FIES velocity stability

Jeppe Sinkbæk Thomsen
Supervisors: John Telting & Frank Grundahl

April 13, 2023

1 Preamble

This document summarizes the 10 month student project on FIES by Jeppe Sinkbæk Thomsen with supervision from John Telting and Frank Grundahl, and in collaboration with Graham Cox and Joonas Viuho. The report gives measurements on the current short term velocity stability of the spectrograph, exclusively from Thorium Argon lamp (ThAr) spectra.

2 Spectrograph introduction

FIES (FIbre-fed Echelle Spectrograph) is a cross-dispersed echelle spectrograph at the Nordic Optical Telescope, located at Roque de Los Muchachos on La Palma. The most recent paper describing the instrument is Telting et al. (2014). Its highest resolution mode is offered by a 1.3 arcsec fibre with an exit slit, having an effective spectral resolution of R=67000. Since 2014, several upgrades have been carried out to improve the radial velocity precision of the spectrograph. The current fiber bundle (D3) uses octagonal instead of circular fibers, in order to scramble the incoming light. The grating was put in a sealed tank filled with neon gas in 2019. Tests were done using a Fabry-Perot etalon for simultaneous calibration light through a separate fiber, but the etalon is not on la Palma any longer. The fiber layout on the CCD is tightly packed, which limits the use of simultaneous ThAr calibration due to order overlap with the science fiber. The current recommended setup for high-precision radial velocity work is therefore to not use a simultaneous calibration, but rather take a ThAr exposure through the science fiber before and after an observation and interpolate a correction (the ”sandwich” method).

The spectrograph is in a separate and temperature controlled building, and several air insulation layers are present: Building front room (door, aircondition, cryotiger, computers and electronics), instrument room, instrument insulating white box with internal fans and active heaters, spectrograph black box, grating tank filled with neon gas.

3 Project introduction

To investigate and potentially improve the short term precision of the spectrograph, a large amount of ThAr observation series have been taken in the last year. This was done at times where the instrument was not used for regular observing. A description of the standard setup and reduction procedures is also available in appendix A. FIEStool was updated as a part of this project to be able to process ThAr spectra like stellar spectra, without updating the wavelength solution.

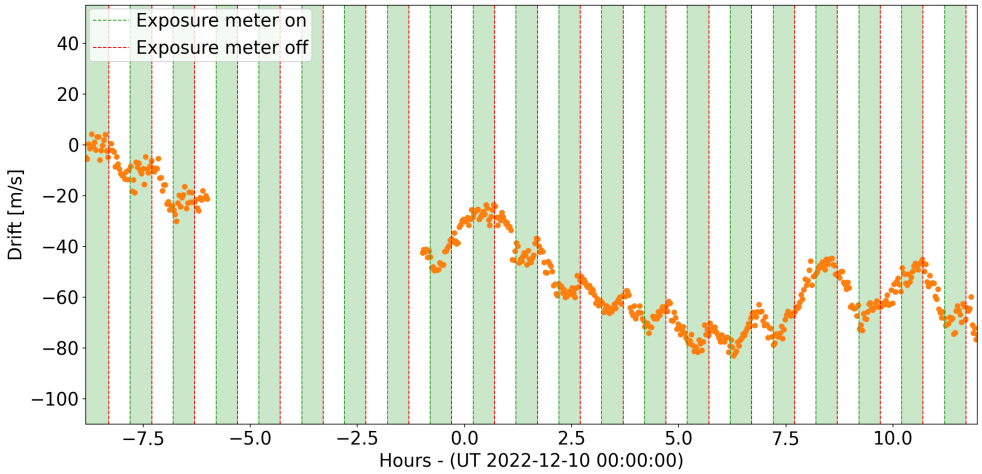

Figure 1: Instrument drift measured with ThAr spectra from 2022/12/09, indicated with orange markers. Overlaid are times exposure meter is switched on (green stripes) and off (red stripes). The green shaded areas indicate when the exposure meter is counting.

The next sections will cover the primary findings of the tests done the last year. In Sect. 4, exposure meter induced drifts are covered. Next, in Sect. 5 the order dependency on measured drift is investigated, possible causes are discussed, and further tests are proposed. After, Sect. 6 covers front room coupling and subzero temperatures. In Sect. 8, an approximation for the sandwich drift correction method is simulated using only ThAr spectra and some basic statistics are shown for most of the ThAr series. Finally, the findings are summarized and concluded upon in Sect. 9.

4 Exposure meter induced drift

In the current setup, the exposure meter induces a significant instrument drift when powered on and counting. On 2022/12/09, a test was performed where the exposure meter was alternated between an on (and counting) state for 30 minutes, and subsequently off for 30 minutes. In Fig. 1, simultaneous drift measurements with ThAr spectra are seen. The green and red stripes indicate the times when the exposure meter is switched on and off, respectively. And the green shaded areas show the time periods when the exposure meter is on and counting. Despite the instrument drifting from a myriad of reasons, it is seen that almost every time the exposure meter is switched on, the instrument starts drifting in positive direction soon after. And when the meter is switched off, the opposite is seen.

On 2023/01/30, another test was performed to see if this effect was reproducible with a different frequency (1h on, 1h off). Plots with and without a 2h moving average subtracted are available in Fig. 2. As a helper plot, the black box temperature is also shown. Winter weather (see Sect. 6) made it difficult to verify from a direct comparison. However, by subtracting a 2h moving average, the effect of the exposure meter is clearly still present in both the ThAr drifts and in the black box center temperature sensor. Around the time the exposure meter is turned on or off, a sharp change in drift occurs. This kind of drift is extremely undesirable, since it is very difficult to account for in long exposure time stellar spectra by using only a sandwiched ThAr correction.

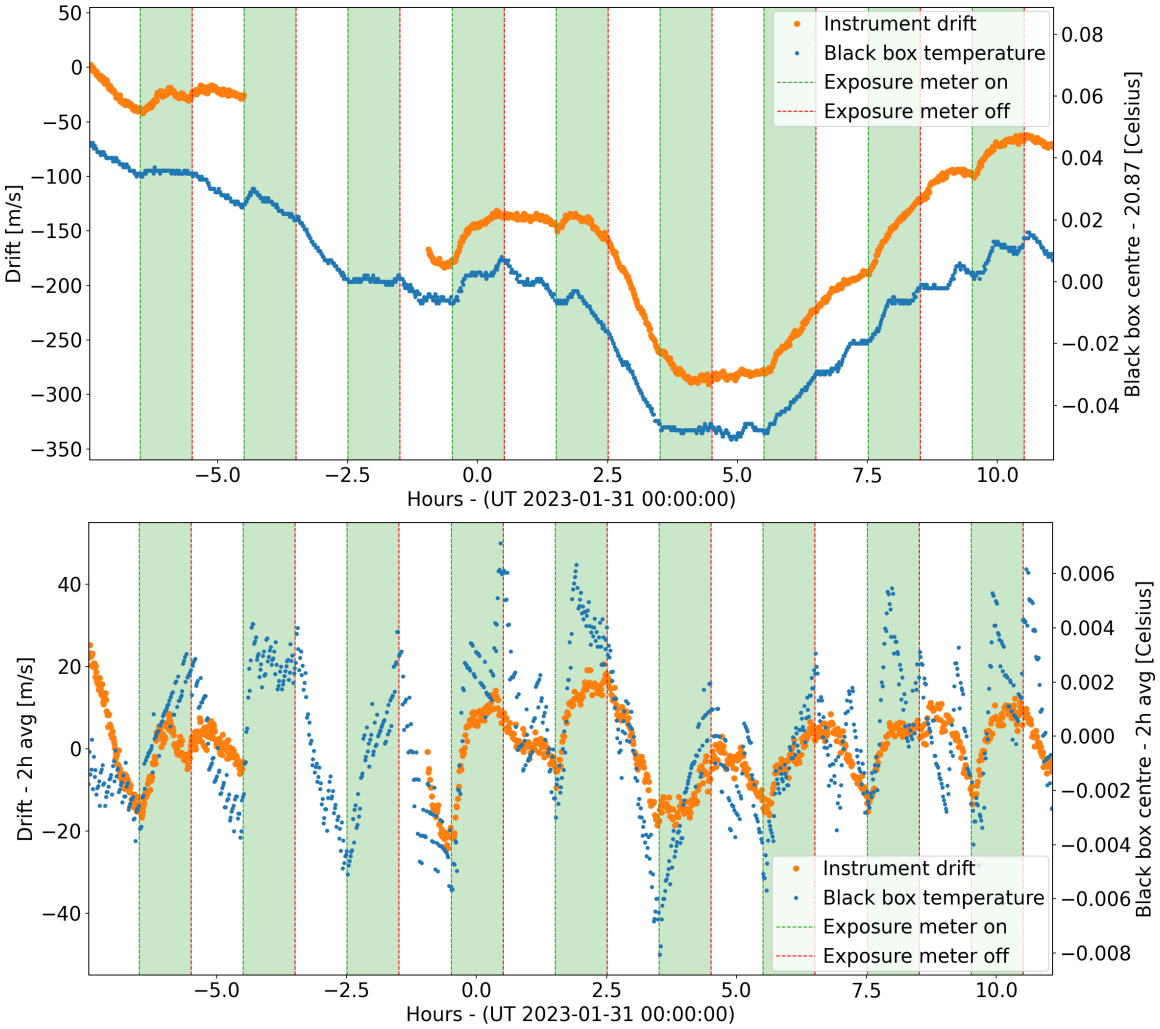

Figure 2: Two plots similar to Fig. 1 for 30/01/2023. Top plot is the unchanged version, while the bottom has had a 2h moving average subtracted. Both plots also include the black box center temperature sensor measurements (blue). The variations in the black box temperature shown in the bottom plot are very close to the sensor precision limit.

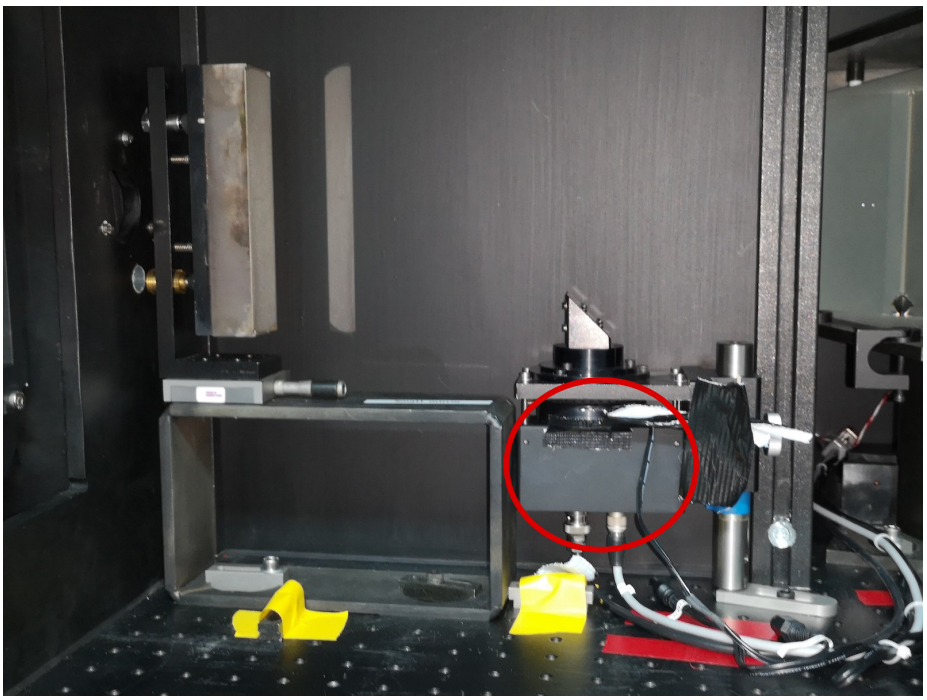

Figure 3: A picture inside the black box, showing the exposure meter (circled in red) proximity to nearby optical elements.

This is evidence both that the exposure meter has a large impact on the instrument drift, and that the act of turning the exposure meter on and off all the time has a significant detrimental effect to the short-term instrument stability. At least when it is not counting.

A picture inside the black box of the exposure meter is shown in Fig. 3. Here, the exposure meter is shown to be located in very close proximity to the mounting of spectrograph optical elements. Heating of this, and through its own mount to the optical table, is the primary suspect for the observed effect on the spectrograph drift.

The exposure meter works by picking up light in the red that does not hit the spectrograph CCD. The currently proposed solution is to move it out of the black box and use a light guide to transport light from its pickup-mirror to the detector. A temporary solution is to keep the exposure meter always on.

The working theory after these results was that the exposure meter heats up the nearby mount to such a degree that there is a time-dependent heat gradient across the mount while the exposure meter is warming up or cooling down. This causes a change in tilt of the mirror due to deformation of the mount, which manifests as a spectrograph pixel shift on the detector. To test this theory, on 2023/03/15 the spectrograph box was opened and temperature sensors 519 (right) and 525 (left) were clamped on both ends of the mount, connected with thermal paste to the metal, and secured with tape. The setup is shown in Fig 4. The day after, the exposure meter was powered off around UT 14:30-14:40 after it had been on (but not counting) continuously for >14 hours. The temperature measurements for both sensors are available in Fig. 5. Before it was powered off, there was a very stable temperature offset of 0.1 Celsius between the sensors. When powered off, the nearest sensor (519) rapidly decreased in temperature, and the other (525) had a much more mild response. This was clear indication that just having the exposure meter powered on gave a strong heat gradient across the mirror mount, and the change in that heat gradient as the

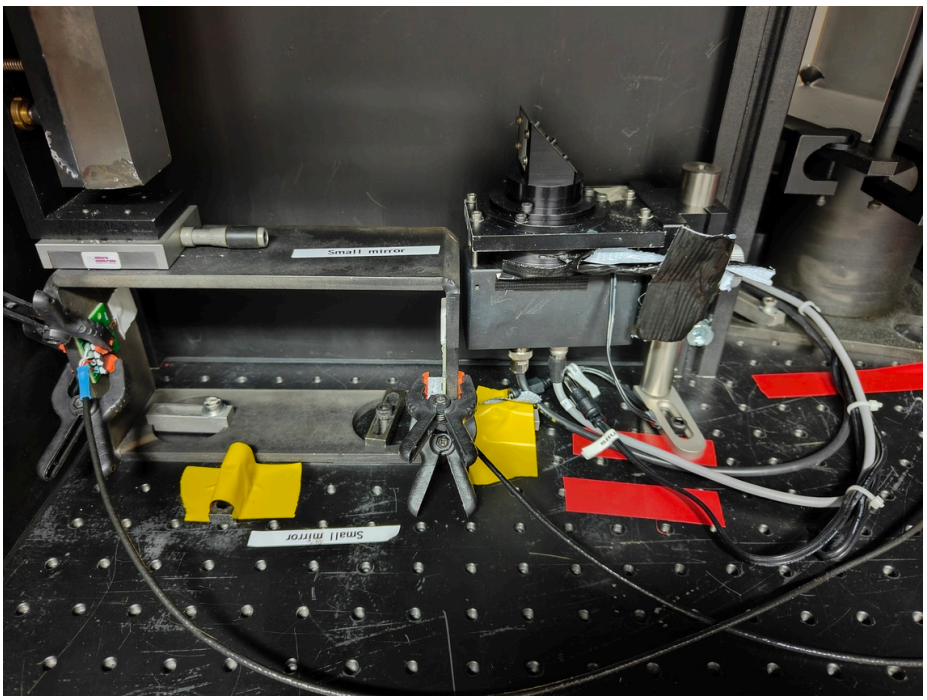

Figure 4: A picture inside the black box, same as Fig. 3, but showing the testing setup for measuring a time-dependent heat gradient across the mirror mount due to the exposure meter powering on or off. The temperature sensors 525 (left) and 519 (right) are clamped on either side of the mirror mount.

on/off state is switched was the likely cause of the shifts.

The test on 2023/03/15-16 also helped show that when powered on continuously, the exposure meter and its surroundings reach an equilibrium temperature eventually.

On 2023/03/19, ThAr spectra were captured while the exposure meter was on and equilibrium had been reached in the temperature sensors. The meter was alternated between counting for 30 minutes, and not counting for 30 minutes. Then, on 2023/03/21, ThAr spectra were captured while the exposure meter was off instead (also in equilibrium). Comparing the two in Fig. 6, it was not possible to measure significant drifts induced by the act of counting/not counting with the exposure meter. The residuals against a 30 min moving average was noted to have RMS 1.5 m/s for 03/19 and 1.8 m/s for 03/21. For the temperature sensors, it was possible to see a small variation in temperature when counting vs. not counting (on in both cases), which are the small periodic variations in Fig. 7. This seems to be far below the measurement precision of the ThAr spectra, compared with the fairly large temperature change when powered off at 21:30 or on at 01:45. The temperature fluctuations from counting occurred regardless of the amount of light hitting the detector.

5 Order dependency on measured drift

In the ideal case, any instrument drift due to either; a pixel shift or; a refractive index change in the gas within the grating tank, should be the same over the whole spectrograph when measured in velocity units (Frank Grundahl & John Telting, private comm.). A pixel shift is a small change in the physical location on the CCD for the incident light in the dispersion direction, while a change in the refractive index due to temperature/pressure variations leads to a direct change in the observed wavelength of the ThAr emission lines.

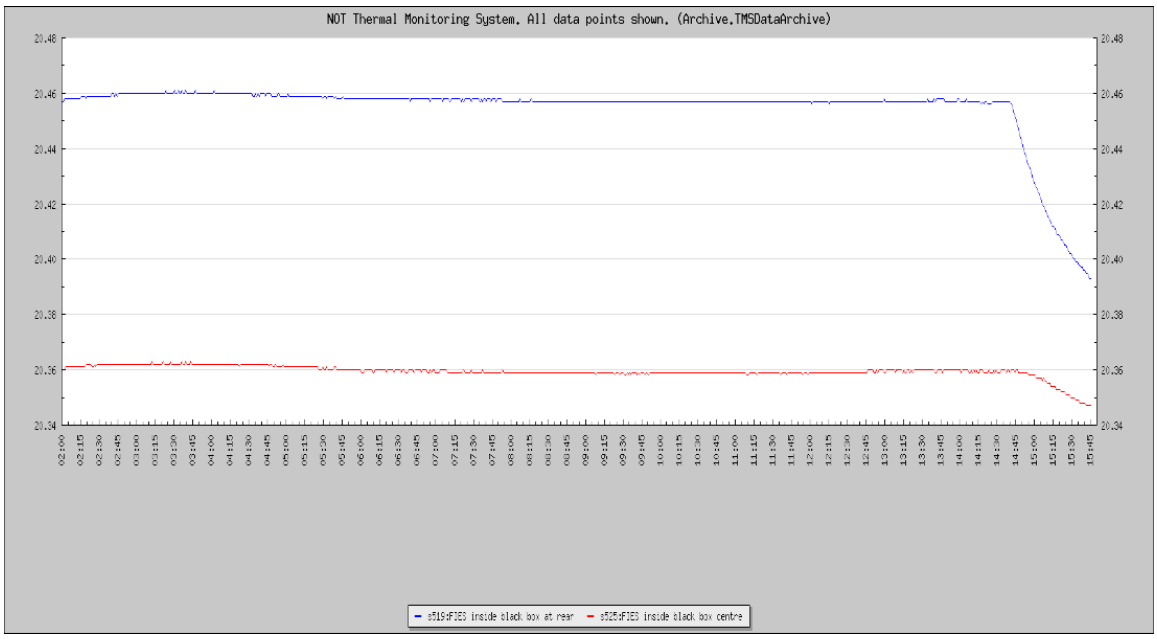

Figure 5: Temperature sensors 519 (blue) and 525 (red) shown in Fig. 4 when the exposure meter is switched off around UT 14:30-14:40 on 2023/03/16 after having been on for >14 hours.

If an order dependence is observed, meaning that it matters whether we measure a velocity drift in the blue or the red part of the spectrograph, that can be an indication of an inaccurate wavelength calibration. The line profile of the instrument is decidedly asymmetric and varies by quite a lot over the spectrograph, which a general order-dependency could likely be related to. Since we fix the wavelength calibration for our tests, then in the ideal case this should show up as an observable but constant order dependence.

In Fig. 8, a test of this assumption has been done for drifts measured on the night of the 19th of February 2023. To measure the order dependency with limited S/N, a rolling 30 order measurements were averaged to produce 40 different observations per ThAr spectrum. The figure shows this as a heat-map going from blue in the bluest 30 orders, to red in the reddest (while not using the 21 reddest orders). During the night, a clear but relatively constant order dependency was found, in line with expectations. Around the early morning hours, the order dependency changes slightly. Time/temperature dependent changes in the order-dependency could indicate a focus or trace change. But it has not been investigated in detail.

5.1 Applying drifts to stellar spectra

When (time-dependent) order-dependency is present in the drifts measured by ThAr spectra, it makes it difficult to apply a derived correction to stellar spectra, since the signal to noise distribution is different for stellar RV measurements. Especially spectra taken on different nights. This makes the wavelength calibration critical for high precision work. On Fig. 9, ThAr measurements are shown where the wavelength solution has been re-fit with FIEStool (IRAF ecreidentify) for each of them. On a single night, re-fitting the wavelength solution shows a clear worsening in the overall precision compared with fixing the wavelength solution, see Fig. 10. In our ThAr reductions, we achieve high-precision measurements since we can constrain ourselves to a single night and a single wavelength solution each time. Given that FIES is not a vacuum spectrograph, it is assumed that it

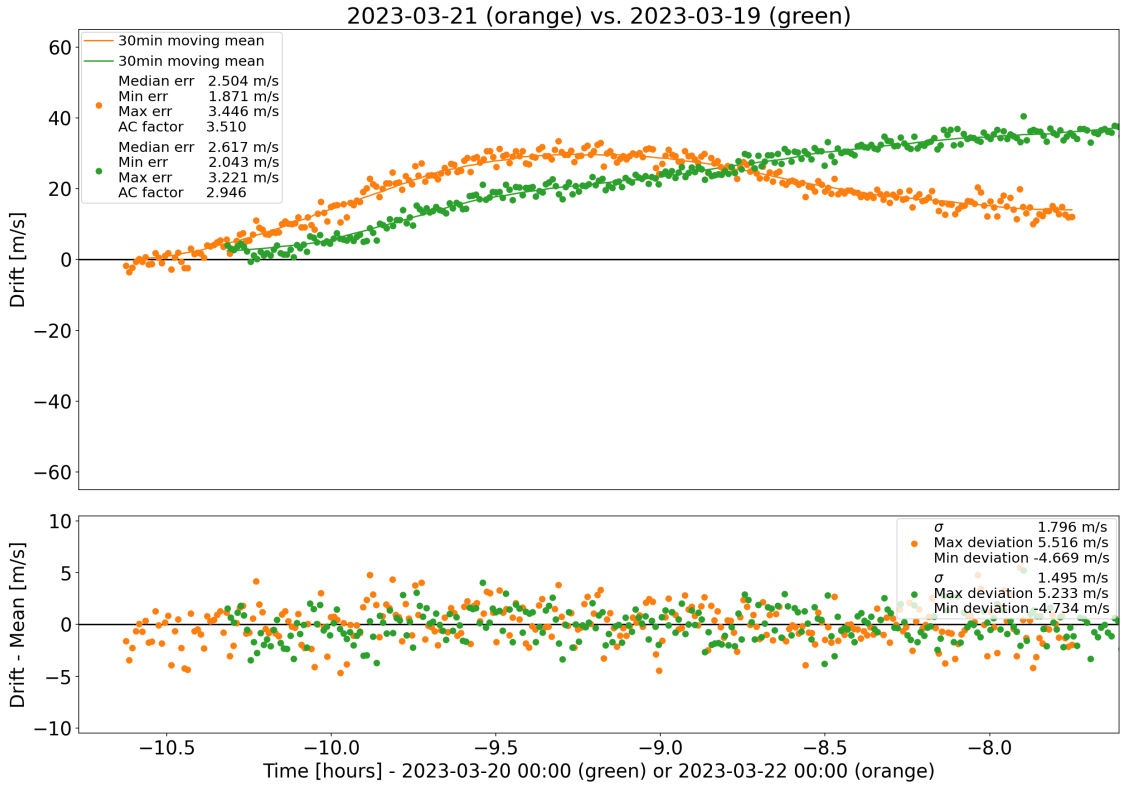

Figure 6: Comparison between the ThAr drifts measured on 2023/03/19 (green) and 2023/03/21 (orange). Top panel shows the drift from each night, measured separately, as well as a 30min moving mean for each. Legend includes median order-to-order error, minimum and maximum error, and anti-correlation (AC) factor. AC factor is a simple root-mean-square (RMS) of each value vs. the mean of the previous and next value. Bottom panel shows the residuals after subtracting the 30min moving mean. Legend gives RMS of the residuals (σ), as well as maximum and minimum deviation.

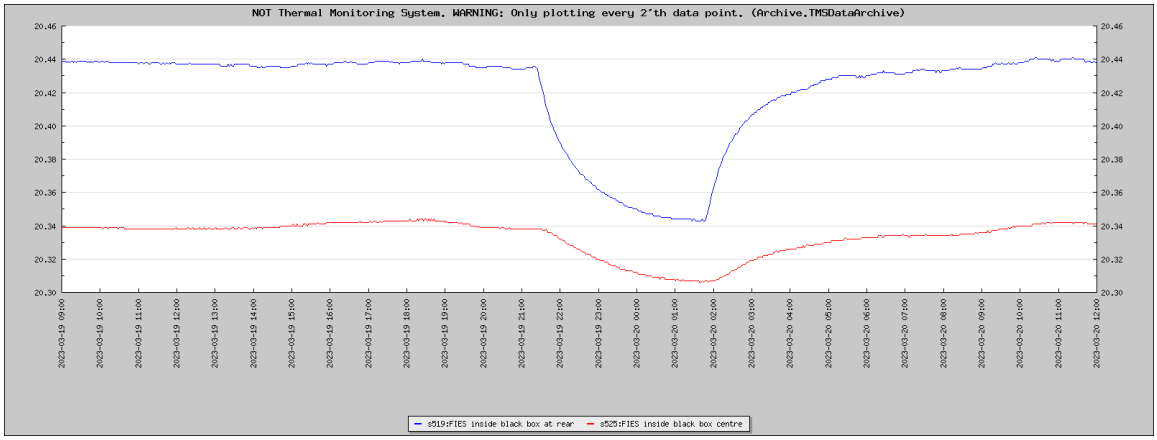

Figure 7: Temperature sensors 519 (blue) and 525 (red) shown in Fig. 4. The exposure meter is switched off at ∼21:30 and on again at ∼01:45. Outside of that period, from UT 13 to UT 12 the next day the exposure meter is alternated between counting and not counting, for 30 minutes each. Before 21:30, a ThAr lamp is on and light is detected on the exposure meter. After 01:45, only background light and dark current is detected.

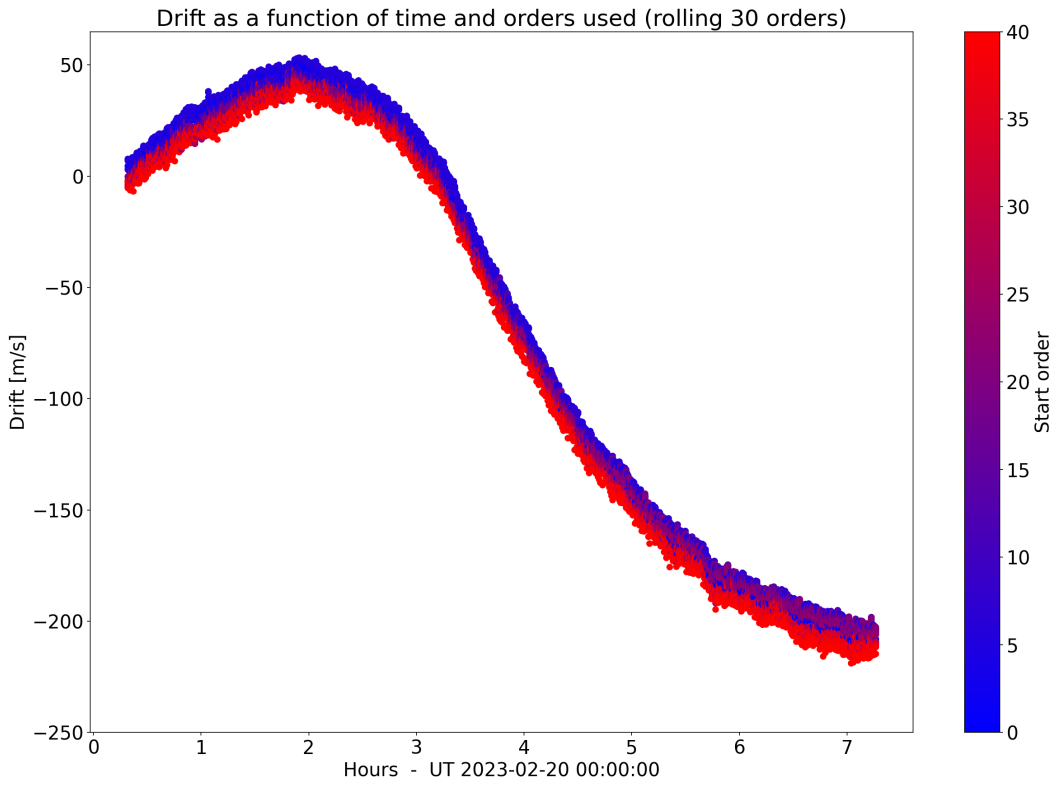

Figure 8: Rolling RV measurements. Heatmap indicates starting aperture order used. For example, with starting order 0, orders used are 0-29.

has to be re-calibrated every day to ensure a stable extracted spectrum and wavelength solution. This puts considerable stress on the reduction technique for any work desiring <5-10 m/s precision on stellar RVs.

It might be possible to correct for an order dependency after-the-fact by measuring the ThAr drift independently for each order. Then, to correct the stellar RVs, they should be extracted on an order-by-order basis. After, each order can be corrected separately, or the average ThAr drift can be weighted by the stellar order RV errors. The key goal is to account for the difference in S/N weighting between the ThAr drift and the stellar RV.

5.2 FIEStool wavelength solution

On Fig. 11, the long term stability of the FIEStool wavelength solution is investigated, by comparing all ThAr series captured with fiber 4 throughout this 10 month period. These all have independent wavelength solution fits with FIEStool. The drift template is the average spectrum from 2022/08/24, which was the first ThAr series taken after ThAr lamp replacement on the same day. However, the master solution used for this observation is from the old lamp, which could explain the offset from later nights. The figure indicates an offset in instrumental drift due to lamp degradation/change. As also found in Fig. 9, the spread is higher when refitting the wavelength solution between each frame. Additionally, there are night-to-night offsets of $\sim$ 10 m/s over time, as can be more clearly seen by comparing only the measurements taken after lamp change. In Fig. 12, a zoomed-in view of 24h shows that even on short timescales the wavelength solution fit can obtain very significant, correlated deviations. These are likely from nights where the number of accepted lines vary by a lot in the wavelength solution fit.

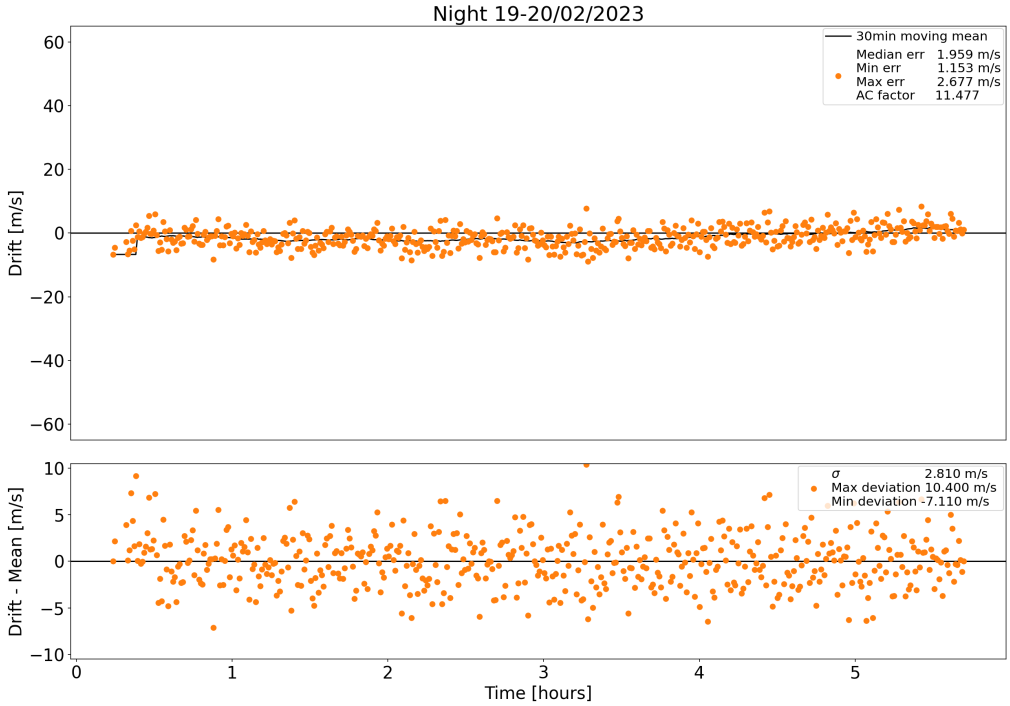

Figure 9: The fiber 5 ThAr measurements from the night of February 19-20 2023, in a similar plotting style as Fig. 6, but where the wavelength solution has been refit for each spectrum independently. A 30min moving mean is shown as a black line in the top panel, and the bottom shows the variation from this mean.

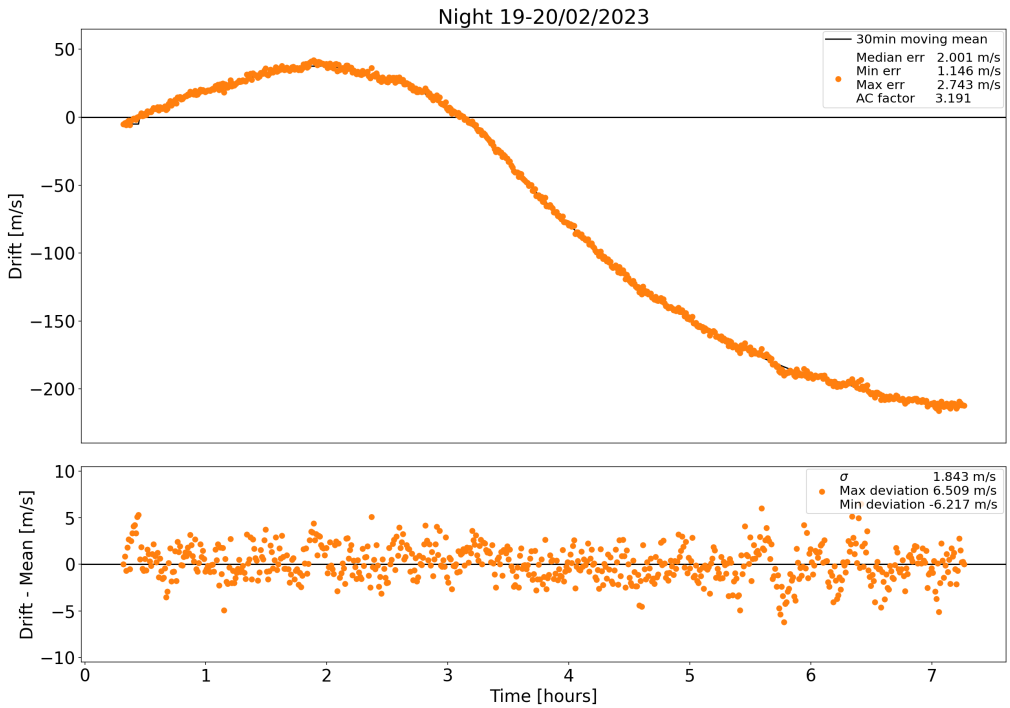

Figure 10: Same plot as Fig. 9, but where the first ThAr defines the wavelength solution for all subsequent spectra.

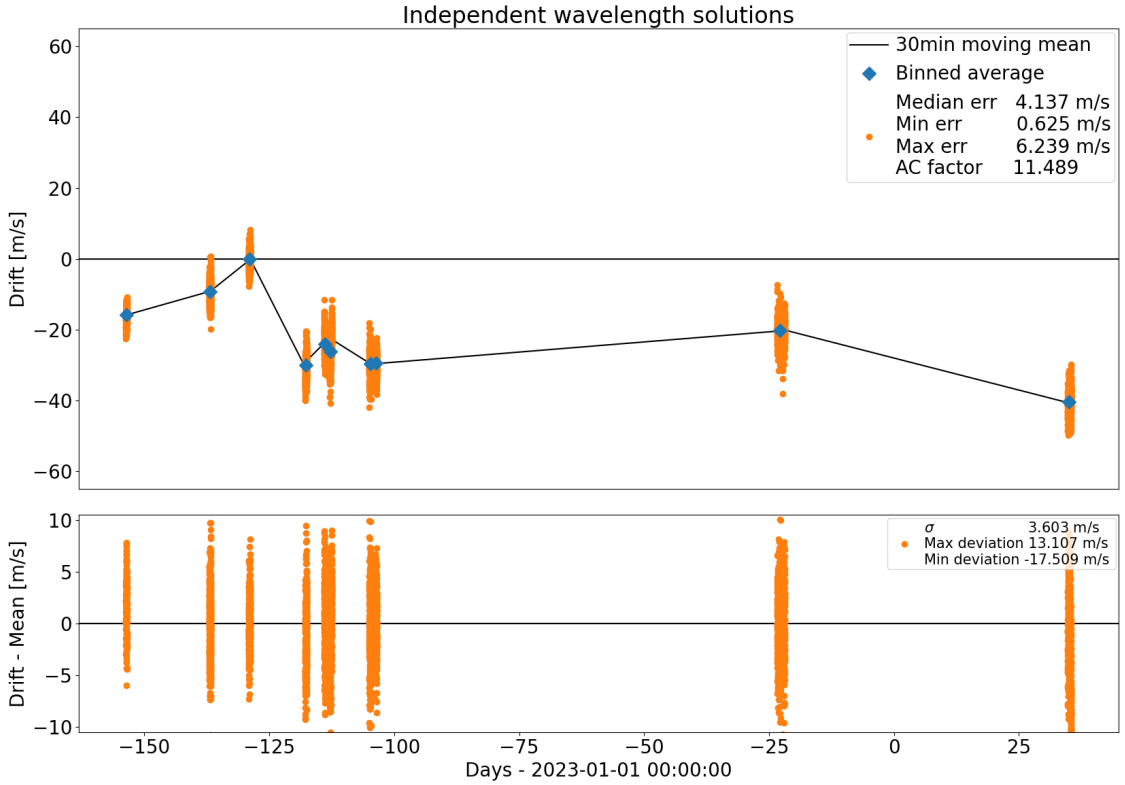

Figure 11: All ThAr series with fiber 4, similar plotting style as Fig. 6, but where the wavelength solution has been refit for each spectrum independently. x-axis has Days instead of Hours. Included in blue is shown the binned average of each ThAr series. Drift template only has spectra from 2022/08/24 (third group from the left). Immediately before the 2022/08/24 series, the top ThAr lamp was replaced.

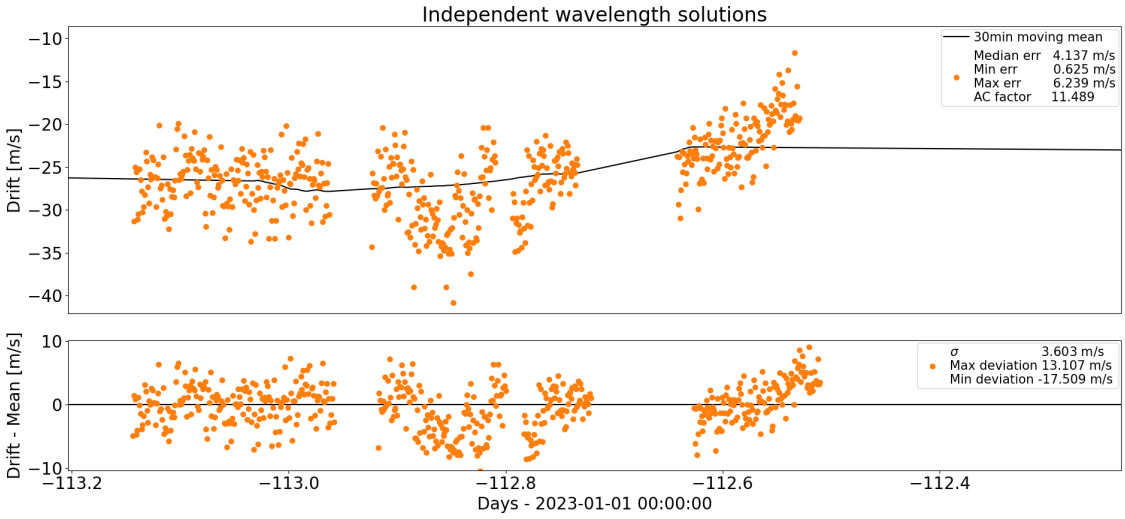

Figure 12: Zoomed in view of Fig. 11 showing only 24h.

The recommended way for vacuum spectrographs of ensuring a consistent wavelength solution while still accounting for night-to-night variations is to only refit a constant offset or linear term in the 2D polynomial fitting pixels and orders, while keeping the higher degree terms fixed (Frank Grundahl, private communication). It could be investigated using RV standard star observations if applying this practice offers improvements for FIES spectra taken on different nights. The danish IDA summer school took several sandwiched 180s spectra of σ Draconis (HD185144) specifically for this purpose, in the time period 2022/08/04 – 2022/11/12.

5.3 Proposed tests

If the order-to-order dependency is not caused by the wavelength solution, but rather a temperature-induced focus change or shift on the CCD, it could be measurable in a 2D fashion using the full-frame images. Potential analysis methods include 2D cross-correlation, or even by application of standard techniques used for stellar image analysis (like SExtractor) to fiber 5 ThAr frames as suggested by John.

A dedicated project to test the wavelength calibration on stellar spectra could involve wavelength calibration and data reduction of standard star observations with and without using IRAF. Here it could be investigated if RV improvements are seen when only refitting a constant offset and/or a linear term for the wavelength calibration between separate nights.

6 Weather sensitivity and front room heat coupling

This winter, whenever outside temperatures dropped below freezing, the front room temperatures would oscillate strongly. This was seen quite consistently by telemetry monitoring. It was also discovered that when these large oscillations occurred, it was possible to see simultaneous, but smaller, variations in the instrument temperature sensors, but not the instrument room itself. An example from 2023/02/04 is shown in Fig. 13. Shown sensor telemetry is black box centre (blue), grating tank housing (green), inside grating tank under grating (red), front room temperature (purple). Also included is simultaneously captured ThAr drift (orange). It is not directly evident that these oscillations have a measurable impact on the ThAr drift in the short term, but it has also not been ruled out due to measurement precision. The reason for the apparent lack of large oscillations in the ThAr drift is likely the very minimal amount of variation in the temperature inside the grating tank. This is clear evidence that even just the extra layer of air insulation offered by sealing the grating in a tank has a beneficial effect to the stability of the spectrograph, regardless of the gas inside.

Overall, the stability of the instrument was much worse whenever temperatures approached freezing, and the winter period had lower precision on ThAr drift measurements than the (few) observations taken in July and August (with a different lamp and master solution). While the cause of the oscillations in the front room is not clear, they have helped illustrate that there is a measurable heat coupling between the front room and the instrument table, bypassing the instrument room. The only exposed metal connecting the two is the cryotiger tubes, which are therefore the primary suspects.

At the moment it is still debated what can be improved, and if improvements are necessary. Since there is no obvious solution to reduce the heat coupling between the front room and instrument, any improvements would be limited to stabilizing the front room temperature near the cryotiger.

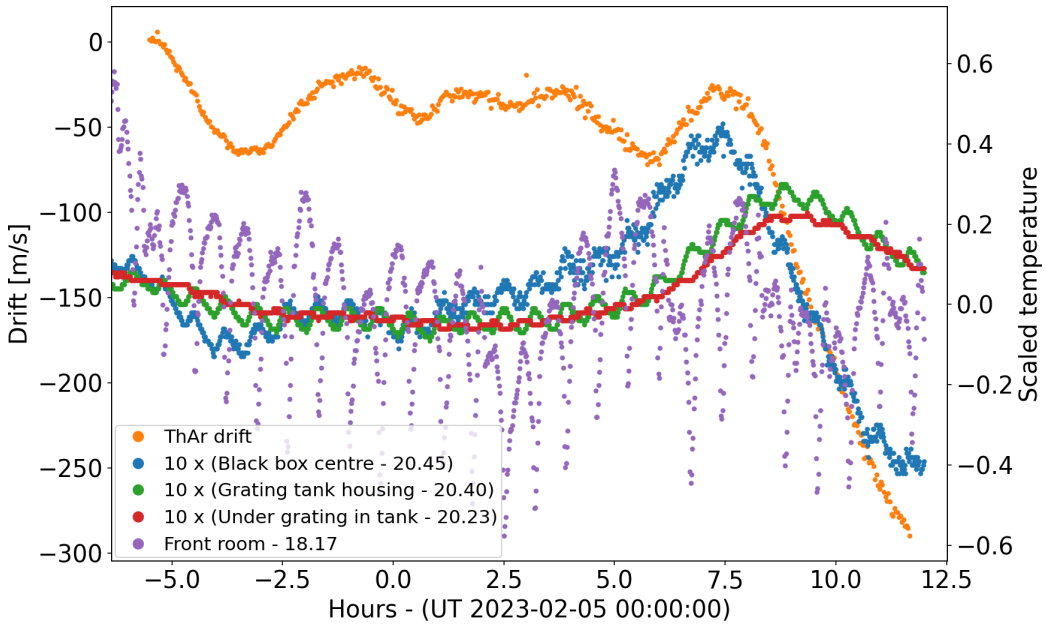

Figure 13: ThAr drift from 2023/02/04 overlain with temperature telemetry. All temperature measurements have been corrected for their mean value. The front room unit is in Celsius, while the other three (black box, grating tank housing, in grating tank) are in decigrade Celsius.

7 Heaters off

On the night of 24th of August 2022, ThAr measurements were done after the heaters in the white box had been switched off (and temperature equilibrium had been reached). Results can be seen in Fig. 14. This measurement overall shows low variation, low errors, but also not a lot of correlation between the different ThAr spectra. The ThAr lamp had just been changed before this, but the master wavelength solution was not redefined yet.

It was difficult to compare this run with the nearest ones. On 2022/07/30, the order-to-order precision (1.2 m/s vs. 1.2 m/s) was similar but the spread against a moving mean was worse (2.2 m/s vs. 1.9 m/s). The 2022/07/30 run was also shorter. However, on the subsequent runs the master solution had been changed and the internal error was higher than previously.

On 2023/03/24-26, the heaters off test was re-performed with fiber 5 instead, as seen in Fig. 15. Since the setup on the heaters off tests on 2023/03/21 and 2023/03/19 (Fig. 6) are similar, they can be used as comparison. On Fig. 16, 17, and 18 the 2023/03/24-26 has been split into sub-groups for easier viewing.

The primary conclusion is that there is no apparent improvement in the scatter compared with Fig. 6. In fact, it appears that point-to-point scatter is larger than before the heaters were switched off. There is also indication that the precision is decreasing over the run. The calibration frames were taken in the middle of the run, so if it was caused by distance to the calibration (in time) then the middle group should be the most stable. So the cause can either be weather/lamp/master-calib, or it can be that leaving the heaters off slowly degrades the precision on longer timescales.

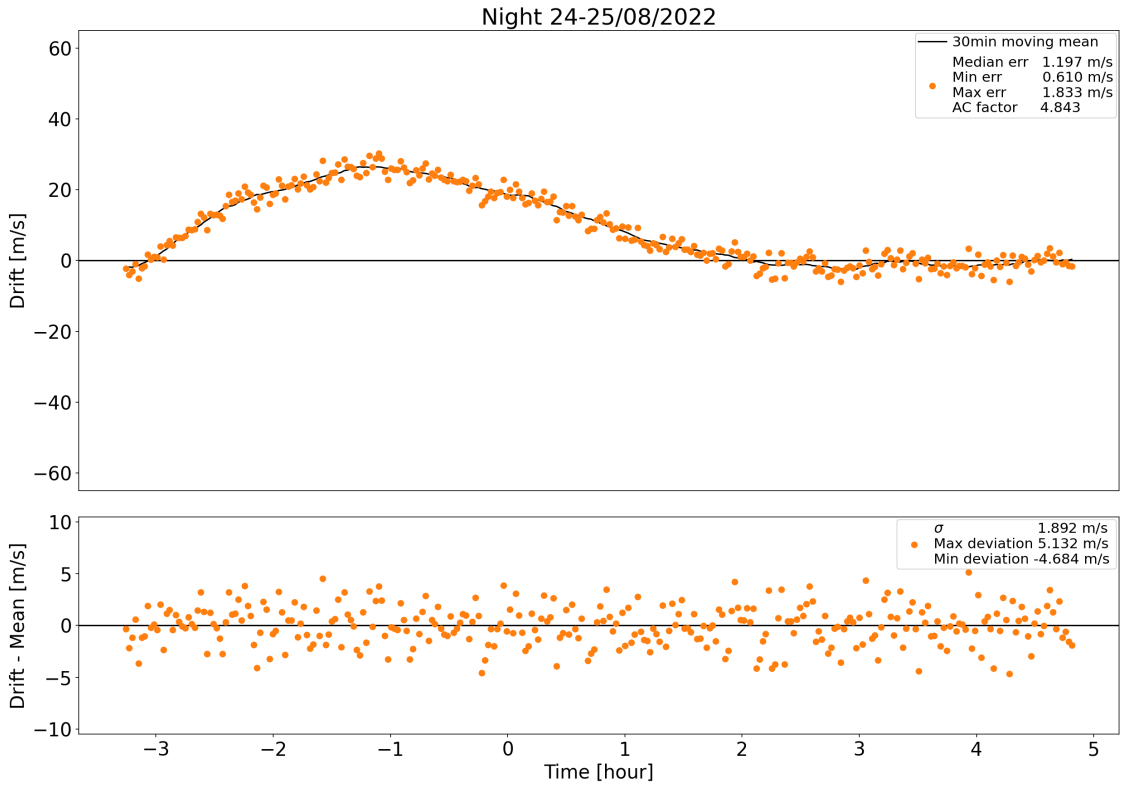

Figure 14: Drifts measured for the heaters off test on the night of the 24th of August 2022. Plotting style is similar to Fig. 6.

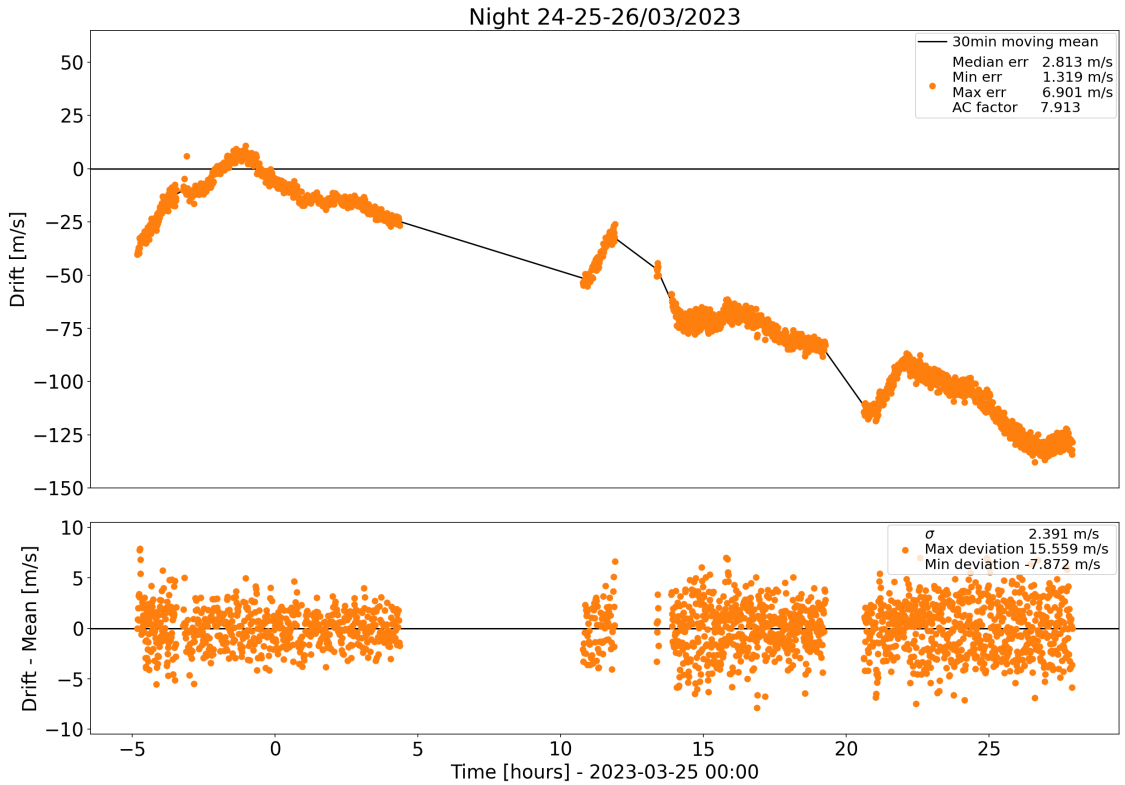

Figure 15: All the ThAr measurements from the heaters off test around 2023/03/25.

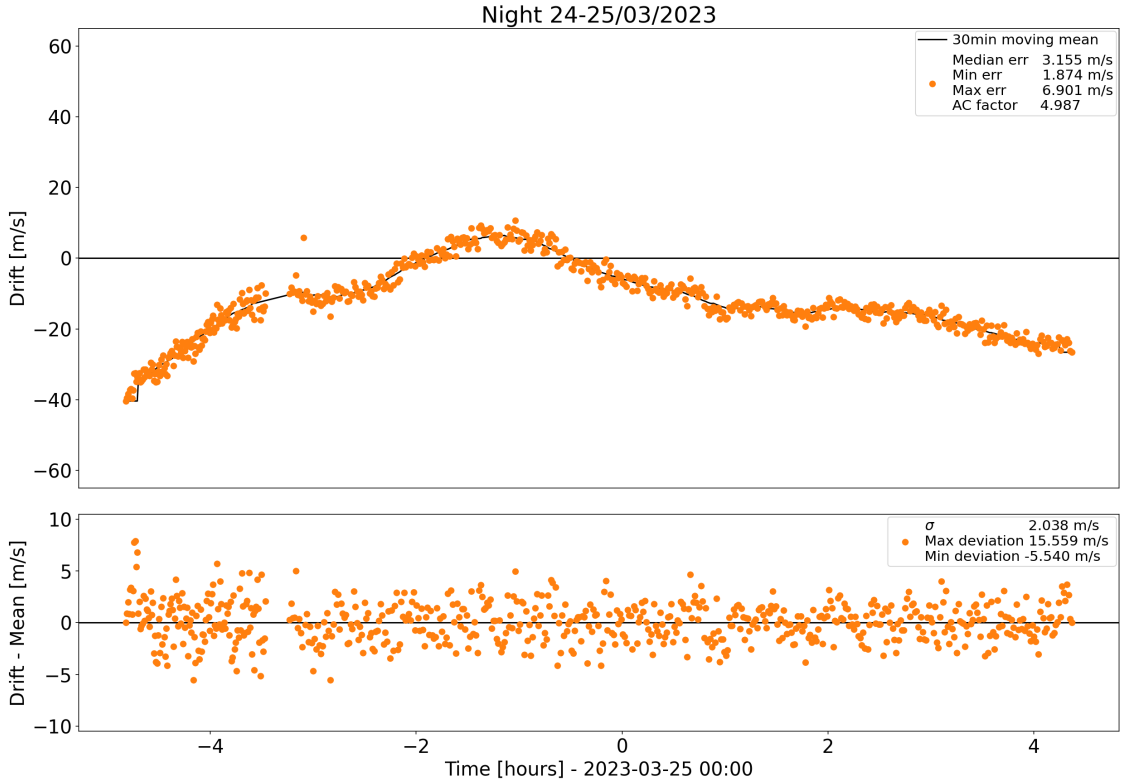

Figure 16: View of the first partition of Fig. 15, showing statistics from only this part.

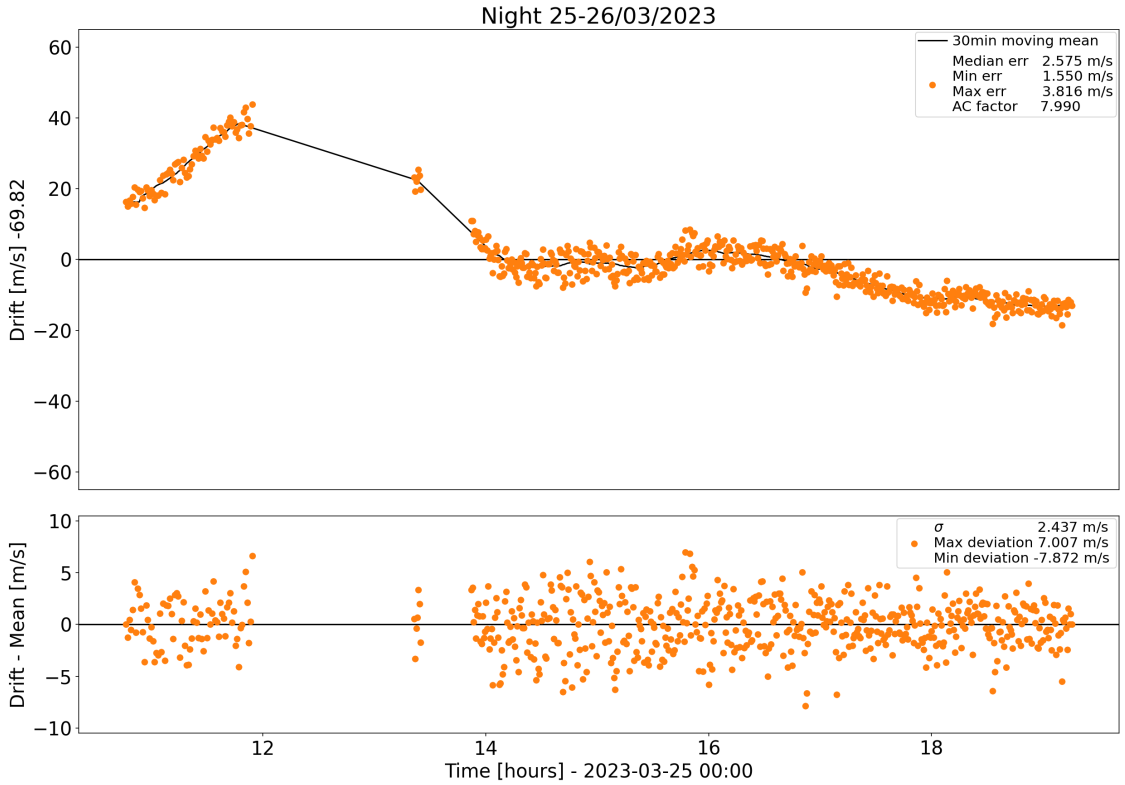

Figure 17: View of the second partition of Fig. 15, showing statistics from only this part.

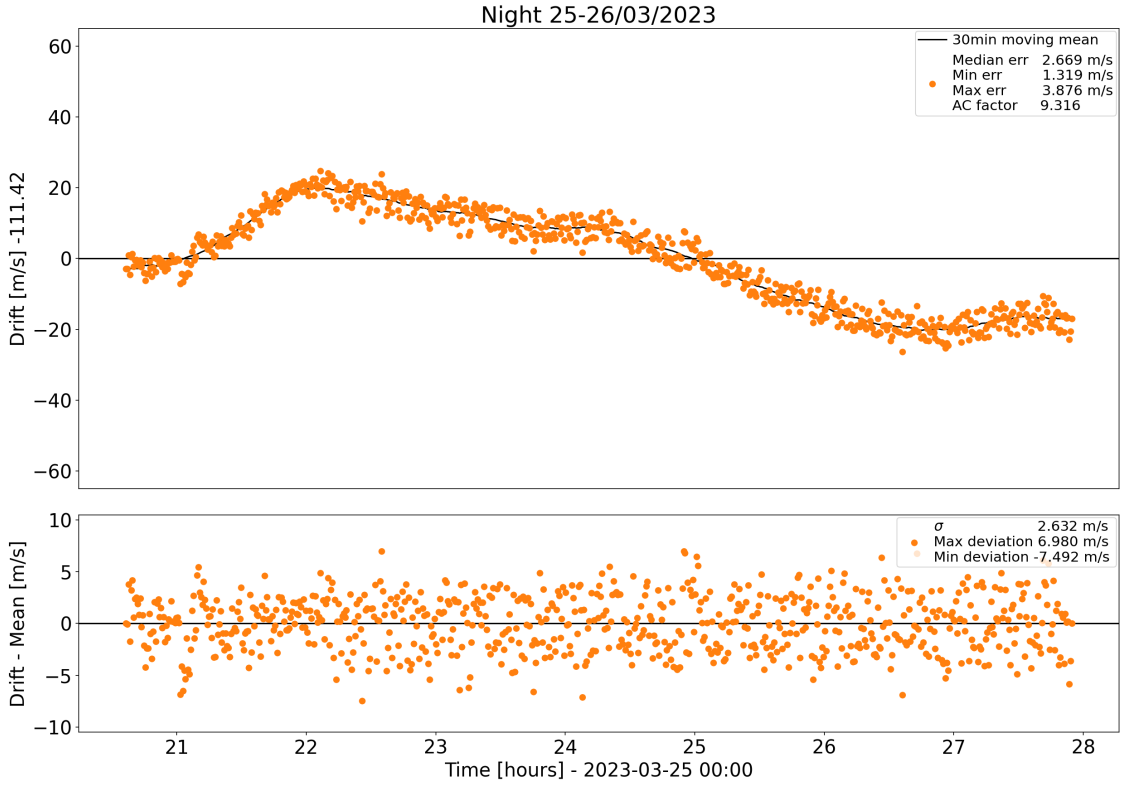

Figure 18: View of the third partition of Fig. 15, showing statistics from only this part.

8 Simulated stellar observations

With the spectrograph design of FIES, high-precision RV work necessitates use of the "sandwich" method, in which a ThAr measurement is taken immediately before and after the stellar spectrum. A drift correction is then linearly interpolated from these two spectra.

To give estimates on the short term precision of the spectrograph, it is possible to simulate this behavior using the ThAr series by averaging several ThAr measurements to produce one observation covering a longer exposure time. By using the ThAr spectra immediately before/after this average, we can perform the sandwich correction and measure the left-over residual. This will indicate the short-term precision available when not using simultaneous ThAr measurements.

As an example, Fig. 19 shows simulated residuals after sandwich correction from a night with a high degree of short-term variations (the exposure meter was alternated between on and off every 30 minutes). On the other end, Fig. 20 shows performance on a calm day without any significant short-term variations.

In the presence of significant short-term variation, only observations with shorter exposure time can accurately correct for it. Still, even the 60min average in Fig. 19 shows a precision of 4.4 m/s, which is pretty decent.

In Table 1, these statistics are recorded for most of the ThAr series taken. Note on the Table: The small upper/lower numbers on the STD and SW columns are the 0.16 and 0.84 quantiles of the absolute deviations. They are not proper measures of scale, which is why they are not symmetric around the RMS. To convert a 0.5 quantile (median) of the absolute deviation to a measure of scale (σ), it has to be multiplied by ~ 1.4826 when assuming a

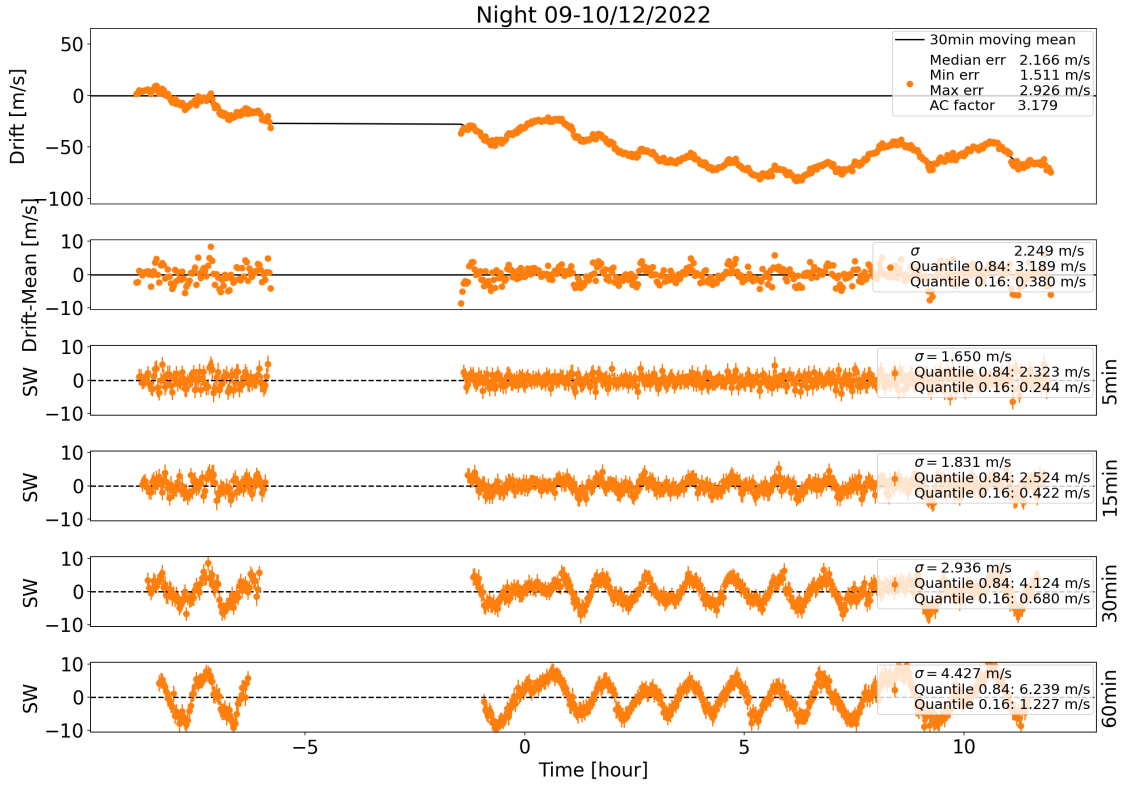

Figure 19: Drift series 2022/12/09, residuals vs. 30 min mean, and simulated long-exposure time (5, 15, 30, 60 min) observations corrected for drift with sandwich method.

normal distribution.

When the instrument/weather is stable, exposure times of 15-30 minutes show slightly improved precision over 5min. This can be explained as a balancing act between the averaging done by a longer exposure, and the accuracy of the preceding/succeeding ThArs used in the sandwich correction. When temperatures are very stable, an average short term precision of 1.1-1.3 m/s can be attained using the sandwich method (depending on exposure time). Otherwise, it might be more reasonable to expect an average short term precision of 1.5-2.0 m/s for 15-30min exposures.

Statistics on ThAr series							
Date	n	Error	STD	SW (5m)	SW (15m)	SW (30m)	SW (60m)
2022/07/30	175	$1.17^{2.01}_{0.73}$	$2.25^{3.30}_{0.36}$	$1.98^{2.64}_{0.33}$	$1.69^{2.42}_{0.25}$	$1.70^{2.24}_{0.41}$	$2.31^{3.04}_{0.38}$
2022/08/24	303	$1.20^{1.83}_{0.61}$	$1.89^{2.75}_{0.35}$	$1.80^{2.59}_{0.32}$	$1.50^{2.14}_{0.30}$	$1.57^{2.21}_{0.37}$	$1.72^{2.45}_{0.39}$
2022/09/04	358*	$1.76^{2.89}_{0.92}$	$2.59^{2.70}_{0.34}$	$1.29^{1.89}_{0.25}$	$1.31^{1.44}_{0.25}$	$1.44^{2.08}_{0.35}$	$2.52^{3.66}_{0.37}$
2022/09/08	420	$1.90^{2.65}_{0.79}$	$2.08^{2.94}_{0.35}$	$1.65^{2.37}_{0.25}$	$1.55^{2.23}_{0.35}$	$1.68^{2.40}_{0.31}$	$2.28^{3.20}_{0.50}$
2022/09/18	399	$1.96^{3.24}_{1.13}$	$2.14^{2.90}_{0.49}$	$1.85^{2.54}_{0.29}$	$1.70^{2.28}_{0.35}$	$1.74^{2.47}_{0.32}$	$1.83^{2.48}_{0.40}$
2022/12/08	628	$1.86^{2.71}_{1.17}$	$3.44^{4.89}_{0.70}$	$3.16^{4.35}_{0.59}$	$2.88^{4.11}_{0.62}$	$2.73^{3.76}_{0.60}$	$2.96^{4.07}_{0.70}$
2022/12/09	530*	$2.17^{2.93}_{1.51}$	$2.25^{3.19}_{0.38}$	$1.65^{2.32}_{0.24}$	$1.83^{2.52}_{0.42}$	$2.94^{4.12}_{0.68}$	$4.43^{6.34}_{1.23}$
2023/01/29	402	$2.10^{3.76}_{1.13}$	$2.34^{3.03}_{0.47}$	$2.00^{2.71}_{0.37}$	$1.98^{2.80}_{0.42}$	$2.09^{2.87}_{0.38}$	$3.57^{4.75}_{0.74}$
2023/01/30	1037*	$1.63^{2.33}_{0.90}$	$2.51^{3.32}_{0.44}$	$1.25^{1.71}_{0.24}$	$1.70^{2.37}_{0.33}$	$2.83^{3.97}_{0.55}$	$5.44^{7.93}_{1.23}$
2023/02/04	402	$2.25^{3.78}_{1.36}$	$2.25^{3.03}_{0.44}$	$1.89^{2.61}_{0.39}$	$1.88^{2.52}_{0.42}$	$2.04^{2.87}_{0.28}$	$3.48^{4.84}_{0.68}$
2023/02/19	669	$2.00^{2.74}_{1.15}$	$1.84^{2.51}_{0.36}$	$1.14^{1.67}_{0.24}$	$1.40^{1.91}_{0.26}$	$1.70^{2.35}_{0.41}$	$3.60^{5.11}_{1.06}$
2023/03/19	484	$2.62^{3.22}_{2.04}$	$1.51^{2.05}_{0.32}$	$1.10^{1.58}_{0.25}$	$1.09^{1.54}_{0.22}$	$1.21^{1.81}_{0.22}$	$1.32^{1.89}_{0.26}$
2023/03/21	292	$2.50^{3.45}_{1.87}$	$1.78^{2.55}_{0.33}$	$1.37^{1.93}_{0.29}$	$1.31^{1.86}_{0.25}$	$1.50^{2.08}_{0.30}$	$3.11^{4.32}_{0.81}$
2023/03/24	637	$3.16^{6.90}_{1.87}$	$2.04^{2.69}_{0.41}$	$1.50^{2.04}_{0.28}$	$1.45^{1.92}_{0.27}$	$1.47^{2.02}_{0.31}$	$1.90^{2.61}_{0.35}$
2023/03/25 (1)	623*	$2.58^{3.82}_{1.55}$	$2.44^{3.46}_{0.53}$	$1.86^{2.64}_{0.34}$	$1.87^{2.61}_{0.37}$	$1.87^{2.58}_{0.35}$	$3.00^{3.75}_{0.44}$
2023/03/25 (2)	729	$2.67^{3.88}_{1.32}$	$2.63^{3.70}_{0.58}$	$2.00^{2.85}_{0.37}$	$1.99^{2.83}_{0.47}$	$1.99^{2.89}_{0.40}$	$2.40^{3.34}_{0.48}$

Table 1: Statistical measures for reduced ThAr series. Columns show (in order left to right): Date of initial observation, number of spectra n in series (stars * mark data-sets where large gaps are present). "Error" refers to median internal (order-to-order) error, and upper/lower values indicate maximum and minimum error in the data. "STD" refers to the RMS vs. a moving 30 minute mean, and upper/lower numbers are the absolute deviation quantiles 0.84 and 0.16, respectively. "SW" refers to simulated sandwich correction observations where all measurements within a time-interval (indicated in the parenthesis) are averaged, after which a correction is made using the ThAr spectrum immediately before and after. The RMS of the residuals are then indicated in the table, along with absolute deviation quantiles 0.84 and 0.16 as the upper and lower values.

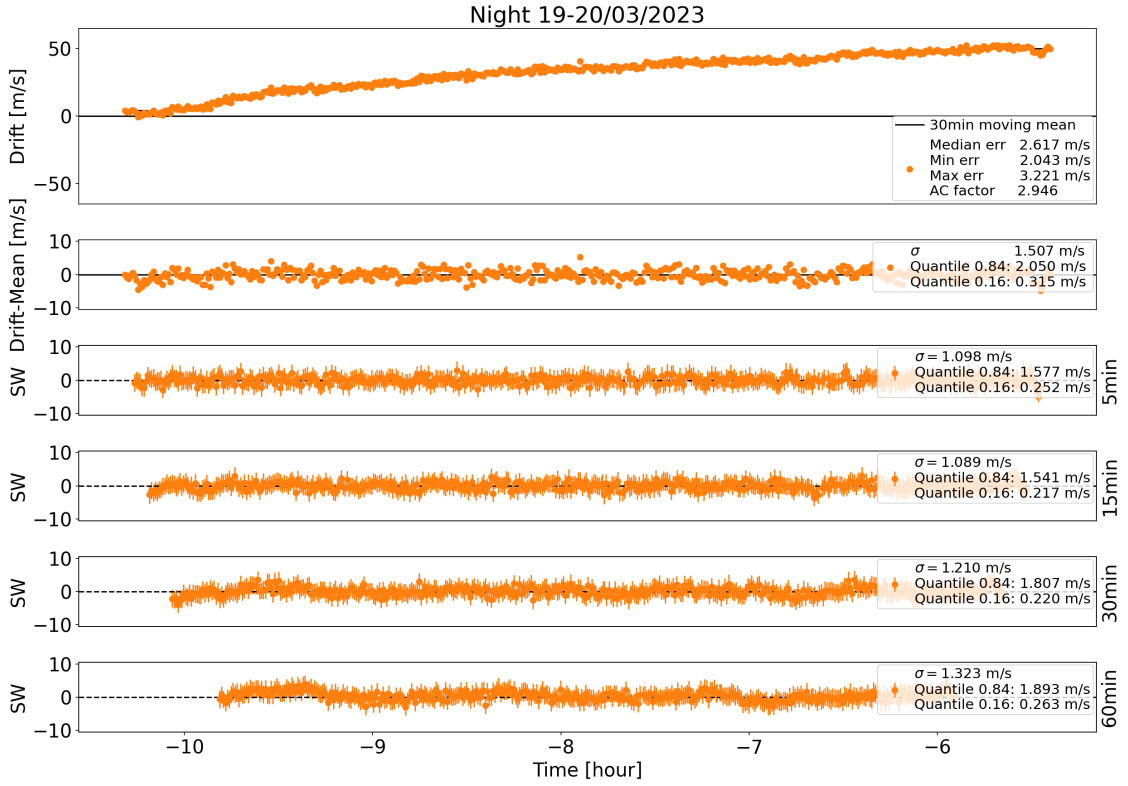

Figure 20: Drift series 2023/03/19, residuals vs. 30 min mean, and simulated long-exposure time (5, 15, 30, 60 min) observations corrected for drift with sandwich method.

9 Discussion and conclusion

It was clearly determined that the exposure meter has significant influence on the spectrograph drift. It does this by producing a heat gradient across a nearby mirror mount, effectively tilting the mirror. Keeping the exposure meter in an "always on" state reduces the effect to non-measurable amounts, since it is the time-dependent change in heat (warming up, cooling down) that produces a measurable drift. But the heat gradient is still there, and tiny temperature fluctuations are measurable on sensors whenever the exposure meter is counting/not counting.

During winter, large oscillations in the front room temperature was measured. These were also directly measurable (with reduced amplitude) in the instrument temperature sensors, bypassing the instrument room. This was a strong indication that a temperature coupling exists between somewhere in the front room and the instrument box. The suspected cause is the cryotiger cooling tubes, which are the only metal pieces connecting the two rooms at this time. While the scatter in drift measurements is higher in the winter period, it was not possible to measure these oscillations clearly in the drift from ThAr spectra. This is also evidence that the extra temperature insulation of the grating afforded by the pneumatically sealed grating tank goes a long way in ensuring stability of the instrument.

Re-fitting the wavelength solution for each spectrum with FIEStool showed a clear reduction in short-term precision by $\sim$ 1 m/s or more. It was also evident that separate nights had significant deviations in the wavelength solution with variations on the order of 10 m/s between different nights. Especially changing the ThAr lamp and/or redefining the FIEStool master wavelength solution had a large impact on the measured drifts.

By averaging drifts measured for several ThAr spectra, long exposure time observations were simulated. Taking the two ThAr spectra immediately preceding/succeeding such a simulated observation, it was possible to make a "sandwich method" style correction to the average drift. In turn, the residuals show how well the method could be expected to perform in removing short-term drifts from stellar spectra. With this, it would be reasonable to expect a short term precision of 1.5-2 m/s on the average night for 15-30 minute exposures, and 1.1-1.3 m/s on nights with very stable weather.

Taking the project further, there are two avenues that makes sense to explore. The first focuses on the long-term stability: To verify the long-term wavelength calibration, verify the spectrograph against standard stars, verify the ThAr lamp stability, and investigate if focus-changes in the spectrograph has a significant effect on the measured drifts over time. The other avenue would involve re-analysis of the measured drifts with the explicit purpose of doing Fourier analysis on the residuals to search for possible periodic signals that can be improved upon from an instrumental point of view, thereby attempting to improve the short-term stability of the spectrograph even further.

References

Anglada-Escudé G., Butler R. P., 2012, The Astrophysical Journal Supplement Series, 200, 15

M. Zechmeister et al. 2018, Astronomy & Astrophysics, 609, A12

Telting J. H., et al., 2014, Astronomische Nachrichten, 335, 41

A Standard setup and reduction procedure

The ThAr capture and reduction to drift measurements in m/s is a multi-step process. Scripts used for performing the actual data acquisition are available on the instrument computer, in folder *~/scripts/65-901/jet_scripts* (user *obs*).

For reduction of the 2D full frame images to order-by-pixel wavelength calibrated spectra, FIEStool (v. 1.5.2) was used, which allows reduction of ThAr spectra as-a-star. The following reduction tasks were used to reduce the spectra for velocity drift measurement (as labelled in the GUI): *Preprocess frame*, *Subtract BIAS*, *Divide by 2-D flat*, *Extract spectrum*, *Add wavelengths*. The wavelength solution is kept fixed, and defined from the first spectrum in the series. For generating a template to measure drift against, the previous tasks were checked, as well as *Check for new ThAr solution* which refits the wavelength solution for each spectrum. This means that cosmic ray mitigation, scattered light modelling, blaze correction, and order merging is not performed. Blaze correction is avoided specifically because it increases the weight of low-signal lines in the drift measurement.

The actual drift measurement is done with a python program called *tharli*, which has been distributed to John. It is a modified version of the code *lyskryds* by Emil Knudstrup, Marcus Marcussen, and Jeppe Sinkbæk Thomsen. The modification was done to measure drifts in pure emission line spectra (ThAr spectra) and remove some program features only relevant to stellar spectra. The drift measurement is by unweighted least-squares fitting of a template to an observed spectrum using a single free parameter, a velocity shift. The χ^2 is sampled by direct stepping, with default step size 6.67 m/s and sampling limits [-0.8, 0.8] km/s. The sampled minimum is used to select the 11 nearest sample points, of which a second order polynomial is fitted. The polynomial is then measures the true minimum and produces a relative error estimate from its curvature. This is done separately for each order, with 67 orders used by default (aperture orders, not grating orders) [2, 68] going from bluest (0) to reddest (91). A weighted average and weighted standard error is produced from the individual measurements, which is the final drift. The template is generated by performing a uniform basis spline fit (M. Zechmeister et al., 2018) to all (separately reduced) spectra that have individual wavelength solutions measured by FIEStool, effectively an average spectrum for each order. To avoid edge effects, the nearest 3 orders are used when producing the template for one order.

The (unweighted) least squares fitting (LSF) method is equivalent to the cross correlation function method (CCF). When using unweighted least squares, the primary difference between the CCF and the LSF is the evaluation parameter. In the CCF, the evaluation parameter is a likelihood measured as $\sum \text{spec} \cdot \text{template(dv)}$, where the template is shifted by a trial velocity dv. For the LSF, the evaluation parameter is instead $\chi^2 = \sum(\text{spec} - \text{template(dv)})^2$. In both cases, direct stepping in velocity space is performed to sample the function, and it is followed by a fit with a model function to find the extremum. For the CCF, a Gaussian model is frequently used, while for LSF the model is a second order polynomial. This also means that on paper there is no particular advantage to using the LSF method for ThAr spectra over the CCF. An example of the measured χ^2 value and subsequent polynomial fit for a single order can be seen in Fig. 21.

Weighted, rather than unweighted, least squares fitting has been used previously for radial velocity measurement of stellar spectra. Particularly for the HARPS-TERRA project (Anglada-Escudé & Butler, 2012) and the Carmenes survey (M. Zechmeister et al., 2018) with measured precision advantages over the CCF method. The stellar spectra require flux uncertainty weighting due to the continuum, which is where the difference between the two methods likely appear. For the ThAr spectra, as long as they are not blaze corrected, I

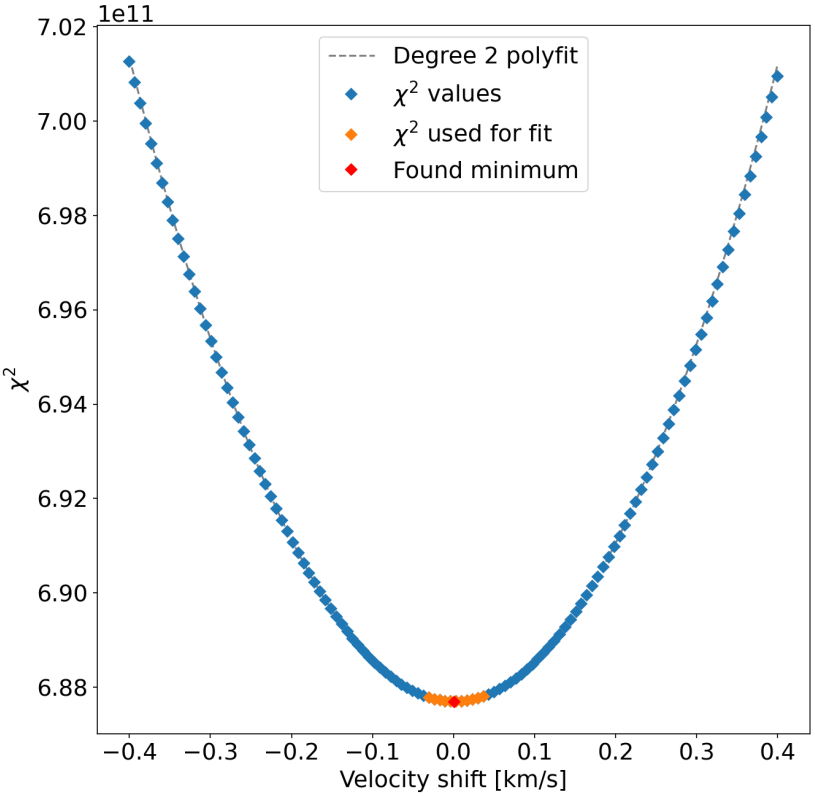

Figure 21: Example of the χ^2 evaluation and polynomial fit for a single order of a ThAr spectrum, which is used to measure the velocity drift of the spectrograph. The red marker indicates the minimum found by the polynomial fit.

see little reason to assume that the (unweighted) LSF is particularly better suited than the CCF, if both methods are properly implemented. The choice of LSF over CCF for this project was purely convenience, as a working code had already been written, was accessible, and could be modified for the use-case.