Ten reasons why accurate pointing is non-trivial

Patrick Wallace[1]

2016 September 20

The basic principles of computer-corrected telescope pointing have been understood since the 1960s if not earlier. See for example the 1968 paper[2] by Meeks *et al.* about the Haystack dish, and note that when Peter Stumpff did similar work on the Effelsberg 100m dish he did not regard it as sufficiently novel to bother writing it up in English. My own early work on the pointing of the 3.9m AAT involved little more than an extension of well-established techniques for dealing with errors in transit circles and apart from one or two conference papers was never written up.

The pace of hardware advances since the 1970s means that any limitations related to processing resources were swept aside decades ago. Efficient algorithms in numerical methods and (admittedly to a lesser extent) in fundamental astronomy were readily accessible by the 1980s. Why, then, is delivering the full pointing potential of a new telescope still a challenge, and why is it so seldom achieved? Here are ten areas in which telescope control system developers sometimes underestimate what is involved.

1 ASTROMETRY

Telescope engineers and users alike tend to regard the pointing challenge as identifying defects in the machinery and optics and applying appropriate corrections. They assume that the other part of the process, namely predicting the *observed place* – the direction from which the radiation from the science target is received – is a solved problem. But in practice many telescope control system (TCS) designs fall at this first hurdle. Why?

1.1 Precession and the rest

The first challenge is knowing what effects to take into account. Precession obviously, as it is an inexorable change in the spatial attitude of the Earth's axis that over the course of centuries amounts to tens of degrees. But what about nutation? And aberration? How important is proper motion? Does parallax matter? Polar motion? Earth tides? Gravitational deflection of starlight? And what part does timekeeping play in the calculation? What happens when there's a leap second? And we haven't even begun to discuss atmospheric refraction, a tricky problem in its own right.

---

[1] `ptw@tpointsw.uk`

[2] M.L.Meeks, J.A.Ball & A.B.Hull, *The pointing calibration of the Haystack antenna*, IEEE Transactions on Antennas and Propagation, **16**, 746-751 (1968).

Once the TCS developer has collected together all the astrometric algorithms that will be needed, and has worked out in what order they have to be applied, how does he/she know the answers are correct? You tell your program the star coordinates, say where the observatory is, pick a time, specify various other things including weather readings, and it spits out an $[\,Az, El\,]$ prediction. How do you verify your numbers are correct, and to what accuracy? And even if you've found authoritative examples on the web or in the literature, how confident are you that they're in the same coordinate system as yours? Worse, when you discover that your predictions disagree with them, how do you know where to start looking?

The problem is difficult enough for some developers to hope that any deficiencies in their astrometric calculations will be mopped up by the pointing model. This will work to a limited extent – for example an empirical $\tan\zeta$ vertical "flexure" term will fudge out imperfections in the refraction model – but some effects, such as aberration, will be locked to the rotating sky rather than the local coordinates that the pointing model deals with. It is incomparably better to get the astrometry right by dead reckoning and only then start to evaluate the telescope pointing errors.

Three real-life examples demonstrate that these difficulties do exist:

- At an important radio observatory where extensive work had been done on modeling the pointing of a large dish, it turned out the TCS was neglecting nutation, an effect of 10 arcsecond magnitude.
- The developer of a TCS for a big optical telescope was using as the basis of his astrometric predictions a popular freeware planetarium system that turned out to be omitting both nutation and aberration. In fact its astrometric algorithms were fundamentally powerless to deal with aberration because it had been assumed that all reference systems differ only in relative orientation.
- A subcontractor to a space agency spent many expensive months trying to discover why the front-end computer for a deep-space antenna appeared to disagree by several millidegrees with apparent star coordinates from a USNO web tool. He was unaware that aberration applies only to stars, whereas spacecraft predictions deal with the equivalent phenomenon using light time.

1.2 Refraction

Atmospheric refraction is a particularly fruitful source of difficulty, partly because no fixed-formula models exist that work for all zenith distances. Wavelength dependence is another danger area: optical users seldom appreciate how important atmospheric dispersion is, and that to get the desired 1 arcsecond pointing accuracy they need to specify the color, whereas some radio TCS designers are unaware that radio and optical refraction are markedly different, and that humidity really matters. A particularly common mistake by amateur observers is to use electronic pressure gauges that display QNH rather than QFE, so that even at high sites the numbers are always of order 1000 hPa, rather than the much lower numbers a mercury barometer would read and which are the values needed by the refraction formula. For radio astronomers operating beyond Ka band there are of course much finer details to get right, involving rapid changes of refractive index in the vicinity of water and oxygen bands.

A final difficulty with refraction is that a TCS has to react to ever-changing atmospheric conditions. If readings from a weather station are fed straight into the tracking algorithm in real time, and there is millibar-level noise, this will manifest itself as sub-arcsecond tracking noise. The readings must be low-pass filtered, or at the very least sampled only once for each new slew.

1.3 Time

Time is a notorious source of difficulty. TCS developers usually start with a belief that an accurate source of UTC is all that is required. The assumption is that fundamental astronomy was perfected in the 19th century, and there must be formulas that link UT to Earth rotation and also to the positions of planets. Unfortunately it's more complicated than that, principally because the time it takes for the Earth to rotate once isn't exactly 86400 SI seconds and is somewhat irregular.

To predict telescope pointing you need to know the Earth Rotation Angle or equivalently Sidereal Time, and to calculate these quantities you need to know UT1 (more or less the equivalent of the old GMT) and Terrestrial Time TT (more or less the equivalent of the old Ephemeris Time).

The predicted difference between UT1 and UTC is available from the International Earth Rotation and Reference Systems Service and can be up to 0.9 seconds, affecting pointing by up to about 13 arcseconds. Professional telescope control systems must take UT1−UTC properly into account, which includes operational procedures for maintaining an up-to-date value and for dealing with leap seconds. And a TCS that hopes to track right through a leap second needs to use not UTC but TAI as its timebase – which is in any case the right thing to do.

A further common blunder is using Greenwich *mean* ST (which is what the standard formulas yield) when Greenwich *apparent* ST is what is needed. The difference is called *the Equation of the Equinoxes*, and neglecting it can introduce pointing errors of over 16 arcseconds. Again, this is commonly neglected for equatorial mounts and fudged out using an *ad hoc* change in hour angle zero point (known as a "sync", of which more later), whereas altazimuths must do it properly.

It is also incorrect to use UTC when computing ephemerides for moving bodies such as the Moon and planets: TT (or more precisely TDB) is needed (also for part of the GMST formulation strictly speaking). This has a much smaller effect on pointing than neglecting UT1−UTC, but the difference TT−UTC, nowadays over a minute and growing each time there is a leap second, matters for fast-moving targets such as the Moon and near-Earth asteroids.

2 POINTING TESTS

There are two tried-and-tested ways to run a pointing test. The older way is to point at a known target, such as a star, center or peak up on it, and log the current encoder readings. The more modern way (and arguably better, for optical/IR telescopes at least) is to track nominal celestial coordinates, take an exposure of the field, and perform an automatic plate-solve to record where the telescope was actually pointing. In both cases the encoder readings (or preferably demands to take account of systematic servo effects) are sampled at a recorded instant, and logged along with the circumstances, namely catalog data, time, site location and weather readings. The job of the pointing-analysis tool is then to determine the relationship between where the telescope was told to point and where it actually ended up pointing.

Simple enough. But what is actually done in practice almost always departs from this, in ways which the perpetrators insist don't matter but that in practice cause trouble.

The most common mistake is to log not a timestamped set of encoder demands but *how far out the pointing was.* Doing it properly requires specific real-time actions (*i.e.* read the clock, interpolate the stream of pointing demands) whereas logging how far the telescope had to be nudged is much easier, the line of least resistance. But there are disadvantages. For a start, the details of the operational pointing model in use at the time have to be recorded, and the pointing analysis has to be able to reconstruct what the TCS actually did. There is lots to go wrong with that – sign reversals, incorrectly written-down formulas, confusion between angular and on-the-sky distances, use of rigorous and non-rigorous formulations *etc.* The bookkeeping task of remembering the operational state is on its own an opportunity for blunders.

A further set of difficulties arises if the modeling is not of *absolute* samples but of the *residuals.* It sounds easy – the fitted amounts can be added to the operational terms rather than replacing them – but the statistical judgments made by the fitting procedure are undermined. A small correction to a run-out term, for example, might not pass significance tests and so leave the operational value unchanged, whereas an absolute fit would have given a revised value.

Note, however, that for the purposes of the pointing test simply operating without a pointing model is not satisfactory. Although the observations are then by definition absolute, it is usually much better to start from a basis of excellent pointing, and this is obviously the case where the observations involve acquiring single stars. The important thing is for the presence of the operational model not to affect the pointing samples in any way, any more than the operator's choice of T-shirt.

Further problems arise in the case where a camera has been used and the offset of the target star from the center has been measured. Quite apart from the opportunity to confuse angular and on-the-sky distances ($\Delta\alpha$ versus $\Delta\alpha\cos\delta$ for example), the method relies on accurate knowledge of both plate scale and picture orientation. And use of first-order spherical trigonometry formulas can introduce significant errors near the pole and zenith.

An especially pernicious technique, beloved of radio astronomers, is to sample azimuth and elevation separately and fit scalar expressions to each that fail to exploit their interrelationship. In particular, the only satisfactory way to model azimuth axis tilt is to fit vector expressions to vector observations: there is only one set of EW and NS tilts, not separate sets for each of azimuth and elevation. The right way is to peak up in both coordinates and then log the position during the period of on-target tracking that immediately follows.

The pattern of test stars is also something that people often get wrong. A particular favorite is a regular grid of either $[\,\alpha, \delta\,]$ or $[\,Az, El\,]$. Such patterns mean that different areas of the sky, namely around the pole or zenith, are sampled more densely than others, and this will mean the model is unwittingly weighted towards these regions. Another objection is that the residual plots may not be evenly populated, with for example only specific declinations sampled so that what goes on in between the samples is never sampled.

Distinguishing between mechanical and celestial coordinates is vital. For GEM and cross-axis equatorials the "flip state" is an essential component. And where axes have ranges that exceed

360° (and associated cable-wraps) the full range must be recorded.

The path followed is also important, to expose hysteresis and to distinguish between spatial and temporal changes. A pseudo-random and "busy" pattern is needed, that while time-efficient makes the mount meander around and keeps both axes well exercised. A degenerate case is where the investigator believes that targeted pointing observations will prove more incisive than random ones, and for example restricts observations to the meridian or low in the sky. In practice this is never satisfactory, and an all-sky random pattern is much more informative.

To sum up, whatever pointing test techniques are used, they must fulfil the following requirements:

- The observations must cover the full patrol range of the telescope and be evenly spread across the sky. They must be random and not a regular grid of any kind.
- The quantities recorded must be absolute and not relative to the current pointing model. In the case of altazimuth mounts this will require adroit real-time handling.
- The telescope readings must be mechanical rather than celestial, while the star readings must be *observed place*, *i.e.* with everything up to and including refraction allowed for, or enough information to compute observed place.
- The path must keep both axes busy and should allow temporal and spatial effects to be distinguished.
- All observations must be vector ones, *i.e.* both coordinates at once.
- The observations should either be centered or astrometrically solved. Recording pointing errors or adjustments is not acceptable.

3 EFFECTIVE USE OF MULTIPLE POINTING TESTS

Pointing tests on professional telescopes are typically 50-100 samples. This is enough to identify the principal model terms, delivering a model of perhaps a dozen terms. A series of such tests, done a few weeks or months apart, will provide evidence of changes, and it may be possible to identify which terms, and hence which physical effects, are on the move. It is, however, possible to go further.

One important consideration is that the different pointing terms are all to some extent correlated, and how the fit apportions the pointing behavior into the various terms will vary from test to test. Some terms, for example on an altazimuth mount (i) the azimuth zero point, (ii) the nonperpendicularity between the mount axes and (iii) the nonperpendicularity between the telescope and elevation axes, can be very highly correlated indeed, and unraveling them will take a lot of observations over the full elevation range. Putting it another way, a typical 50-100 sample pointing test will generally not deliver trustworthy estimates for highly correlated terms.

Another issue is that small effects will be swamped by noise unless the number of observations is large. This is not a problem if everything is stable – the multiple runs can simply be concatenated and a single model then generated – but where something is known to change, due perhaps to temperature gradients, that won't work: the resulting noisy fit will prevent smaller effects being seen.

What is needed is a more sophisticated modeling technique that can simultaneously fit (i) a set of core terms assumed to be common to all runs and (ii) individual sets of terms that belong only to a specified run. This will allow small effects to emerge from the noise, very accurately determine the global effects such as tilt, and monitor changes in pointing offsets *etc.* This capability is not commonly implemented.

4 THE BOTTOM LINE

Claims are often made that a given telescope points to (say) 1 arcsecond. On further investigation it turns out this was simply the post-fit *a posteriori* RMS, in some cases of a carefully selected pointing test. The quoted RMS will not necessarily reflect the operational pointing performance, and experience shows that the achieved performance is typically two or three times worse than the single-pointing-test result.

To make an honest assessment, there is no substitute for applying the model from one pointing test to a whole new set of observations not used in the model fitting. Not only will this produce an RMS that means something, but the pattern of the residuals may be a good guide to what is changing with time.

A further consideration is how consistently any given model term is calculated (a) when fitting models and (b) in the control system. Quite apart from any gross errors such as sign differences or other programming errors, there may well be a lack of harmony over the use of approximations. An example would be where the fitting uses first-order expressions for mount tilt whereas the control system implements a rigorous 3D rotation. The only safe way is to use identical code in the fitting tool and the control system – *i.e.* linked from the same libraries – in order to achieve tautological agreement.

5 OUTLIERS

Naïve model fitting usually involves rejecting samples that lie more than (say) 3σ from the main error distribution. This is very dubious, especially when (as is always the case) assuming normal (Gaussian) statistics has no *a priori* justification. Normal statistics are assumed not because they correspond to physical reality but because a multitude of neat mathematical formulas are available to deal with that particular case.

More rigorous statistical tests are available, for example that take into account not just the size of the residual but the influence on the model of that particular point. And bootstrap sampling may generate more reliable confidence estimates than simply assuming normal statistics.

6 WHICH TERMS?

It is often claimed that correcting pointing and tracking is simply a matter of applying polynomials or spherical harmonics of sufficiently high order. Orthogonal polynomials are a favorite recommendation from math experts unaware that the orthogonal property applies only in an analytical context and not to irregularly spaced samples. Others, with more justification, advocate look-up tables and a simple interpolation scheme. However, in practice the most successful systems are ones where terms are specifically designed to map onto known geometrical or mechanical effects.

This is not to say that empirical functions and look-up tables should be avoided. There will always be effects that though neither predicted nor convincingly diagnosed are nevertheless rather well represented by smoothly varying *ad hoc* functions. And some mechanical effects, such as irregularities in azimuth tracks, are by their nature best represented by actual measurements. Metrology is of course the next step, where the measurements are made in real time.

The best of all worlds is achieved by basing the pointing model on physical/geometrical terms, applying any look-up tables or metrology inputs, and then mopping up with judiciously selected empirical terms such as polynomials and harmonics. But the tricky part is knowing how far to go: what is the criterion for including, or not, any particular term in the model?

Except in cases of ill-conditioning, adding additional terms will *always* reduce the RMS pointing error, the most common figure of merit. Consequently, a reduced RMS is not on its own evidence of model improvement: better statistics are required. It is particularly important to take account of the number of degrees of freedom, so that the number of samples remains large compared with the number of model terms.

An important objective during modeling is to base the predicted performance of the model not on the samples themselves but on hypothetical future samples drawn from the same population. This can be done either by correcting the RMS by a factor that depends on the relative numbers of terms and samples, or by resampling techniques, for example jackknife or bootstrap. Analogous techniques can be applied to the terms themselves, where the effect on the resulting RMS of removing terms can be assessed and only the most successful terms retained. These techniques necessarily involve multiple fits (often running into thousands), and computer resources can become an issue. Fortunately much of the processing is per-sample, and consequently multithreading techniques can readily be employed.

A particular aspect of modeling that is often ignored is whether to treat the two axes of the mount as independent or not. As was mentioned earlier, radio astronomers often choose the former, so that for an altazimuth mount a pointing sample is a scalar quantity $\Delta Az \cos El$ or ΔEl and each axis is modeled separately. But the better way is to treat it as a vector problem, so that a pointing sample is $[\,\Delta Az \cos El, \Delta El\,]$ and the single model returns a vector correction. The reason the vector method is better is that for some terms, for example the azimuth axis tilt components, there is an inter-relationship between the corrections in the two axes. If each axis is treated separately, the model will include two estimates of the same physical quantity, which is a potential loss of information and an unnecessary degrees-of-freedom depletion.

7 HANDLING ZERO POINTS PROPERLY

Control systems for amateur telescopes commonly sweep various hour angle offsets under the carpet in a procedure called “synching”. A sync in effect makes *ad hoc* adjustments to the encoder zero-points to eliminate the pointing error for the current star, the hope being that this will improve the all-sky pointing. Apart from canceling out unmodeled pointing errors, the procedure also deals with (i) neglecting the equation of the equinoxes, enabling mean sidereal time to be used instead of apparent, and (ii) neglecting UT1−UTC. It works because both of these are offsets purely in hour angle.

It would be easy to assume that the same applies to altazimuth mounts, and that merely resetting the encoder zero points to match the current star would be a good first-order correction to the all-sky pointing. But this is not the case: compensating for an offset in hour angle would require adjustments both to the azimuth zero point and to the east-west component of the azimuth axis tilt, proportional to sine and cosine latitude respectively. So with altazimuths there can be no such thing as a "clock star": dead reckoning is the only way to go.

8 ROTATORS

In practice, on telescopes that have instrument or image rotators these are usually ignored for pointing-test purposes. This has its dangers. If the pointing reference (for example the central pixel of a camera) is not coincident with the rotator axis, and if the camera is rotating, then rotator-angle-dependent pointing changes will occur.

This is likely to be a consideration on an altazimuth because of the need to stabilize the field for the test exposure. It can even be a problem on equatorials, especially for asymmetric designs such as GEMs and cross-axis, if the rotator flips to maintain north on the picture and the analysis fails to take proper account of this.

An even trickier problem arises if the pointing reference is off-axis, which often happens on optical telescopes if an autoguider on the edge of the science field is being used as the pointing reference, or a special feed horn on radio telescopes. In this case transforming the fitted model to refer instead to the rotator axis is less straightforward than for a pointing reference fixed in the focal plane. The most sophisticated way to handle the problem is to encourage rotation during the pointing test and to log the rotator angle as part of the pointing observation. It is then possible to deduce the [x,y] coordinates on the rotator of the pointing reference by including them as two terms in the pointing model.

A common omission when dealing with the off-axis pointing reference case is taking account of the pointing corrections needed to place the target image in the right place. In the case of an altazimuth (and equivalent arguments to apply to equatorials), as the zenith is approached, a pointing boresight that is not perpendicular to the elevation axis will require ever-increasing azimuth corrections, as well as the need for adjustments also in elevation. The azimuth corrections will have the secondary effect of rotating the field of view, which must be taken into account when the pointing is re-referenced to the rotator axis. The effects can grow arbitrarily large, and indeed a stage is reached when the pointing calibrator cannot be acquired at all, at any azimuth setting. A related effect, which many designers of altazimuth mounts are blissfully unaware of, is that the angle to which a rotator needs to be set is not exactly the parallactic angle. The angle that is needed depends on some pointing model terms, and is best calculated in real time by what amounts to ray tracing.

Some of these complications arise even when the telescope has no rotator. In particular, deriving a pointing sample from measurements of a calibration star image in a camera exposure requires knowledge of the field orientation and the $[\,x, y\,]$ of the pointing reference relative to where the image needs to be for science observations; in short, an offset in x does not simply turn into an azimuth change of $x \cos El$ and no change to elevation. Failure to deal rigorously with the condition will work fine over most of the sky but produce rapidly increasing errors for targets near the pole of the mounting.

9 TRACKING

Optical/IR astronomers tend to treat pointing and tracking as two separate aspects of performance, namely (i) target acquisition (plus associated maneuvers such as finding guide stars) and (ii) maintaining alignment during the course of the science observation. Radio astronomers tend to regard pointing as covering both aspects, which is better. The danger with the optical/IR mindset is that the techniques used to accomplish both aspects might be different and inconsistent, perhaps even the work of different software developers.

One obvious danger is that unless pointing and tracking are working from the same pointing model, either acquisition or tracking will be compromised, or both. The extreme example is the bulk of amateur mounts (there are one or two exceptions) where a pointing model helps to drop targets onto the camera chip but tracking is plain sidereal, *i.e.* 15 arcseconds per sidereal second in hour angle and zero in declination. Changing refraction and flexure then conspire to cause images to drift during exposures, making autoguiding mandatory. Only if the model is used to modulate the tracking rates will unguided exposures be possible.

Even where efforts are made to take account of the pointing model when tracking, developers typically focus on tracking velocities, as if they are the end of the story. They aren't: what actually matters is maintaining *position*, and the only safe way to do this is to calculate the required pointing and to manipulate velocity demands to steer the axis encoders to the correct places.

Even with a position-based scheme, there are still pitfalls. One is to calculate the tracking rates very accurately and to send them to the servos essentially as analog quantities. Doing so means that control over rounding errors has been lost, and after some unknown period of tracking the rounding errors will build up into significant position errors. Extra uncertainty can be introduced if the rates are obtained by evaluating analytical derivatives of the pointing terms, which sounds like a good idea until you realize that it involves not only the telescope model but all the astrometry (precession-nutation, aberration, refraction and so on) as well. The only sure way is to obtain the rates (*i.e.* the pointing changes over some time interval) by numerically differencing the pointing demands themselves.

This leads on to the question of servo requirements. The best servo designs involve feedforward of dead-reckoning quantities, in particular velocities. This means that the servo loop is there only to mop up small calibration errors and is not responsible for generating the tracking rates themselves. It is perfectly reasonable for the pointing kernel to evaluate the pointing model multiple times per iteration, and to generate velocities (and accelerations *etc.*) of the required smoothness by parabolic interpolation, spline functions or whatever. And of course it goes without saying that the pointing updates should not rely on exact timing, and the algorithms should be tolerant and self-healing in the presence of lost or repeated demands, and perhaps even if the demands arrive in the wrong order.

10 OBSERVATORY MANAGEMENT

Voltaire said something that translates to "The best is the enemy of the good". The opposite tends to be true of telescope pointing. What usually happens in practice is the following:

1. The specification for a new optical/IR telescope stipulates 1 arcsecond RMS pointing.
2. The designers take this very seriously and do their best to provide hardware that can deliver this level of performance, and fierce battles are fought over how the 1 arcsecond error budget will be shared out among the different engineers.
3. Horrendous things happen during commissioning (often about trying to get elaborate metrology systems into operation) that means pointing accuracy comes way down the priority list.
4. Nature takes its course, and default line-of-least-resistance measures manage to deliver (say) 5-10 arcseconds RMS.
5. Providing there are no showstoppers such as not being able to find guide stars, the astronomers dig out their finding charts and happily observe in the time-honoured fashion.
6. With a scientifically productive telescope in full flow, management is disinclined to award dedicated observing time for pointing investigations, and things just stay as they are.

There are exceptions of course. The Anglo-Australian Telescope met or exceeded its pointing specification almost from the start. Still competitive even now, it represented such a quantum leap at the time that it was seen by management as good publicity, and worth spending time on. A more recent example from the radio is ALMA, where in order to do the science, blind pointing to 2 arcseconds, and an even more challenging offsetting accuracy, were essential, and measures were taken to ensure these requirements were met operationally.

The way for the future is to exploit the astrometric accuracy of science observations. A few decades ago it was commonplace for an astronomer to provide target coordinates where there was some doubt about the reference system being used, or that came from telescope readouts logged by previous observers. Nowadays very accurate and dense star catalogs are available, and observers are required to submit accurate coordinates just as they would for time on a space observatory. Therefore if the science instrument has been properly calibrated, so that the focal-plane coordinates of the detector or entrance aperture are known, as well as the effective color, there is no reason why the science observation cannot deliver a trustworthy pointing sample *en passant*. It is remarkable that no plans to do this on the large survey telescopes VISTA and LSST have been reported. Dedicated pointing tests should be a thing of the past on such facilities, which continuously acquire vast amounts of super-accurate astrometric data. The all-sky pointing model should be updated in real time and without human intervention, so that not only are exposures accurately placed but the open-loop tracking performance is always in a high state of tune.
