant temperature gradients. These gradients then change in response to ambient temperature. The oven consisted of a thermally isolated section of the PCB ringed by heater resistors and containing substantial amounts of copper on all 4 layers. Use of so much copper, it was hoped, would reduce gradients to insignificant levels. The proximity of the current setting resistors (see R58,59,61,62 in Figure 11) to the servo heater resistors probably doesn't help.

Since the preamplifier was mounted in its own oven it was straightforward to measure its temperature coefficients. The temperature was varied between 302 and 309K as illustrated in Figure 8. The servo loop controlling the oven was first programmed to have maximum slopes of 0.1K/minute. This was later modified to 5K/minute. The response of CH2 (diode) and CH4 (Pt100) were rather different. The diode channel coefficient was -0.8mK/K whereas the Pt100 channel gave +0.6mK/K. This is consistent with the current sources used to energise the sensors having a small +ve temperature coefficient. It also gives further evidence that the current sources are the major source of error. One further observation was that the diode channel was particularly sensitive to the derivative of the temperature changes whereas the Pt100 channel was only sensitive to the absolute temperature. Sudden changes in temperature are likely to mechanically strain the PCB components.

Later versions of the preamplifier will aim to reduce thermal gradients using a metal heat-spreader and also by use of through-hole rather than surface mount resistors. The former are likely to be less prone to mechanical strains since these will be absorbed by the compliance of the wire leads.

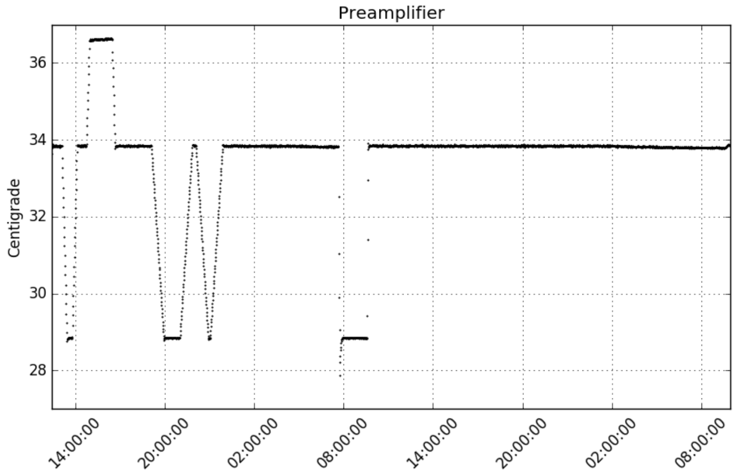

Figure 8. Programmed changes in preamplifier temperature.

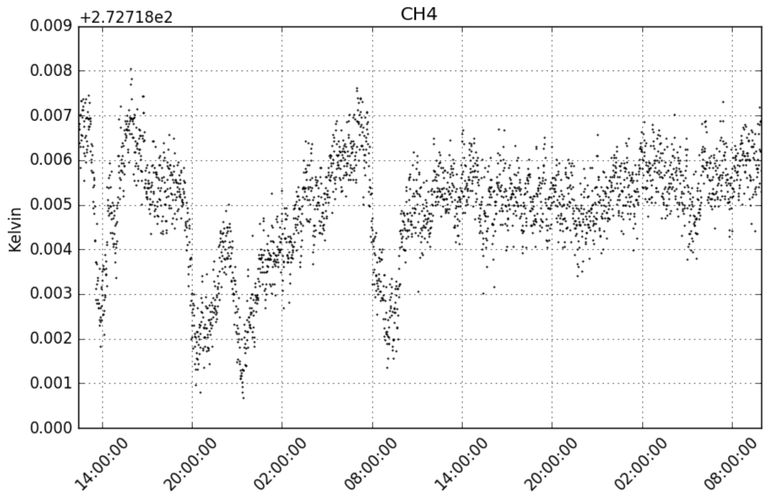

Figure 9. Pt100 sensor channel response showing changes due to the variations in preamplifier temperature illustrated in Figure 8.

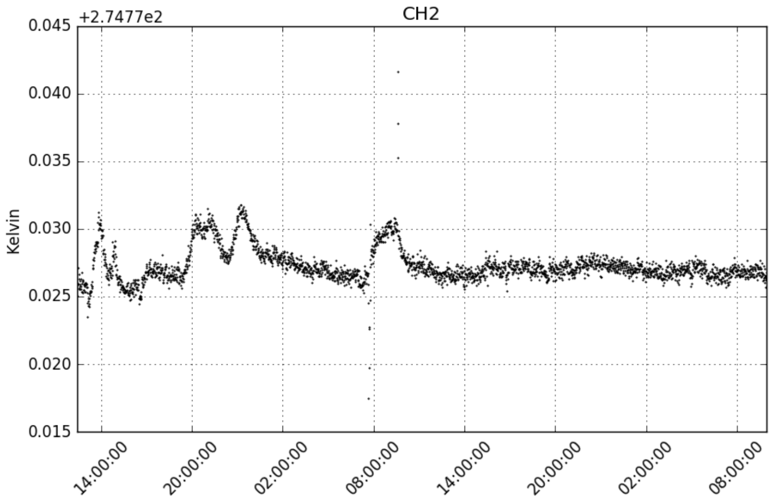

Figure 10. Diode sensor channel response showing changes due to variations in preamplifier temperature illustrated in Figure 8.

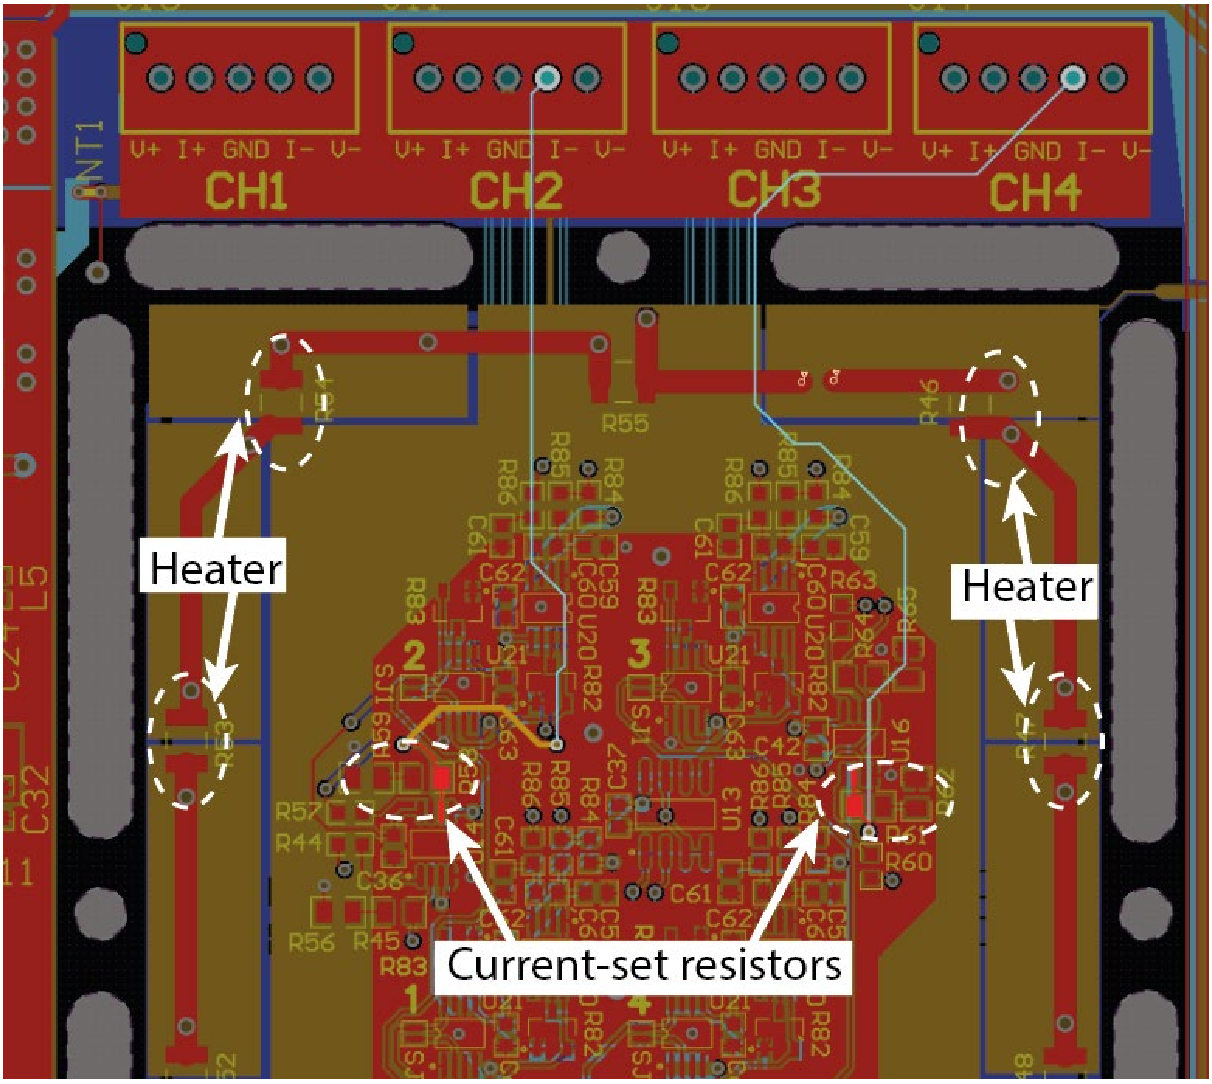

Figure 11. PCB design file artwork showing central part of controller board containing the preamplifier oven. The large regions of copper pour, shown colour-coded on the various layers, aim to minimise gradients within the oven.

7. Conclusions

Measuring true temperature stability over extended periods is extremely challenging. It really requires the use of an absolute temperature reference so that drifts in both sensors and electronics can be measured. Such references are hard to obtain so instead we have developed a way of at least characterising the short-term stability of our measurement electronics. We have not measured intrinsic drifts in the sensors themselves.

The controller has demonstrated sub-mK stability over 100 hours using both types of sensor. It seems highly likely that the true stability of the diode sensor channel is actually much better than we have been able to demonstrate using our technique.

Further revisions of the preamplifier PCB will focus on improvements to the current sources and mechanical changes to make the preamplifier oven more isothermal.

The preference of diodes over Pt100 is now cast into doubt due to the difficulty of demonstrating the true stability of the diode-sensor electronics.

---*End of Document*--