Working with weakly magnetized samples on the OPRF cryomagnetometer

OPRF booklet 6
Pontus Lurcock
27 April 2009

> But now… how turn ye again to the weak and beggarly elements, whereunto ye desire again to be in bondage? – *Galatians 4:9*

Contents

1 Introduction

At the OPRF we routinely work with very weakly magnetized samples whose remanences are at the limit of the 2G cryomagnetometer's range. Under these circumstances it is obviously desirable to determine and document a set of best practices to produce the most precise measurements possible with the equipment we have.

1.1 On precision and accuracy

This project is mainly concerned with maximizing the *precision* of our measurements – that is, the repeatability of measurements under identical conditions. It is less concerned with *accuracy*, the relation of those measurements to the true physical values being measured (but see notes on the tray correction below). For us, precision is probably a more pressing concern: we are more often interested in demagnetization trends than absolute values, so systematic bias is a less serious problem than it might at first appear. Conveniently enough, precision is also somewhat easier to work on than accuracy, since it can be investigated simply by remeasuring the same samples repeatedly under different conditions.

2 First set of tests

For this study, I focused on four of the most familiar settings offered by the 2G control software:

- **Number of measurements**

  The measurement queue editor allows the user to specify how many measurements are made per sample. the 2G manual states that they are 'averaged at each sample point' (page 6;4) and the data file format refers to the result as a 'mean magnetic moment' (page 10:4), which seems to imply that an arithmetic mean is taken of the measurements. However, the manual does not make this explicit, and the results of this investigation (see Section 4.1.2) imply that some other procedure is being used.

- **Drift correction**

  This correction involves measuring the background moment or susceptibility before and after measuring the samples or long core, and interpolating linearly in time between the measurements to produce an estimate of the background value at each point in time at which a sample measurement was taken. This estimated background is subtracted from the measured value.

- **Delay after moving tray**

  'After the core has moved to a new position, Long Core will wait for a period equal to the Delay Time before latching the SQUID's analog and count values. This delay is designed to allow the VRM to decay to a stable value.' (page 6:5)

- **Tray correction**

  'When Tray Corrected is set to Yes, Long Core will subtract the magnetic moments or susceptibility of an empty tray from the measured values.' (page 6:5). Thus, this setting should affect accuracy rather than precision, unless the tray is remeasured repeatedly.

2.1 Procedure

The basic experimental plan was to vary each of the settings mentioned above while holding the others constant, in order to determine which settings have the greatest effect on precision of measurement.

No. of measurements	Drift corrected?	Tray corrected?	Delay (seconds)
1	no	no	1
1	no	yes	1
1	yes	no	1
1	no	no	5
1	no	no	10
3	no	no	1
5	no	no	1
10	no	no	1
5	no	no	5

Table 1 Settings for precision tests

For each of the combinations of settings, ten successive measurements were made. Note that the final set of measurements differs slightly from the stated protocol, in that two settings (number of measurements and delay time) are simultaneously varied.

2.2 Sample selection

For these tests I selected four samples of limestone collected at Limestone Point on Campbell Island in November 2007. Their initially weak remanences had been further eroded by thermal demagnetization to 500 °C. The previously measured intensities of these samples were as follows.

Sample number	Sample code	Intensity (G)
1	LP2469.1	1.16×10^{-7}
2	LP2774.1	1.33×10^{-8}
3	LP2775.1	1.93×10^{-8}
4	LP2877.1	1.30×10^{-8}

Table 2 Samples and nominal intensities for precision tests

Sample 1 has an intensity an order of magnitude greater than the others; it was deliberately chosen to provide data on the behaviour of the machine with slightly more magnetized samples.

2.3 Empty sample positions

Samples 1, 2, 3, and 4 were placed in sample positions 1, 3, 5, and 7 respectively on the discrete measurement tray. Positions 2, 4, 6, and 8 were left empty in the hope that variations in the measured values at these positions could be used to measure the current background and correct for it.

3 Results

3.1 Analysis technique

I used standard deviation (across the ten runs with each configuration) as a proxy for precision, and wrote a small Python script to calculate, for each experimental run, the sum of the standard deviations for each of the three SQUIDs. (Initial data exploration indicated that, while there are differences in SQUID behaviour, their response to measurement settings is fairly similar.)

The raw data is on the OPRF RAID at `/cryomag/pont/weak-tests`.

3.2 Incidental observations: the warm-up effect

Initial perusal of the data revealed a factor outside the initial remit of this project: whenever a run of successive, identical measurements is being made, there appears to be more noise on the initial measurements. I will hereinafter refer to this phenomenon as the *warm-up effect*, and discuss it more fully in Section 4.1.

3.3 Effect of number of measurements

The effects of averaging more than one measurement proved dramatic: as can be seen in Figure 1, increasing the number of measurements to three decreases the standard deviation by two orders of magnitude for the weakly magnetized samples. The reduction is less pronounced for the more strongly magnetized Sample 1; this may be partly due to the stronger magnetization, and partly due to the sample being in position 1 on the tray.

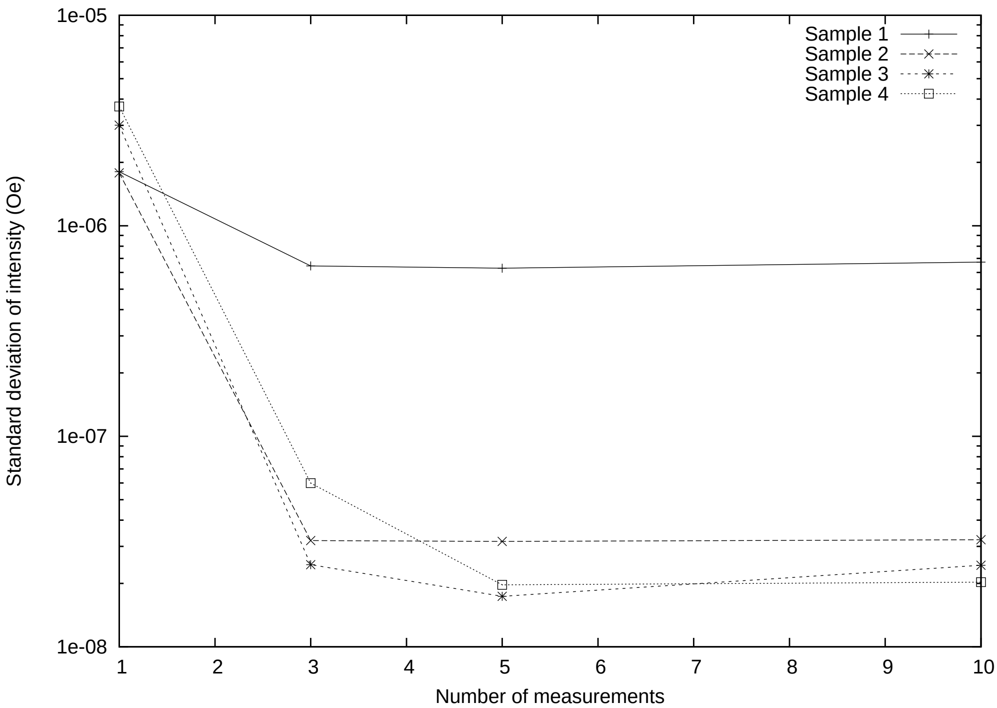

Figure 1 Effect of number of measurements. Note the logarithmic scale.

3.4 Effect of delay time

The effect of delay time is a little surprising: as Figure 2 shows, there is no noticeable correlation between delay time and precision. It is possible that the effect is too small to be seen under the conditions of the experiment: the plotted data are all with the number of measurements set to 1 in the 2G magnetometer, and a similar test with 3 or more measurements might show a clearer trend. It is in any case clear that the effect of delay time, if any, is dwarfed by the effect of increasing the number of measurements.

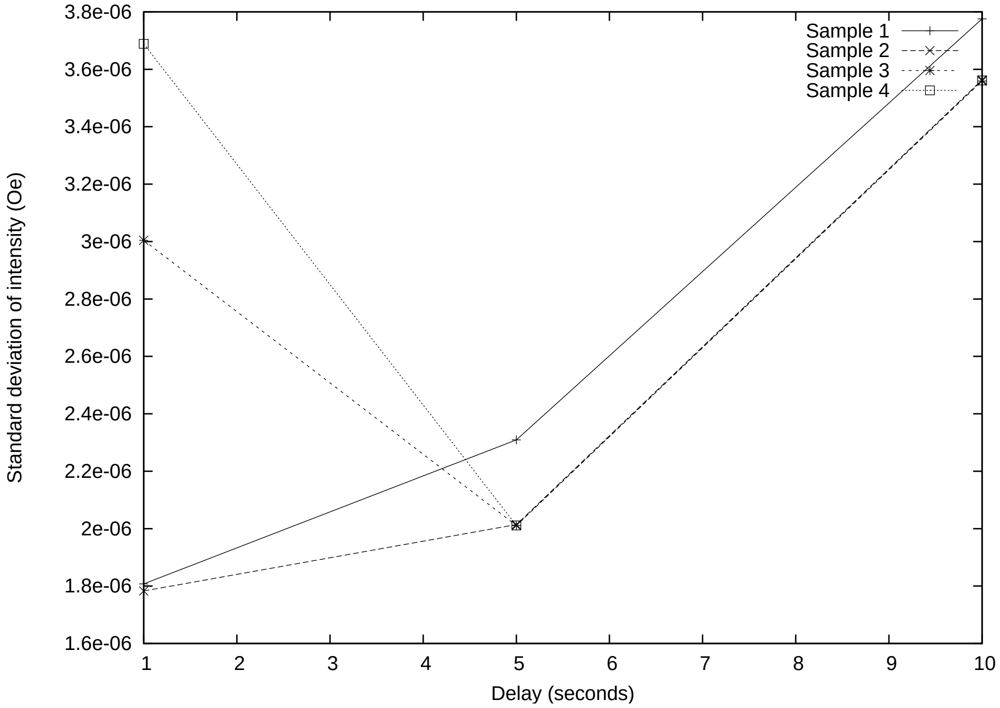

Figure 2 Effect of delay time.

3.5 Effect of drift correction

The results for drift correction are promising but inconclusive, as can be seen from Table 3. There was a slight improvement in precision for sample 2 and a significant improvement for samples 3 and 4, but for the more strongly magnetized sample 1 precision *decreased* slightly. These results definitely warrant further investigation of drift correction with a larger number of measurements per sample.

	σ (μG)		
Sample number	uncorrected	drift corrected	tray corrected
1	1.8	2.3	2.7
2	1.8	1.3	2.1
3	3.0	.90	4.4
4	3.7	.59	4.7

Table 3 Effect of drift and tray corrections on precision

3.6 Tray correction

As expected, applying the tray correction did not improve precision: in fact, the result appeared to be a slight reduction.

4 Second set of tests

It is clear from the first round of tests that the easiest way to improve precision significantly is to increase the number of measurements per sample. The second round of tests consisted of three sets of repeated measurements using an average of three measurements per sample, with the goal of investigating the warm-up effect and the effect of subtracting the reading from an empty sample position.

4.1 The warm-up effect

Keeping the number of measurements per sample at 3, I performed three test runs of measurements on the samples used in the first round: two sets of 10 measurements with a delay of around half an hour between them; and a set of 20 measurements, performed after the machine had not been used for around 14 hours.

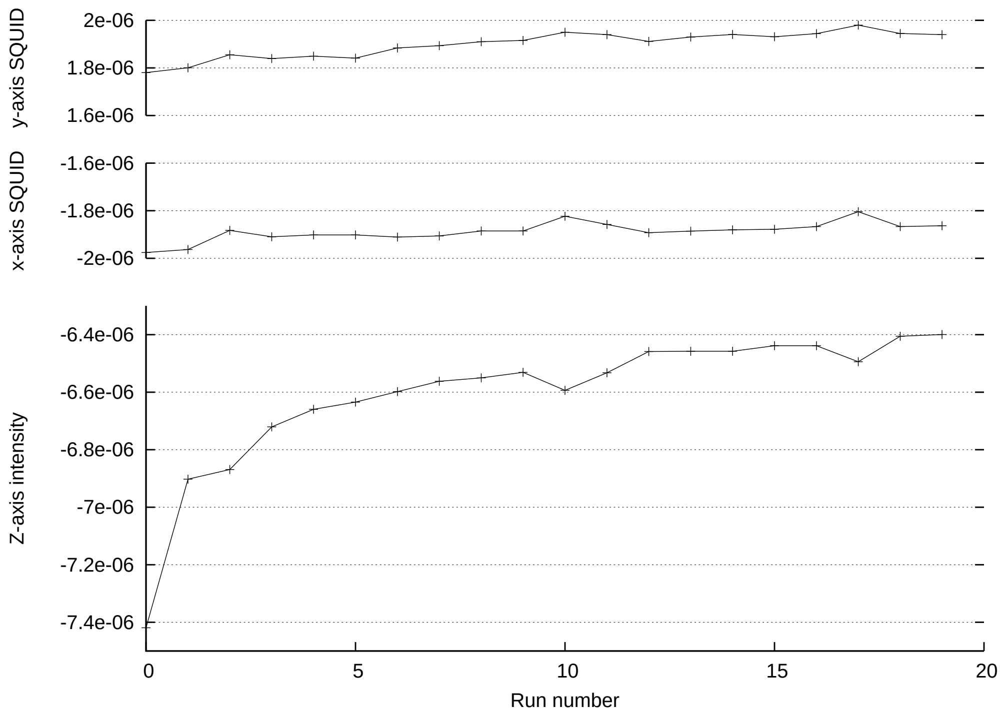

Figure 3 20 successive measurements of sample 1; output from the three SQUIDs plotted separately.

As Figure 3 shows, the warm-up effect is worryingly pronounced for the more strongly magnetized Sample 1. Even after 20 measurements, the readings for all the SQUIDs still seem to be increasing.

The readings from the empty sample slots cannot be used to correct for this effect: there is no significant correlation, and the variations in the empty slot readings are far lower (see Figure 4).

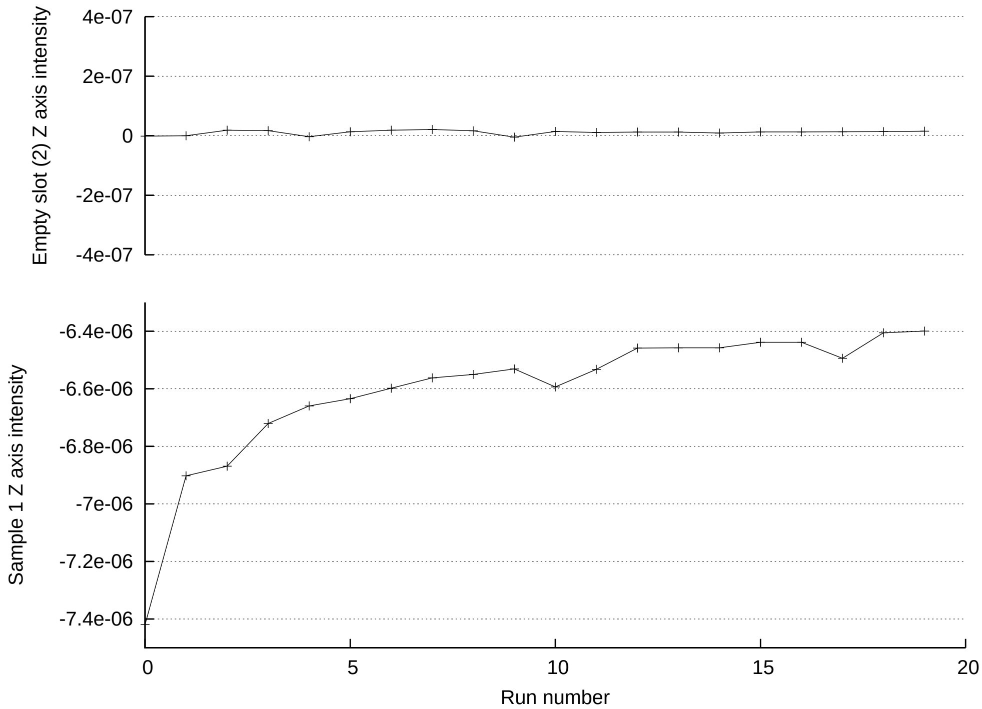

Figure 4 Successive x-axis SQUID measurements: sample 1 and adjacent empty slot.

These variations give noticeable but not disastrous variations in the measured orientation of the magnetization vector: over five degrees of declination and three degrees of inclination, as seen in Figure 5.

4.1.1 The warm-up effect: how fast is the 'cool-down'?

Figure 6 shows that recent use of the machine hardly mitigates the effect: the second run was undertaken about half an hour after the first (with no intervening use of the machine), and shows roughly the same variation in measurement.

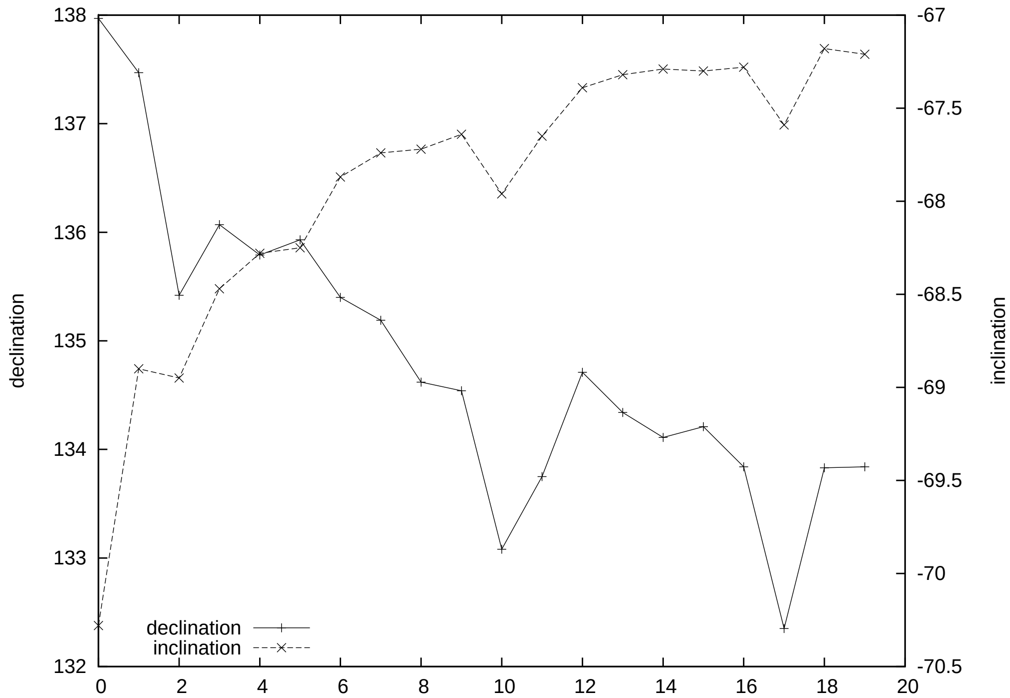

Figure 5 20 successive measurements of sample 1; orientation.

4.1.2 Number of measurements versus the warm-up effect

Increasing the number of measurements does not alleviate the warm-up effect: in fact, reducing the number of measurements to 1 seems to dispel the effect entirely – but at the cost, of course, of a huge increase in random noise.

This behaviour can be clearly seen in Figure 7, which also demonstrates another strange and pervasive effect of increasing the number of measurements: all the averaged values are significantly lower than any of the non-averaged readings. This strongly suggests that (despite suggestions to the contrary in the Long Core manual) these values are unlikely to be the result of a straightforward arithmetic mean calculation, or indeed of anything commonly denoted by the term ‘average’.

4.1.3 The warm-up effect on very weak samples

Fortunately, the effect seems almost absent on the more weakly magnetized samples; furthermore, much of the variation in measurement of these samples is correlated with variation in the measurements of the empty positions on the sample tray (see Figure 8).

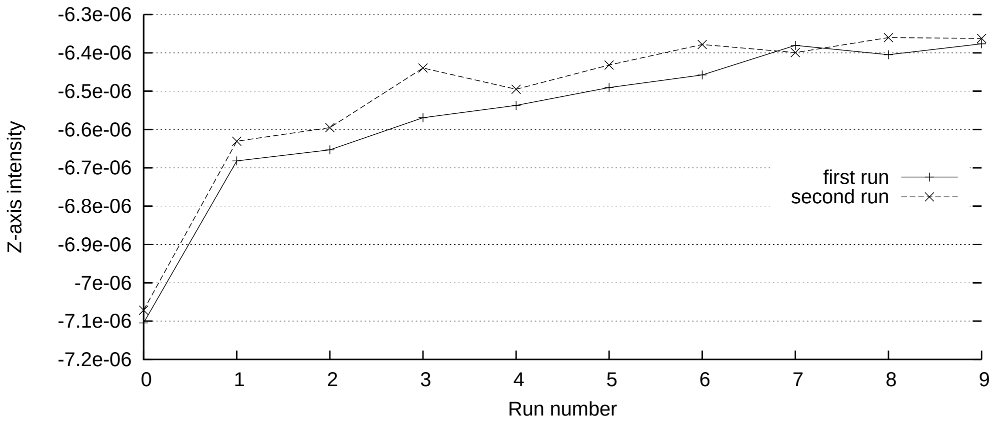

Figure 6 10 successive measurements of sample 1 on two different runs.

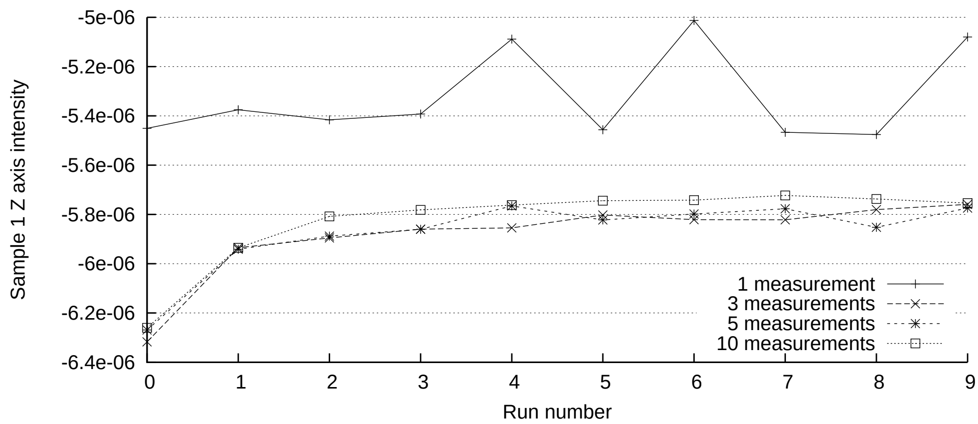

Figure 7 Successive measurements of Sample 1, with various measurements-per-reading settings.

4.1.4 Conclusions on the warm-up effect

It's worrying that there is such systematic variation in measurements for slightly stronger samples, but the effect on direction isn't too dramatic. The obvious ways to mitigate it are

1. Make several measurement runs and throw the results away before making the 'real' measurement. This would involve manually hitting the 'remeasure' button several times, or using 'autosave' for the throw-away runs. The dummy runs would then have to be filtered out of your data before further processing.

2. Reduce the number of averaged measurements to one, significantly reducing the overall precision.

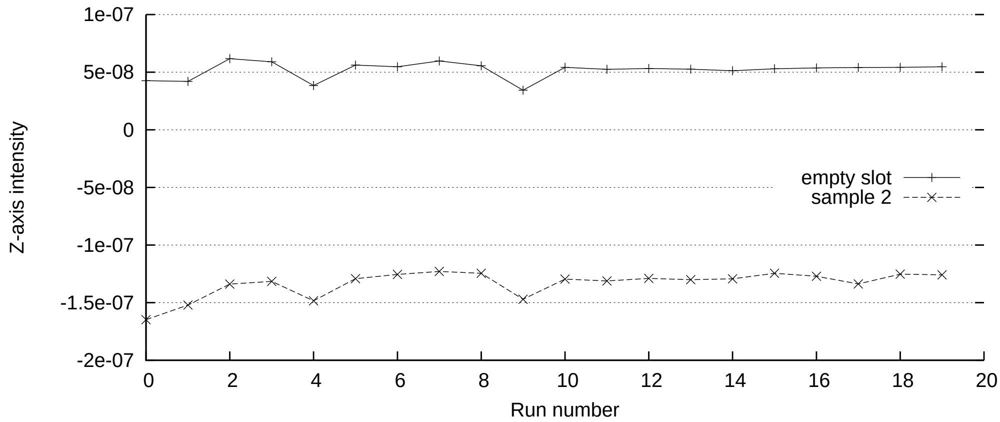

Figure 8 Successive measurements of Sample 2 (X component).

(1) would clearly be the best way to maximize precision, but at a significant cost in time and effort. In practice, the practical solution for now is to ignore it and hope (with some justification) that the errors are not too great. Since my study included only one sample with magnetization above 10^{-7} G, a more thorough investigation should be done before any major effort is put into mitigating the effect.

4.2 Empty slot correction

Using a sequence of 20 consecutive measurements, I instrumented PuffinPlot to subtract the SQUID readings for one of the empty slots from the corresponsing values for the samples. As Figure 9 shows, the results were encouraging: before applying the empty-slot correction, the measured magnetization directions ranged over 18 degrees of declination and 11 degrees of inclination; with the correction, these ranges are around 6 and 4 degrees respectively.

Brief qualitative analysis indicated that, as might be expected, an empty-slot correction seems more effective the closer the empty slot is to the measured sample. Thus, best results would be obtained by leaving every other slot empty, at the cost of halving the machine's capacity and complicating data processing. If, more practically, only one empty slot is used, it should be near the centre of the sample tray (position 3 or 4) to ensure minimal average distance from the measured samples.

5 Conclusions

If you are measuring very weak discrete samples on the cryomagnetometer, you can easily obtain a dramatic improvement in precision by taking the following two steps.

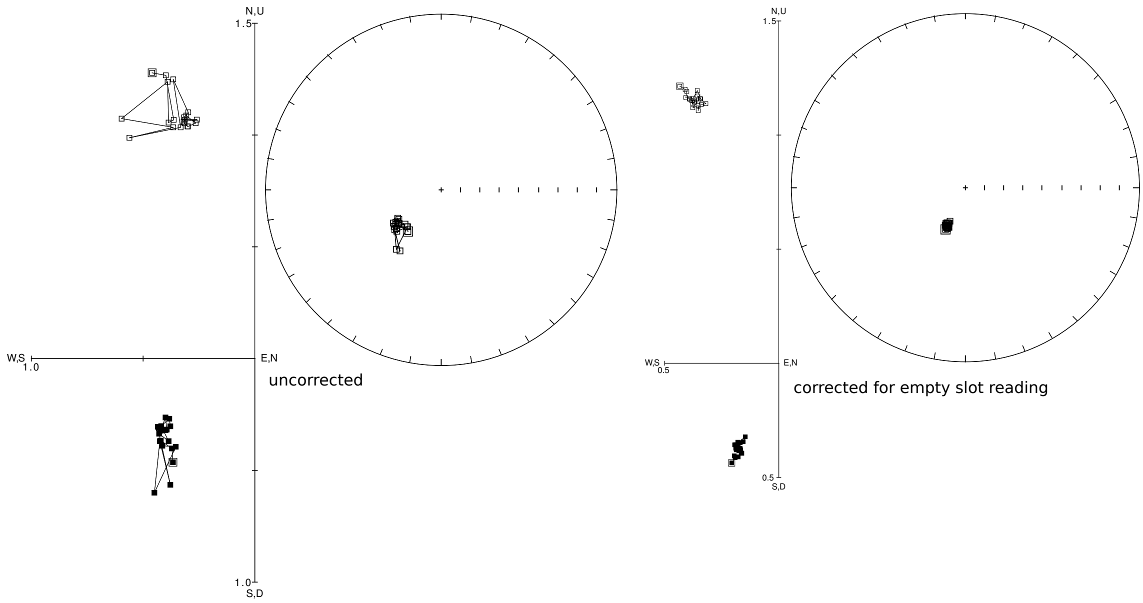

Figure 9 Effects of subtracting measured magnetization of empty slot (position 7) from that of Sample 4 (position 8).

1. Set the number of averaged measurements to at least 3.

2. Keep position 3 or 4 empty as a control on machine noise, and subtract its readings from those of your samples. PuffinPlot can do this for you.

It is likely that there are other ways to further increase the machine's precision. By applying the two practices mentioned above and varying other parameters of the magnetometer's operation, a future investigation might find effects from, say, delay time or drift correction which were too slight to be noticeable in the tests undertaken for this report.