The NIST Silicon Lattice Comparator Upgrade

Marcus H. Mendenhall, James P. Cline, Csilla I. Szabo, and Albert Henins

National Institute of Standards and Technology, U.S. Department of Commerce,

100 Bureau Dr., Gaithersburg, MD, 20899, USA[*]

Abstract

The NIST Silicon lattice comparator has been in service, in various forms, since the 1970s. It is capable of measuring the difference in lattice spacing between specimens of high-quality float-zone silicon to $\Delta d/d \approx 6{\times}10^{-9}$. It has recently undergone a thorough update of its control systems and mechanics. These upgrades result in the ability to collect data with improved stability and less settling time of the instrument, and with less operator intervention.

[*] marcus.mendenhall@nist.gov

I. BACKGROUND

The concept of using a non-dispersive double-crystal diffractometer as a lattice comparator dates at least to the work of Hart [1, 2]. The NIST lattice comparator, which operates with Laue diffraction using AgKα_1 radiation, and uses pendellösung fringes to enhance its resolution [3], has measured lattice constants with a relative error $\Delta d/d \approx 6{\times}10^{-9}$. After the "Avogadro" silicon kilogram project [4, 5] was completed, it was used for various other lattice measurements, such as those of Vaudin et al. [6]. At the end of those projects, it had a break in demand, during which a major upgrade to its control electronics and some of its mechanical systems was undertaken. This work documents those changes, along with their effect on the performance of the instrument.

The current configuration of the instrument is illustrated in figure 1. The top panel is a photograph of the instrument. The total length of the table it sits on is about 2 m. The bottom schematic (from Kessler et al. [7]) shows the relative location of the various components.

The central transfer crystal (C1) on the instrument has a thickness of 455 μm hence, to get the most intense central pendellösung fringe, the optimal thickness of a specimen in position C2 is the same. Because this thickness would result in a fairly flexible crystal, most specimens are cut with a thick base. The typical geometry is shown in figure 2; this also labels the axes as they will be described in the text below, and shows the orientation of the 220 planes from which the diffraction occurs. The stippled appearance of this crystal is the result of its etch in 80° C KOH in an ultrasonic bath to remove polishing defects and strain from the surface. This process seems to work as well as the HF/HNO_3 used on many samples, without requiring HF. It is also quite a bit slower, with better control of material removed.

II. MECHANICAL UPGRADES

The mechanical upgrades of the system include improvements to the stepper motor drives which operate the various shutters, the replacement of piezoelectric tilt (roll) system of the crystals with stepper motors driving a fine-pitch screw, and the specimen crystal mounts.

The stages on which the specimens are mounted have been significantly changed. In previous versions of the system, the crystals were mounted on a flexure which had a manual screw adjustment to coarsely adjust the roll axis of the crystal, and a piezoelectric mount under the crystal which controlled the fine adjustment. The coarse adjustment took considerable time to carry

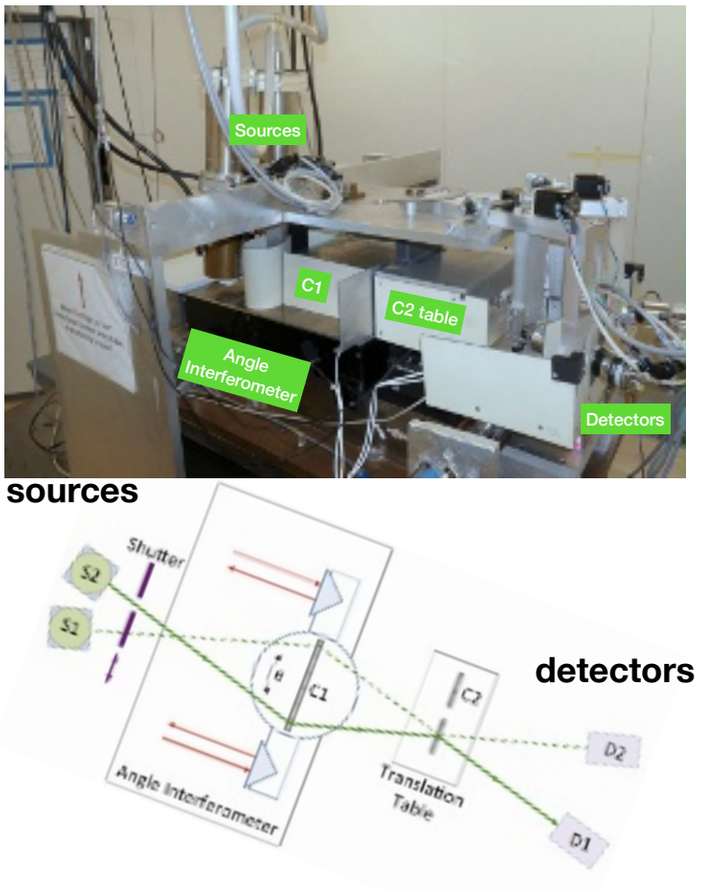

Figure 1. Top: photograph of the NIST lattice comparator. Bottom: Schematic. C1 is transfer standard crystal. Two crystals at C2 are the reference and actual specimen being measured. (non-copyrighted figure from Kessler et al., J. Res. NIST 122 (2017), altered)

out, and was generally performed using an autocollimator and a mirror on the stage, since the adjustment screws did not have very high resolution, and the piezo stage had a very small adjustment range. Furthermore, the piezoelectric stage drifted, as all such stages do when maintained at constant voltage, so it required long settling and checking to stabilize.

This system was replaced with a screw adjuster (Kozak Micro Adjusters, Randolph, NJ, USA)[8]) with 50 μm per turn (508 threads per inch) pitch. These were attached to geared stepper motors (Faulhaber MicroMo, Schönaich, Germany) with 20 steps per turn, and 16:1 gearing. This system provides a single point adjustment with a vertical lift range of many millimeters, and resolution in full-step mode for the motor of $50\,\mu\mathrm{m}/(20\,\mathrm{steps} \times 16:1\,\mathrm{gearing}) \approx 156\,\mathrm{nm}$. Since the distance between the flexure and the screw is approximately 75 mm, this corresponds

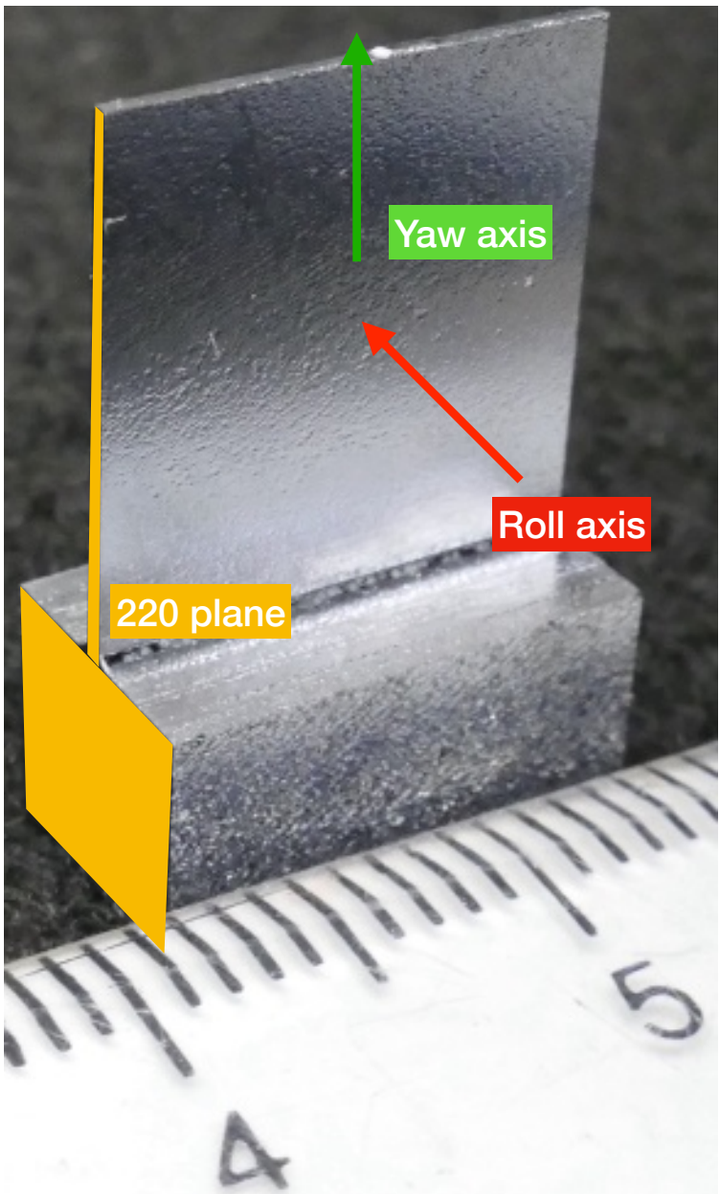

Figure 2. Typical specimen. Scale in cm.

to an angular resolution for the roll axis of 2.1 μrad. This is sufficient resolution that there is no need for any finer alignment with a piezo stage. It has been determined that this system is stable over time at a level which is sufficient for lattice comparisons ($< 10\,\mu$rad drift over one day). See figure 3. It must be noted that these steppers produce significant heat when they are operated, so the electronics driving them are set up to turn off all power shortly after a move. Under typical operating conditions, running the motors long enough to align the crystal causes a temperature increase of a few$\times 10^{-2}\,^{\circ}$C at the specimen that is allowed to cool off for at least 10 minutes before

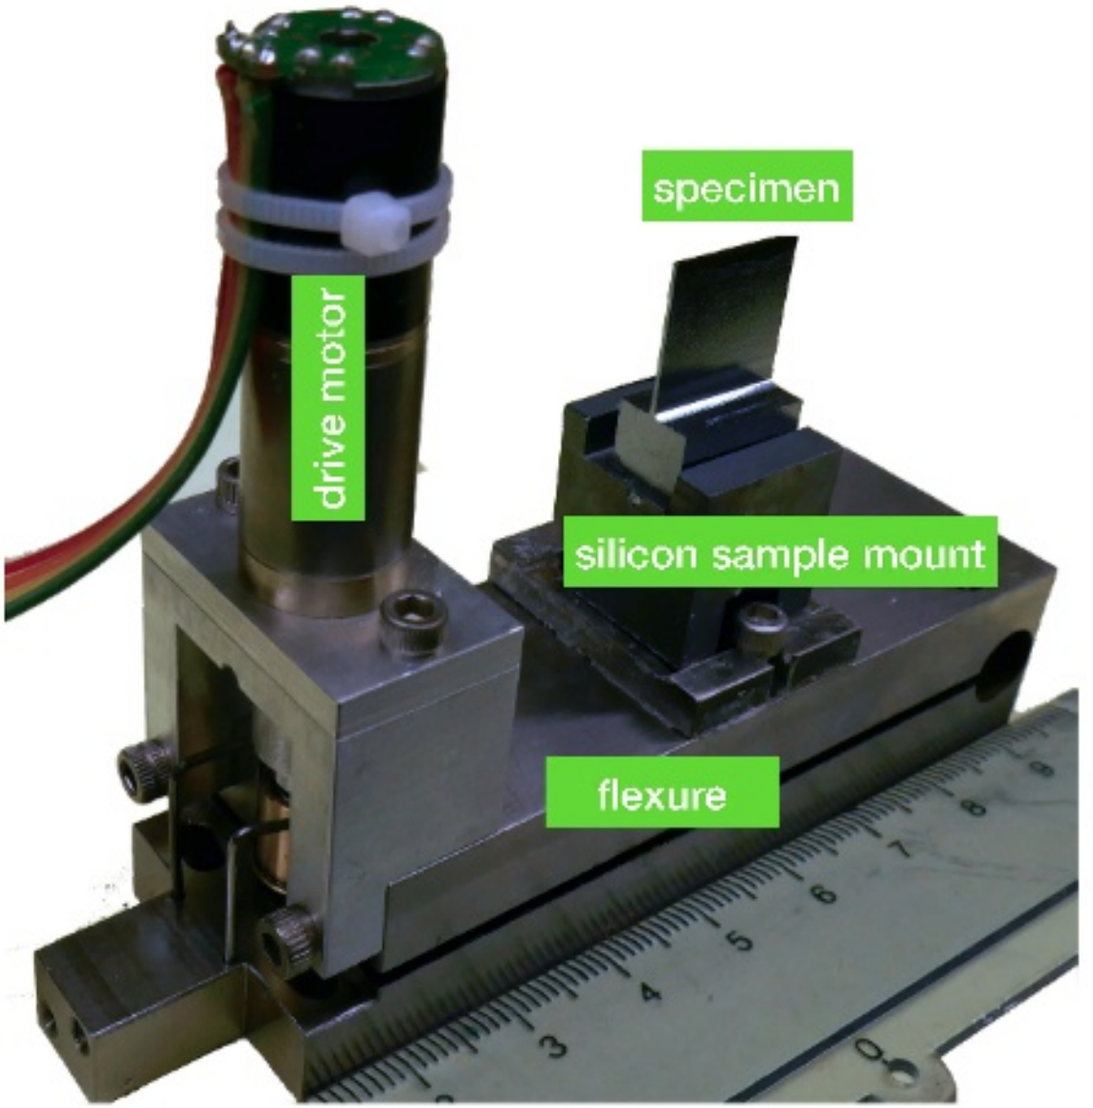

Figure 3. Flexure mount for specimen base with fine adjustment screw and stepper motor. Scale in cm.

production measurements are made.

The flexure mount is held down to the specimen translation stage with rare-earth magnets set in the stage. This allows its coarse yaw angle to be adjusted by sliding so that it pivots around a mounting screw, with the magnets providing sufficient frictional force that once the coarse alignment is complete, no tightening of screws is needed. The coarse alignment of the flexure is probably the most time-consuming part of mounting a specimen for the system. When a new crystal is mounted, it must be rotated around its vertical (yaw) axis so that the 220 planes are near the center

of the beam diffracted from C1. The total useful width of this beam is about 4000 interferometer fringes, which is 3 mrad $\approx 0.18°$. Thus, the flexure base has to be repeatedly nudged in increments smaller than this until the reflection is found, often requiring as much as 30 minutes of time.

Two major changes have been made to the mounting of the specimens on their stages. In the previous version of the system, the specimens were attached with wax to a silicon base to hold them in place without introducing strain between the specimen and the base. Now, a newly designed base, also cut from silicon, has a channel in which the specimen sits and is very gently constrained to not slide around in the channel via a bent brass shim-stock spring which presses the thick base of the crystal against one edge of the channel. This mount is shown in figure 4. This also shows the high-thermal-conductivity graphite shell around the specimen which is discussed below.

The other major change is the enclosure of the crystals in an isothermal box which assures that the thermistor which measures the temperature of the crystal is at the same temperature as the blade of the crystal. These isothermal boxes consist of three layers of very high-thermal-conductivity graphite, separated by air insulation, with an aerogel blanket on the outside of thickness 3 mm. The assembly is supported by a 3D-printed plastic frame which holds the graphite with the air gaps. Figure 5 illustrates this. The two outer layers of graphite are Panasonic EYG-S091210, which is a flexible expanded graphite with a density of approximately $1\,\mathrm{g/cm^3}$ and a thickness of $100\,\mu$m. It is specified as having an in-plane thermal conductivity in excess of $600\,\mathrm{Wm^{-1}K^{-1}}$. The inner layer of graphite is Wurth Electronics WE-TGS, which is a self-adhesive graphite bonded directly to the crystal base, such that both the specimen and the thermistor lie inside it. This material is a full-density graphite foil on a polyester base; the graphite component has a thickness of $17\,\mu$m and the polyester component has a thickness of $50\,\mu$m. The in-plane thermal conductivity is specified as exceeding $1500\,\mathrm{Wm^{-1}K^{-1}}$. Because of the low density and very low x-ray absorption of graphite, this multilayer structure does not interfere with the transmission of x-rays through the system. This structure was modeled in detail using COMSOL Multiphysics as part of a future project which will require one of the crystals to be heated or cooled; this modeling determined that even if a crystal in this enclosure differs from the ambient temperature by 70 °C, the temperature difference between the thermistor inserted in the well in the base and the blade of the crystal will be less than 0.002 °C. Details of this model will be presented in the paper on that experiment, since it is only weakly relevant for this work.

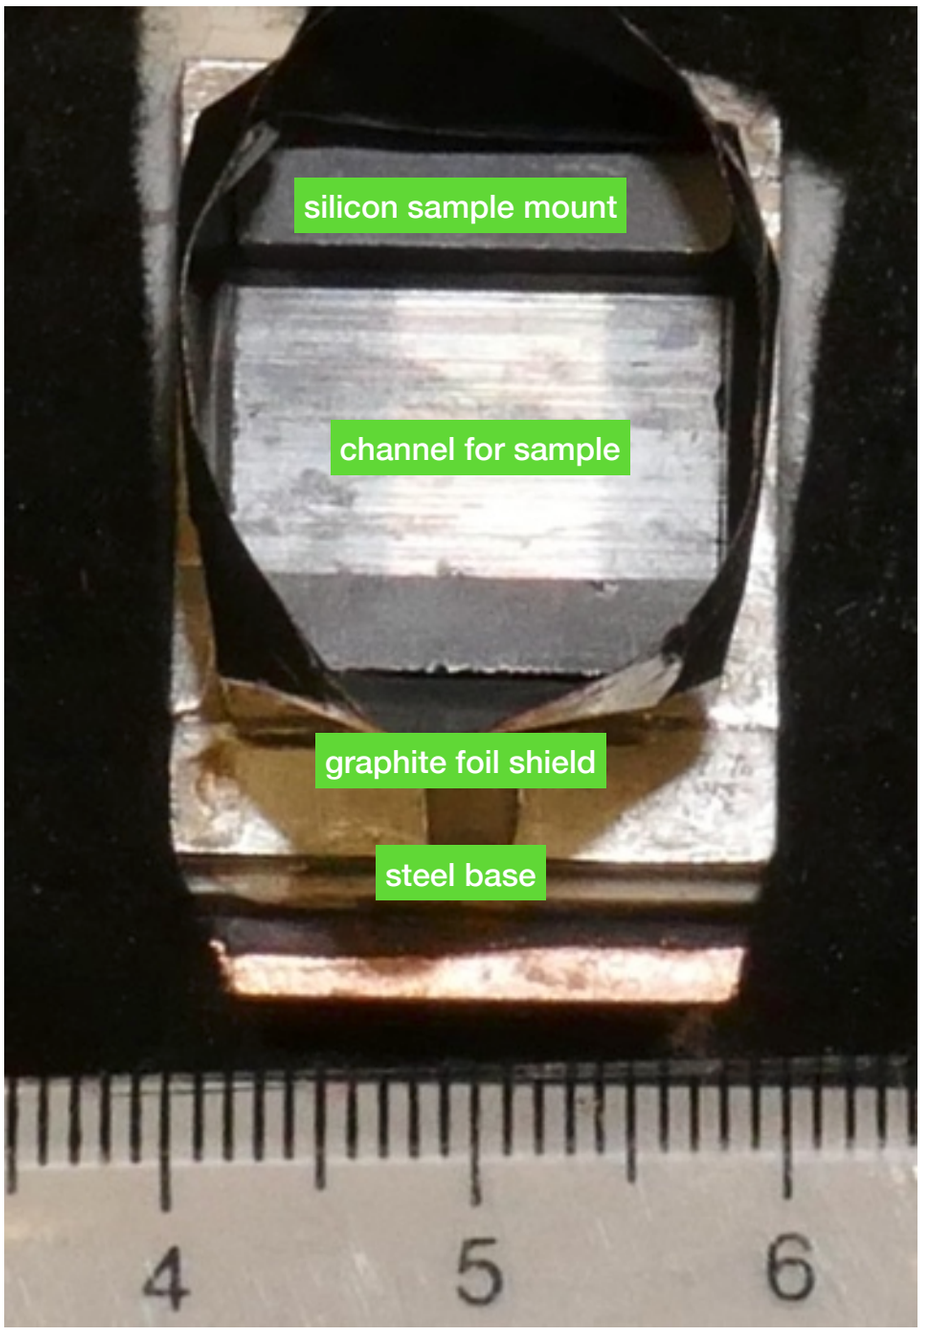

Figure 4. Silicon channel specimen mount with graphite thermal shield. Scale in cm.

III. ELECTRONICS UPGRADES

Except for the bias supplies and amplifiers for the phototubes, nearly all of the electronics on the system were replaced in this upgrade. They will be briefly enumerated here, and where

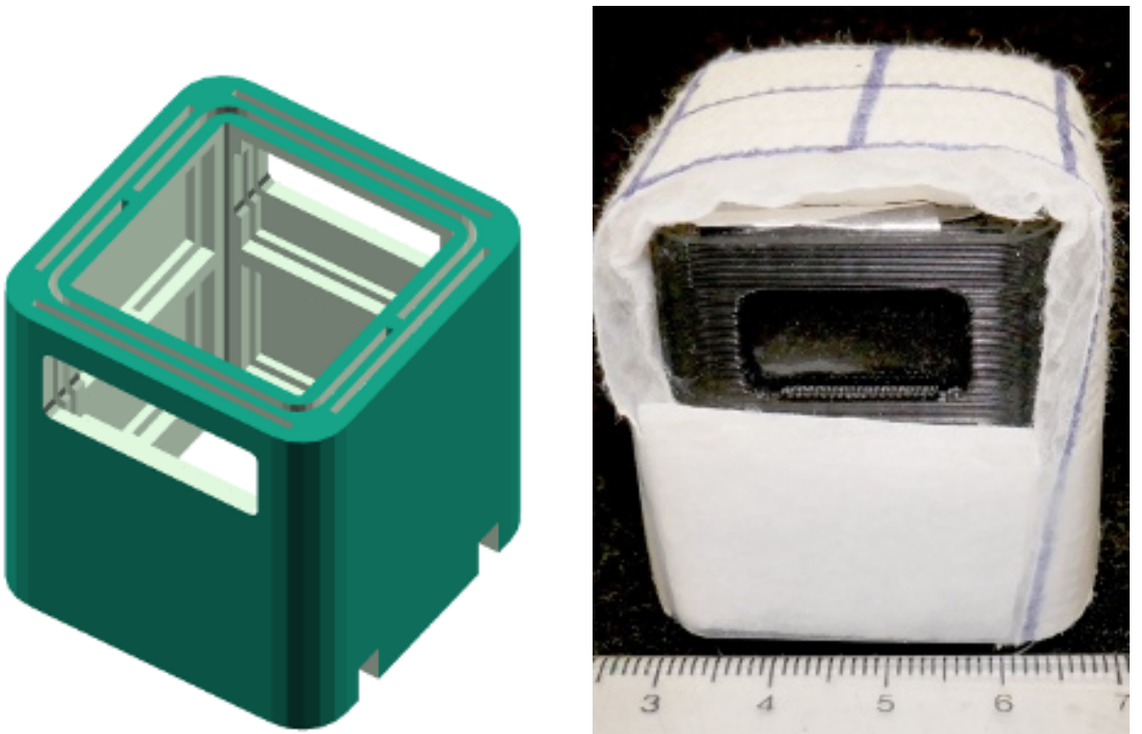

Figure 5. Isothermal specimen box. Left: 3d-printable design; right: front view with aerogel blanket. Scale in cm.

necessary discussed in detail below.

- The CAMAC-based motion control system was replaced with modern, Ethernet-interfaced stepper controllers; this does not need further discussion.

- The piezoelectric drive for the main axis rotation was replaced by an Ultravolt 1 kV bipolar power supply, driven by an Ethernet-connected analog output module.

- The laser angular interferometer system was interfaced to a purpose-built analog quadrature phase detector which converts the square-wave output of the Keysight 10780C detector modules to ± 1 V sinusoids, which are then digitized by a Heidenhain EIB741 encoder digitizer; this also is critical to the whole functioning and will be discussed below. This replaces both the PCI-card digitizer which was used for monitoring the angle in the previous versions, and the hardware phase locking electronics which regulated the axis piezo voltage to hold it at a defined position. The circuit diagram is in appendix A. This was then placed under software proportional-integral-differential (PID) control to regulate the motion of the

primary axis; this will be discussed in much more detail, since it affects the entire scanning capability of the system.

- The environmental monitoring for the system was placed under control of a stand-alone computer which does nothing but monitor the Hart (now Fluke) BlackStack temperature logger (via a Keysight E5810B Ethernet IEEE-488 interface), a USB analog input module (Measurement Computing USB-1608G) which monitors the humidity sensor and ambient light sensor, and a USB frequency counter (Measurement Computing USB-CTR04), which monitors the frequency outputs of the quartz barometric pressure sensors. This computer logs the environmental information to an sqlite database, via a Python interface, which is then queried by the computer running the rest of the system for current values. The frequency outputs of the quartz barometers is measured by counting the time for 1000 cycles to elapse, instead of counting cycles in a fixed time window, as was done previously. Since the counter's clock has a period of 21 ns, this results in a determination of the frequencies between 30 kHz and 100 kHz to $\Delta f/f \approx 2\times10^{-6}$ in < 30 ms.

A. Angular interferometer electronics

The Keysight 10780C (and equivalent) interferometer detectors produce an approximately 2 MHz square wave whose timing depends on the phase shift between the heterodyne laser signal and the signal from a fixed detector (usually within the laser). The phase detector circuit we use is very similar to that of the original lattice comparator (see figure 11 of Kessler et al. [3]), except with significantly modernized electronics. It consists of filter amplifiers to convert the square wave to a sinusoid, an IQ demodulator to generate in-phase (I) and quadrature (Q) sinusoids which track the phase shift of the laser signal and further filter and scaling amplifiers to produce clean 1 V peak-to-peak signals. This output is digitized with a Heidenhain EIB741 quad encoder digitizer.

Aside on angular measurements: because fringes are the actual quantity reported by the hardware on this system, we report many internal measurements in the following sections in fringes. Also, because of the history of the use of seconds of plane angle in high-resolution crystallography, we report and plot some data in these units, which will always be designated by the symbol " ″ ". Analyzed results will be reported according to GUM [9, 10] procedure. However, the definition of a fringe has changed from that used in previous work on this system. Previous, the unit called a fringe was one zero crossing in either the I or Q channel. The new system and software considers

a fringe to be a complete set of four quadrature crossings. Thus, papers on the old system called a fringe approximately 194.69 nrad (about $0.04''$); the new fringe is 778.76 nrad (about $0.16''$) at zero degrees. Note that the conversion of fringes to angles is not linear, since the angle interferometer measures $\sin\theta$. Most measurements are made at very low interferometer angles (< 5 mrad), so for the general discussions below, the assumption will be that $\sin\theta \approx \theta$. The precision data analysis does not make this assumption.

To obtain the necessary resolution and repeatability for the system, this fringe count must be linearly related to the angle to approximately 1 nrad (0.001 fringe). To provide this level of linearity of interpolation either requires extremely low distortion of the output sinusoids, or a method of correcting for it. We follow the method of Mendenhall et al. [11] with minor modifications for a non-periodic system. To measure the error, the digitizer is operated continuously at a rate of 10 kHz while the stage is rotated as uniformly as possible with its drive screw. Figure 6 shows the results of this measurement. The top pane shows the observed fringe count on the horizontal axis, and the same with a linear fit subtracted on the vertical axis. The large deviations are due to various effects in the drive mechanism. The bottom pane shows a subset of this result (blue curve, "raw offset", shifted to zero vertically), and also the same with a low-pass filer applied (green curve, "drive error"). The orange curve labeled "compensation" is the Fourier-windowed difference, with 10 harmonics included. It is produced by the "direct" method from section 3.5 of Mendenhall et al.. It has an amplitude of about 0.01 fringe, and is corrected during analysis. Note that if this correction is not made, the apparent lattice constant develops quite a bit of scatter, since the left-hand and right-hand beam measurements are made at different angles, and therefore sit at different positions on this curve. The center point of the diffraction pattern slowly drifts, so the offset wanders around depending on where on this compensation curve the measurement is made. In our normal analysis procedure for data, we only use harmonics 1, 2, and 4, since all the rest are within statistical noise of zero amplitude.

B. Angular interferometer calibration

The response of the angular interferometer is $\sin\theta = \zeta n$, where n is the quadrature count and ζ is the calibration constant. The constant ζ has been determined repeatedly since 2004, with the interferometer in its current configuration. This calibration is set by the spacing between the apexes of the two corner cubes on the interferometer arm. This arm is a super-invar bar, and is

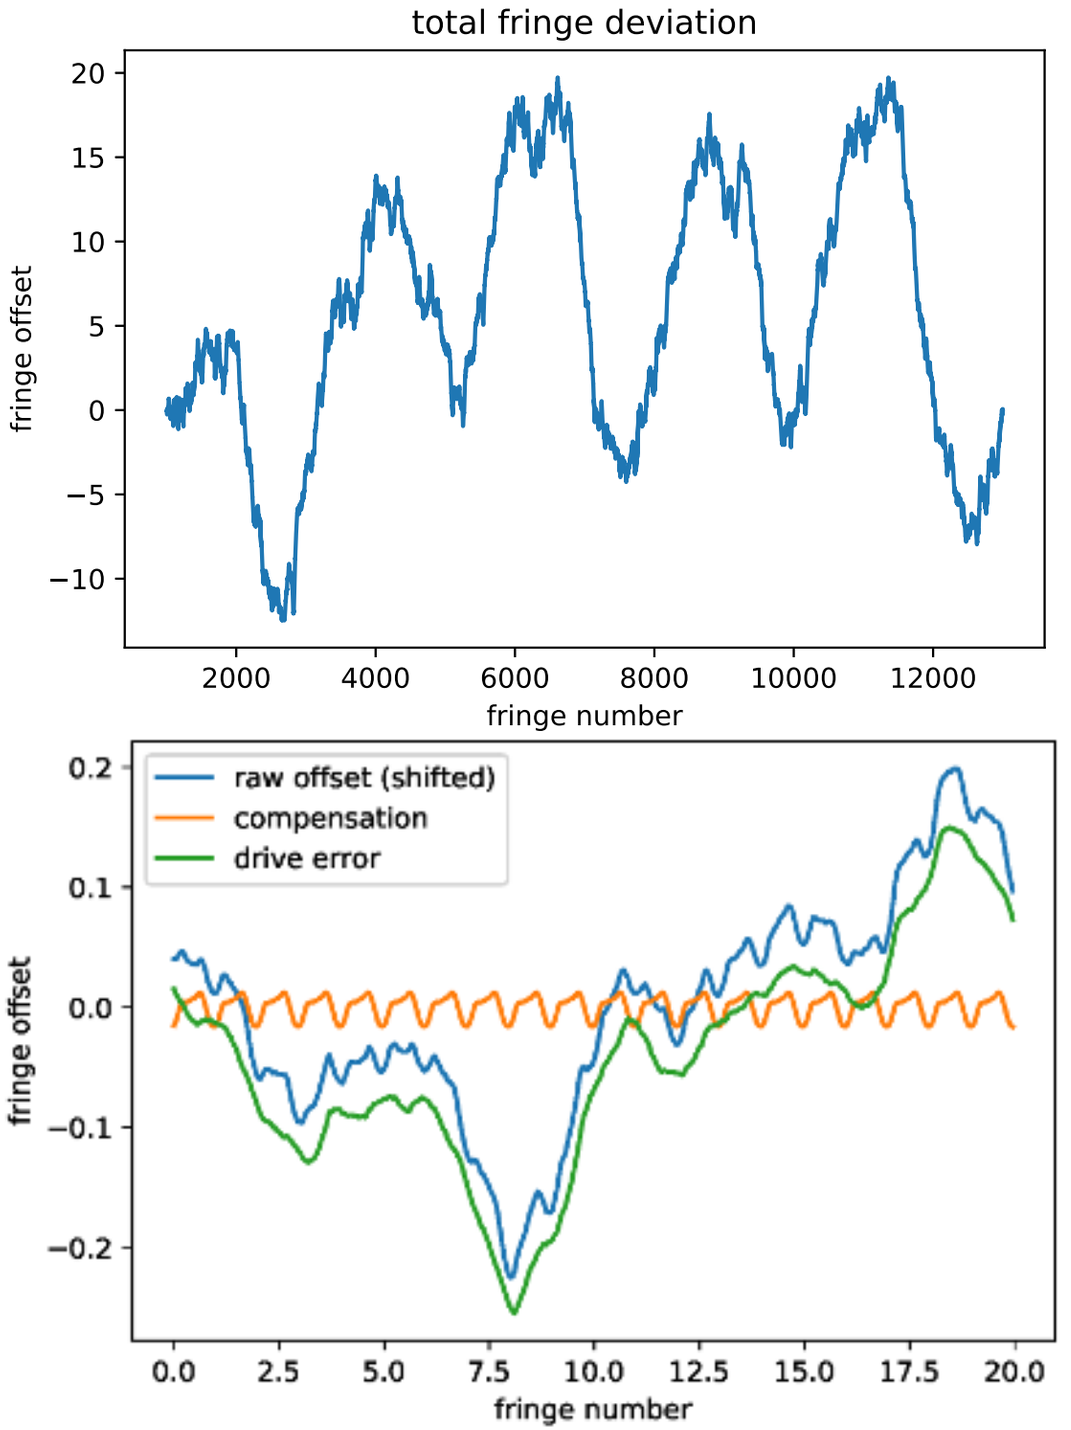

Figure 6. Data from short-period compensation measurements. Top: total fringe error; bottom: subset of data with smooth part subtracted, showing error correction signal.

stable, as shown below. The calibration procedure has been documented in Kessler et al. [3], where this constant is called K. For the purposes of this paper, all of the data sets which are available since 2004 have been reanalyzed. These were collected by Kessler and his collaborators, and have not been previously published. These use the extracted zero crossings ("Z-cross" column, see Supplementary Materials), which are corrected for atmospheric conditions, so the measurements are effectively computed at zero pressure. The code and data for this are in the supplementary data. Figure 7 shows the measurements grouped by pairs of faces on the polygon, including data from some incomplete runs. The variability between face pairs in this is the result of the slight deviations from the ideal 15° face angles on the polygon; these angles have been measured previously (see table 2 of Kinnane et al. [12]), and the errors are of the order of 15 μrad (3″). The clusters at some points in the plots are regions where repeat runs were collected; the size may be indicative of the actual random uncertainty of each measurement. A few runs from 2010 were rejected, as they had large outliers which appeared to be a result of mis-indexing the polygon faces. All the runs from 2016 include at least 2 measurements, one moving forward and one backwards, but their values are not distinguishable at the scale of this graph. There is a hint of an actual change of calibration over the time interval, but the possible drift is not enough to affect any measurements. Since the sum of all face angles must be 360°, the average over the face-by-face datasets contains no bias due to face errors, and yields a measurement of ζ independent of the characteristics of the polygon. Combining the data from 2010 and 2016, the $k = 2$ expanded type A statistical value $1/\zeta = 5\,136\,322(3)\,\text{quadratures}/\Delta\sin\theta$ (*in vacuo*) is adopted for data analysis. For comparison, Kessler et al. (1994) obtained a value of $1/\zeta = 5\,138\,552$; in between the 1994 measurement and these, the machine was repeatedly moved and rebuilt, including the replacement of a stainless steel interferometer arm with the super-invar one. As discussed in that paper, for the near-null measurements used for lattice comparison, this constant is not that critical; it will become so in a future experiment with large temperature differences and, hence, large lattice spacing differences.

IV. SOFTWARE UPGRADES

The main control computer runs a suite of interfaces written in National Instruments LabVIEW. The interfaces provide the following primary functions:

- Keep track of all the stepper motor positions;

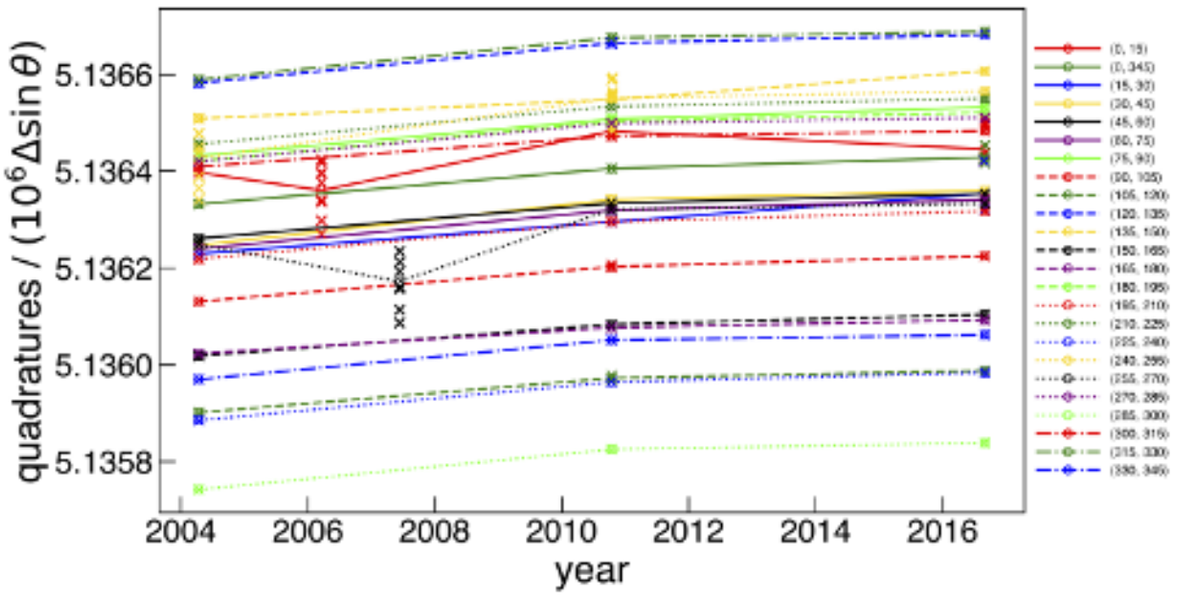

Figure 7. Calibration constant measurements, grouped by polygon face pairs. Lines connect means of groups, and are to guide the eye. Markers are individual measurements.

- Run a PID control loop to keep the main rotation axis locked to its requested position, by adjusting the voltage on the main axis piezo, based on data from the laser angular interferometer via the EIB741;
- Carry out rapid, open-loop (PID disabled) scans of the main axis to find the diffraction peak from a newly-mounted specimen;
- Communicate with environmental monitor computer to extract needed parameters;
- Scan a specimen around its roll axis while monitoring the offset of the peak position between a beam tilted above center, and one below center, to align the specimen to the C1 crystal;
- Carry out the right beam-left beam fine scans over the central pendellösung fringe, to measure the difference between the specimen's lattice spacing and that of C1;
- Carry out the right beam-left beam scans repeatedly, and switch between the specimen crystal and a reference crystal, to compare the specimen's lattice spacing to that of the reference crystal.

The PID loop and the right-left scan characteristics are the critical operations that affect the final lattice spacing measurements; these will be discussed in detail below.

A. Control of angle

The rotation angle θ_1 of the C1 crystal is the measurand most critical to the operation of this system. It must be extremely stable and repeatable, at a level of around 1 nrad to not contribute significantly to the error budget of the measurement. The coarse control of θ_1 is provided by a stepper motor rotating a shaft with a helical bearing on it. This is a system which provides quite fine control for a mechanical drive, at the level of a few μrad (a few fringes). It is somewhat non-deterministic, in that the helical bearing does not have positive engagement with a screw, but it has quite low backlash. On the other hand, the piezo pusher on this axis has a range of about ten fringes, very high resolution, and zero backlash, but it drifts when held at constant voltage. To set the system to a well-defined angle requires two processes, then.

- If the target angle is not within the scanning range of the piezo, the piezo voltage is set to zero and the stepper drive is iteratively moved until the interferometer is within a few fringes of the desired target. The decreasing series of steps in a fixed direction towards the target is necessary because of the non-deterministic nature of the mechanical drive. Once the mechanical drive has stopped and is within range, the system switches to piezo control. Figure 9 (discussed below) shows an example of this happening.

- If the angle is within the piezo scanning range, a software PID loop monitors the interferometer and controls the piezo voltage to move to the new target point. With the PID parameters in use, this results in settling to 0.001 fringe in about 0.5 s if the step is small ($\lesssim 0.01$ fringe). If the step is much bigger, it takes up to 2 s to fully settle, in addition to the time the motors spend making a move, if needed. The code used to request the angle change also monitors for stable PID output and angle readback, and delays its return until stability is verified. The very fast response for small steps is in part due to the stability check not seeing large variations that cause further delays. The data for the PID controller comes from a continuous stream of 1 kHz samples from the Heidenhain EIB741 digitizer, operating in its low-bandwidth mode. These samples are grouped in blocks of 50 samples, thus averaging over exactly 3 cycles of the 60 Hz power line, to minimize line interference. This results in a processed stream of 20 measurements each second. The PID loop integral averaging time is 0.3 s, so the stream provides sufficient bandwidth for response.

Figures 8 through 11, show the behavior of the PID loop under various conditions. These are direct

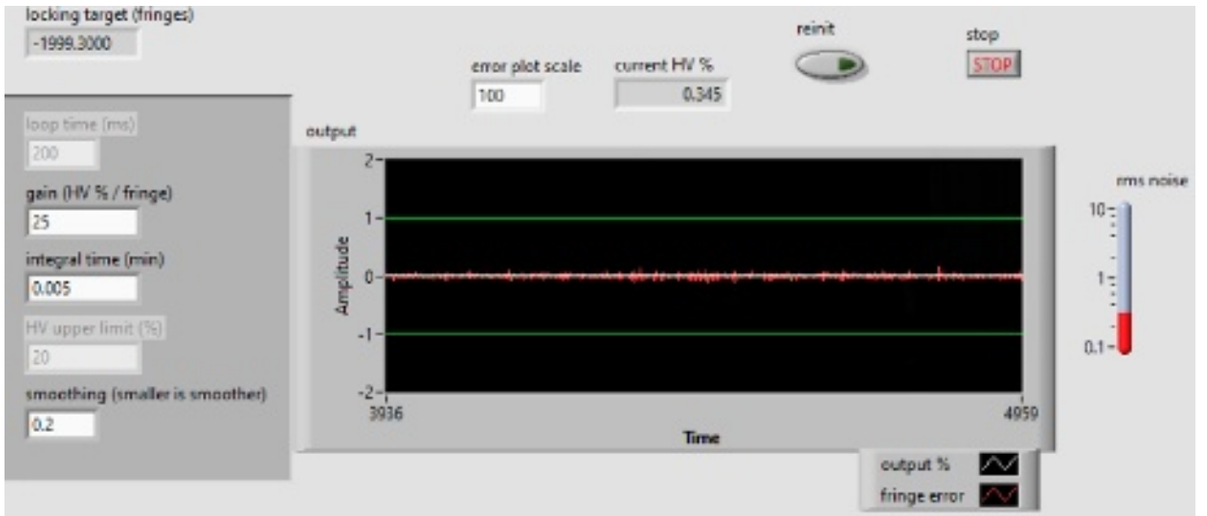

Figure 8. The PID controller when it is holding at a fixed angle. Red curve: deviation from setpoint; white curve: piezo drive amplitude. Vertical scale on red plot: 1 unit = 0.01 fringe; white plot is arbitrary percentage. Approx. 1 minute on the horizontal axis. "Thermometer": RMS deviation from last 2.5 s, 1 unit = 1×10^{-3} fringe.

screen captures of the LabVIEW panel which controls the PID. The red curve is the error signal; immediately after a step is commanded, it jumps to the full value of the step, which is off-scale on the plot except for the case of the smallest, 0.01 fringe steps. The PID then servos it back to zero via the piezo voltage. The white curve is the output to the piezo, in arbitrarily scaled units. The green lines are the limits beyond which the piezo is shut down and the motor moved. The PID parameters on the system have been tuned for the fastest settling on 0.01 fringe steps, since that is the mode the system usually runs in when scanning. Large steps are not optimized. Figure 9 shows an automatic reset of the motor position, at the place where the white (piezo output) curve jumps in the negative direction.

B. Scanning strategy

Theory

The method used to optimally scan the angle of C1 must be guided by the sources of error in the system, such that one can collect data efficiently while controlling the error to lie within necessary bounds. There are two classes of type A errors [9] which come into this system. The first is the

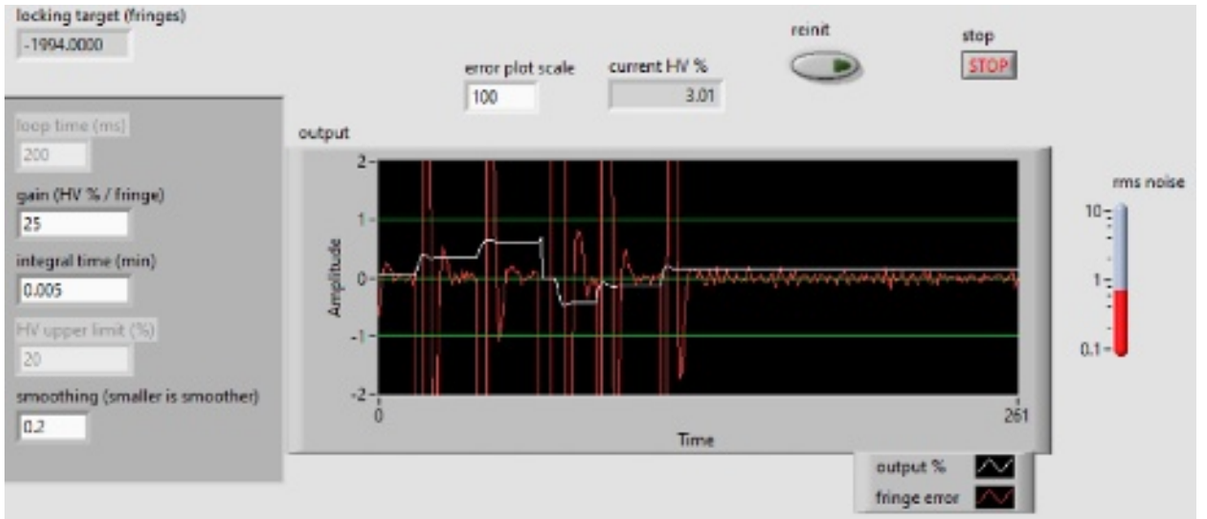

Figure 9. The PID controller when it is being commanded to step one fringe at a time at 5 s intervals. See figure 8 for details.

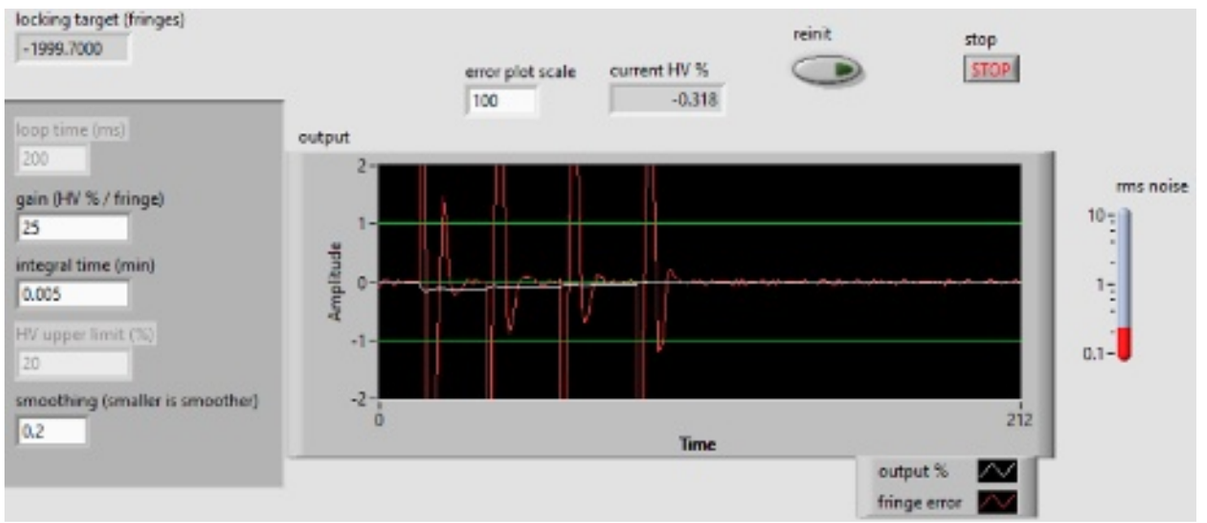

Figure 10. The PID controller when it is being commanded to step 0.1 fringe at 5 s intervals. See figure 8 for details.

Poisson counting noise from the process of detecting the signal in the first place. The only way to directly reduce this is to increase the counting time. The second source of errors is drift in the machine, resulting from temperature instability, source instability, and any mechanical drift. Since the desired outcome of a measurement is the position of the central fringe in the non-dispersive diffraction pattern, and this is obtained by a least-squares fit to the pattern, one has to evaluate the

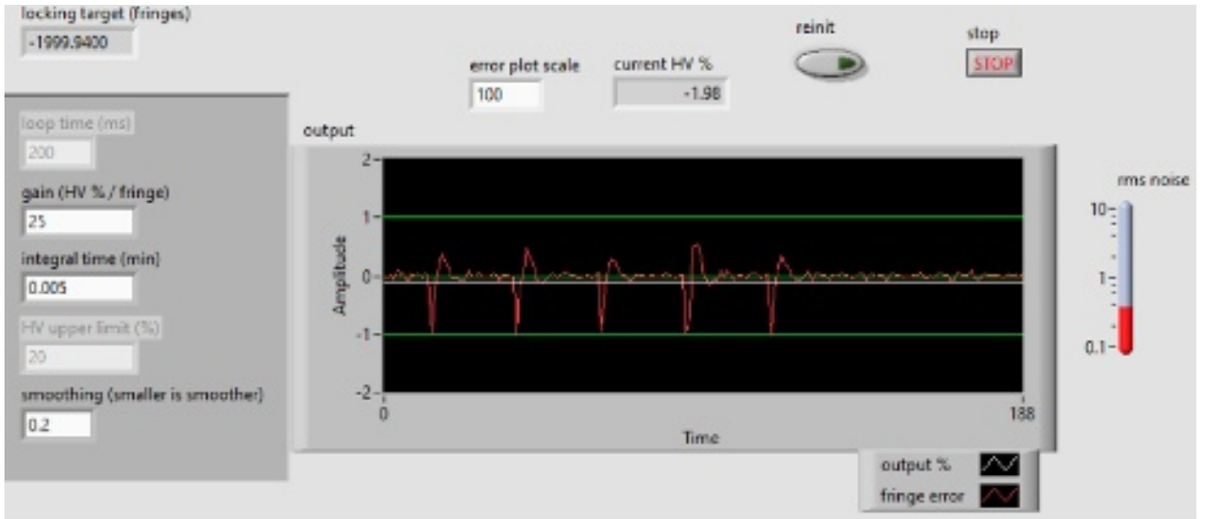

Figure 11. The PID controller when it is being commanded to step 0.01 fringe at 5 s intervals. See figure 8 for details.

effect of these source of variation on the fitting process.

To accomplish this, it is necessary to look a bit into the information theory that lies behind the processing of these signals. First, note that the important parameter to be determined is the center of the peak; the shape has nothing to do with the analysis, as long as a single careful scan has shown that the peak is well formed.

To fit a function and estimate whether the fit is well-centered on the data, one can conceptually use Taylor's theorem on $F(\vec{a};\theta+\Delta\theta)$, and fit

$$y_n(\theta_n) = F(\vec{a};\theta_n) + (\Delta\theta)\left.\frac{\partial F}{\partial\theta}\right|_{\theta=\theta_n} \quad , \tag{1}$$

where $y_n(\theta_n)$ is the measurand, $\vec{a}$ is the set of parameters defining the shape of the fit function F, and $\Delta\theta$ is a shift in the position in θ space of the best fit of the function. Now, in the least-squares process, one is attempting to adjust parameters $\vec{a}$ to minimize the sum of the squares of the residuals $r_n(\theta_n)$ and, if the function is perfectly centered, to drive $\Delta\theta$ to zero. Note, though, that the sum of the squares of the residuals χ^2 is:

$$\chi^2 = \sum\left((y_n(\theta_n) - F(\vec{a};\theta_n)) + (\Delta\theta)\left.\frac{\partial F}{\partial\theta}\right|_{\theta=\theta_n}\right)^2 \quad . \tag{2}$$

To minimize χ^2 one expands this with the binomial theorem and differentiates with respect to $\Delta\theta$

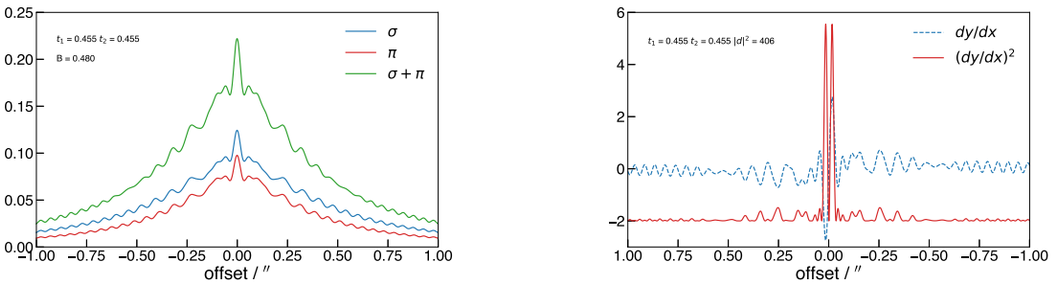

Figure 12. Theoretical diffraction patterns from the lattice comparator. Left: diffracted intensity; right: derivative. In the right-hand panel, the square of the derivative is offset vertically for clarity.

to get the extrapolated $\Delta\theta$:

$$\Delta\theta_{\text{estimate}} = -2\frac{\sum (y_n(\theta_n) - F(\overrightarrow{a};\theta_n))\left.\frac{\partial F}{\partial\theta}\right|_{\theta=\theta_n}}{\sum \left.\frac{\partial F}{\partial\theta}\right|^2_{\theta=\theta_n}} \quad . \tag{3}$$

Note that the denominator implies that working in regions of large $(\partial F/\partial\theta)^2$ yields the ability to measure very small $\Delta\theta$. Collecting data in regions in which the square of the derivative isn't large provides very little information about the position of the center of the function. In such regions, one is collecting mostly noise and instrumental error that may pull the apparent center away from its true position, and not useful information, resulting in loss of fidelity in the center of the peak.

To graphically demonstrate how this affects the typical diffraction patterns, a computation of the Laue pattern for an ideal specimen is shown in figure 12. The left-hand panel shows the contribution from the two polarization components, and the total diffracted intensity. The broad peak is the usual Laue non-dispersive pattern. The ripples are the pendellösung fringes resulting from interference of the fast and slow normal modes of the X-ray propagation. The right-hand panel shows the derivative of this function (blue dashes), and the square of the derivative (red solid). As can be seen, almost all of the information lies within the central fringe (in particular, on the two shoulders of that fringe).

Implementation

Previous iterations of the system scanned the diffraction pattern of interest to the lattice measurement by setting the instrument to a specific angle, and then measuring the intensity of the

diffraction signal from the left-hand and right-hand beams, and then stepping to the next angle. This allowed the left and right beams to be compared at very close to the same time, but resulted in fairly slow scanning due to the time for the shutters to actuate during each angular step. If the reference crystal C1 happened to have the same lattice constant as the specimen C2 (at the current temperature difference between them), the left and right peaks would be at the same angle, and the system might have reduced sensitivity to drift. There are two problems with this strategy. The first is the low data collection rate, due to dead time while the shutters are opening and closing. The second stems from the fact the C1 and C2 are almost never at the same temperature, so the peaks are significantly displaced. Then, the slow scanning means it may be quite some time (many minutes) between when the two peaks are seen. This problem will be further aggravated in an upcoming experiment in which the C2 crystal will be at a very different temperature than C1, to measure the thermal expansion.

For these two reasons, the new system uses a somewhat different strategy to scan. This strategy is to scan the entirety of the needed pattern on the left beam, then switch beams, and then scan the pattern for that beam. As discussed above, the part of the pattern that contains most of the information lies entirely inside a region of width less than $0.8\,\mu$rad ($0.16''$ or 1 fringe) around the center. Thus, the instrument can be scanned over one fringe, while collecting data every 0.01 fringe, with one second dwell time per point, in about 180 s, including the time for the instrument to step and settle. This is carried out first for the left-hand beam and then the right-hand beam. The time between the acquisition of the peaks on each side is thus less than 180 s (since the the full acquisition time includes points taken out onto the shoulder). As long as the instrument is stable during this time, it will get a correct comparison of the lattice of C2 to that of C1. Results of this process are shown in detail in section V. This strategy will also allow a rapid jump to quite different angles between the peaks, when C1 and C2 are at very different temperatures.

V. EXAMPLE DATA

A complete lattice comparison consists of a series of measurements on the specimen crystal at C2 interleaved with a series of measurements of a well-known reference crystal at C2, using C1 as the transfer standard. Typically, between 6 and 10 measurements are made of the specimen, and then the same number made of the reference. This pattern is repeated a few times. The whole process takes about 8 hours. The following figures show the report that is generated for this

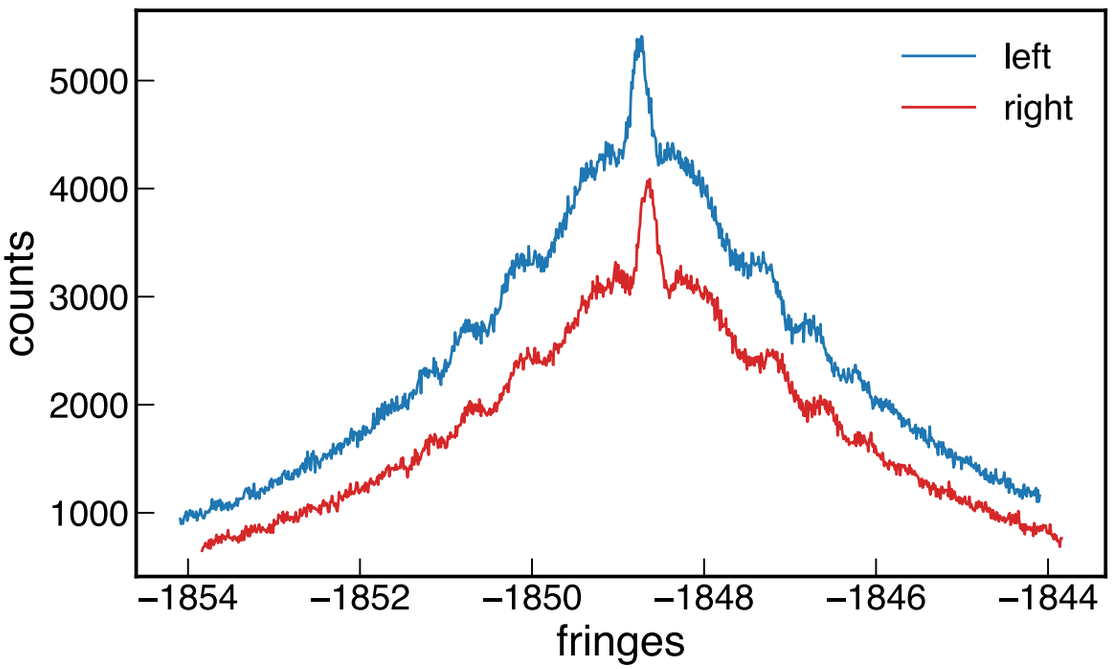

Figure 13. Typical left-hand and right-hand scans of most of the Laue pattern.

process, and how the data analysis is carried out. The series used for this was intentionally taken from a cold start on the instrument, to stress-test the ability of the temperature and drift corrections to be made. This example was run for longer than typical, to demonstrate long-term stability.

Figure 13 shows a wide scan of the diffraction pattern from the two beams. These closely match the computed pattern in figure 12. Figure 14 shows a comparison of the two vertically offset beams, which are used to determine if the C2 crystal planes have a vertical axis parallel to those of C1. In some cases, as shown in the example below, one or the other of these scans may show poor pendellösung contrast; nonetheless one can verify that the axes are aligned. It is presumed that the poor contrast results from one or the other beam being too close to the edge of the crystal, where the thickness may vary from the ideal value.

The result of a typical left-beam and right-beam scan is shown in figure 15. Note that the peaks are offset from the center because the C1 crystal is typically about 0.1 °C warmer than C2. A symmetric quartic polynomial of the form $a_0 + a_2 (x - x_0)^2 + a_4 (x - x_0)^4$ is fit to the peak. The extremum of this quartic is used as the estimate of the center of the peak. The calculation is iterated so that a first quartic fit uses a data window centered on the channel with the highest count rate, and the center of that quartic, which can be pulled around by the counting statistics in its determination

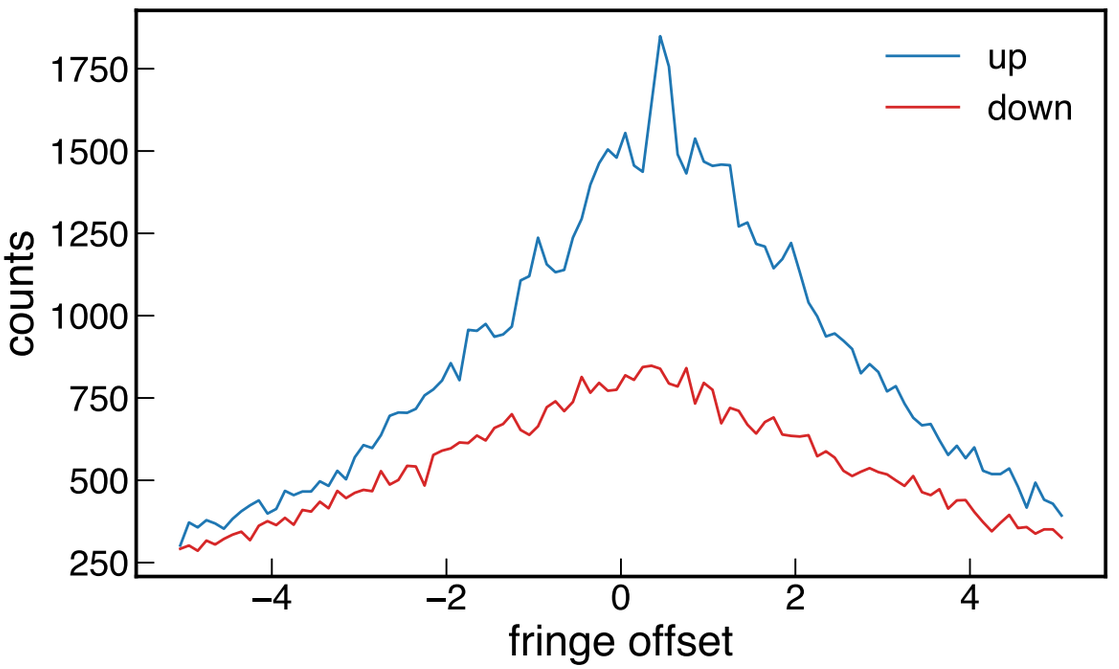

Figure 14. Slightly defective, but usable, above-center and below-center beam scans for crystal roll-axis alignment.

of the fitting window, is used to select the fitting window for the second quartic. This assures that a uniform subset of data is used for the fit which is centered over the peak. No attempt is made to propagate statistics through this part from the counting statistics, for two reasons. First, the brightest part of the pattern is not that much different in intensity than the dimmest part, so the statistics are fairly uniform. Second, the statistics will be measured directly from the scatter of repeated measurements, thus including all sources of variation.

Figure 16 shows the temperature at various points in the instrument during this acquisition series. When this instrument is fully stabilized, drift rates of about $1\times10^{-3}\,^{\circ}$C/hour can be achieved, but this requires very long settling times (days).

Figure 17 shows a critical step of the data analysis. It is observed that, over time, the center of the diffraction pattern drifts, due to thermal effects and (probably) due to mechanical drift in the linear stage which is used to place one of the two specimens in the beam. This figure plots the mean position of the left and right measurements of a crystal during a set of repeated scans. A group connected by a smooth curve represent that data taken with the linear stage parked, and then the stage is moved to the other specimen and a set of scans taken, and back again. The crosses are

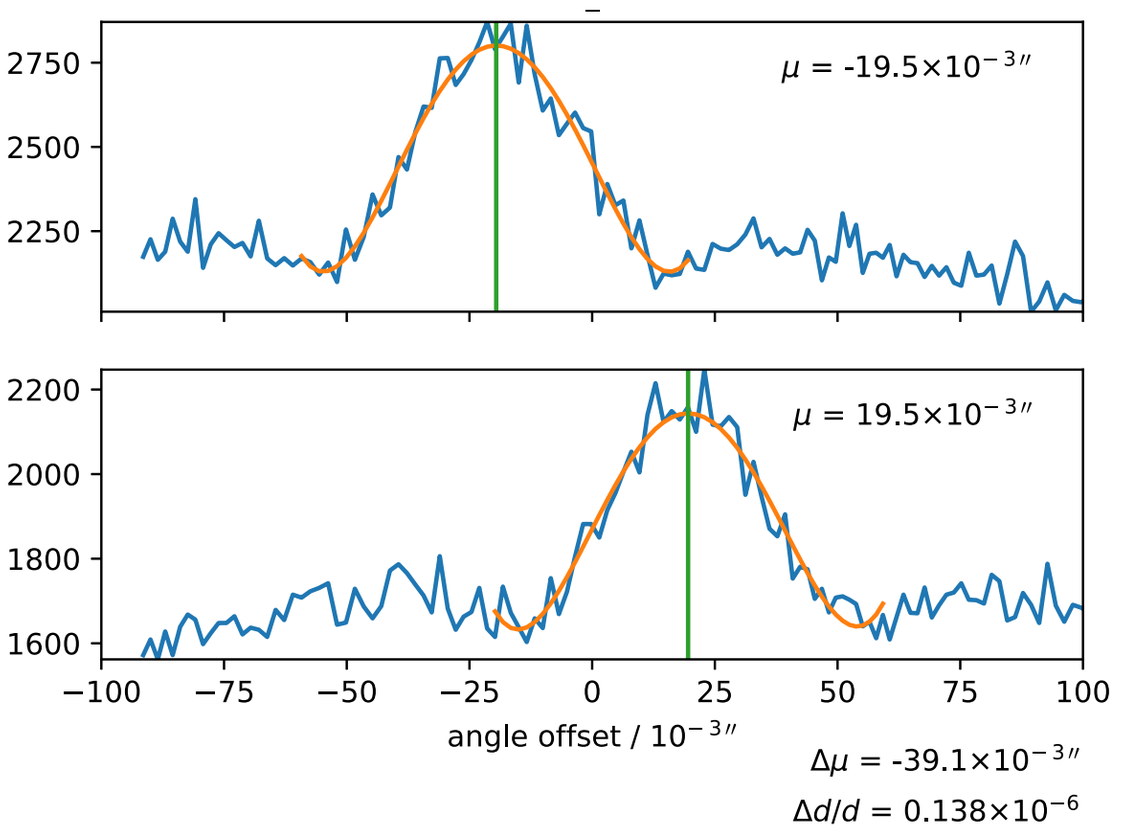

Figure 15. Typical left-hand and right-hand scans of the central fringe of a specimen. Blue: raw data; red: quartic fit; green: extremum of fit.

the measured position relative to the mean of the positions for that set; the smooth curve is a cubic fit. Note that these are very smooth, so the drift rate of the crystal is a well-defined quantity. This then allows one to note that the left-hand and right-hand scan are taken at slightly different times (a few minutes apart), and one can correct their measured position difference based on the shift of the center of the crystal at the time each peak was observed. This capability allows the machine to have a sample loaded and start analysis within a few hours. Previously, the system had to settle to a temperature drift rate of $1\times10^{-3}\,^{\circ}$C/hour for optimal performance; now, it is possible to analyze data at drift rates of $1\times10^{-2}\,^{\circ}$C/hour with minimal degradation.

Figure 18 shows the result of analyzing an entire series of measurements of both the unknown specimen crystal and a reference crystal, over an extended period of time. The blue and red symbols are the data; they can be seen to have very definite time structure at the start of each sequence, which is the result of instrumental settling. The orange curve is a smooth interpolating fit to the reference crystal data, and the green $\star$ plot is the difference between the unknown specimen

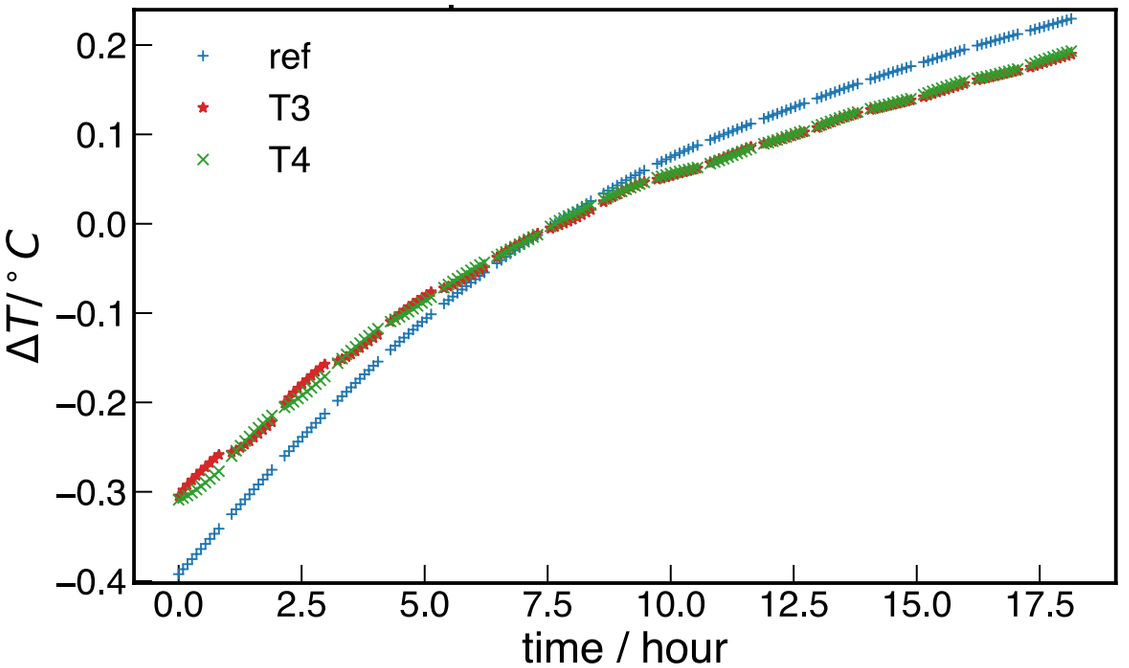

Figure 16. System temperatures during the example series. "ref" is C1; "T3" is the reference; "T4" is the specimen. The temperatures are plotted as offsets from the individual means.

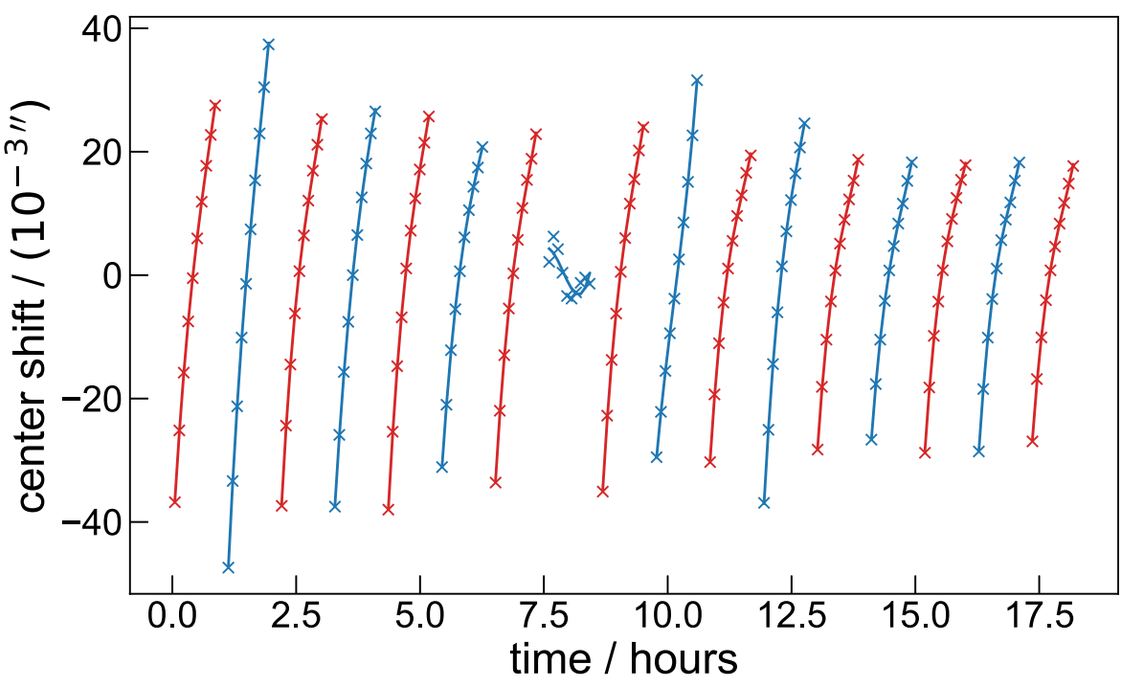

Figure 17. Center position of the pattern vs. time. red: reference scans; blue specimen scans. Crosses: data; smooth line: cubic fit.

measurements and the interpolated result from the reference crystal. The upper pane, which is not corrected for drift, the standard deviation of the measurements is 17×10^{-9} in $\Delta d/d$. The lower pane shows the result when the drift correction is applied. This reduces the standard deviation to 8×10^{-9}; it also almost entirely removes the structure from the scatter of points within each measurement set.

VI. DISCUSSION

The renovation of the NIST lattice comparator has produced an instrument with significantly improved performance, relative to its previous iteration. In particular, the thermal settling time is now of the order of hours from a cold start, instead of days. The PID-controlled piezoelectric angular stage provides a root-mean-square noise floor of 0.5 nrad. The higher scan rate possible allows many repeated samples of the peak to be made, allowing measurement of, and correction for, mechanical drift in the system, along with improved determination of the statistical variations of the result. Measuring the lattice constant of a specimen of high-resistivity FZ silicon gave a point-to-point $k = 1$ prediction interval of $\Delta d/d \approx 8\times10^{-9}$, and in an 18 hour run, 80 individual measurements of both the specimen and the reference crystal were able to be made. This results in a $k = 2$ expanded type A uncertainty of just below 2×10^{-9} on the mean. This is less than the absolute uncertainty in the lattice constant of the reference crystal, so this measurement could have been made in a shorter time, with much less data collected. Nonetheless, given the time it takes to set up a sample, it is likely that for this purpose, one specimen can be completely characterized in a day.

There are a number of upcoming applications for the instrument. It is currently being fitted with a temperature control stage that will allow the specimen to be operated between $-30\,^{\circ}$C and $70\,^{\circ}$C, while the reference crystal is held at room temperature. This will allow a determination of the lattice constant as a function of temperature. Another sample stage is in progress which will be used for the much simpler task of measuring the miscut angle on flat, polished silicon specimens.

VII. SUPPLEMENTARY MATERIAL

All the data and code used to generate the figures in this paper are provided as supplementary materials. A detailed description of these files is in the supplementary file

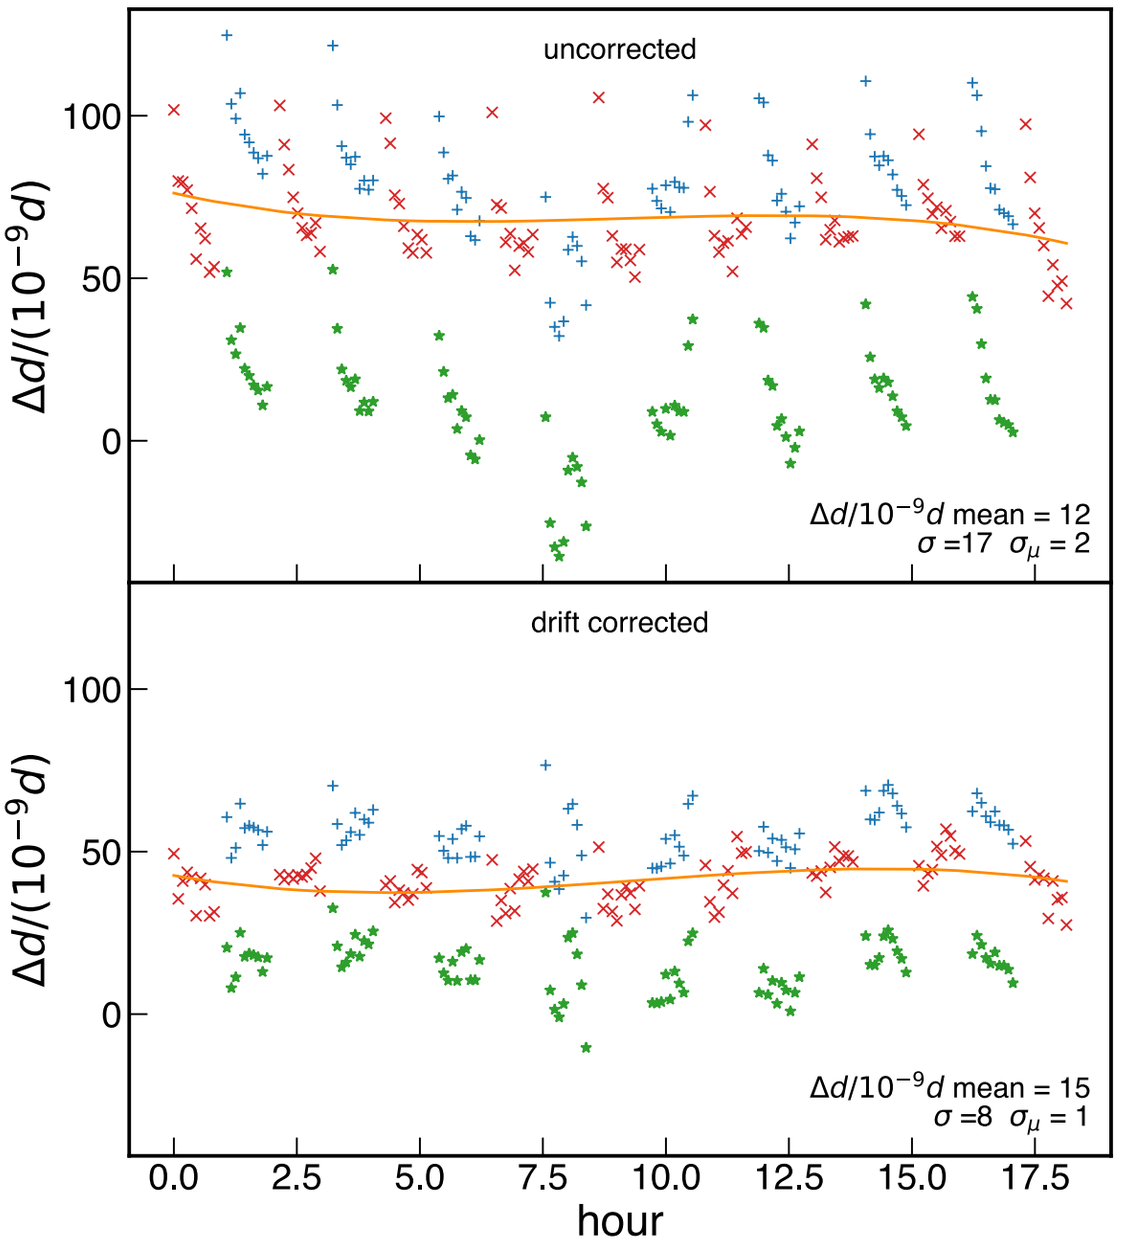

Figure 18. Apparent $\Delta d/d$; Red $\times$: reference; blue $+$: specimen; green $\star$: difference; orange curve: smooth fit to reference. Top: uncorrected for instrumental drift; bottom: corrected.

"README_supplementary_material.txt".

ACKNOWLEDGMENTS

Official contribution of the U.S. Department of Commerce National Institute of Standards and Technology; not subject to copyright in the United States. Funded by internal DoC operations money. Data for the interferometer calibrations was collected by Ernest Kessler and Mark Vaudin. Silicon material preparation was carried out in the NIST Center for Nanoscale Science and Technology (CNST) (`https://www.nist.gov/cnst`).

Appendix A: Phase detector

The 4-channel phase detector circuit diagram is shown in figure 19. As of now, only one channel is in use, but this provides future expansion for a 3-channel tip/tilt/translation sensor which has been used for other experiments. The version in use was designed using the gEDA tools, but a newer version of the schematic and layout was completed using KiCAD, which is significantly more modern in design. The KiCAD project for this is provided in the supplementary data. Figure 20 shows the board as assembled and mounted in its box; the photo is oriented the same way as the schematic. The left-hand side shows the row of coupling transformers that provide full galvanic isolation of the system. The large metallic unit is the quadrature splitter that provides the IQ reference for the two phase detectors on each channel.

DATA AVAILABILITY

All data are available on the NIST MIDAS site at `https://doi.org/10.18434/mds2-2913`. Code to analyze the data should run in any python3 installation newer than 3.9 which supports NumPy (or SciPy) and matplotlib. Data are also included as supplementary data to the journal.

[1] Michael Hart and Frederick Charles Frank. High precision lattice parameter measurements by multiple Bragg reflexion diffractometry. *Proc. Roy. Soc. A: Mathematical and Physical Sciences*, 309(1497): 281–296, 1969. doi:10.1098/rspa.1969.0042.

[2] M. Ando, D. Bailey, and M. Hart. A simple Bragg-spacing comparator. *Acta Cryst. A*, 34(4):484–489, 1978. doi:10.1107/S0567739478001047.

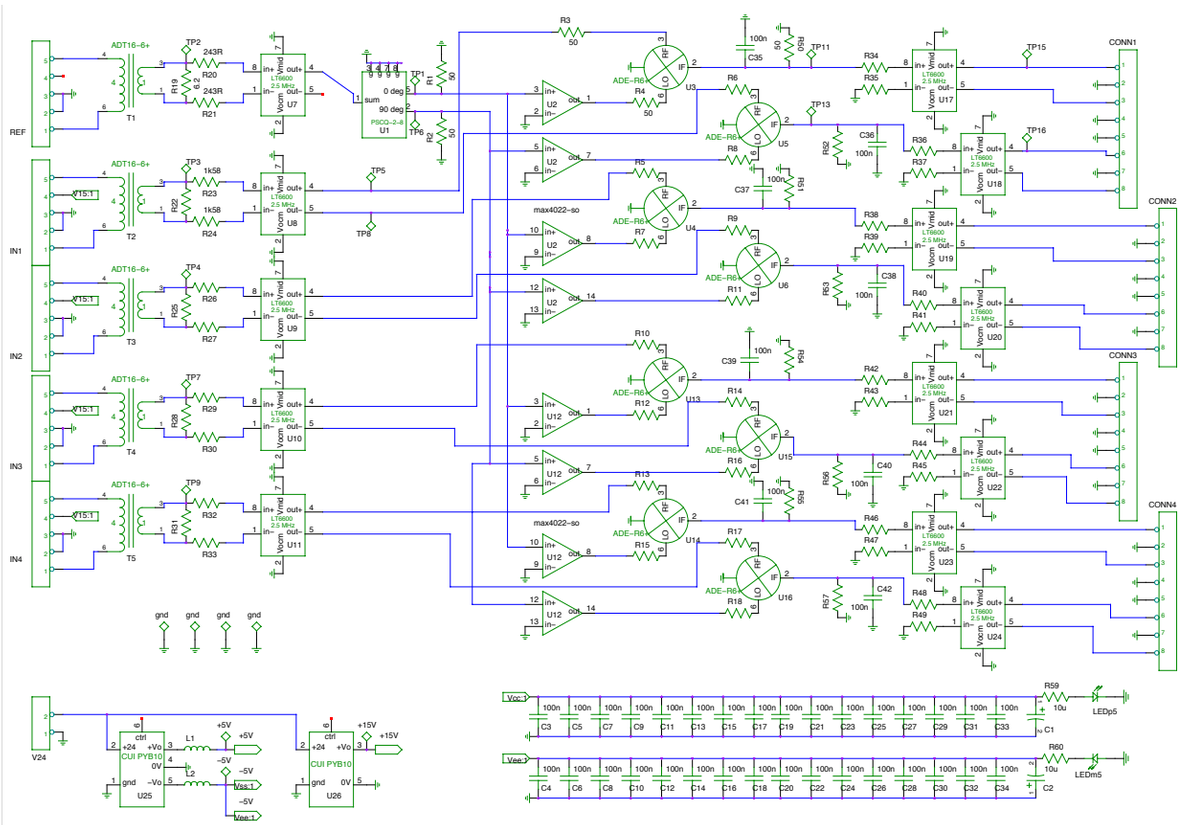

Figure 19. Circuit schematic for 4-channel phase detector.

[3] E. G. Kessler, A. Henins, R. D. Deslattes, L. Nielsen, and M. Arif. Precision comparison of the lattice parameters of silicon monocrystals. *J. Res. NIST*, 99(1):1–18, Jan 1994. doi:10.6028/jres.099.002.

[4] E. G. Kessler, S. M. Owens, A. Henins, and R. D. Deslattes. Silicon lattice comparisons related to the Avogadro project: uniformity of new material and surface preparation effects. *IEEE Trans. Instr. and Meas.*, 48(2):221–224, 1999. doi:10.1109/19.769568.

[5] E Massa, G Mana, L Ferroglio, E G Kessler, D Schiel, and S Zakel. The lattice parameter of the ^{28}Si spheres in the determination of the Avogadro constant. *Metrologia*, 48(2):S44, 2011. doi: 10.1088/0026-1394/48/2/S07.

[6] M. D. Vaudin, E. G. Kessler, and D. M. Owen. Precise silicon die curvature measurements using the NIST lattice comparator: comparisons with coherent gradient sensing interferometry. *Metrologia*, 48 (3):201–211, 2011. doi:10.1088/0026-1394/48/3/015.

[7] E. G. Kessler, C. I. Szabo, J. P. Cline, A. Henins, L. T. Hudson, M. H. Mendenhall, and M. D. Vaudin. The lattice spacing variability of intrinsic float-zone silicon. *J. Res. NIST*, 122, May 2017.

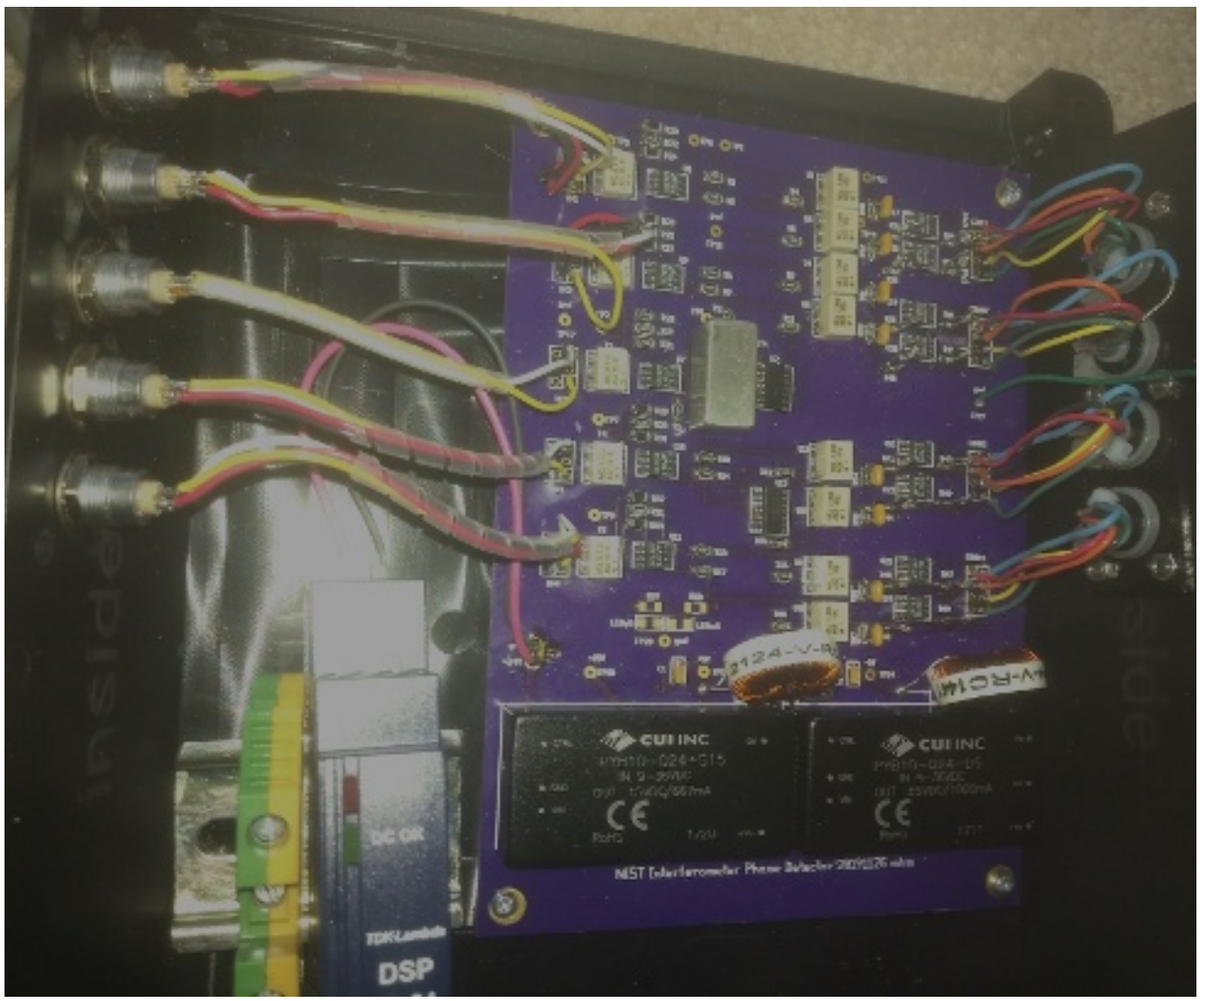

Figure 20. Photo of assembled and mounted phase detector.

doi:10.6028/jres.122.024.

[8] Note1. Certain commercial equipment, instruments, or materials are identified in this paper in order to specify the experimental procedure adequately. Such identification is not intended to imply recommendation or endorsement by the U.S. government, nor is it intended to imply that the materials or equipment identified are necessarily the best available for the purpose.

[9] B.N. Taylor and C.E. Kuyatt. TN1297: Guidelines for evaluating and expressing the uncertainty of NIST measurement results. Technical Report 1297, NIST, 1994. URL `http://physics.nist.gov/cuu/Uncertainty/index.html`.

[10] JCGM. Uncertainty of measurement–part 3: Guide to the expression of uncertainty in measurement (JCGM 100:2008, GUM:1995). Technical Report 100:2008, Joint Committee for Guides in

Metrology, 2008. URL http://www.iso.org/sites/JCGM/GUM-introduction.htm.

[11] Marcus H. Mendenhall, Donald Windover, Albert Henins, and James P. Cline. An algorithm for the compensation of short-period errors in optical encoders. *Metrologia*, 52(5):685–693, September 2015. doi:10.1088/0026-1394/52/5/685.

[12] Mark N. Kinnane, Lawrence T. Hudson, Albert Henins, and Marcus H. Mendenhall. A simple method for high-precision calibration of long-range errors in an angle encoder using an electronic nulling autocollimator. *Metrologia*, 52:244–250, Mar 2015. doi:10.1088/0026-1394/52/2/244.

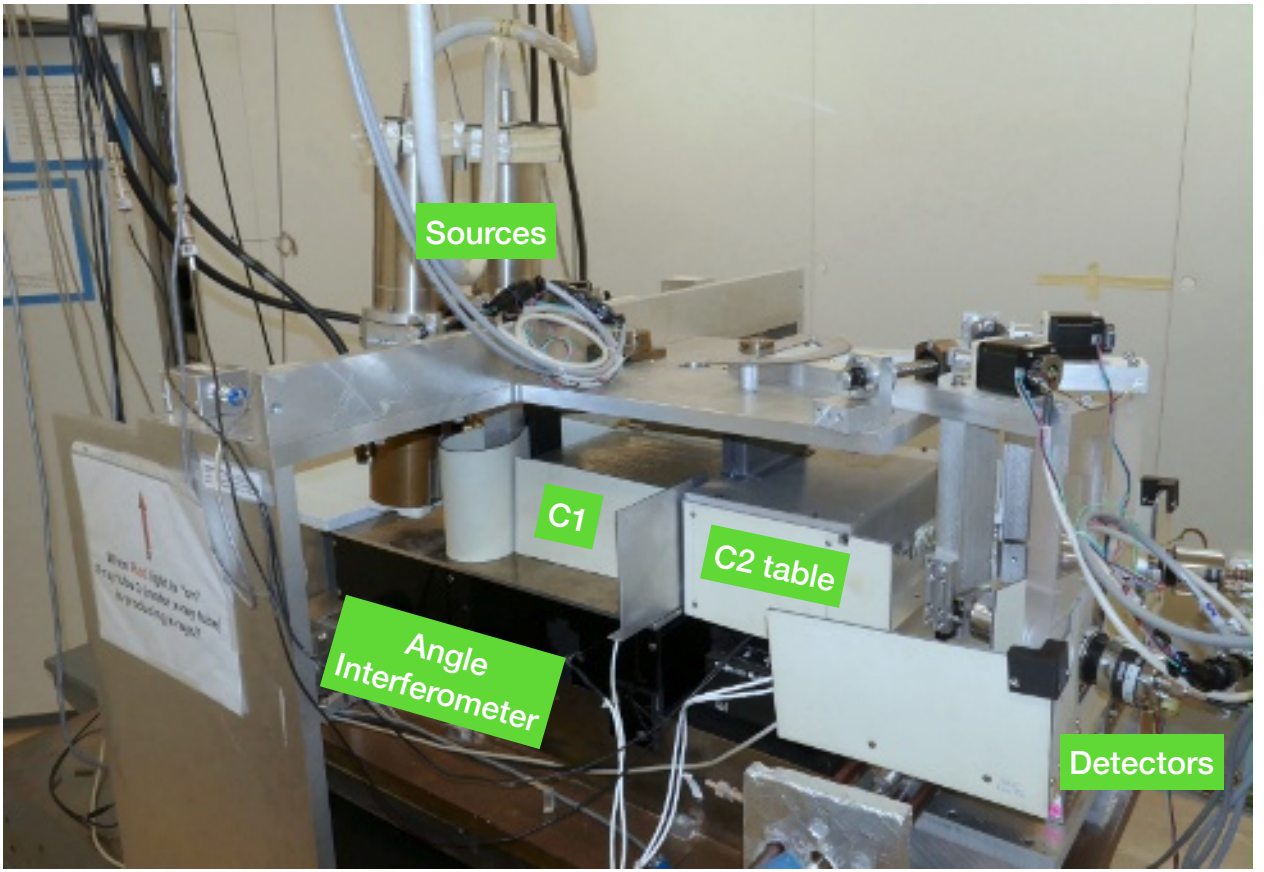
Sources
C1
C2 table
Angle Interferometer
Detectors

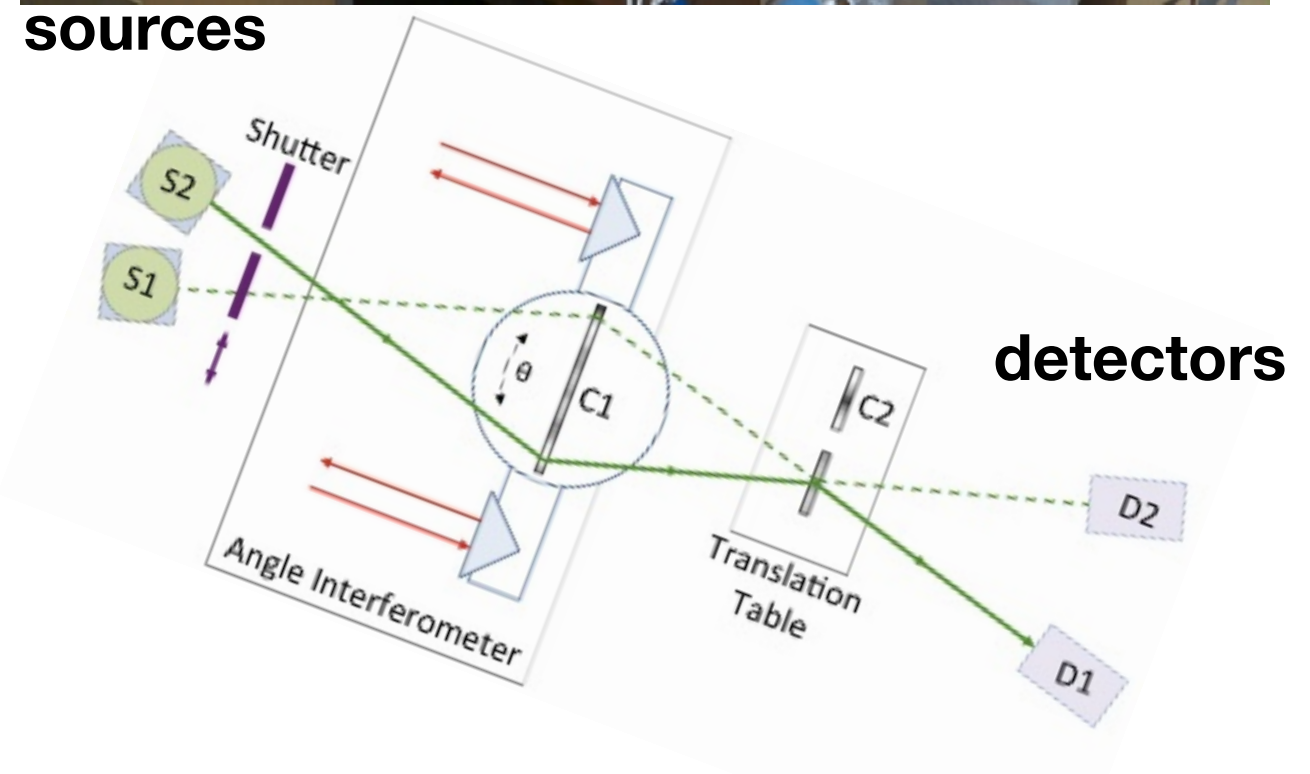
sources
Shutter
S2
S1
θ
C1
C2
detectors
D2
D1
Translation Table
Angle Interferometer

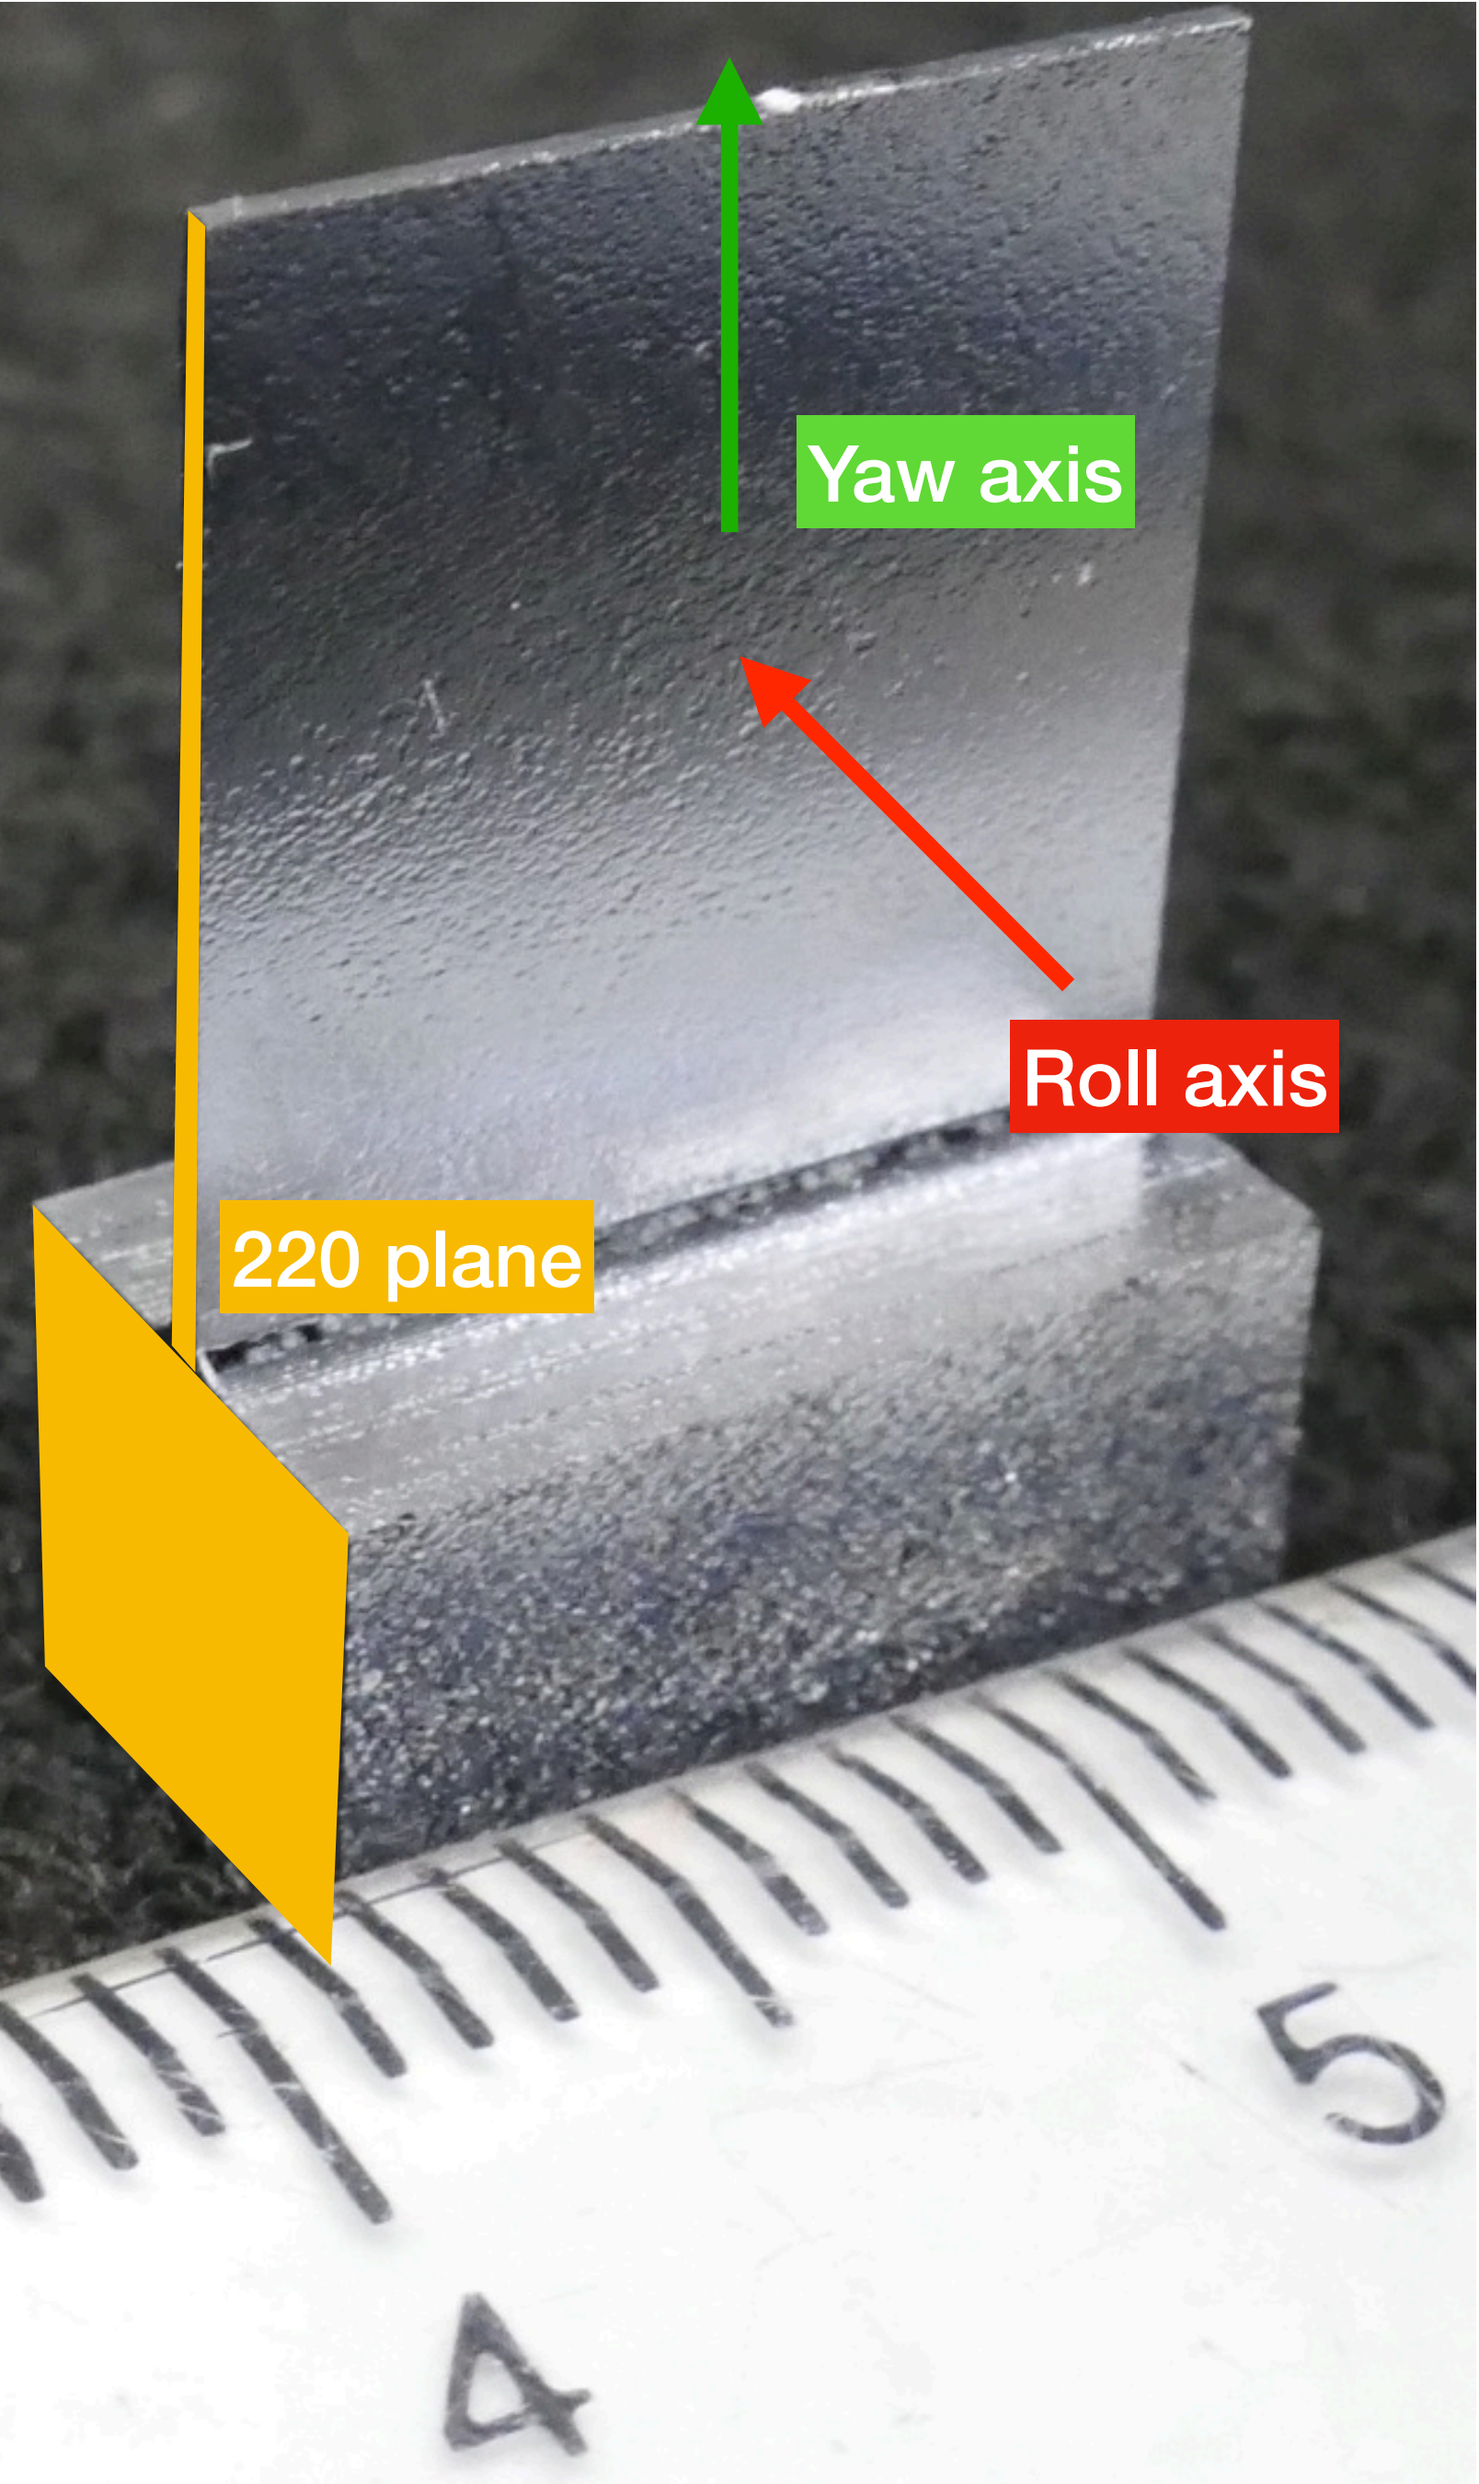
Yaw axis
Roll axis
220 plane
5
4

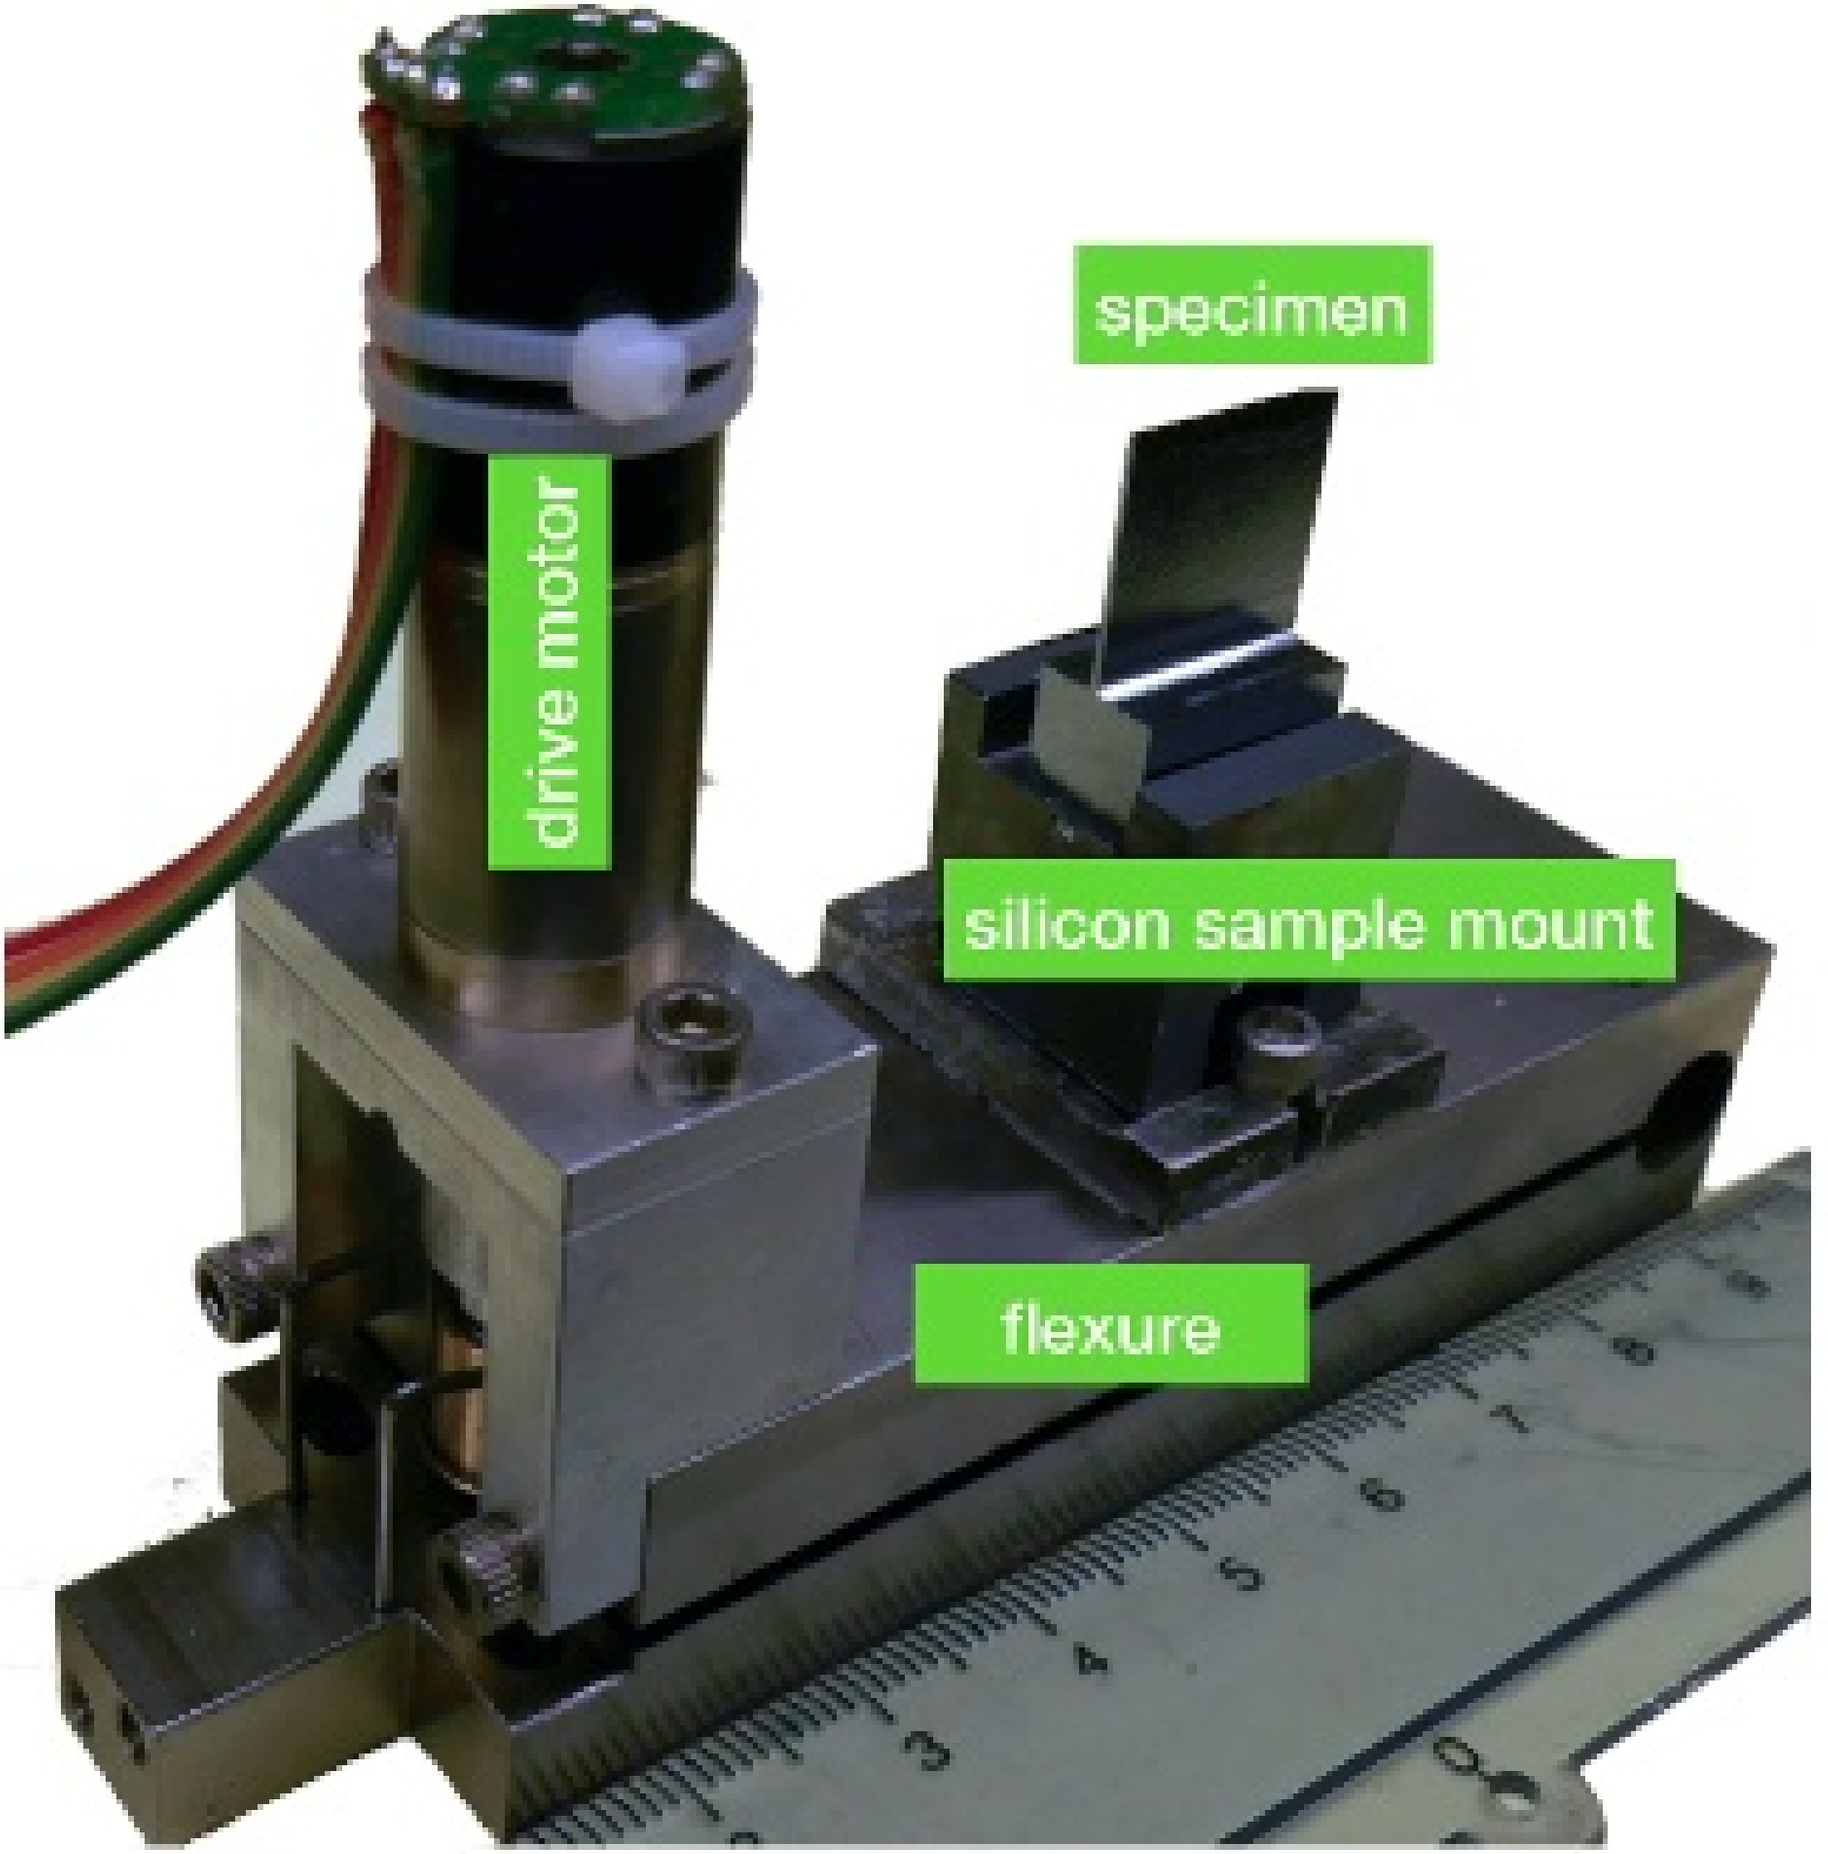
specimen
drive motor
silicon sample mount
flexure

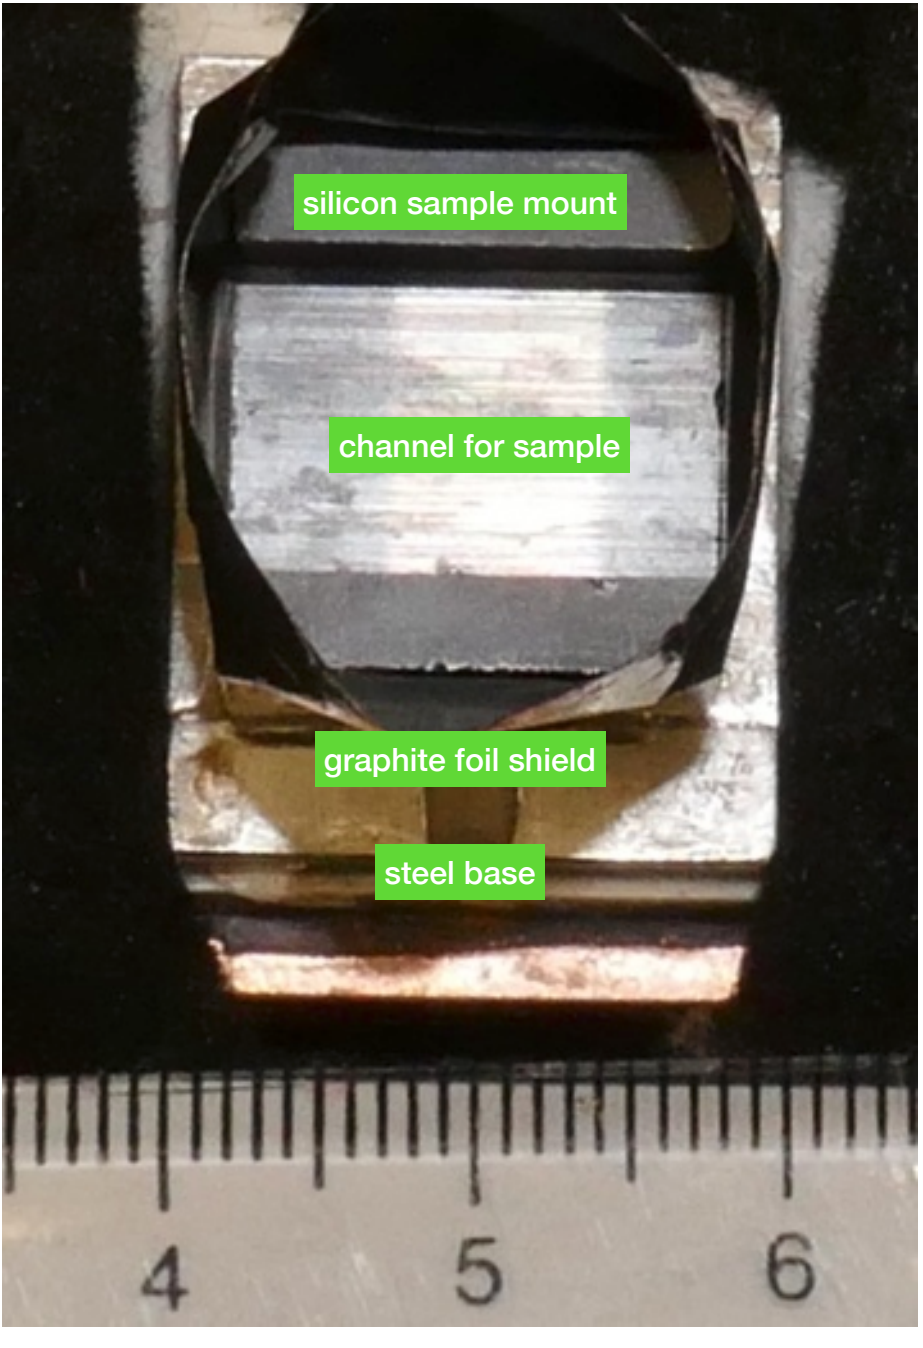
silicon sample mount
channel for sample
graphite foil shield
steel base
4
5
6

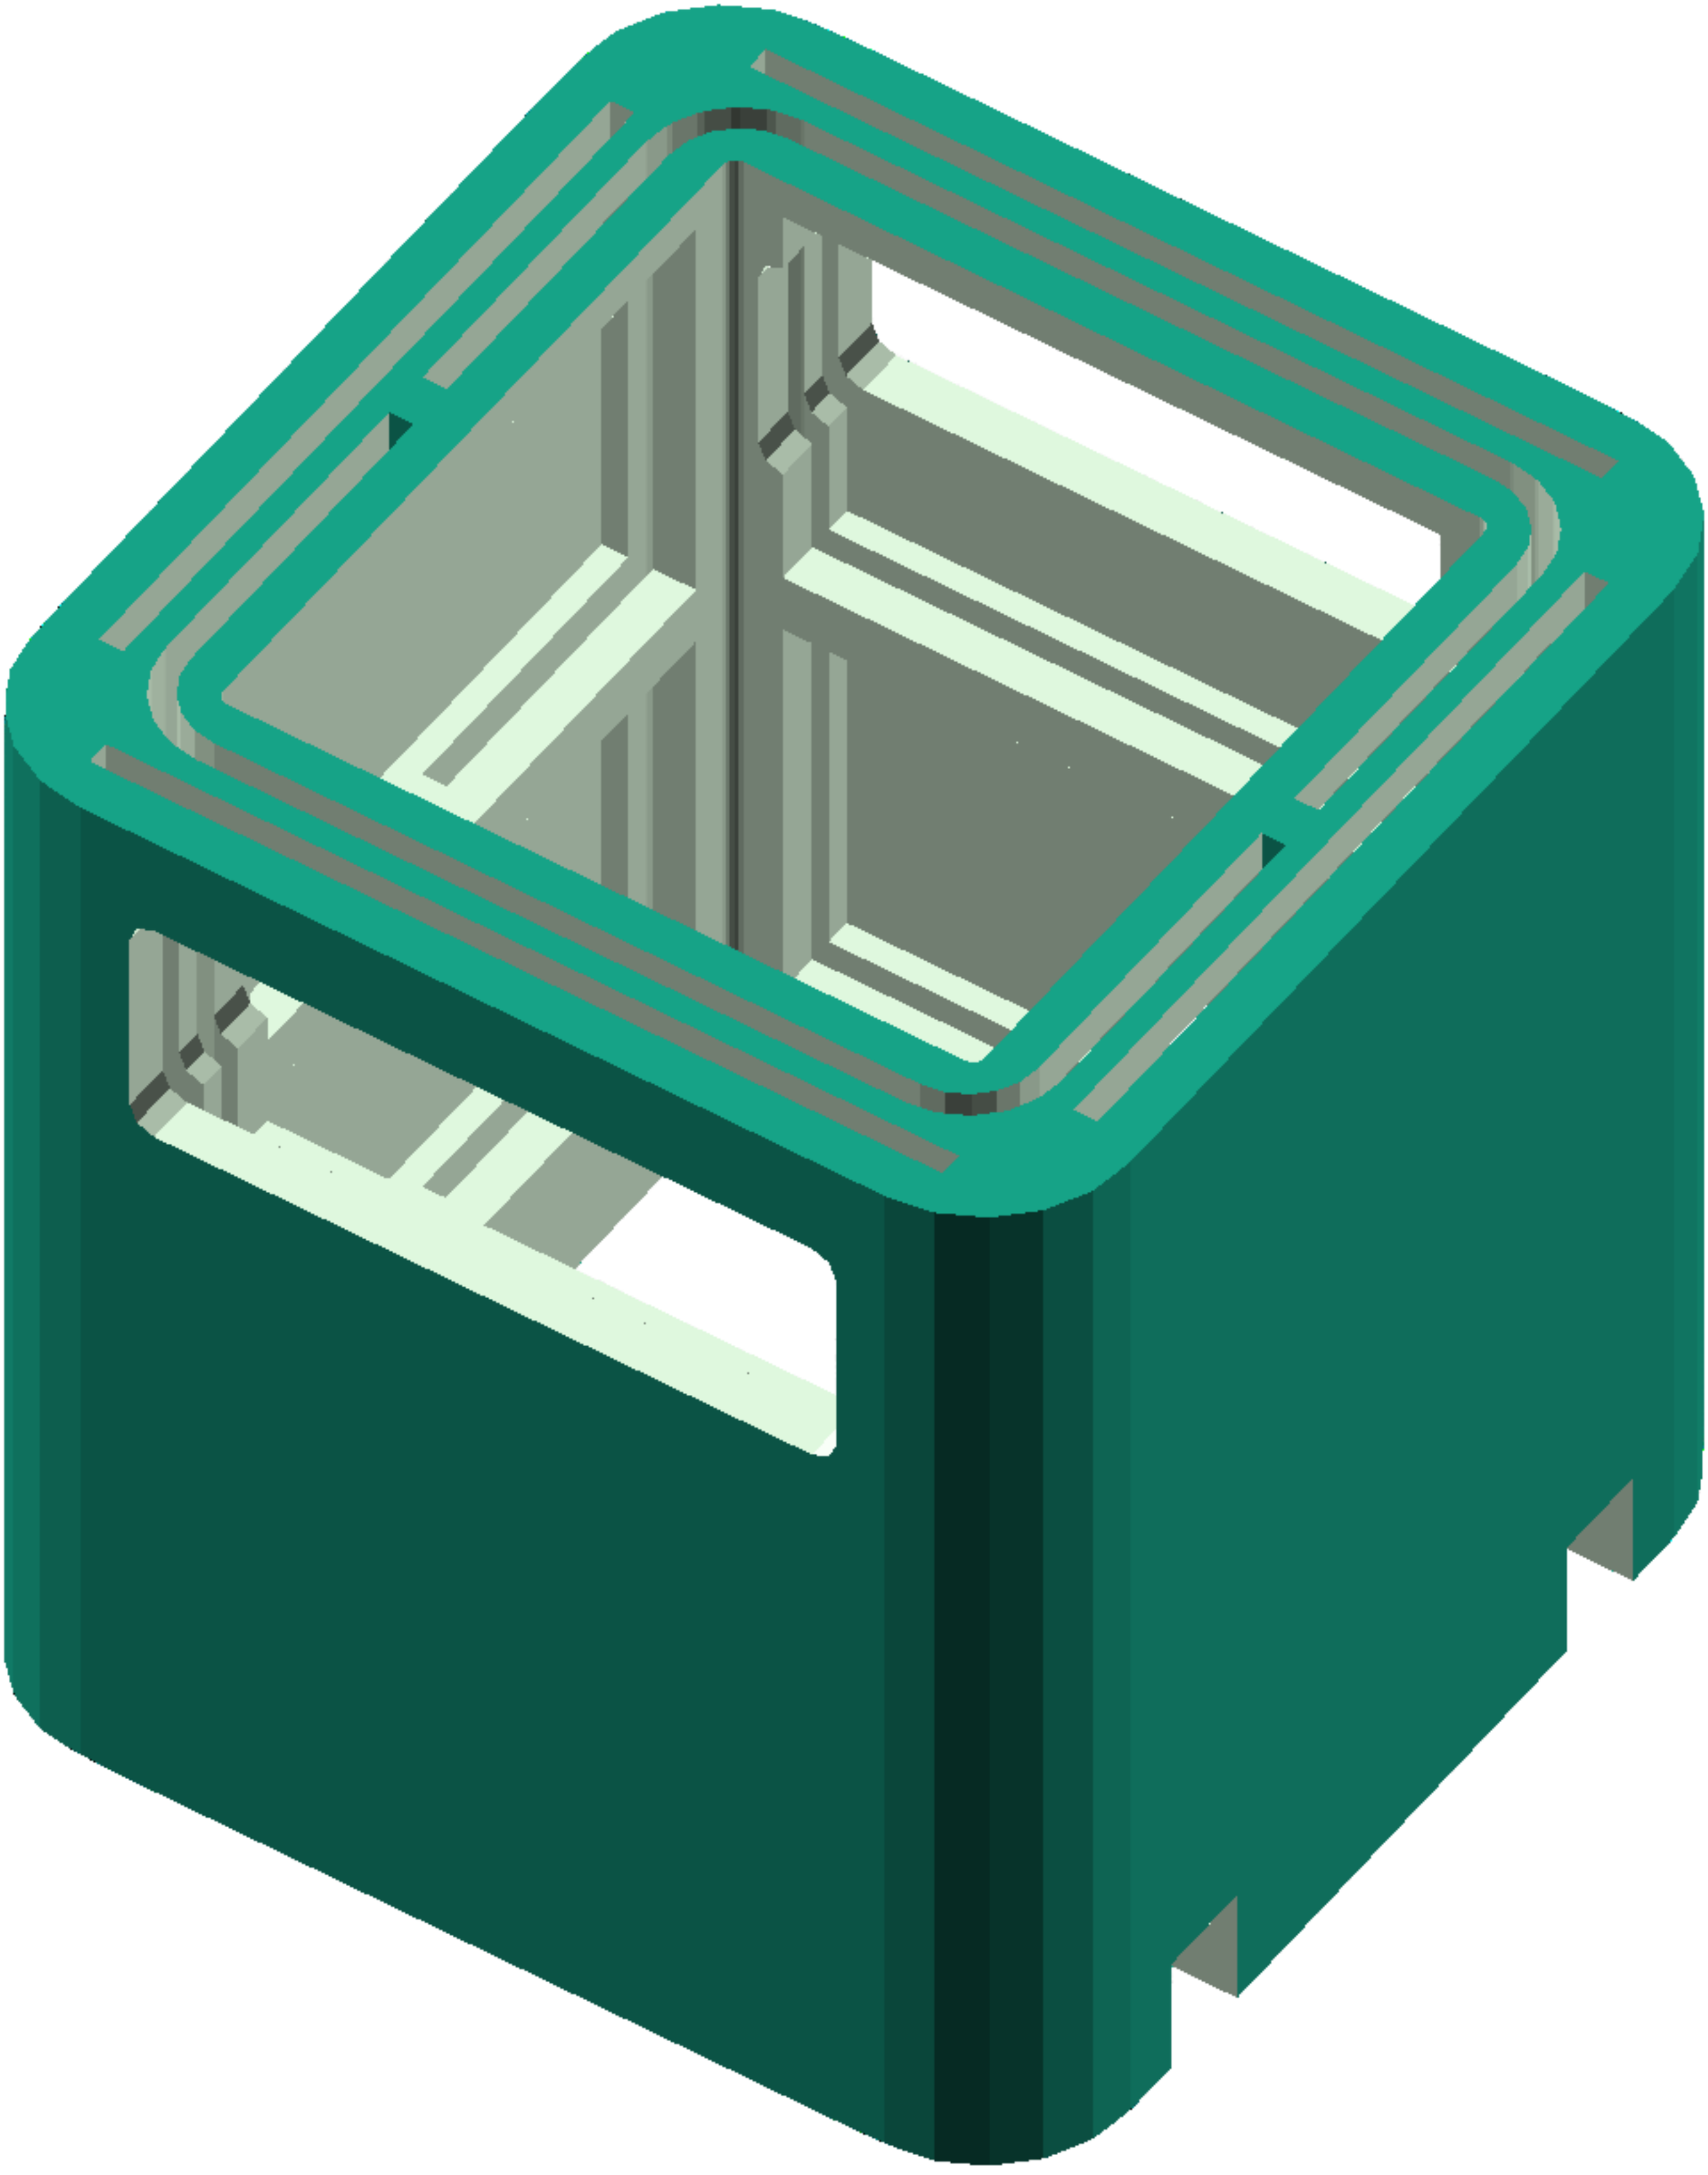

AIP Publishing
Review of Scientific Instruments
ACCEPTED MANUSCRIPT

PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0169355

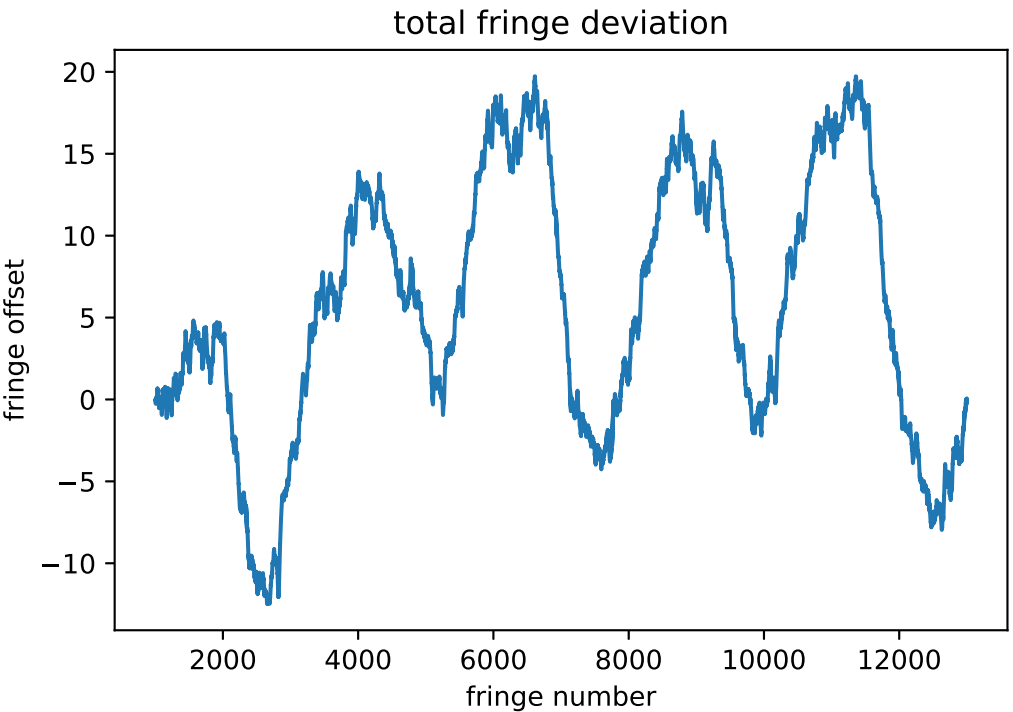
total fringe deviation
fringe offset
fringe number
20
15
10
5
0
−5
−10
2000
4000
6000
8000
10000
12000

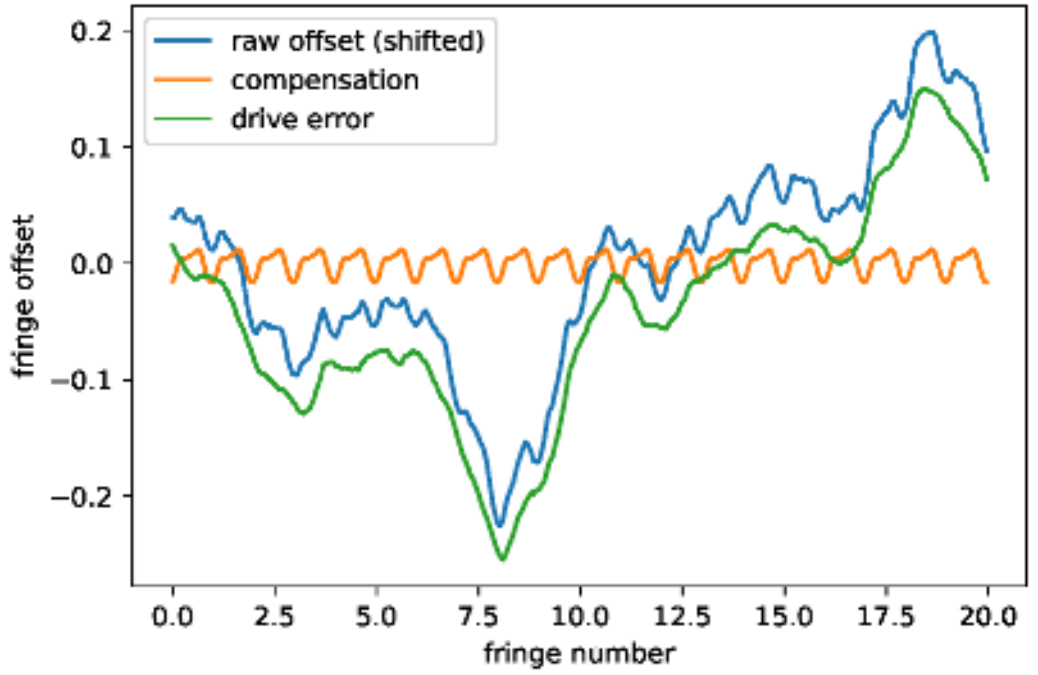
raw offset (shifted)
compensation
drive error
fringe offset
fringe number
0.2
0.1
0.0
−0.1
−0.2
0.0
2.5
5.0
7.5
10.0
12.5
15.0
17.5
20.0

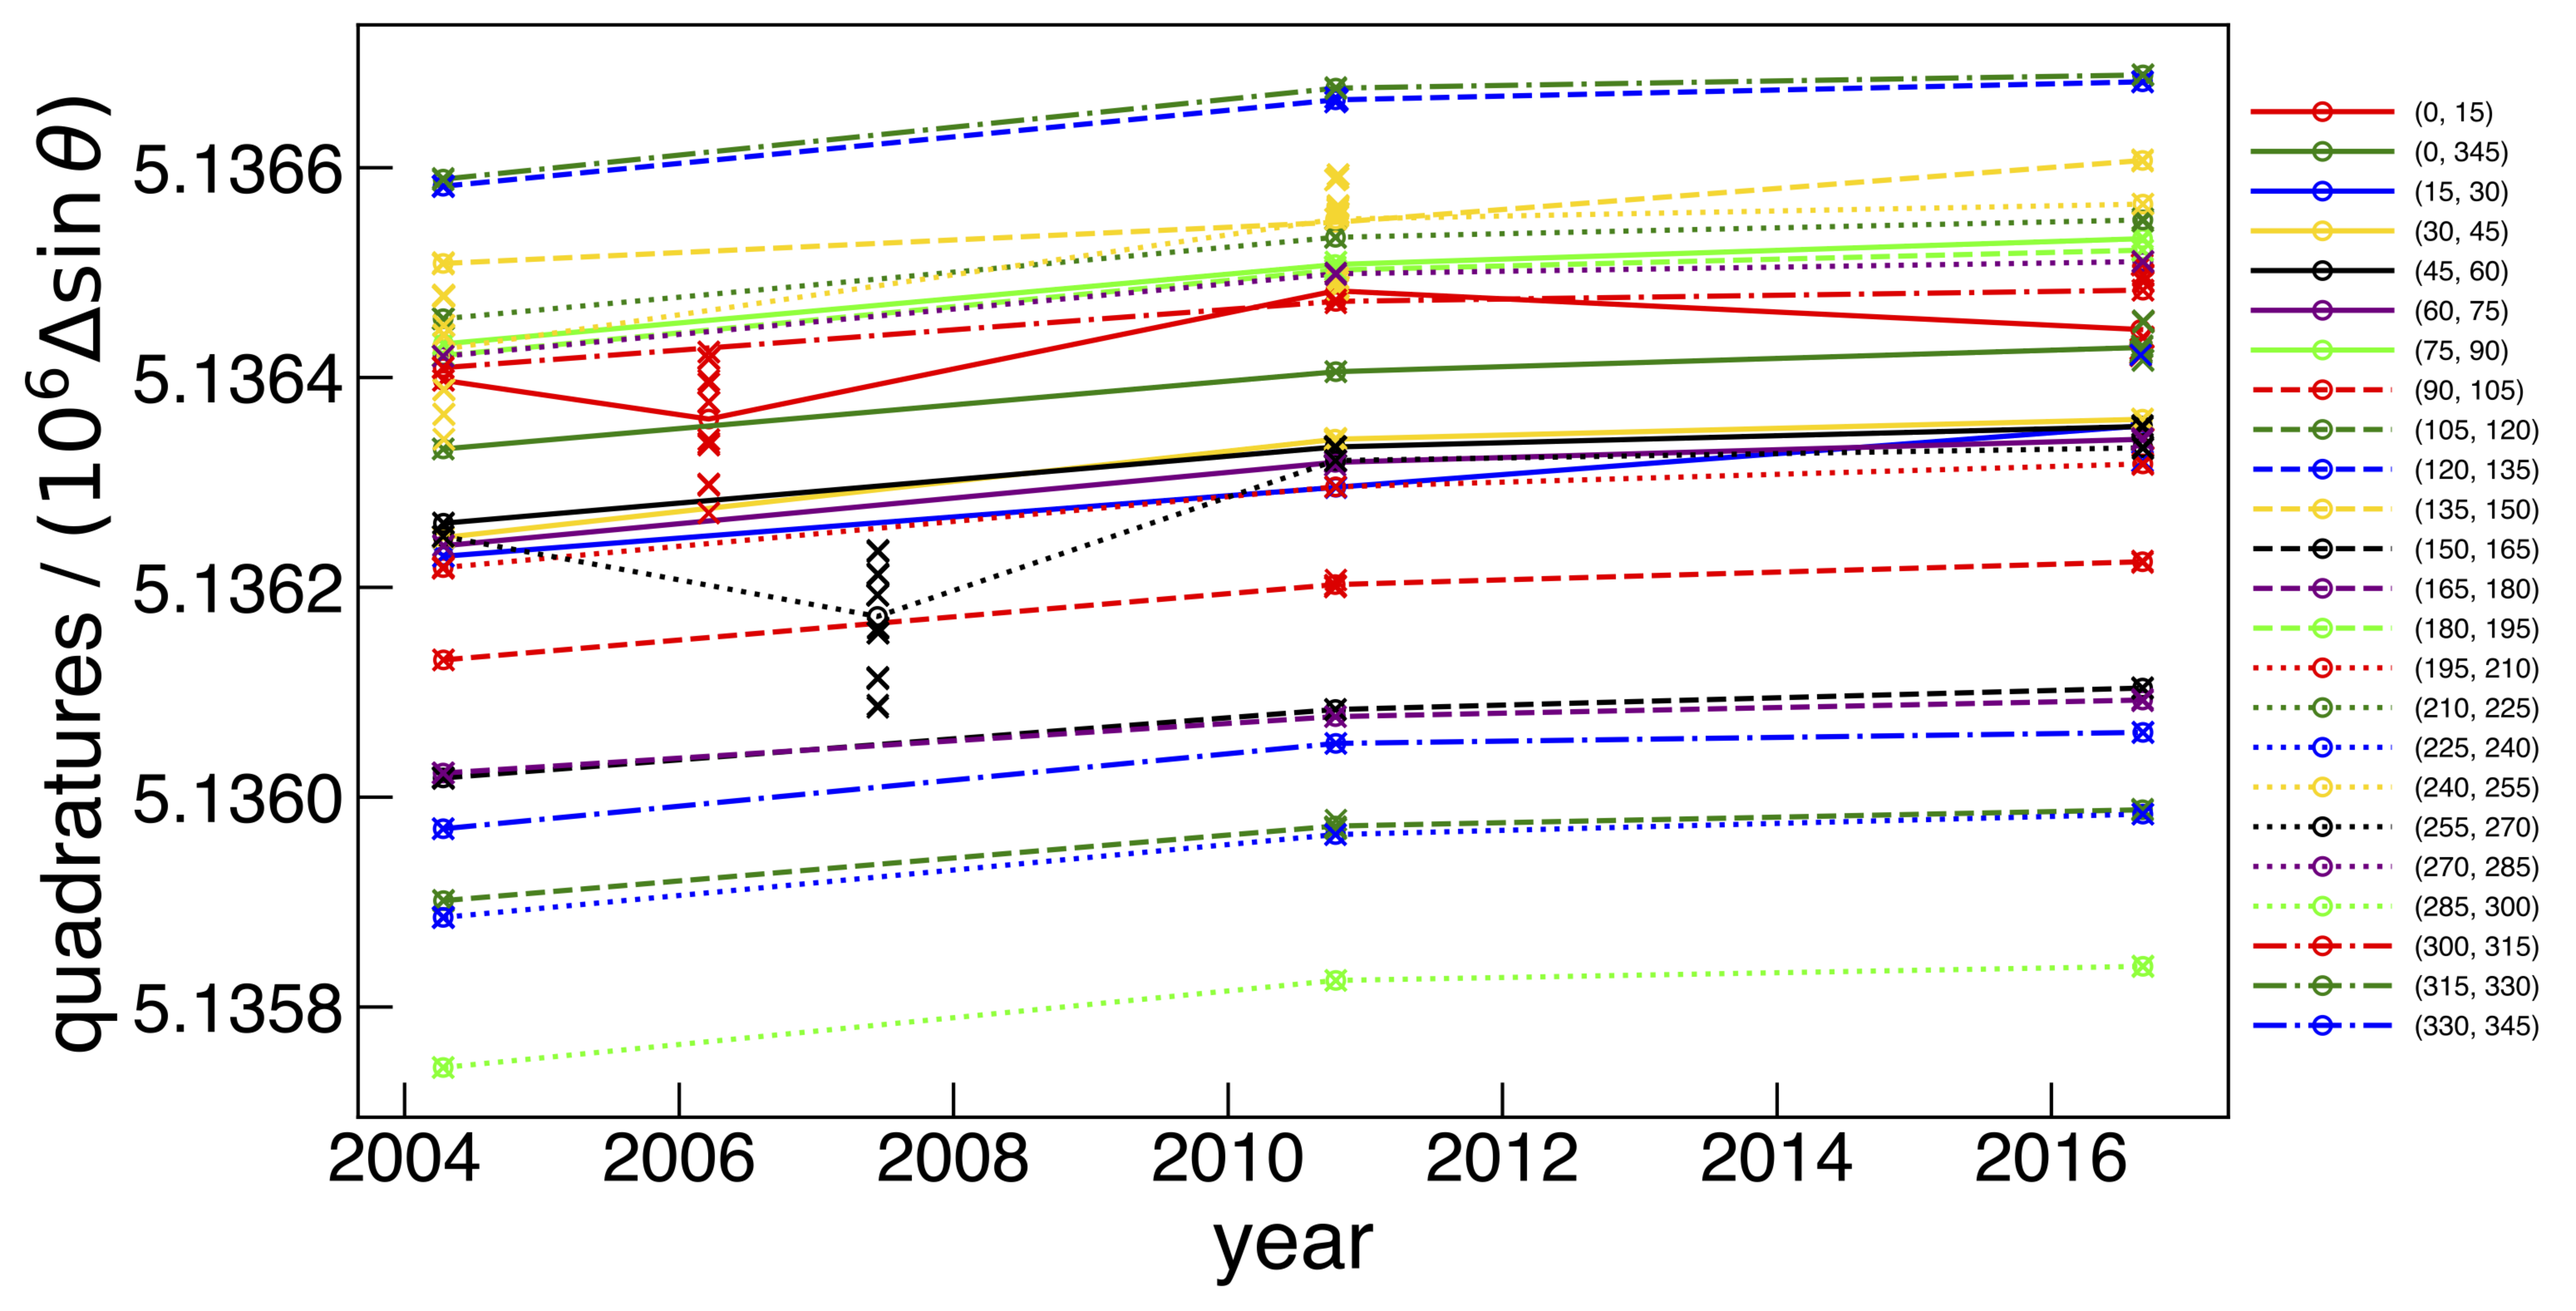
quadratures / ($10^6 \Delta \sin\theta$)
5.1366
5.1364
5.1362
5.1360
5.1358
2004
2006
2008
2010
2012
2014
2016
year
(0, 15)
(0, 345)
(15, 30)
(30, 45)
(45, 60)
(60, 75)
(75, 90)
(90, 105)
(105, 120)
(120, 135)
(135, 150)
(150, 165)
(165, 180)
(180, 195)
(195, 210)
(210, 225)
(225, 240)
(240, 255)
(255, 270)
(270, 285)
(285, 300)
(300, 315)
(315, 330)
(330, 345)

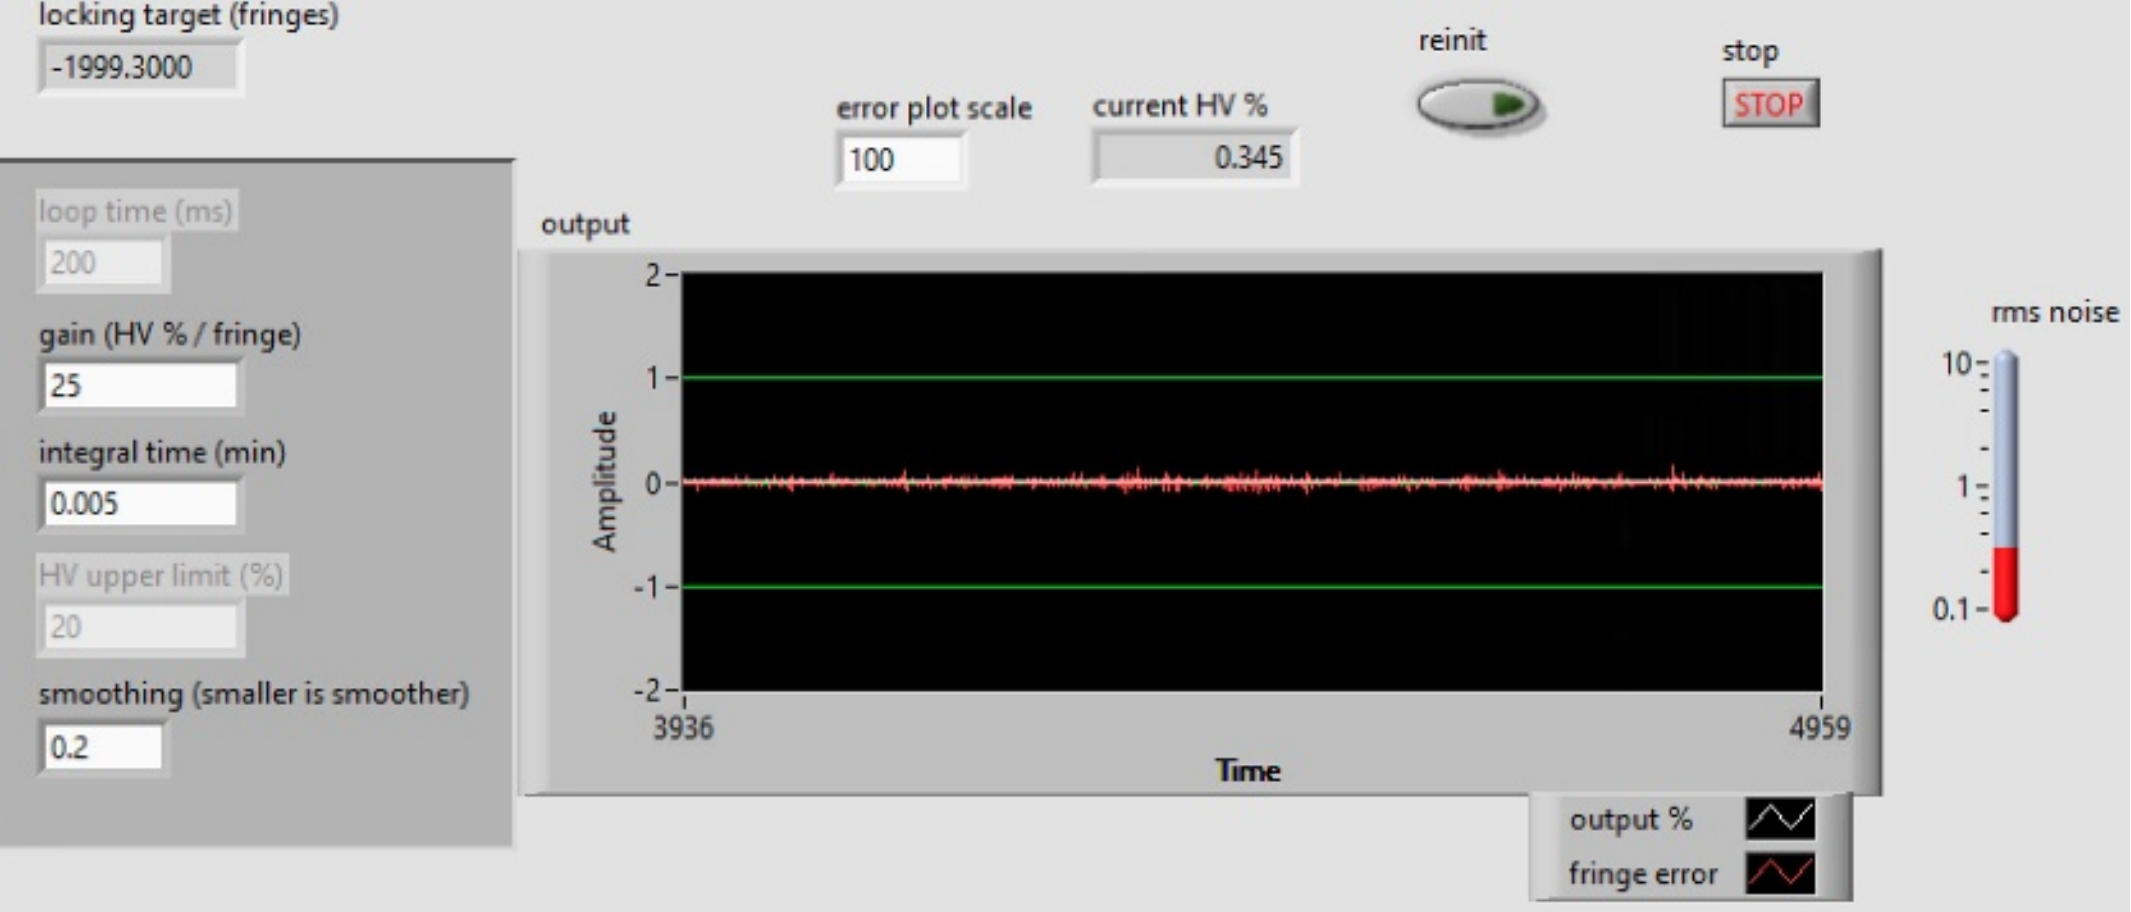
locking target (fringes)
-1999.3000
error plot scale
100
current HV %
0.345
reinit
stop
STOP
loop time (ms)
200
gain (HV % / fringe)
25
integral time (min)
0.005
HV upper limit (%)
20
smoothing (smaller is smoother)
0.2
output
2
1
0
-1
-2
Amplitude
3936
4959
Time
rms noise
10
1
0.1
output %
fringe error

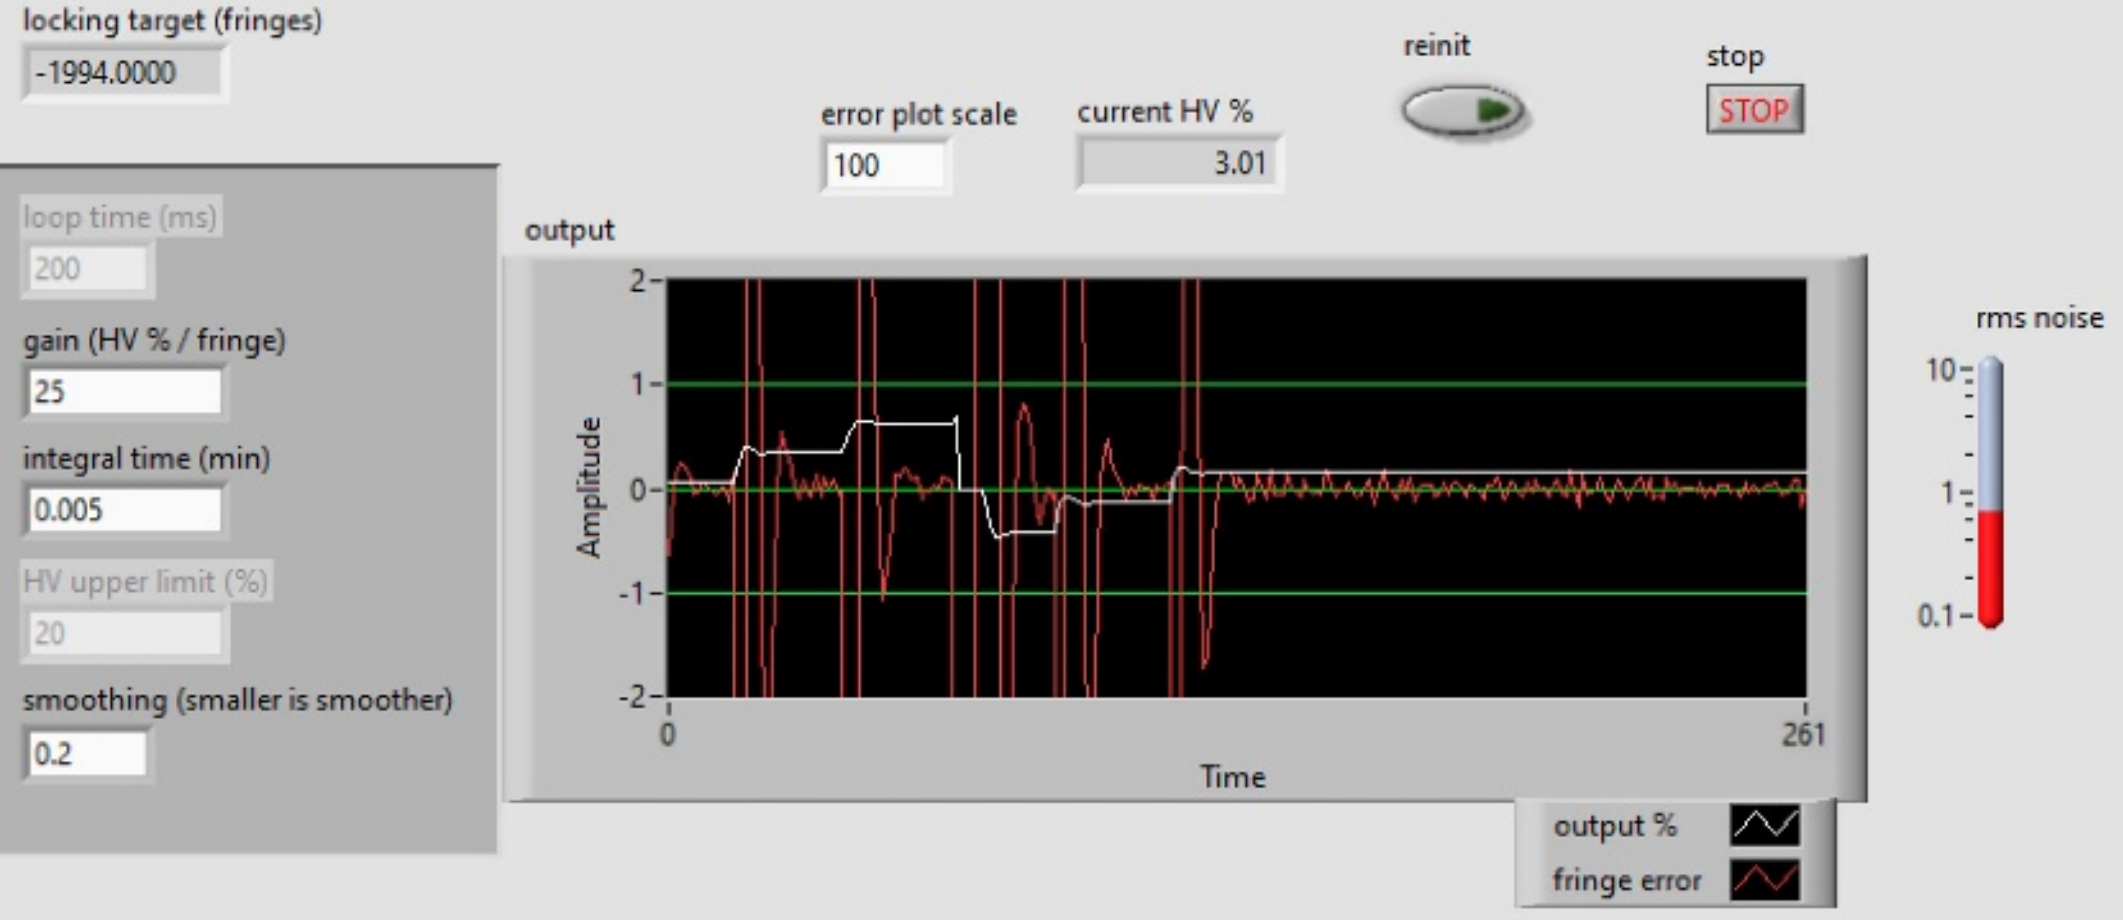
locking target (fringes)
-1994.0000
reinit
stop
STOP
error plot scale
100
current HV %
3.01
loop time (ms)
200
gain (HV % / fringe)
25
integral time (min)
0.005
HV upper limit (%)
20
smoothing (smaller is smoother)
0.2
output
2
1
0
-1
-2
Amplitude
0
261
Time
rms noise
10
1
0.1
output %
fringe error

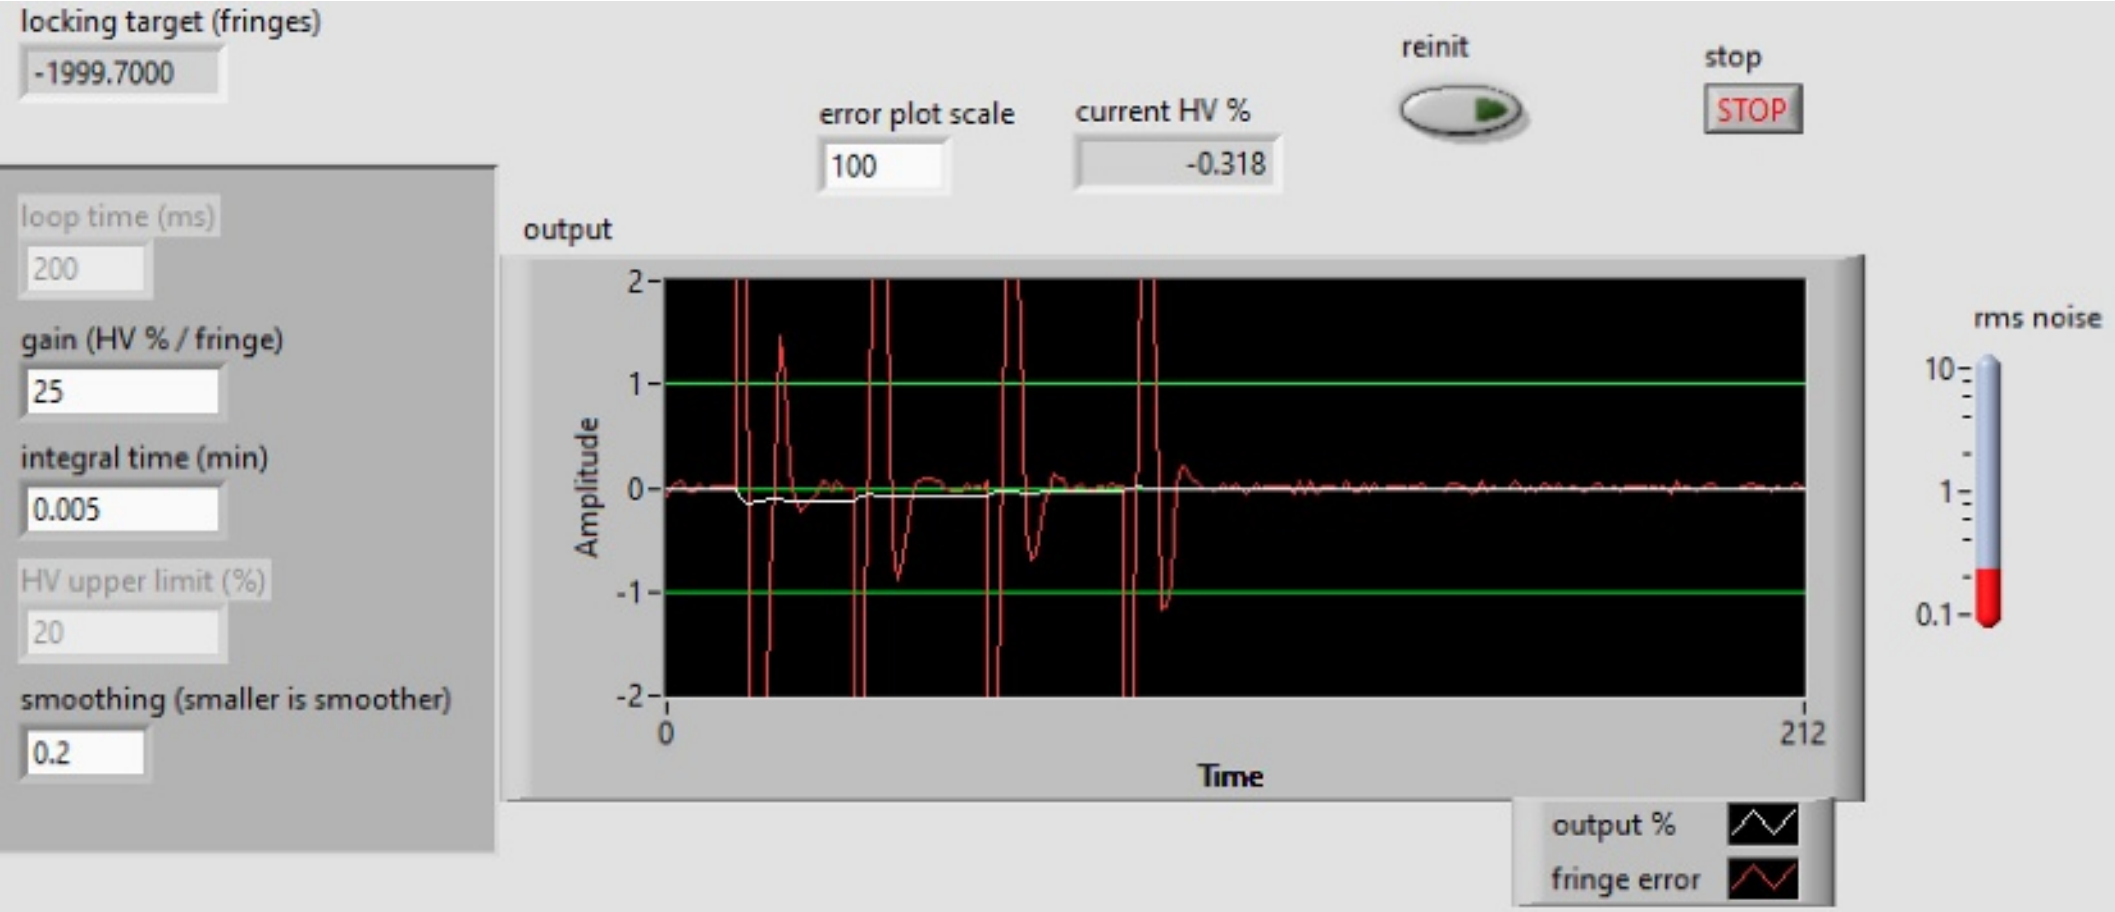
locking target (fringes)
-1999.7000
reinit
stop
STOP
error plot scale
100
current HV %
-0.318
loop time (ms)
200
gain (HV % / fringe)
25
integral time (min)
0.005
HV upper limit (%)
20
smoothing (smaller is smoother)
0.2
output
Amplitude
2
1
0
-1
-2
0
212
Time
rms noise
10
1
0.1
output %
fringe error

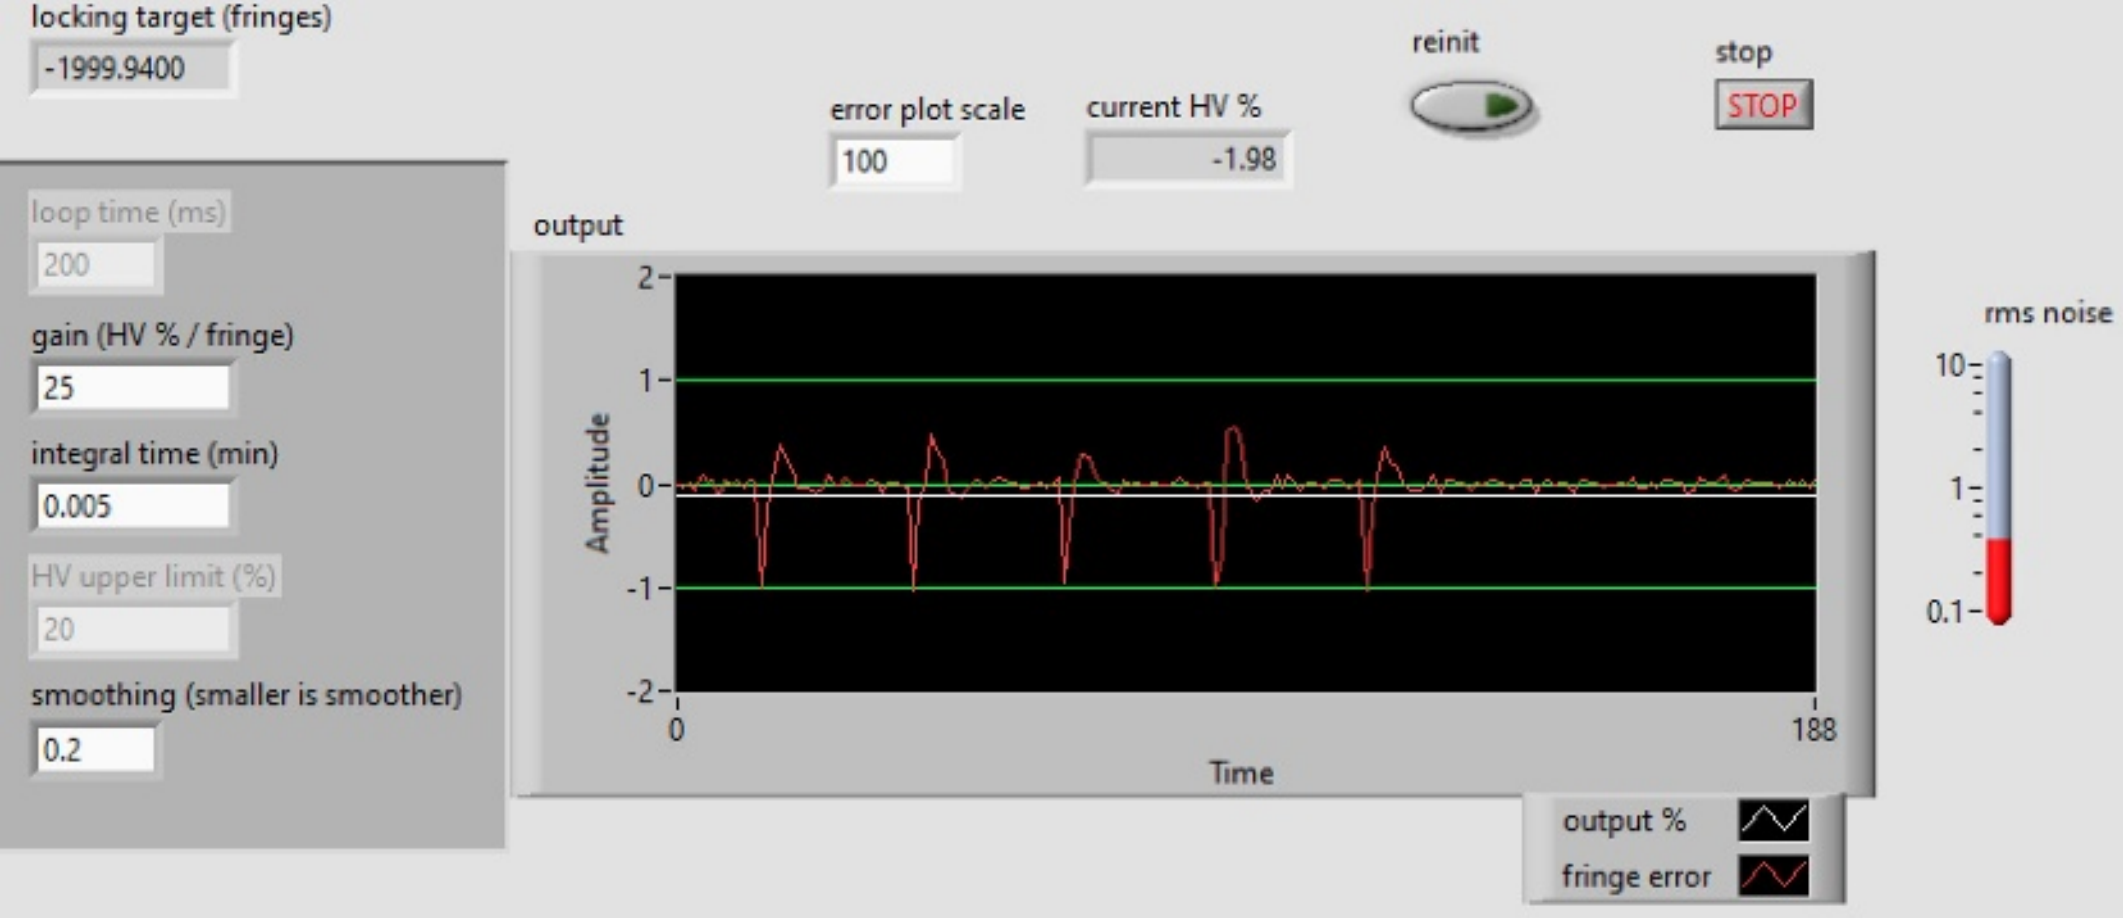
locking target (fringes)
-1999.9400
reinit
stop
STOP
error plot scale
100
current HV %
-1.98
loop time (ms)
200
gain (HV % / fringe)
25
integral time (min)
0.005
HV upper limit (%)
20
smoothing (smaller is smoother)
0.2
output
2
1
0
-1
-2
Amplitude
0
188
Time
rms noise
10
1
0.1
output %
fringe error

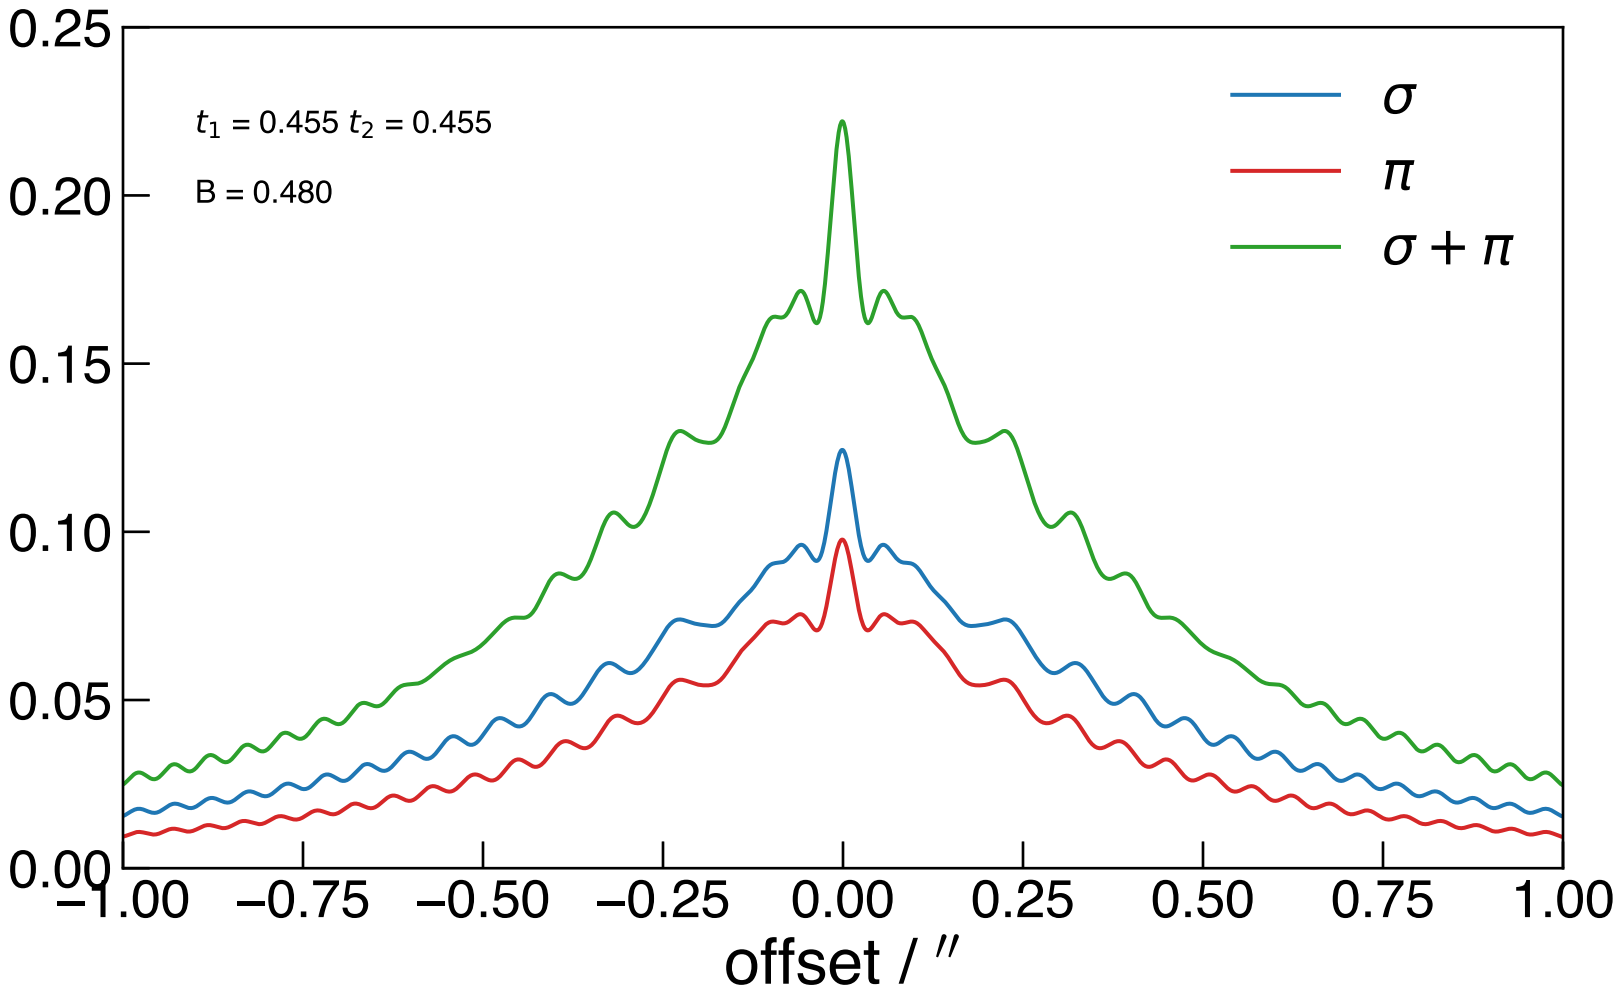
0.25
0.20
0.15
0.10
0.05
0.00
t_1 = 0.455 t_2 = 0.455
B = 0.480
σ
π
$\sigma + \pi$
−1.00
−0.75
−0.50
−0.25
0.00
0.25
0.50
0.75
1.00
offset / ″

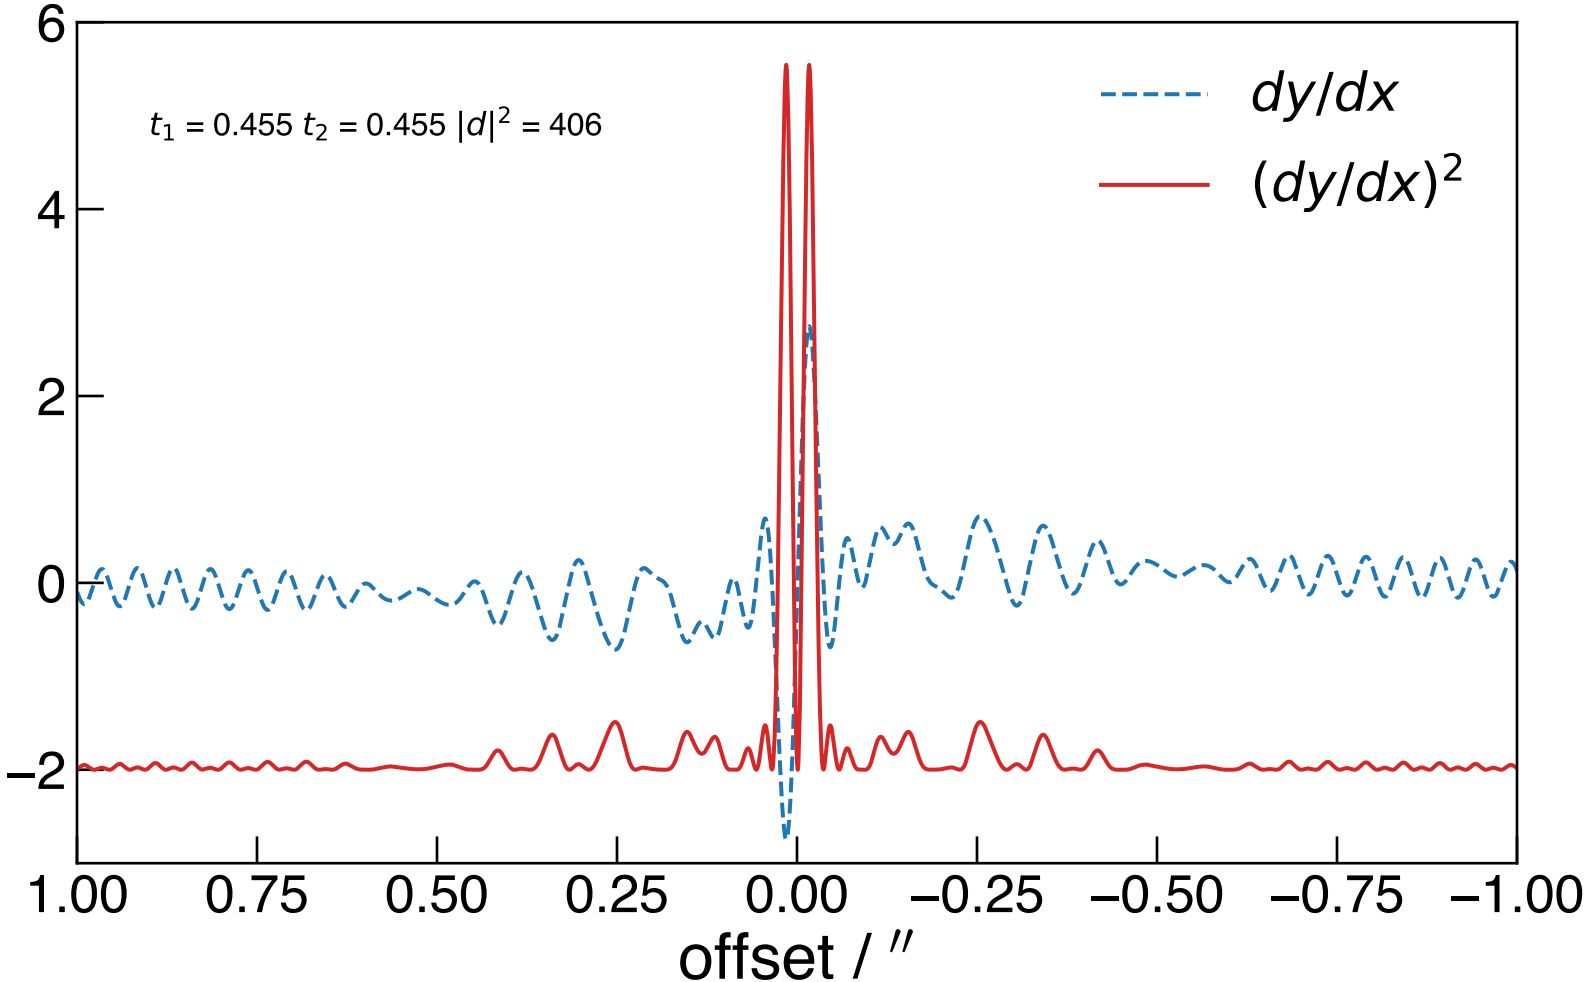
6
4
2
0
−2
t_1 = 0.455 t_2 = 0.455 $|d|^2$ = 406
dy/dx
$(dy/dx)^2$
1.00
0.75
0.50
0.25
0.00
−0.25
−0.50
−0.75
−1.00
offset / ″

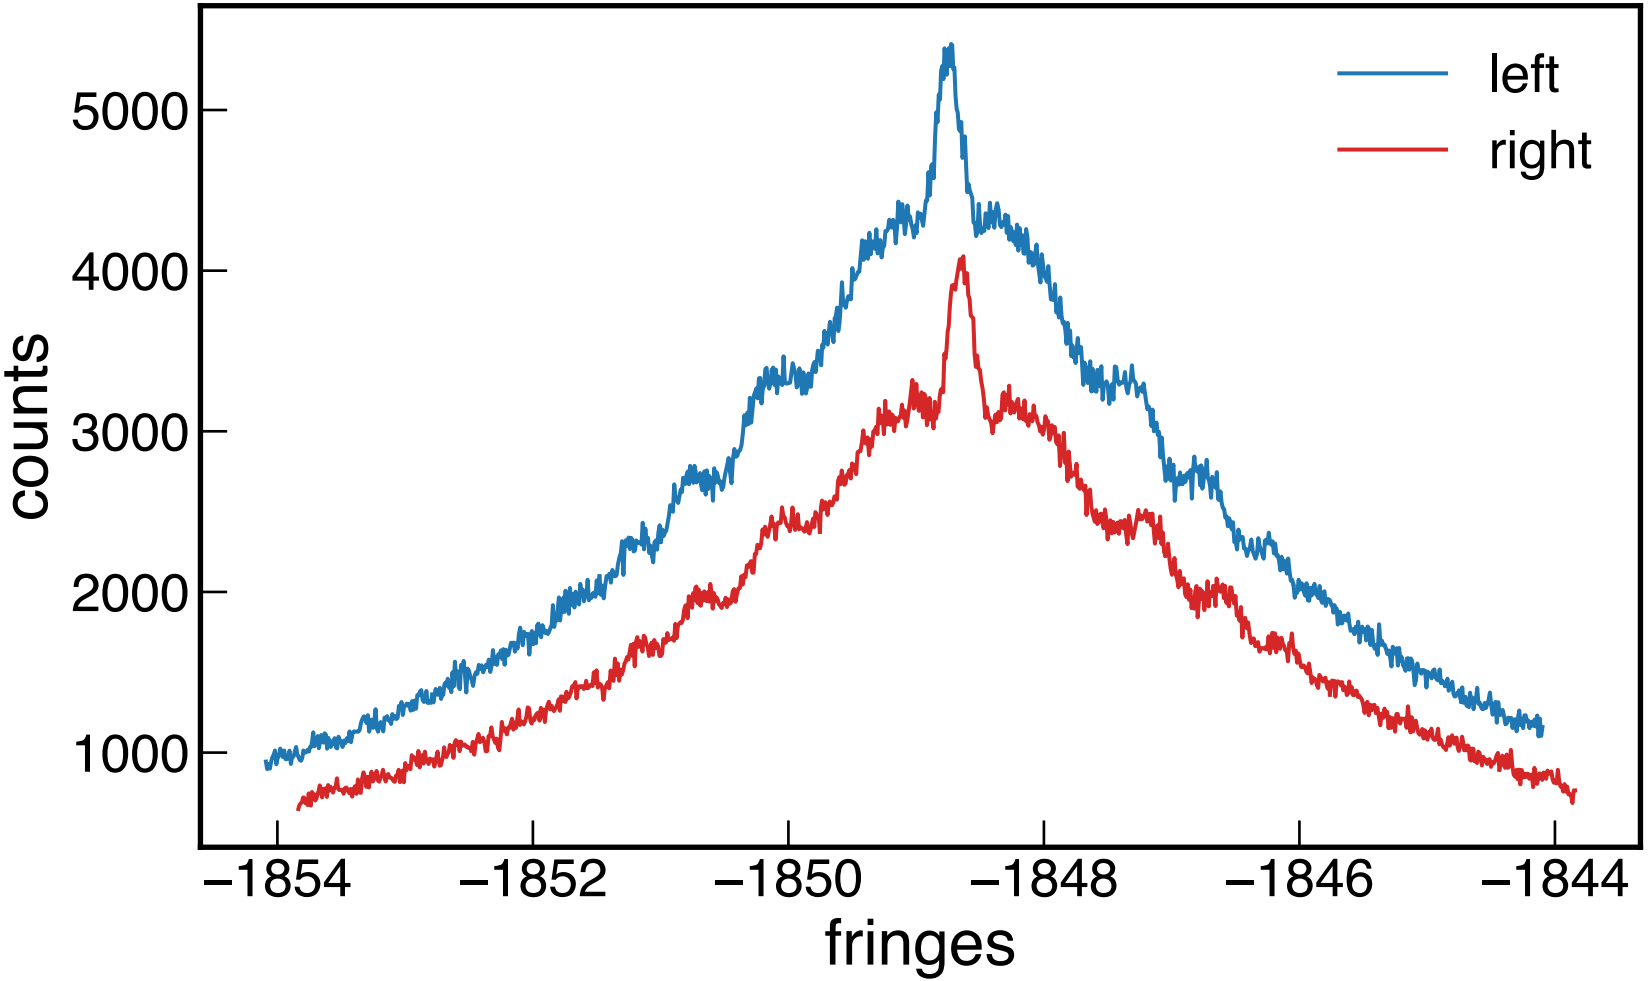
left
right
5000
4000
3000
2000
1000
counts
−1854
−1852
−1850
−1848
−1846
−1844
fringes

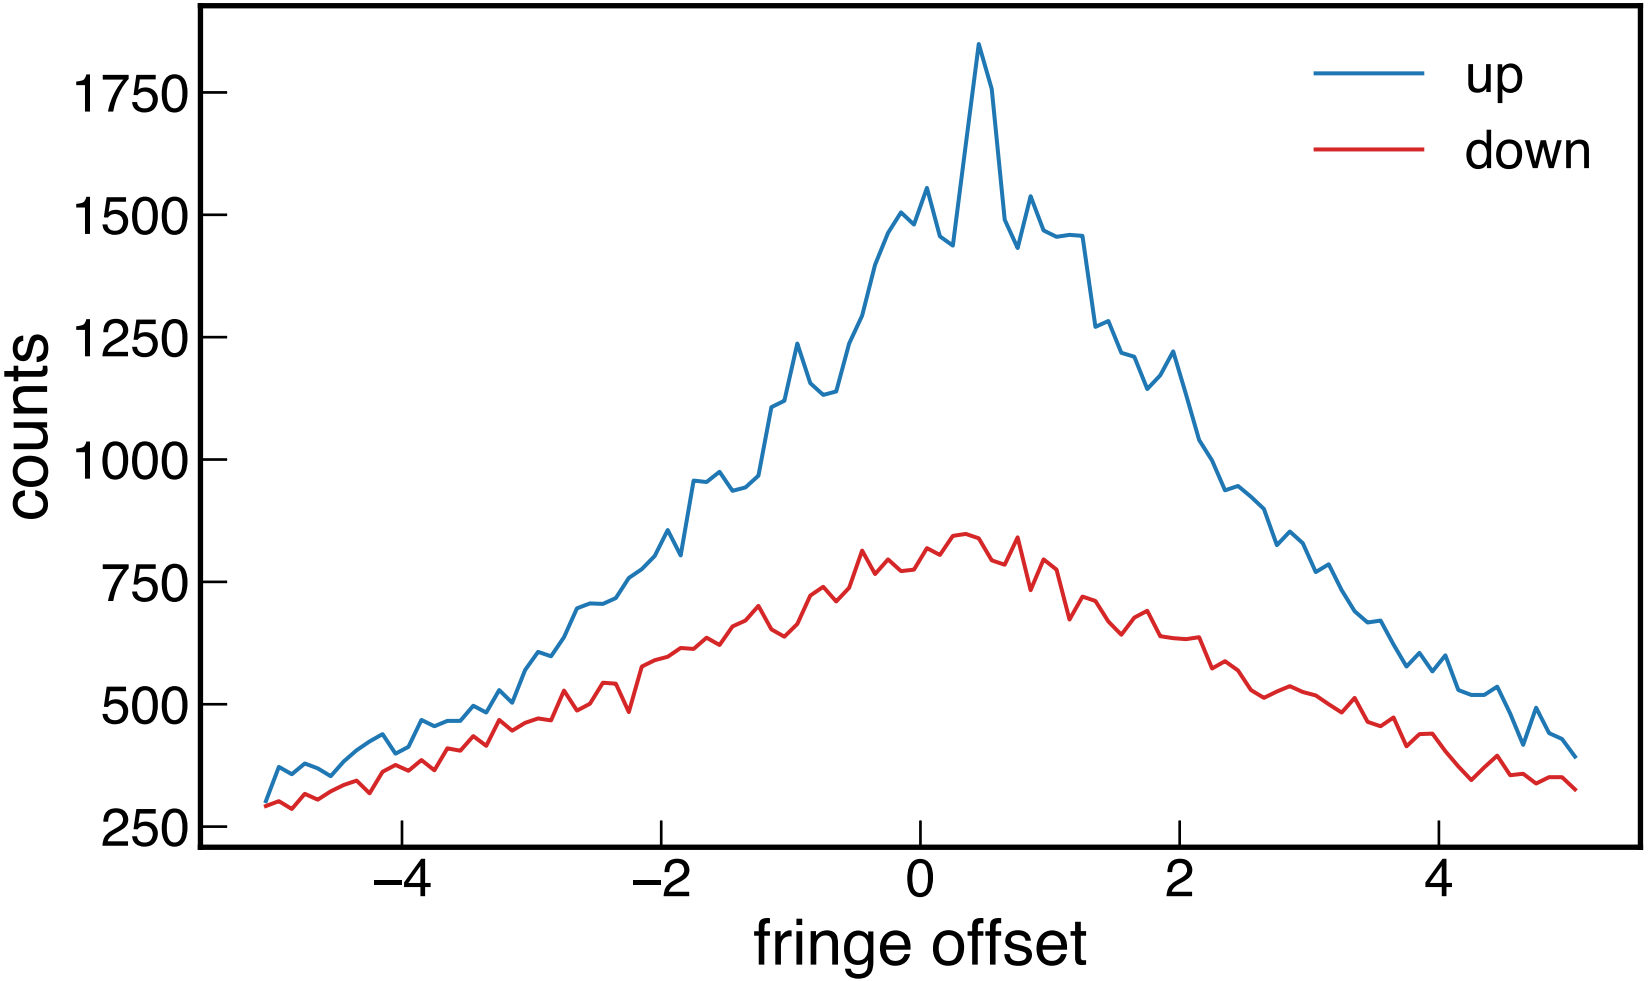
up
down
counts
1750
1500
1250
1000
750
500
250
−4
−2
0
2
4
fringe offset

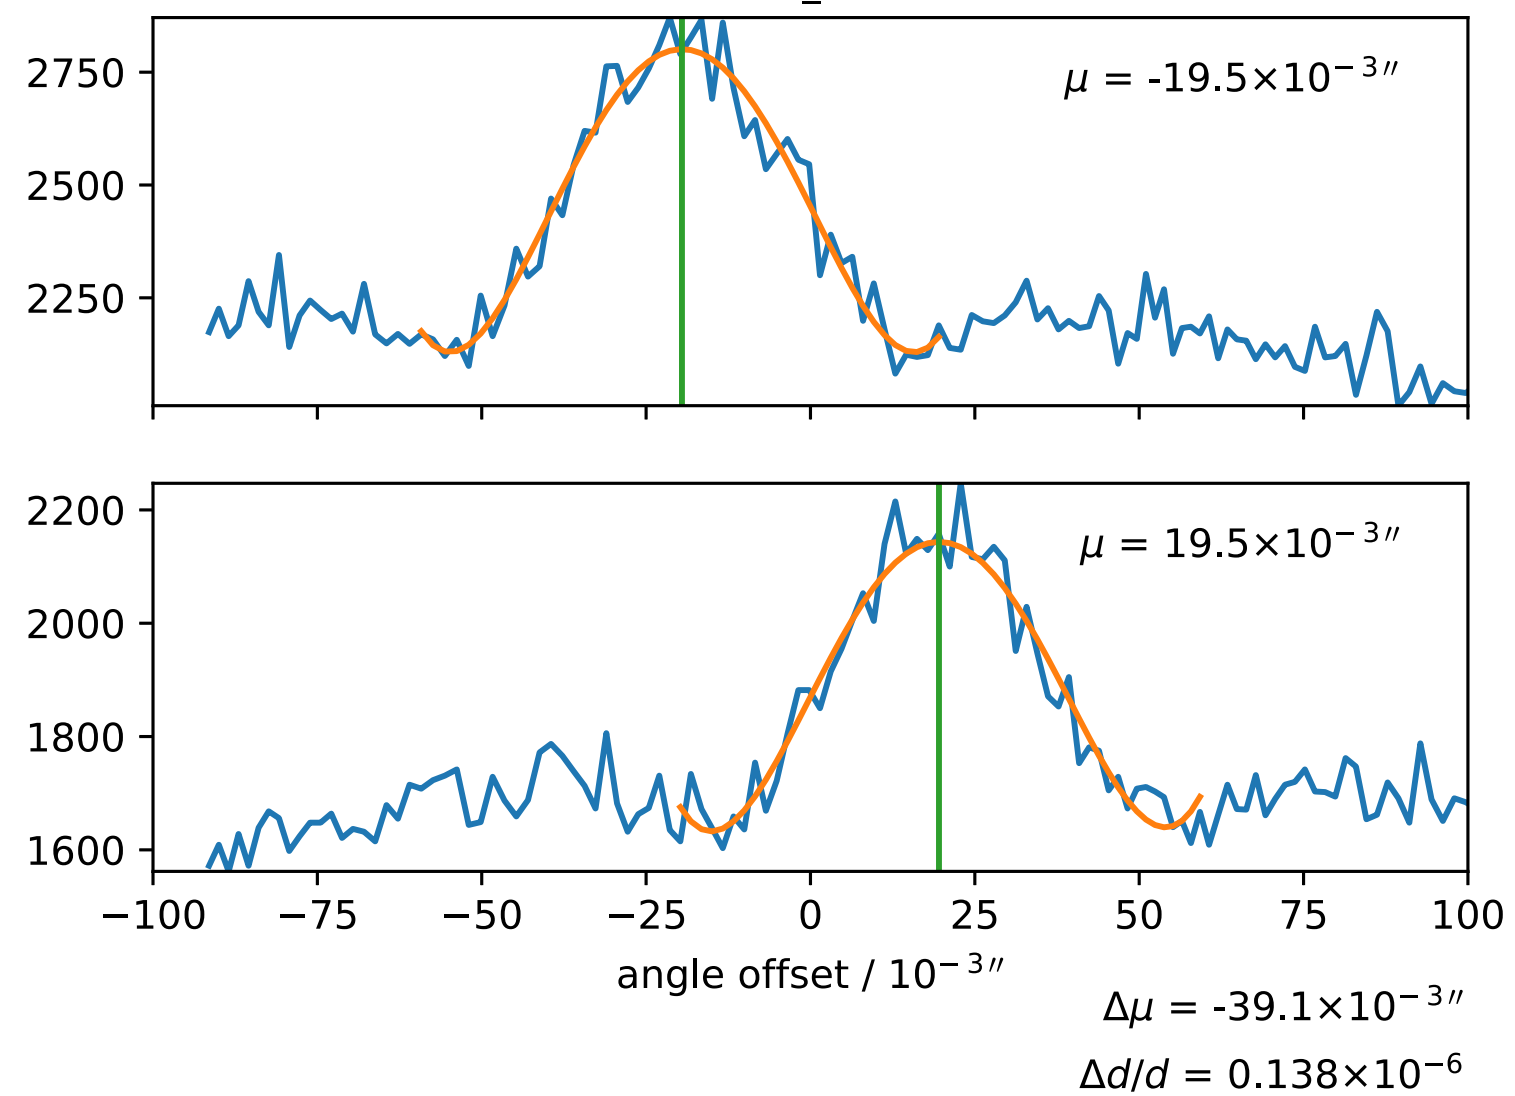
2750
2500
2250
μ = -19.5×10⁻³″
2200
2000
1800
1600
μ = 19.5×10⁻³″
−100
−75
−50
−25
0
25
50
75
100
angle offset / 10⁻³″
Δμ = -39.1×10⁻³″
Δd/d = 0.138×10⁻⁶

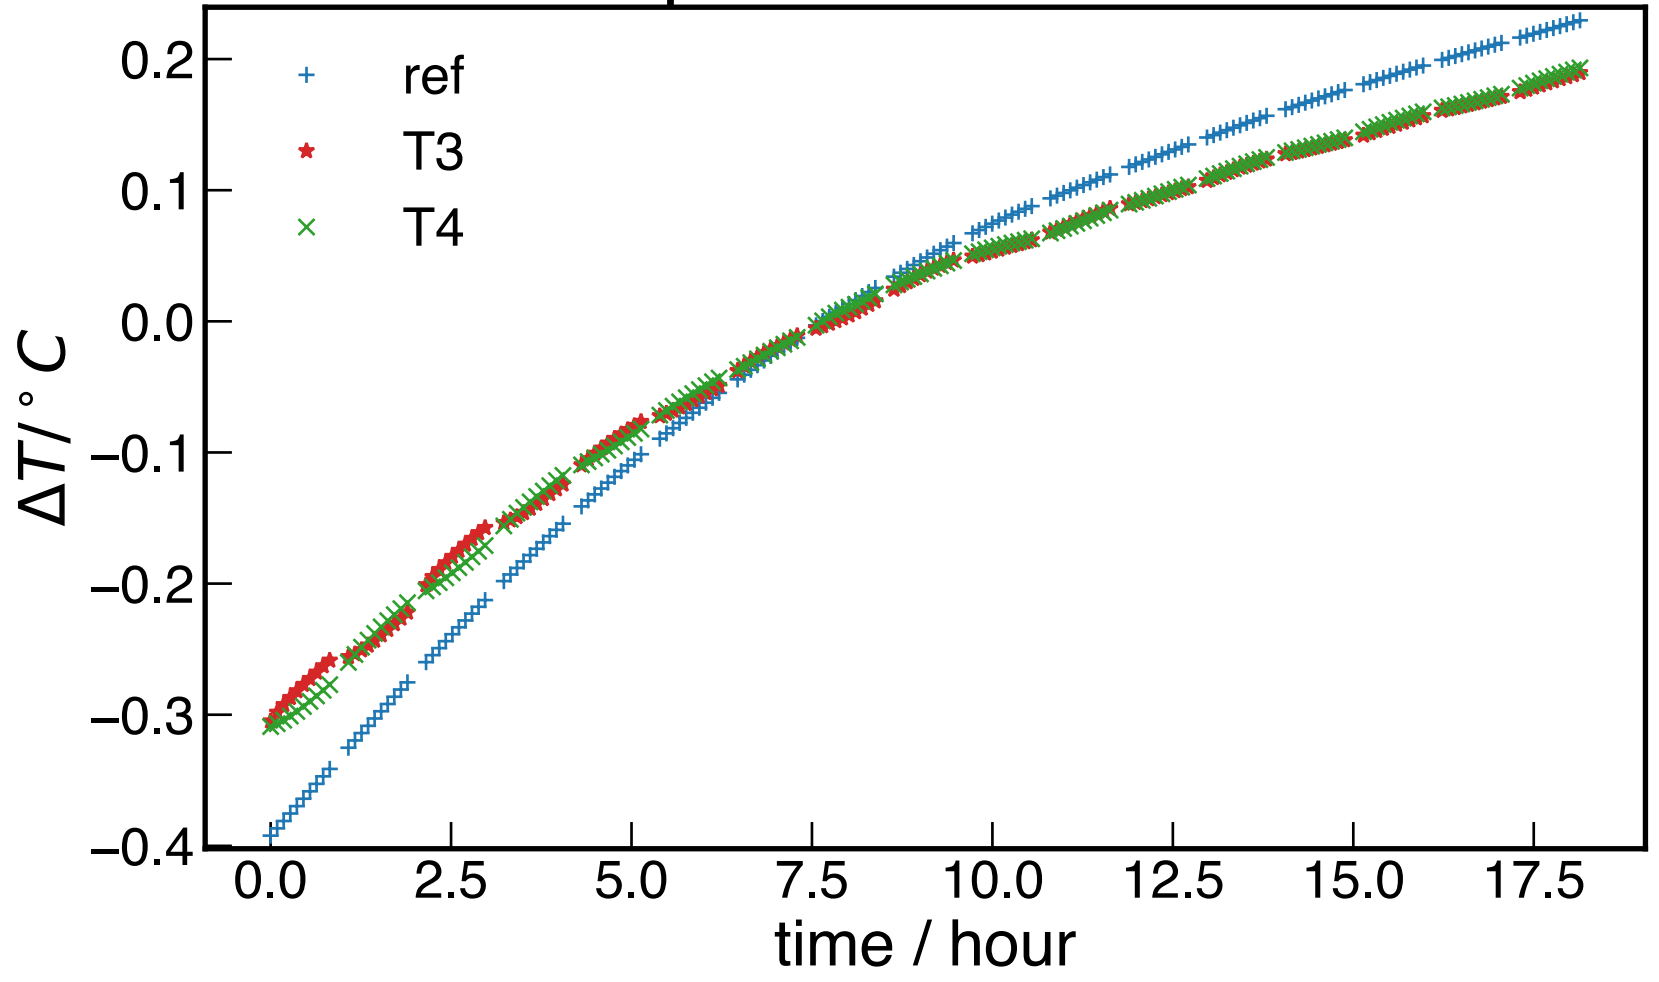
ref
T3
T4
0.2
0.1
0.0
−0.1
−0.2
−0.3
−0.4
ΔT/ ° C
0.0
2.5
5.0
7.5
10.0
12.5
15.0
17.5
time / hour

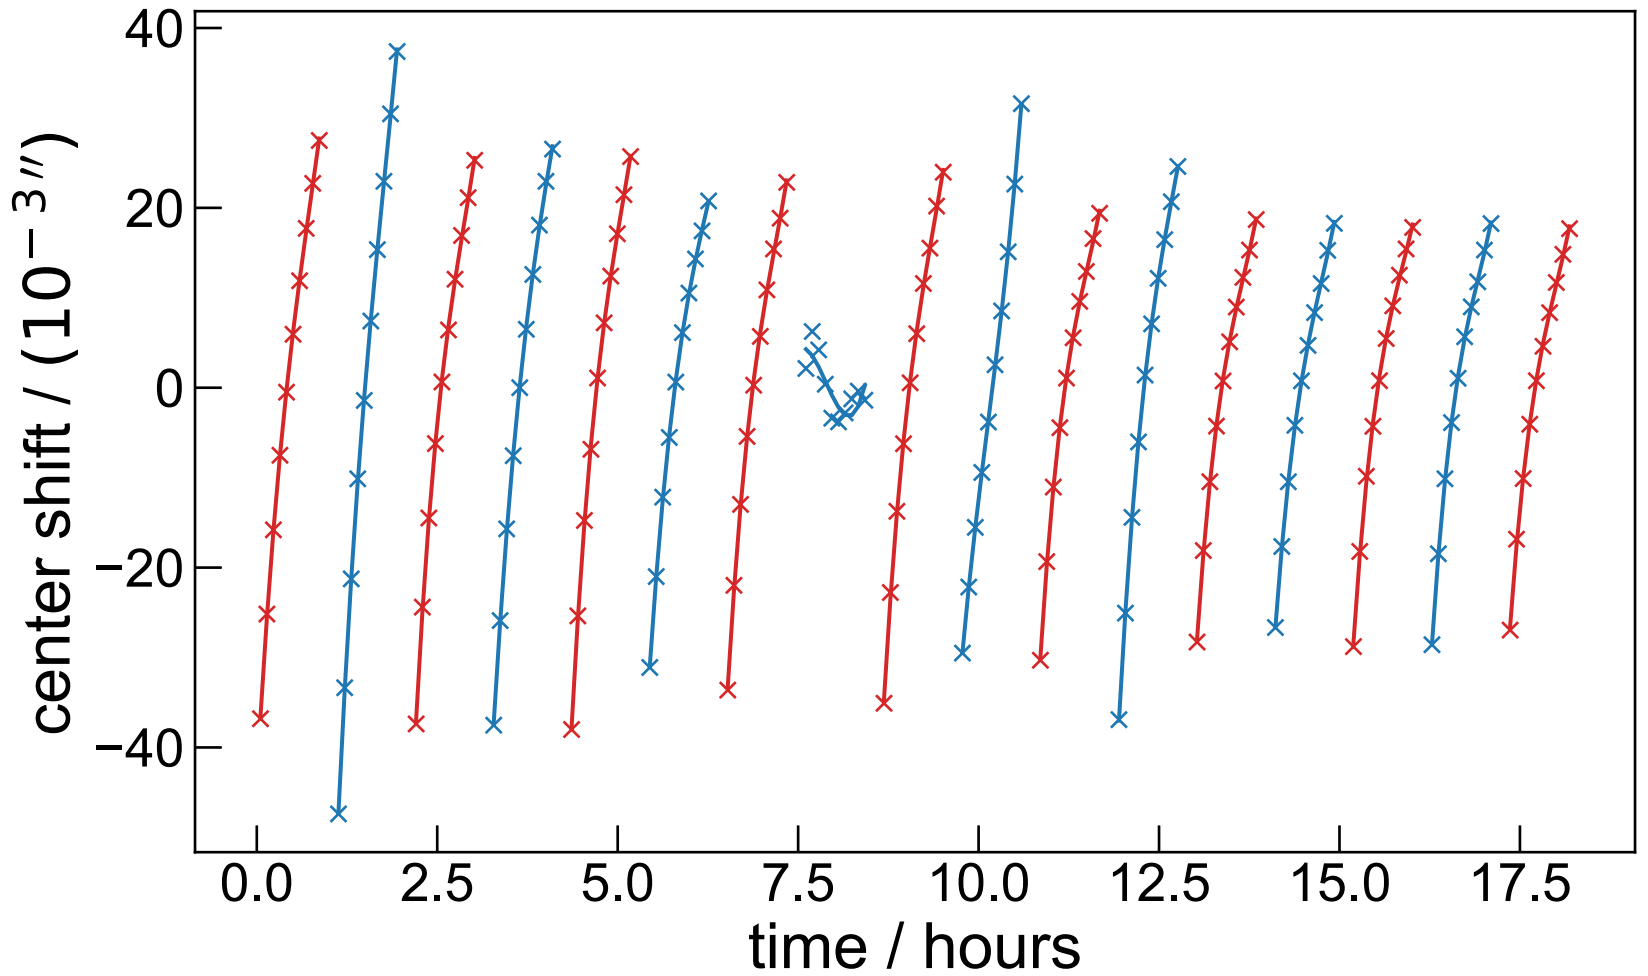
40
20
0
−20
−40
center shift / (10^{-3}″)
0.0
2.5
5.0
7.5
10.0
12.5
15.0
17.5
time / hours

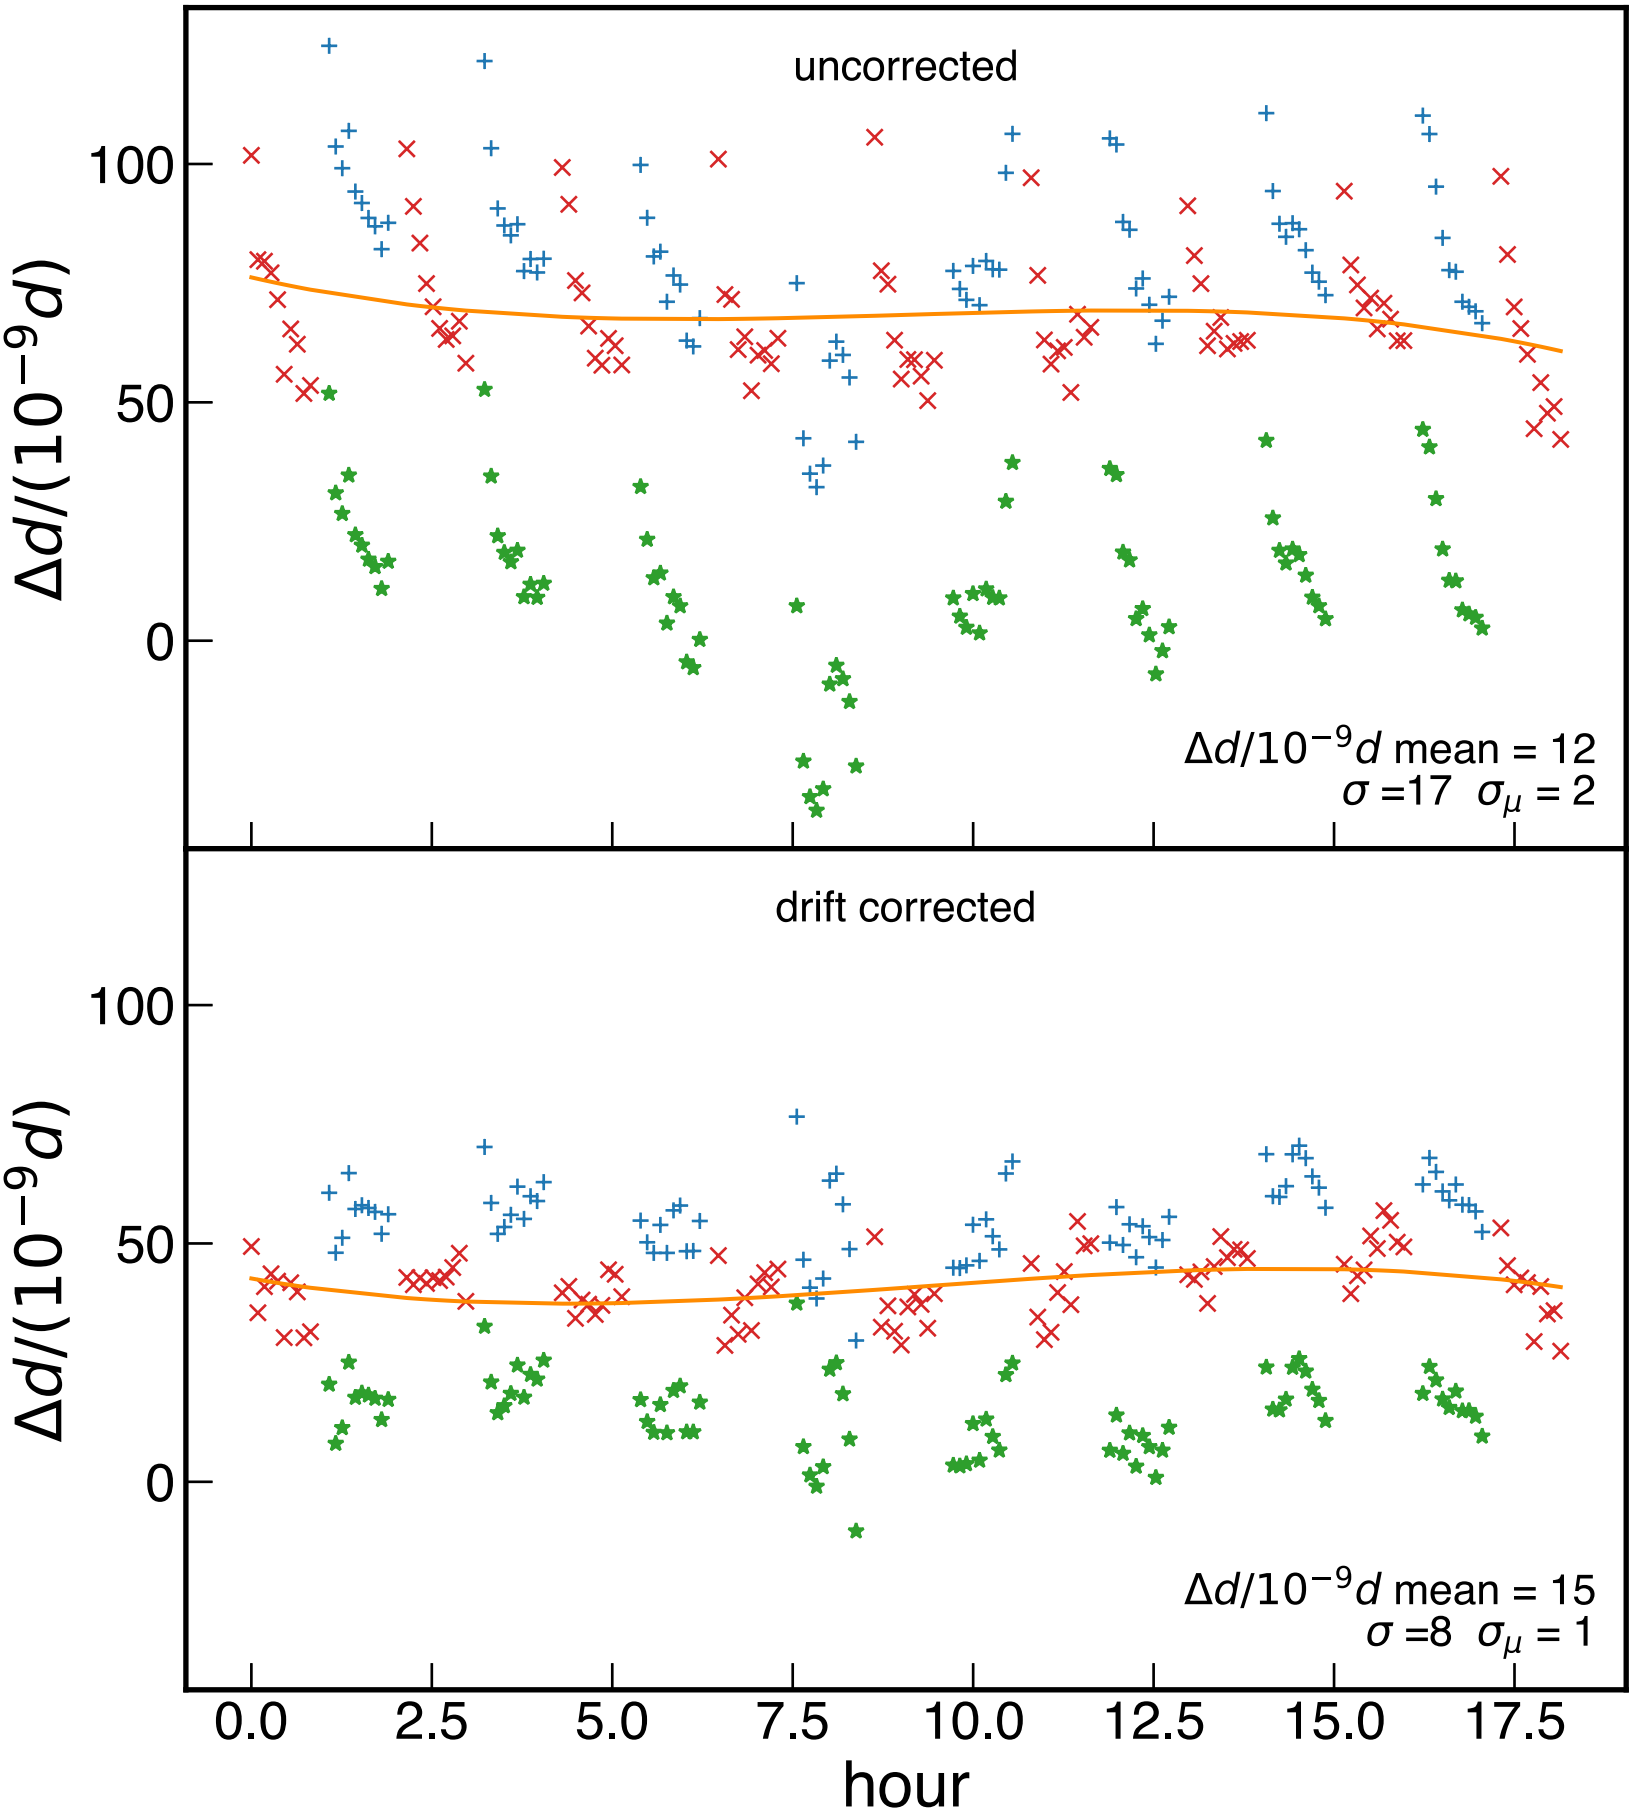
uncorrected
100
50
0
$\Delta d/(10^{-9}d)$
$\Delta d/10^{-9}d$ mean = 12
σ =17 σ_μ = 2
drift corrected
100
50
0
$\Delta d/(10^{-9}d)$
$\Delta d/10^{-9}d$ mean = 15
σ =8 σ_μ = 1
0.0
2.5
5.0
7.5
10.0
12.5
15.0
17.5
hour

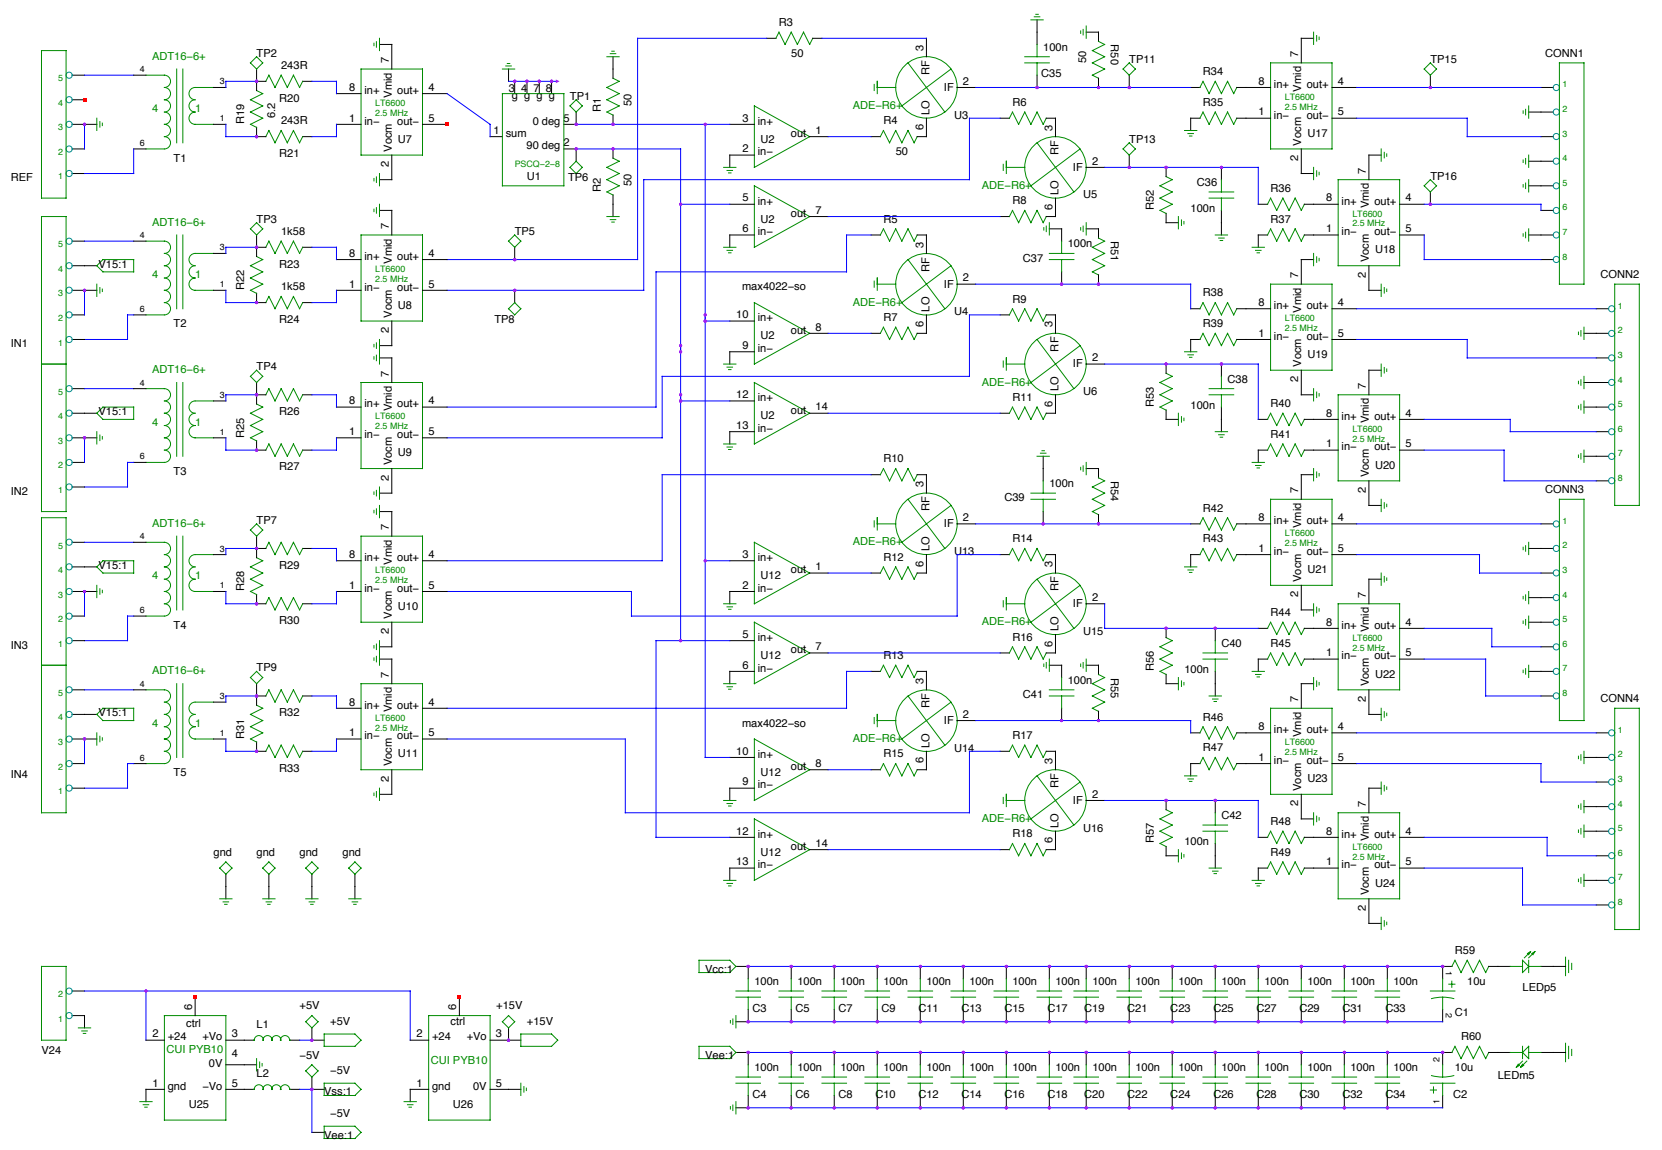
REF
IN1
IN2
IN3
IN4
ADT16-6+
T1
T2
T3
T4
T5
TP2
TP3
TP4
TP7
TP9
243R
R20
R21
R19
6.2
1k58
R23
R24
R22
R26
R27
R25
R29
R30
R28
R32
R33
R31
V15:1
LT6600
2.5 MHz
U7
U8
U9
U10
U11
TP1
TP6
TP5
TP8
R1
R2
50
0 deg
90 deg
sum
PSCQ-2-8
U1
R3
ADE-R6
U3
U4
U5
U6
U13
U14
U15
U16
RF
IF
LO
R4
R5
R6
R7
R8
R9
R10
R11
R12
R13
R14
R16
R17
R18
max4022-so
U2
U12
100n
C35
C36
C37
C38
C39
C40
C41
C42
R50
R51
R52
R53
R54
R55
R56
R57
TP11
TP13
TP15
TP16
R34
R35
R36
R37
R38
R39
R40
R41
R42
R43
R44
R45
R46
R47
R48
R49
U17
U18
U19
U20
U21
U22
U23
U24
CONN1
CONN2
CONN3
CONN4
gnd
V24
CUI PYB10
U25
U26
L1
L2
+5V
-5V
+15V
Vss:1
Vee:1
Vcc:1
C1
C2
C3
C4
C5
C6
C7
C8
C9
C10
C11
C12
C13
C14
C15
C16
C17
C18
C19
C20
C21
C22
C23
C24
C25
C26
C27
C28
C29
C30
C31
C32
C33
C34
R59
R60
10u
LEDp5
LEDm5

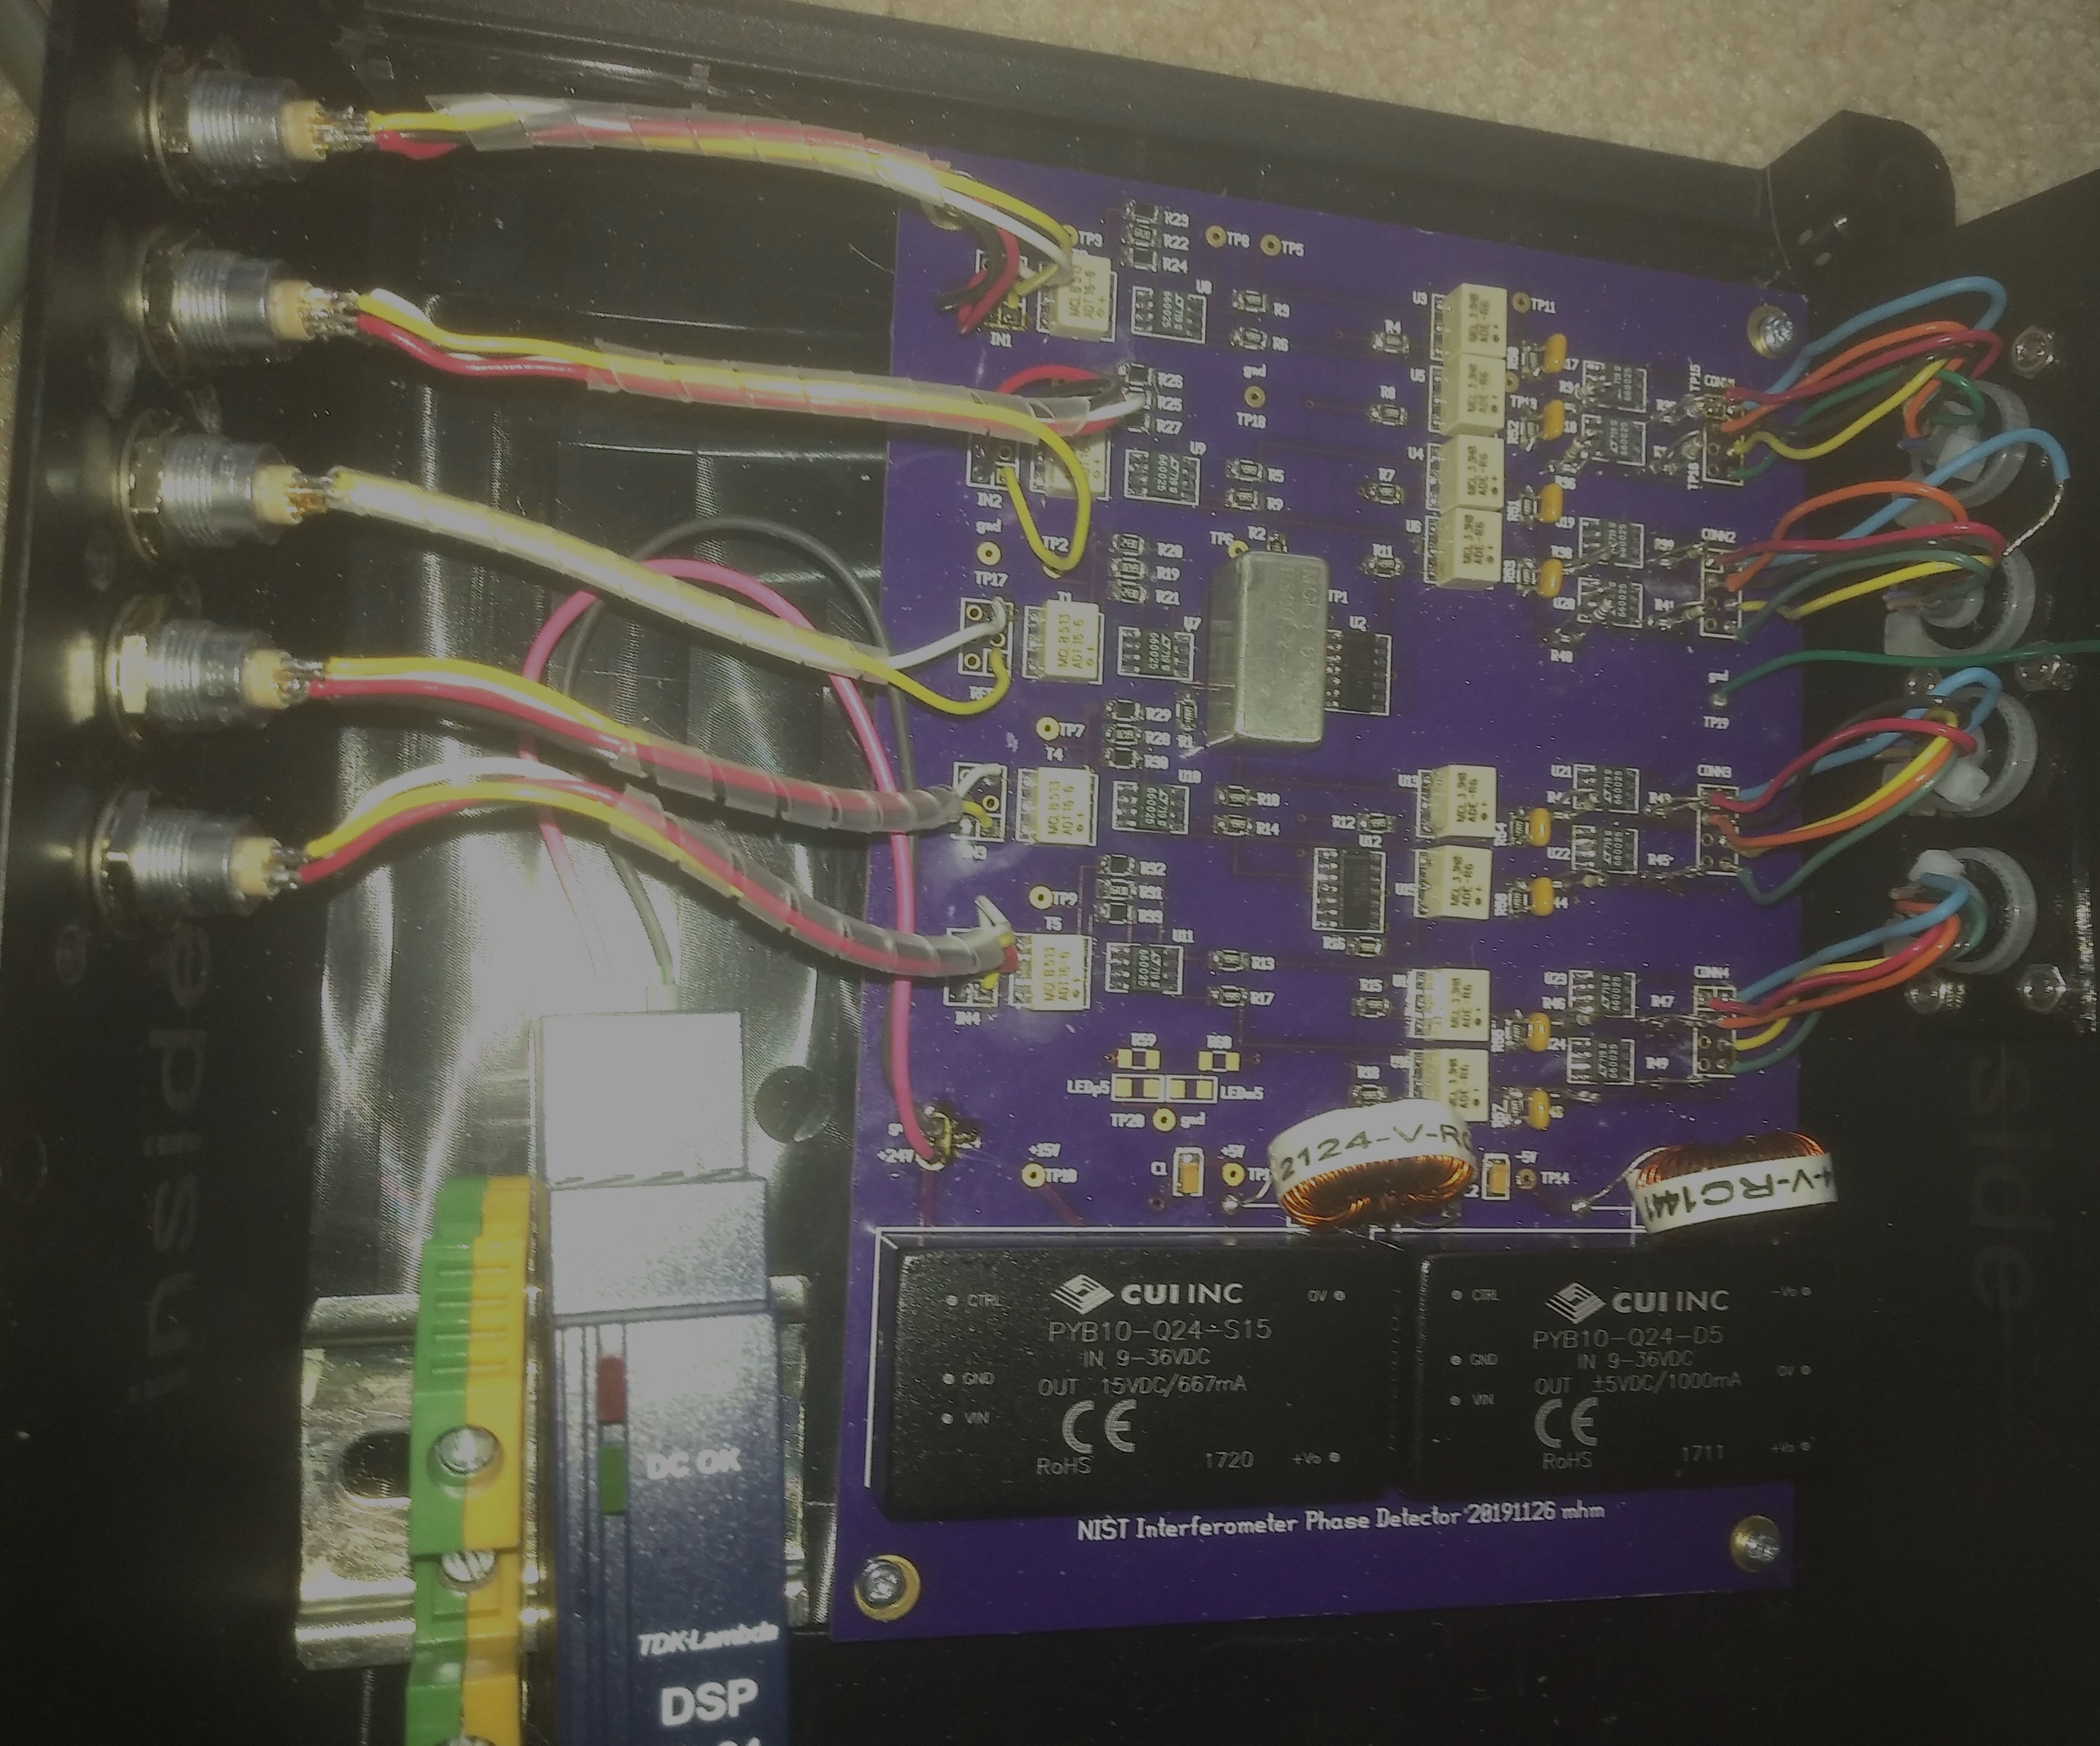
CUI INC
PYB10-Q24-S15
IN 9-36VDC
OUT 15VDC/667mA
RoHS
1720
CUI INC
PYB10-Q24-D5
IN 9-36VDC
OUT ±5VDC/1000mA
RoHS
1711
NIST Interferometer Phase Detector 20191126 mhm
DC OK
TDK·Lambda
DSP