Precise “Astrometry” from Undersampled Data for the *Veloce* Spectrograph

Chris Tinney (UNSW)

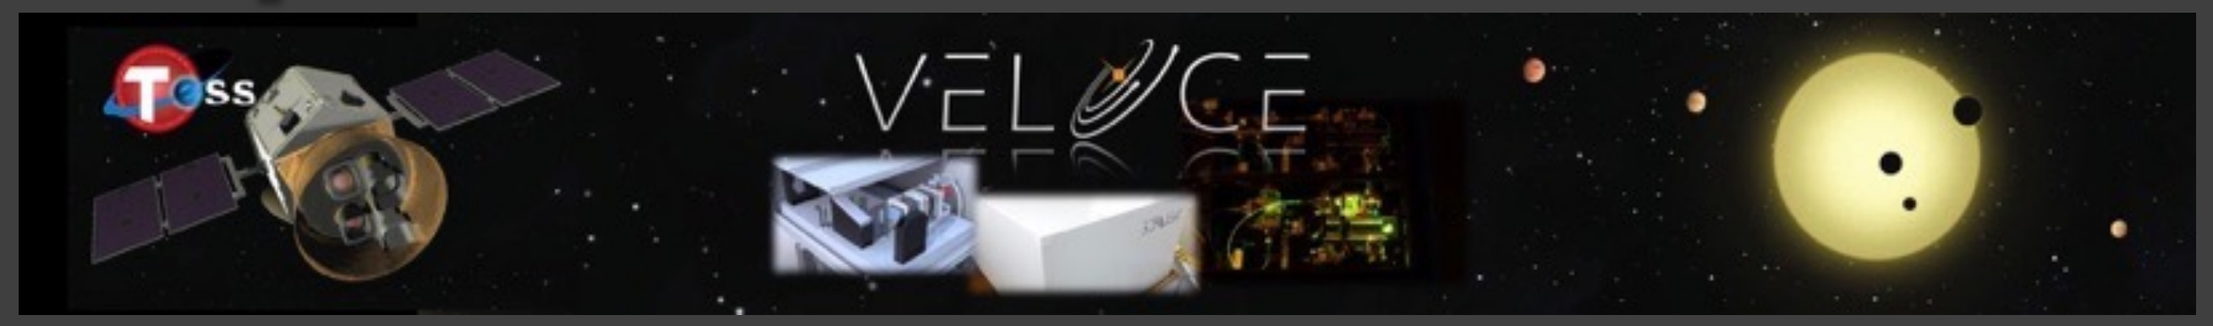

Echelle spectrographs are a fundamental for measuring the dynamical masses of exoplanets. The state of the art is now velocity measurement at the 1m/s (or less) level. Almost all of these instruments are now fed by optical fibres, that present a single light source of almost entirely spatial-information-free spectrum at the entrance to the spectrograph. The resulting spectrum is recorded in 2D as an “echellogram”, with many interference orders slicing the spectrum into multiple wavelength ranges, that must be traced and extracted.

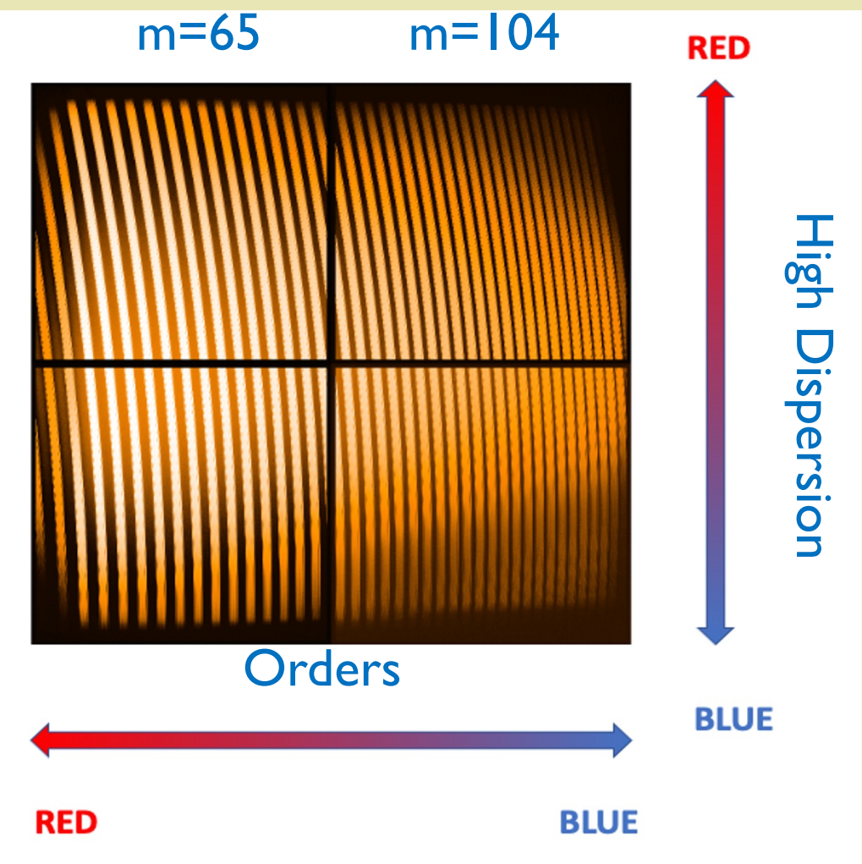

Design trade-offs must be made. Ideally the high-dispersion direction would have resolution so large that it’s close to the 1m/s velocity precision desired (i.e. $\lambda/\Delta\lambda \sim 3\times10^8$). But, as this science is photon-limited, you’d also like as large a wavelength coverage as possible (i.e. ~600nm). Detector size at ~4Kx4K is also a constraint, which limits the number of pixels in the high-dispersion direction in each order, as well as the total number of orders that can be recorded.

Finally, spectrographs are *wavelength interferometers*. Systematic precision is determined by physical control of the scale of the diffraction elements (and the speed of light in the spectrograph). Most systems are therefore stabilised by putting the whole, (several-m-scale) instrument in a large, expensive vacuum enclosure.

The “Standard Model” that has emerged for these instruments is to have a few optical fibres are inject into a spectrograph operating at $\lambda/\Delta\lambda \sim$ 100,000-150,000, designed so they are *substantially* oversampled in both dispersion directions (to trivialize extracting of 1dDspectra from 2D detector data).

This is effective, if very expensive. Instruments like HARPS (Pepe et alo. 2000, doi.org/10.1117/12.395516, ~EUR18m), ESPRESSO (Spano et al. 2010, doi.org/10.1117/12.858096, ~EUR35m) and GCLEF (Ben Ami et al. 2016, doi.org/10.1117/12.2232854, ~USD50m) followed (or will follow) this path.

Veloce on the Anglo-Australian Telescope (AAT)

The funds potentially available for a similar facility on the 4m AAT were much more modest at ~A$5m.

Which made us think about compromises to provide a larger wavelength range and higher throughput, while controlling costs, by requiring more precise simultaneous calibration, allowing under-sampling of the recorded data, and requiring a better understanding our instrument.

Veloce is fed by a small 19-element integral-field unit (IFU) that segments a 2.5” diameter aperture (to collect more photons given our 1.5” median seeing) into a narrow, long pseudo-slit at the spectrograph entrance. The narrow slit delivers $\lambda/\Delta\lambda$ = 80,000 spectral resolution in a compact spectrograph with a 100mm pupil imaged onto ab off-the-shelf 31.6l/mm R4 grating.

The long-pseudo slit requires anamorphic compression and undersanploing in the cross-dispersion direction. The cameras+cross-dispersers are quite small, so we fit three of them in the spectrograph enclosure, providing wavelength coverage of 370-930nm in the echelle orders m=65-155.

The enclosure for Veloce is sealed, but not evacuated. It uses a pressure bladder and off-the-shelf controller to maintain internal pressures at a 925±0.1mbar, plus dual shell temperature control at 25.00±0.01K. We know the temperature of the interior will slowly change at the 0.01K level (or less) and that the pressure will slowly change at the 0.1mbar (or less).

The key innovation in this design is dedicated use of Simultaneous Calibration fibres to inject ThXe and/or LaserComb light alongside each stellar spectrum. This data is used to measure changes in the spacing of the grooves in the gratings, while interior telemetry measures changes in the speed of light (which is a known function of pressure, temperature, humidity and wavelength to high precision).

Layout of the original single channel *Veloce Rosso* system

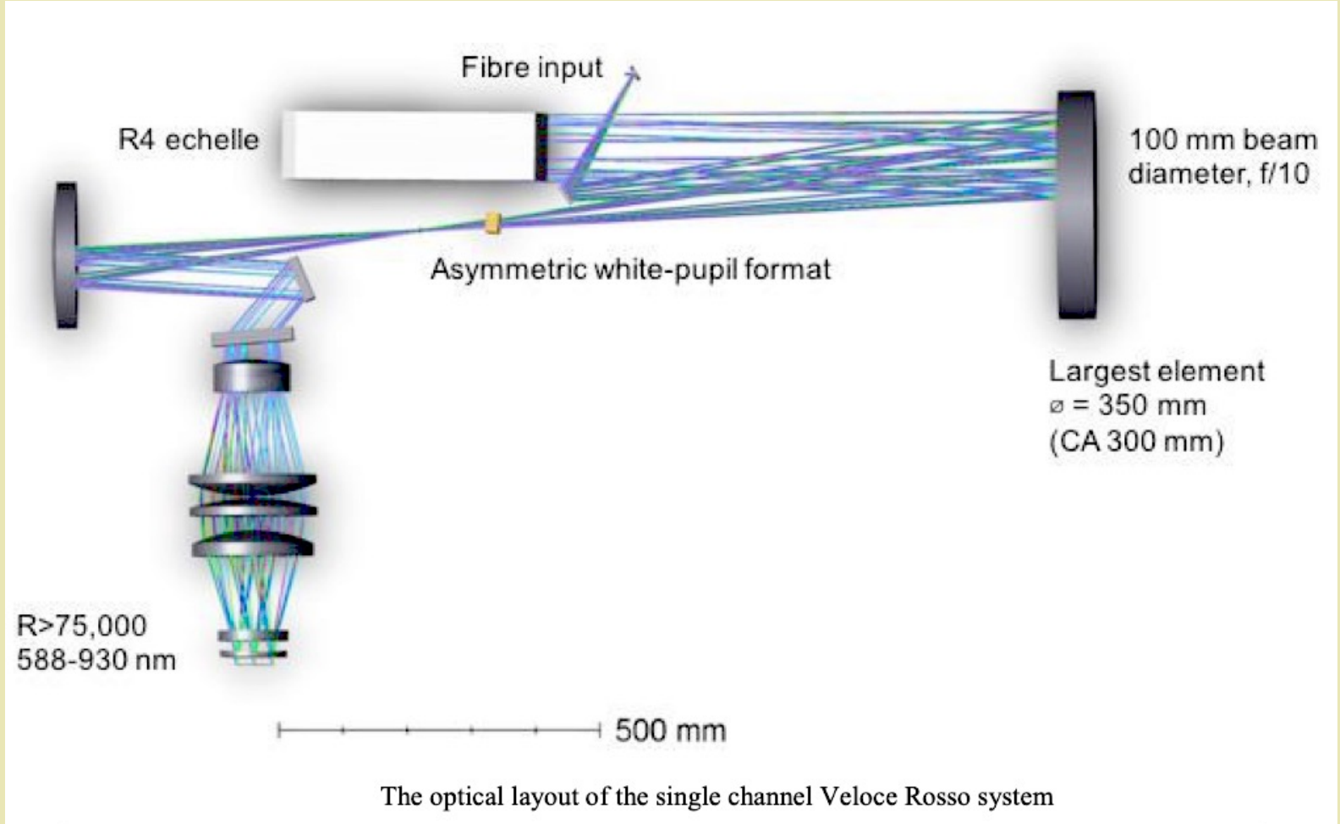

The optical layout of the single channel Veloce Rosso system

Layout of the current three channel *Veloce Rosso+Verde+Azzurro* system

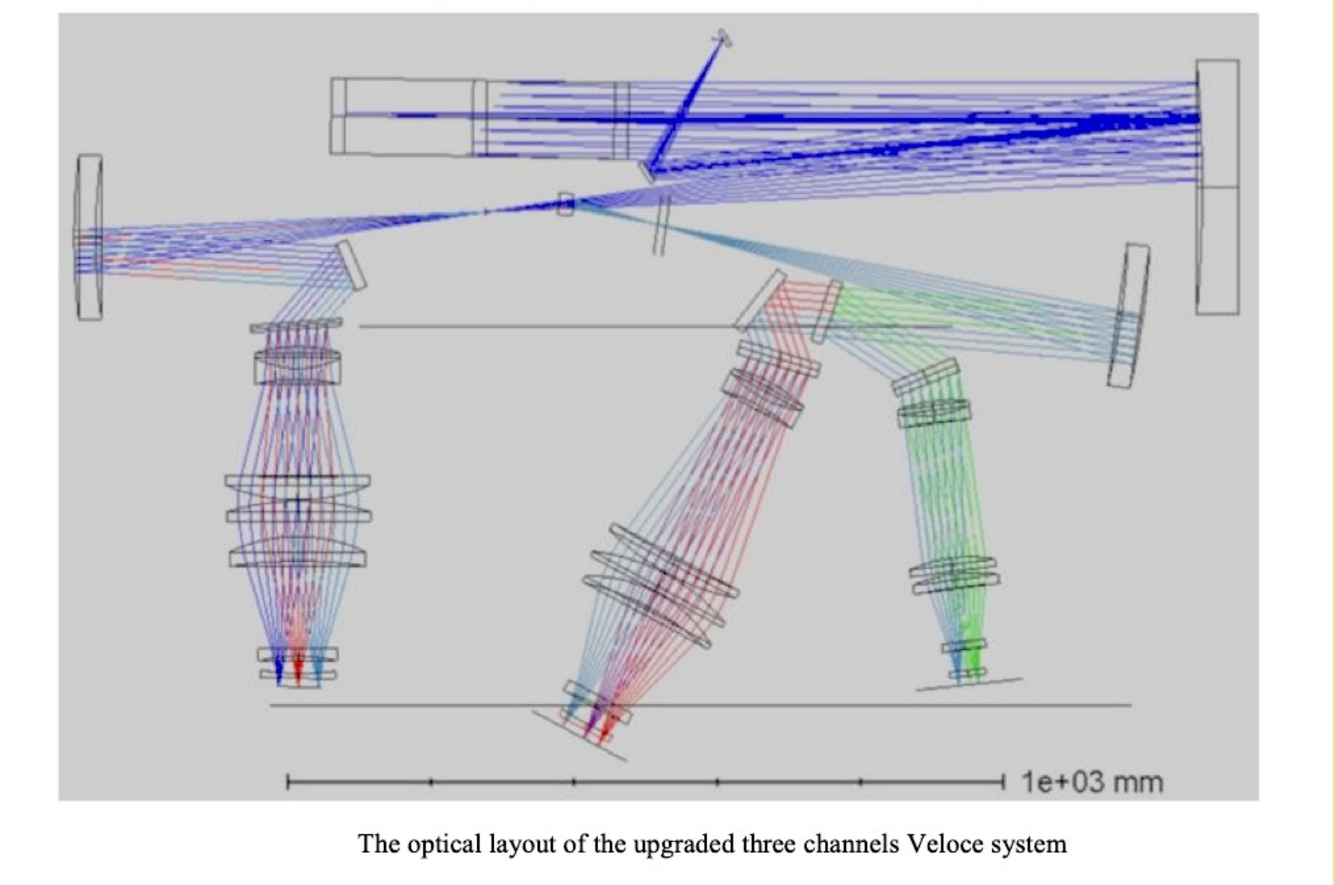

The optical layout of the upgraded three channels Veloce system

A 19-element hexagonal IFU at the telescope focus is reformatted into a long pseudo-slit - along with 5 sky fibres, 2 blank fibres and Simultaneous calibration fibres injecting ThXe and LaserComb (LFC) light

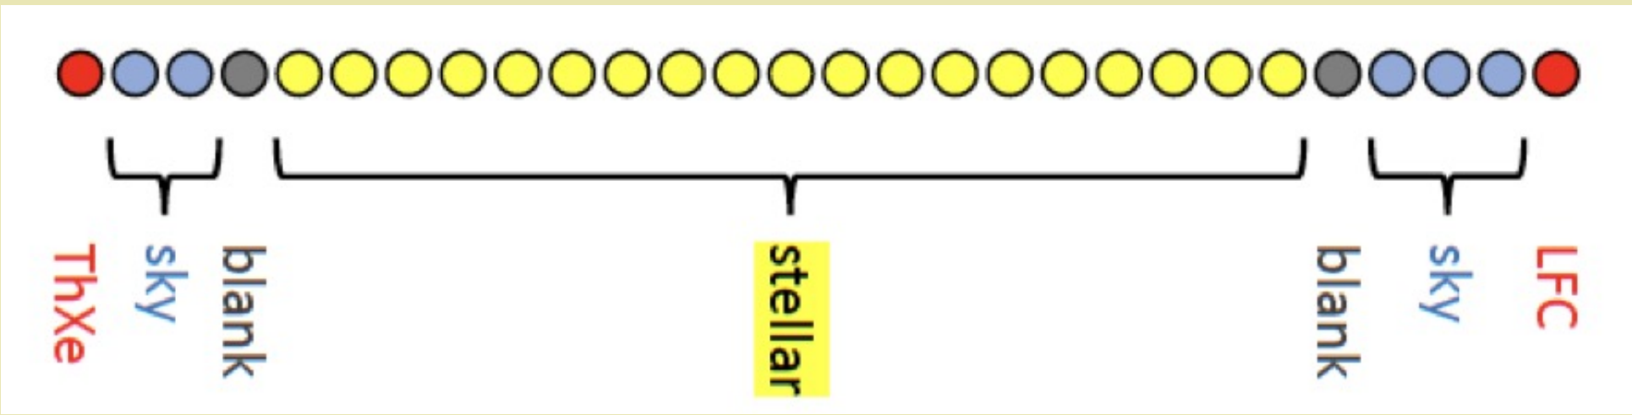

The Menlo Systems Laser Comb is a key element in our calibration. It delivers ~10,000 spectral lines with frequencies precisely determined (i.e. to better than 1-part-in-10^{10}) by a GPS-moderated Quartz clock. LaserComb exposures can be injected alongside any other observation.

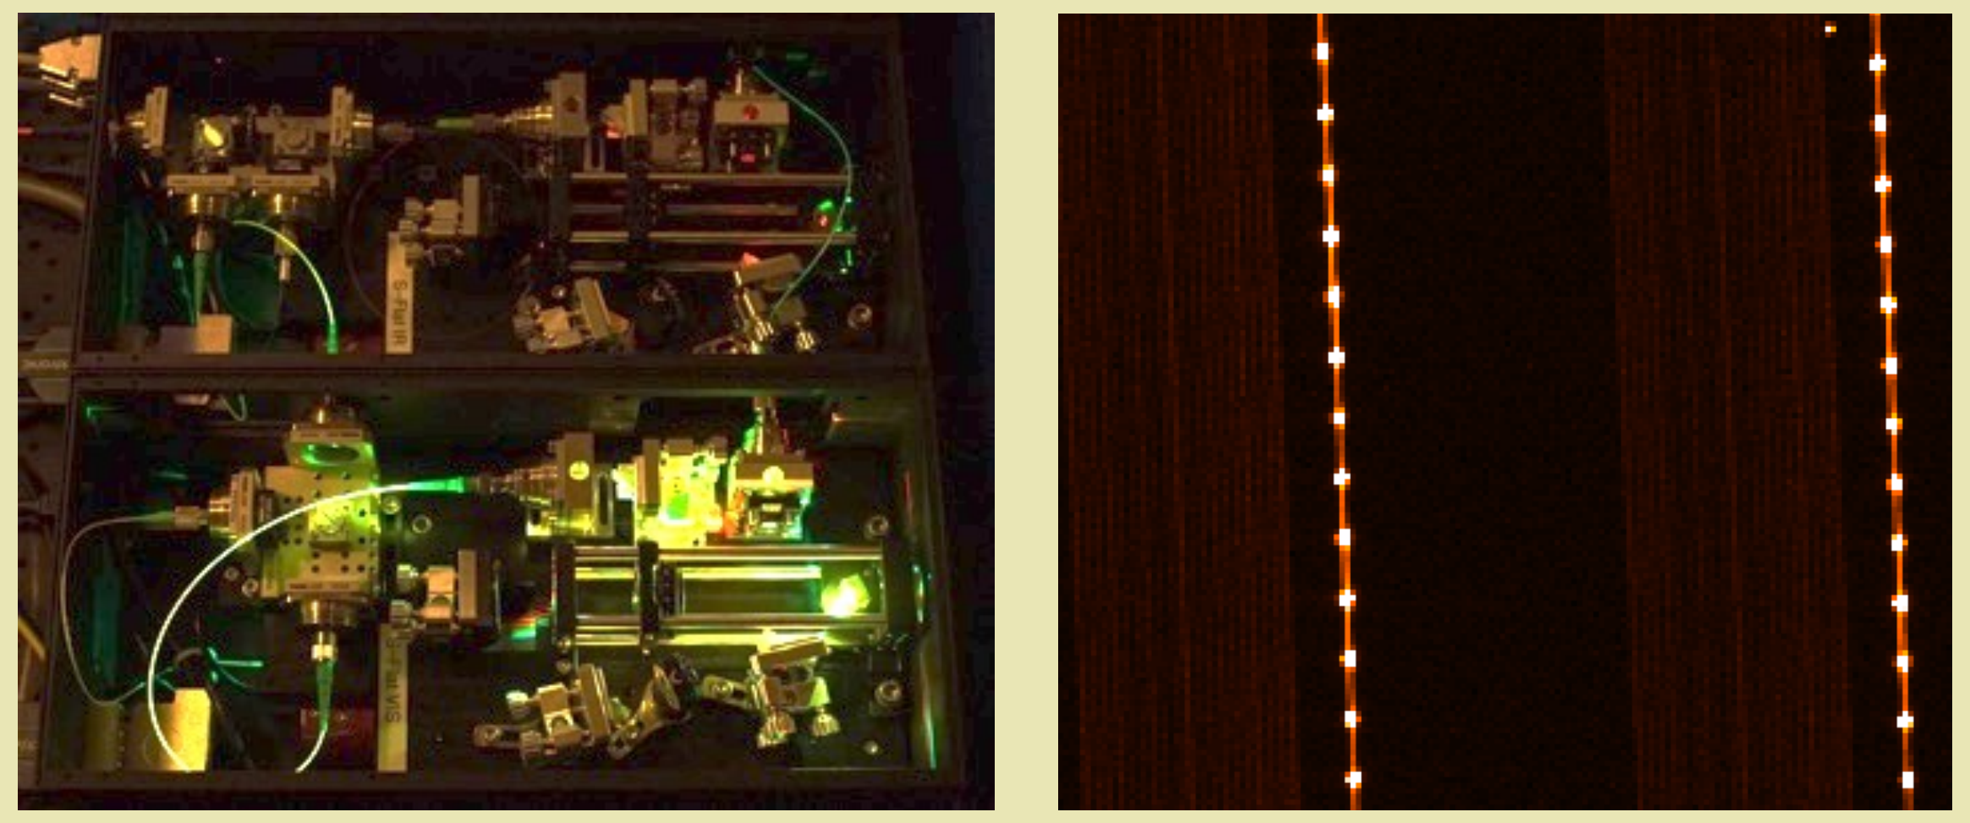

Figure: The ‘lit up’ LaserComb when operating (left) and zoomed-in sample data showing LaserComb lines alongside faint star light, for two orders.

The LaserComb is deliberately injected by a small, endlessly single mode fibre (rather than the 75µm fibres that inject star light). Why? Ask me!

Undersampling - Scourge? Or Interesting Challenge?

To fit our information-dense echellogram onto 4Kx4K detectors, we must run counter to the “your spectrum must always be oversampled” paradigm. *If Kepler/TESS can produce parts-per-million photometry from very under-sampled images, surely we can extract precise spectra from under-sampled spectrograph data?*

Our spectra are both under-sampled and overlapping in the X (cross-dispersion) direction (where every fibre sees the same stellar spectrum) and under-sampled in the Y (high-dispersion) direction for the LaserComb.

This means that most current techniques for extracting 1D spectra from the 2D images (which rely on not needing to understand profiles very well because the data are over-sampled), don’t work. So there’s some software work to do. We want to (1) measure LaserComb line positions (i.e. centroids) so that we can define the tracks of the echellogram in the 2D images, and (2) eventually measure their shapes so that we can precisely carry out a fit to do the 19-fibre optimal extraction.

The reason designers insist on oversampling is that fitting the under sampled images is well-known to produce poor results for the most of the easy fitting techniques. This problem became well known when standard point-spread function (PSF) techniques were used to measure astrometry from the under-sampled images of HST. The resulting astrometry had large (±0.1-0.2pix) errors that depended on where the centroid of the stellar light landed within a pixel – so called “pixel-phase” errors. These are largely due to a combination of (1) an imperfect model PSF, and (2) sub-pixel sensitivity variations in CCD pixels.

We found pretty much *exactly* the same issue in our under-sampled “point sources” in *Veloce* – 0.1-0.3pix errors in measured centroids that are strongly correlated with the pixel phase.

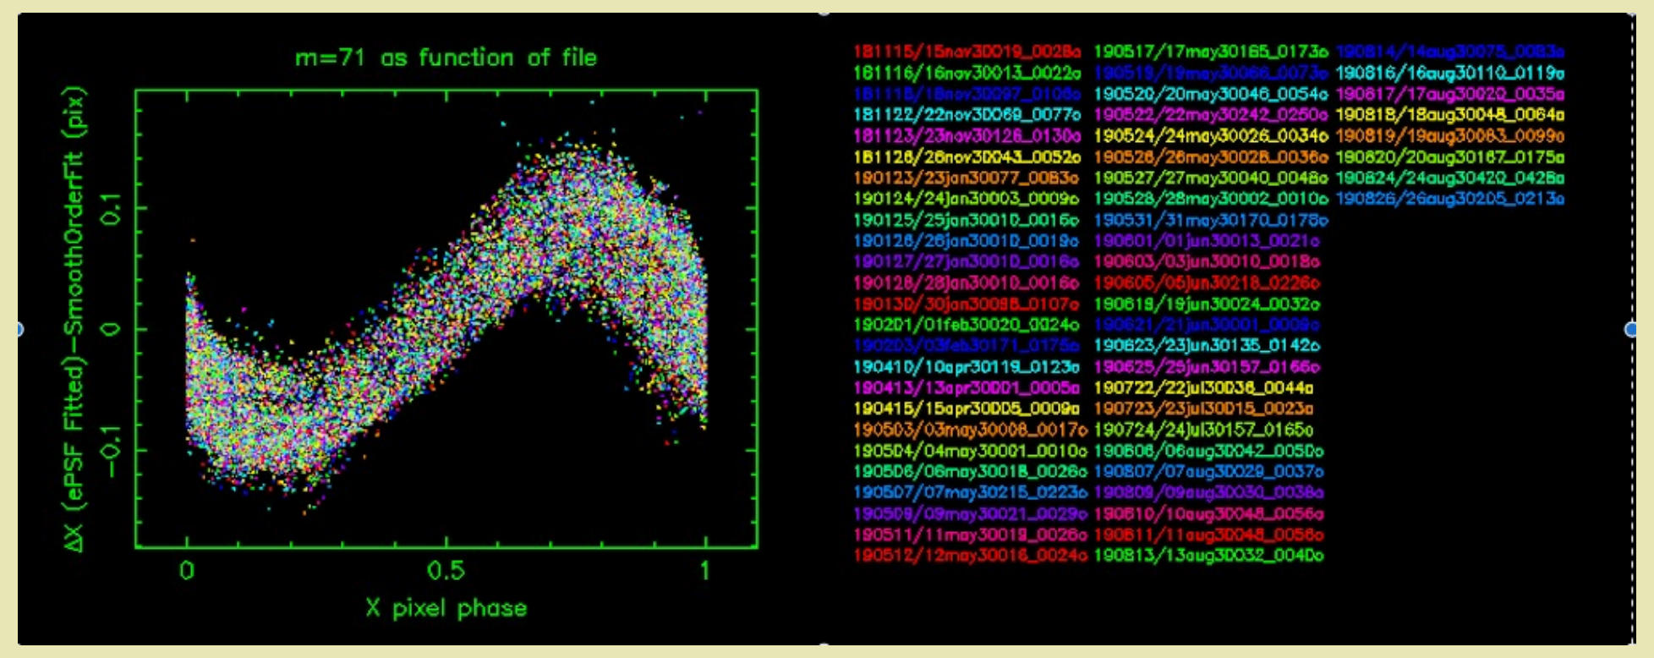

Plot of the deviation of line centroid X positions (from a smooth sit through those line positions) for a single order (m=71) in ~70 LaserComb observations spanning ~9 months. The pixel-phase ‘errors’ produced have a consistent pattern with amplitude > 0.1 pix.

A Solved Problem!

Anderson & King (2000, PASP, 112, 1360) demonstrated that a more sophisticated approach can use the many stars/lines at many pixel-phases observed to create an oversampled empirical “effective PSF” model that varies across the field & includes sub-pixel effects.

We have implemented a first-cut at using ePSF techniques with *Veloce* data using a modified version of the the `astropy` implementation of Anderson’s ePSF algorithms. We produce ePSFs from hundreds of images, containing hundreds of thousands of lines, across broad time periods (i.e. many months).

What if the ePSF varies with time?

A feature of stabilised spectrographs fed by fibres is that the spectrograph’s intrinsic optical PSF doesn’t vary at any meaningful level. The major effect of the small temperature and pressure changes is to stretch and contract the echellogram – PSF changes are much, much smaller by comparison.

In fact, the main changes we have seen in the Veloce ePSF with time have come from the changing electronic performance of our controllers, as we have been forced to change video cards, resulting in changed SCTE..

The echellogram moves, expands and contracts – the PSF doesn’t change.

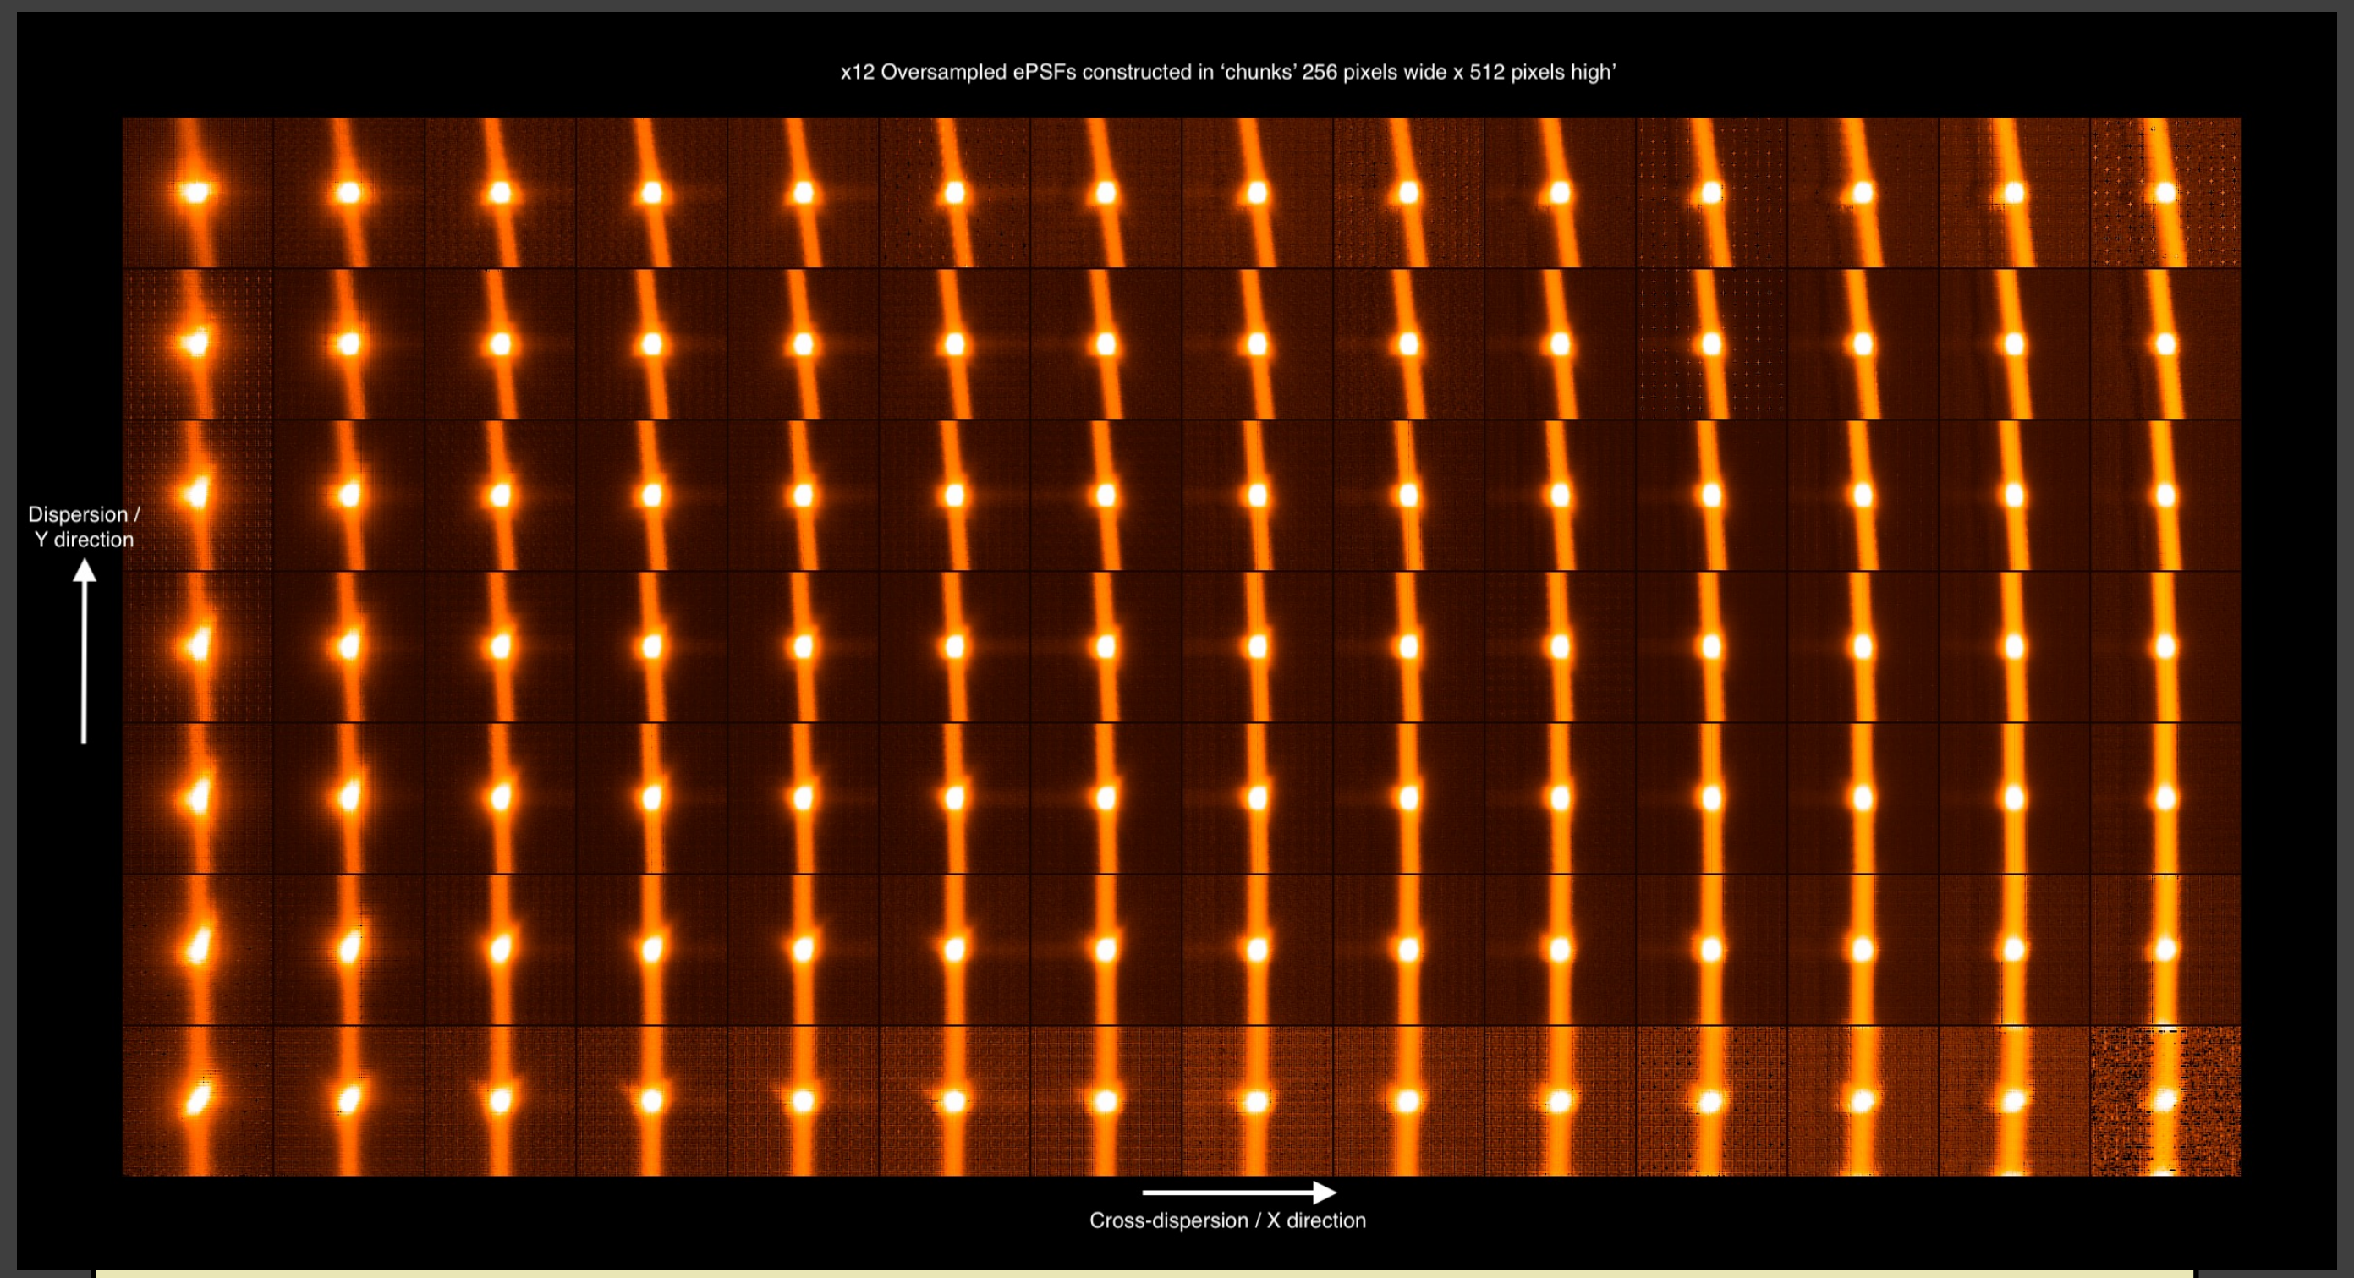

A ‘map’ of x12 oversampled ePSFs (each made for lines in a 256x512 box) across Veloce Rosso’ echellogram. We see the SCTE change direction in the two halves of the detector, and regular ‘noise’ in the background that reflect sub-pixel sensitivity.

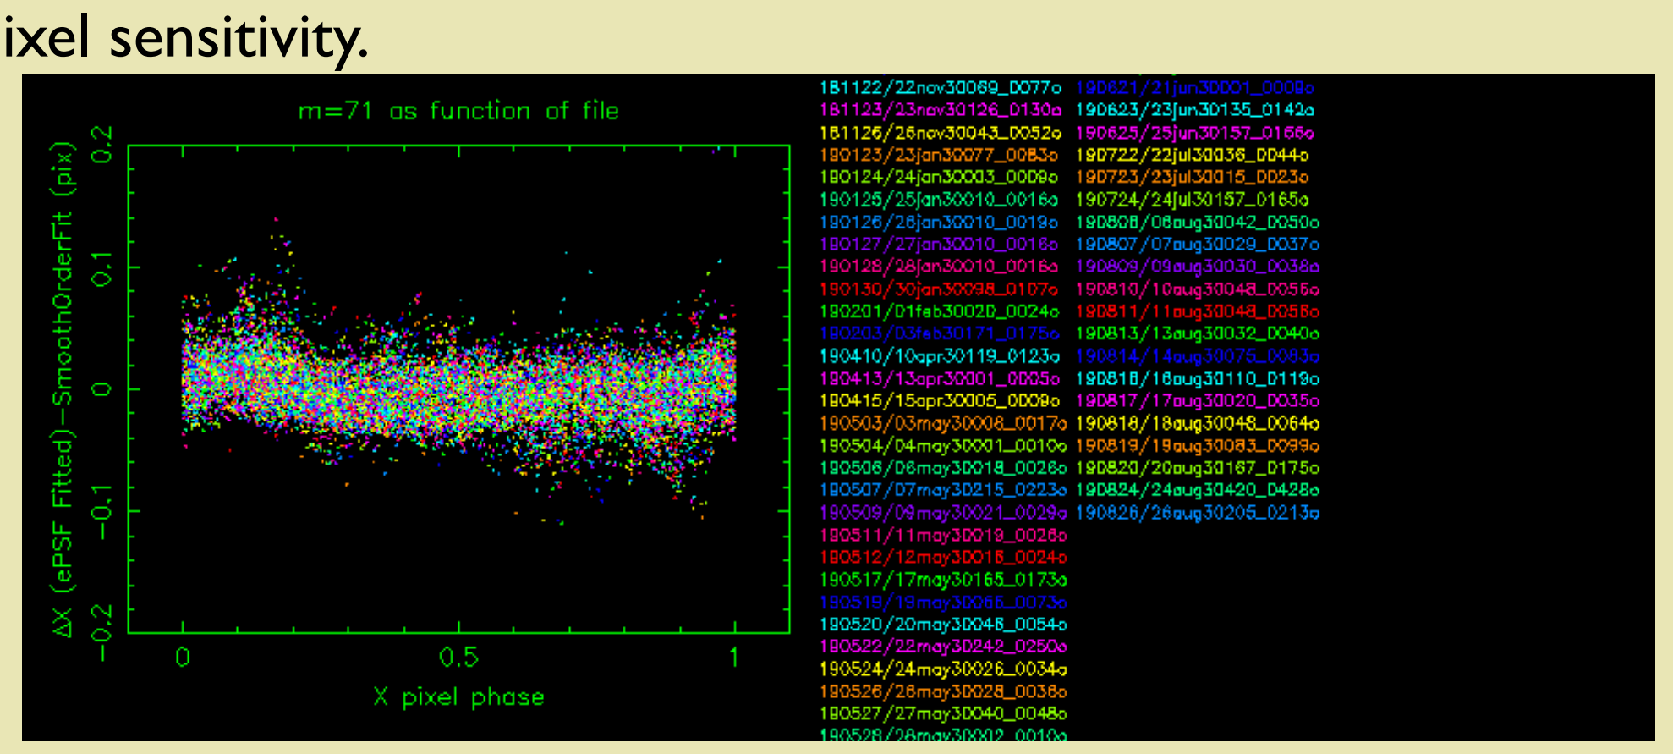

The same line centroid deviation plot as previously (on same scale), but now using ePSF-fitted positions. Pixel-phase issues have gone. The scatter for individual lines is the expected photon counting ~20mpix.

Why go to all this trouble?

A key challenge for Veloce is calibrating the expected variation in the physical spacing of our diffracting elements. To do that, we need a model of the spectrograph that allows us to use the grating spacing (and known speed of light for the Temp, Pressure and Humidity in the spectrograph) to predict the echellogram’s “breathing patterns” (the effects *aren’t* aligned along detector X,Y).

This we have done (at least for an earlier generation of LC and ThXe line astrometry from DAOPHOT (Stetson 1987), where we had to remove pixel-phase effects in an arbitrary and ad hoc manner). We have an analytic model for the spectrograph based on grating equations, that ends up with just 6 free parameters being solved for from ~10,000 LaserComb lines in each simultaneous calibration exposure. This measures the slow change in the effective spacing of the echelle grating over time, with a scatter about an interpolated solution to that spacing of just 84 pm (roughly 3ppm of the nominal grating spacing of 31.6µm).

We have no vaguely constrained and/or degenerate NxM Cartesian polynomials used to track orders and fibres. Instead, we have an analytic model with just a few parameters that correspond to *physical things we expect to vary*, which can be interpolated to any observation MJD to tell us where our fibre tracks lie.

Obviously the more precise the LaserComb astrometry that drives that model, the more precisely the tracks can be modelled and predicted.

The Final Step in this process is to replace the analytic Generalised Gaussian functions used in optimal extraction fits to date, with “effective Line Spread Functions” (eLSFs) that include PSF variation and sub-pixel sensitivity. Since we know the ePSF, finding the corresponding 1D eLSF is an obvious step, and will be the key in getting the correct extraction of spectra from our under-sampled data.

And of course, once we can make this work, the paradigm of needing to always be oversampled is broken – leading the way to cheaper instruments (or more information density for the same cost) in future precision Doppler echelles.