Telescope Pointing

Patrick T. Wallace

What's the problem?

When I want to look at Jupiter with my old 8-inch reflector, I swing the telescope on its mounting until the planet is centered in the Telrad finder, and there it is in the eyepiece. Easy enough. Next, M57: find Lyra, point about a third of the way from β to γ and search around a bit. But what about a deep-sky object? Then I have to peer at a finding chart with a red flashlight, wondering which way up north is in the eyepiece and whether the chart could for some reason be left-right transposed. And even in the comfort of the control-room of a modern large telescope, where worries over dark-adaption or the possibility of hypothermia are reduced, the finding-chart approach, though still used, is unacceptable because of the sheer cost of the wasted time — approaching one US dollar per second on the very largest telescopes!

A more orderly method is to point the telescope by moving it until its read-outs (setting-circles, computer display or whatever) match the coordinates of the object. This is what is done on professional telescopes, and many amateur ones as well these days (and what radio-astronomers have *always* had to do). But what at first sight is a simple idea — just set the dials to the desired RA/Dec — turns out to be a surprisingly complicated problem, so much so that the vast majority of telescopes capable of being set in this way fail to deliver anything like their true potential, and a finding chart is still needed. In this article I look at what is involved in pointing by "dead reckoning" and show how the idiosyncrasies of individual telescopes can be allowed for.

The techniques I am going to describe can lead to startling levels of pointing performance. The best of the giant observatory instruments (and even some sub-mm radio telescopes) can point to 1-2 arcsecond, roughly the diameter of Jupiter's Galilean satellites; it is not uncommon to acquire stars straight into a spectrograph slit. Many telescopes can reliably place the detector on a planetary disk without human intervention, and some can acquire guidestars automatically. Amateur telescopes can reasonably aspire to 30 arcseconds RMS (RMS is short for "root-mean-square"; 30 arcseconds RMS means that the telescope is within 30 arcseconds about 60% of the time), placing objects in the center of even the highest-power eyepiece. The very best amateur mounts, for example the Software Bisque Paramounts, can do considerably better than this, as long as the telescope itself is sufficiently stable.

Where do you think you're looking?

The first problem in pointing a telescope is deciding what you're trying to look at: the target. The usual way to specify the target is by quoting its right ascension and declination, or RA/Dec. As anyone who is reading this knows, RA/Dec is a sort of celestial longitude and latitude; declinations are measured north and south of the celestial equator, while right ascensions are measured eastwards from the equinox, the intersection between the celestial equator and the ecliptic, the plane of the Earth's orbit around the Sun.

Knowing the RA/Dec of the target isn't the end of the story, as we shall see. There is a rather daunting series of positional-astronomy corrections and transformations which must be made before we are ready to point a telescope. Quite apart from knowing what to do, and in what order, a problem experienced by anyone trying to make these calculations is the dearth of test data. You can calculate an apparent place, for example, but how do you know your answer is correct?

Where does your telescope think it's looking?

The second problem in pointing a telescope is how to set it to a specified attitude. This involves measuring the orientations of the two axes of the mounting. The traditional way to do this is to equip each axis with a setting circle, so that, for an equatorial mounting (the case we'll be concentrating on), you can read off the "hour-angle" and declination, or HA/Dec. Ideally, the index for the hour-angle circle is on an adjacent, movable, circle which is driven by a sidereal clock. These two circles together constitute an analog computer which performs the calculation RA = ST−HA (where ST is the local sidereal time), so that the right ascension can be set, or read off, directly. On some professional telescopes of a few decades ago, the same sort of thing was done using synchros, which were an electromechanical way of relaying the telescope coordinates onto the control console. On modern telescopes, the axes of the mounting are equipped with digital encoders and/or stepper motors.

Encoders are either "absolute" or "incremental". Absolute encoders read out the orientation of the axis directly, so that even at switch-on you know where the telescope is pointed. Incremental encoders, which cost much less than absolute encoders for a given resolution, merely keep track of how far the axis has moved. The high cost of absolute encoders means that even on large professional telescopes ones with a sufficiently high resolution to meet the tracking specification may be unaffordable. In such cases there is, as a rule, some form of "zero-set", where passage through one or more places of known absolute position can be detected accurately and the zero-point of the incremental encoder reset appropriately. Or it may be acceptable to establish the zero-point by finding a bright star.

Stepper motors provide both the torque to drive the telescope and the means to determine where it is pointed. The latter function comes about because the electronics providing the pulses knows how many pulses have been sent and hence how far the motor shaft has turned. This correspondence will be reliably maintained as long as the drive has not stalled at any point.

Encoders and stepper-motors are usually coupled to the axis concerned via some form of gearing arrangement. Gears maintain long-term positional accuracy but need to be of high quality to provide tracking of the necessary smoothness. Rollers (or belt-drives) can achieve the required smoothness relatively easily but are prone to drift.

Encoding arrangements — the encoders or stepper-motors themselves, and the gearing to couple them to the telescope axes — are usually the limiting factor in how well a telescope points. Many large telescopes have failed to meet their pointing specifications mainly because encoders of sufficient quality could not be afforded. The problem is an even greater challenge for the amateur telescope maker, because in many respects the problems do not scale down to match the much lower overall cost of the telescope compared with the observatory giants.

What's the point?

Having seen some of the technical hurdles which must be surmounted in order to achieve accurate automatic pointing, we should review the motivations for bothering to do so. Moreover, we must be clear about why it is worthwhile striving for accurate "blind" pointing, when it is possible to achieve even higher accuracies by first finding a bright "reference star" near the object of interest and then offsetting the telescope a short distance. (This technique is used on some popular SCTs.)

One reason to provide accurate absolute pointing is that it will, depending on the way the telescope control system works, improve the unguided tracking. This is because tracking is differential pointing: the two are merely different aspects of the same thing. Another is that studying the absolute pointing properties of a telescope is an important diagnostic tool; at the very least, it will provide a check on the polar axis alignment, and it may expose bearing runout, insecure optics and other shortcomings for which there may be a mechanical remedy. Accurate pointing will also improve operating efficiency, especially important for large observatories: far less time will be wasted searching for the object, there will be no need to spend time setting first on a bright star and there will be no risk of spending hours observing the wrong target. And another advantage is that an accurate position for anything that is observed is available simply by logging the read-outs: instant astrometry! This is useful when archiving CMOS or CCD exposures, which need accurate coordinates to facilitate subsequent access to the data. To someone determined to do everything the hard way, these may not be compelling arguments, but those people who actually use fully-corrected telescopes tend to be enthusiastic proponents of doing the job properly.

Three steps to pointing a telescope

We have the coordinates of our target, and we wish to rotate the axes of our telescope mount so that the target comes into view. It would be convenient if all we had to do was to set the coordinates on the setting-circles of our telescope. Unfortunately, there's more to it than that; we live on a spinning and wobbling planet, in orbit around a star, looking up through an atmosphere and using imperfect machinery: all these factors have to be taken into account. The sequence of transformations and adjustments that we must carry out is shown diagrammatically in Figure 1.

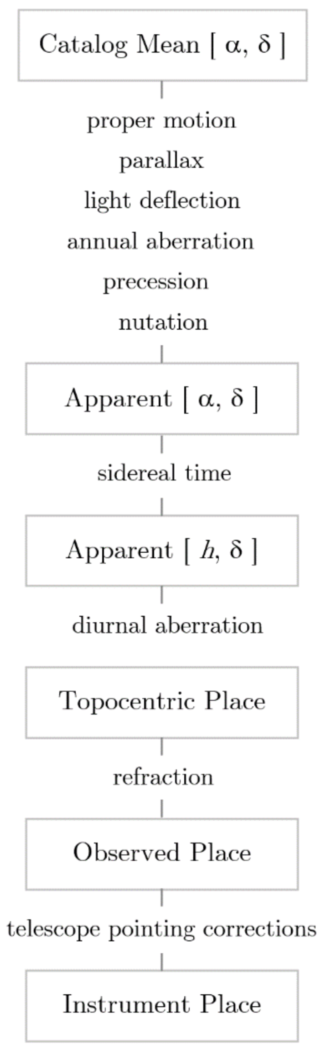

Figure 1.
The sequence of transformations required to convert a star's catalog position into settings for the telescope mount. (The diagram shows the classical, equinox-based, method. A more modern alternative is to start with ICRS coordinates and to compute CIO-based *intermediate places*, with Earth rotation angle replacing sidereal time. All explained here and here (see Figure 1).

The procedure falls into three stages: mean place to apparent place, apparent to observed, observed to instrumental. The first stage, mean to apparent, allows for the fact that the Earth's axis, and hence the celestial equator, is in constant motion due to precession and nutation, and that because of our motion round the Sun the apparent direction of a star is displaced due to annual aberration. The transformation from apparent to observed place involves allowing for Earth rotation and the geographical location of the telescope, and atmospheric refraction. The final stage consists of correcting for instrumental imperfections, virgin territory for most telescope builders.

What does "mean place" mean?

If we wish to observe a planet, we can in most cases look up its "apparent place" from the Astronomical Almanac. An apparent place is the RA/Dec with respect to a celestial equator which is simply the projection onto the sky of the Earth's equator, and an equinox derived from the plane of the Earth's orbit (a slippery concept in its own right). The day-by-day tabulations for a planet, necessary because of its rapid and complex motion, may as well include allowance for precession and the various other effects that lead to apparent place. But this isn't practical for stars, galaxies and so on, where the catalog can afford to provide just one RA/Dec per object, not different coordinates for every day. To get round this difficulty, catalog positions are quoted with respect to the equator and equinox of a certain date, or "epoch"; moreover, the equator and equinox are artificial ones that move steadily and smoothly, unlike the real equator and equinox. This steady movement, the well-known once-every-26,000-years precession of the Earth's pole around the pole of the ecliptic, together with a gradual tilting of the ecliptic, is the long-term component of a complex motion caused by the effects of the gravitation fields of the bodies of the Solar System (principally the Moon) on the distorted and spinning Earth. The steady movement is called "general precession" and produces changes of up to 50 arcseconds a year in the coordinates of stars. The residual wobbles, the largest component of which has a period of 18.6 years, are called "nutation", and affect telescope pointing at the 10 arcsecond level.

The epoch which specifies the mean equator and equinox looks like a year, for example "1950", but often has a ".0" suffix as a warning to the reader that it means more than just a calendar year. For added mystique, the year can be prefixed "B" or "J" (for "Besselian" and "Julian" respectively) for reasons which I am not going into here. (And while we're on the subject, I'm also going to leave out light deflection, annual parallax, diurnal aberration, polar motion and the difference between FK4, FK5 and ICRF, all of which matter if you aspire to 1 arcsecond pointing.) When an RA/Dec is accompanied by, for example, "1950.0", the latter means "with respect to the mean equator and equinox for epoch 1950.0" or "equinox 1950" for short (never "epoch 1950"). Note that (a) the star never actually occupies the given RA/Dec and (b) the mean place isn't, as is popularly supposed "the average place during the given year" — it's closest at the beginning of the year, in fact. Fortunately, a lot of the epoch/equinox confusion is now in the past. Since the 1990s, celestial positions have almost always been given with respect to the International Celestial Reference System (ICRS), which is practically the same thing as "mean J2000". Future star catalogs will stick with ICRS, and J2050 (for example) will never make an appearance.

After allowing for any proper motion, transforming from mean to apparent place consists of applying standard algorithms to allow for precession, nutation and annual aberration. All three effects matter for pointing large telescopes; nutation (up to 10 arcseconds) and even aberration (up to 20 arcseconds) could be omitted for all but the finest amateur telescopes without doing much harm.

The view from your observatory

Next, we need to allow for the fact that the observatory is on the surface of the Earth and beneath an atmosphere. We will need (i) the observatory's latitude, longitude and height above sea level, (ii) the ambient pressure and temperature and (iii) the time.

For Solar-System objects, but not stars etc., we must allow for geocentric parallax and light-time effects. For planets these are quite small, but in the case of the Moon the parallax step is essential, producing a shift of up to a degree.

For reliable blind pointing we need to know the time quite accurately, in the form of Universal Time, UT1. This is not quite the same as Coordinated Universal Time, UTC, which you get from the local civil time by adding or subtracting a whole number of hours (in some places half-hours); to obtain UT1 from UTC you add a correction called ΔUT which allows for the irregular rotation of the Earth. UT predictions from the International Earth Rotation Service are available through the Internet. As ΔUT can grow as large as 0.9 seconds (before a "leap second" is introduced into UTC in order to realign civil and Solar time), which affects pointing by about 10 arcseconds, it is an important effect for big telescopes, and must be allowed for despite the operational complications that it imposes. However, for amateur telescopes it may not be worth worrying about, especially ones without absolute encoders.

From UT1, and using standard algorithms, we can compute the Greenwich Mean Sidereal Time. Adding the (east) longitude we obtain the Local Mean Sidereal Time. Finally we add a nutation term called the "equation of the equinoxes" to obtain the "Local Apparent Sidereal Time". This is ST in the equation HA = ST - RA, giving us the hour angle. From this HA/Dec and the observatory latitude, we can obtain the "topocentric" azimuth and elevation.

(A new scheme, introduced by the IAU in 2000, replaces Sidereal Time with something simpler called "Earth rotation angle". ERA works with a new form of apparent place called "intermediate place", where the RA zero point is almost exactly at ICRS RA zero, avoiding any complications to do with ecliptics and equinoxes.)

The next step is to allow for refraction. The incoming ray from the star is bent downwards as it passes through the atmosphere, so that the object looks higher in the sky than it really is. For a sea-level site, the effect amounts to about an arcminute for an elevation of 45°; there is a rapid increase at lower elevations, enough to squash the setting Sun's disk noticeably. The effect depends mainly on the air pressure and temperature at the telescope; for amateur telescopes, an average pressure and temperature for the site is all that is really needed, but large telescopes may have meteorological sensors that continuously feed readings into the control system. (There are important color effects as well, and the distance between the blue and red parts of an atmospherically-dispersed star image may be several arcseconds.)

What your mounting makes of it

The previous step has given us the "observed place", which is where a perfect telescope, on a perfect mount, perfectly set up, would see the star. But we have a real telescope, which at some level is imperfect in a variety of respects. Its readouts may be offset; the components of the mounting may be out of alignment; the tube may bend under its own weight; the polar axis may not point to the pole. We could take the approach of improving and adjusting the system until it is as good as we can get it; but apart from considerations of time and cost, it may prove difficult to diagnose where the deficiencies are. A more practical plan is to accept the imperfections and correct the star coordinates to take them into account (or apply corrections to the telescope readouts, which comes to the same thing).

The problem of modeling all the distortions and irregularities in a telescope mount may seem intractable. However, it turns out that dramatic improvements in pointing performance can be achieved merely by correcting for a handful of well-understood effects that all telescopes exhibit to some extent. These effects consist of six purely geometrical terms, supplemented by two or three likely flexures. For an equatorial mount, the six geometrical terms are as follows:

Table 1.
The six geometrical terms for an equatorial mount. h and δ are hour angle and declination.

term	*description*	Δh	$\Delta\delta$
IH	h index error	IH	
ID	δ index error		Id
CH	collimation error	CH sec δ	
NP	h/δ nonperpendicularity	NP tan δ	
MA	polar axis left-right misalignment	MA cos h tan δ	MA sin h
ME	polar axis vertical misalignment	ME sin h tan δ	ME cos h

The IH and ID terms are simply the zero-point corrections to the hour angle and declination readouts. The collimation error CH describes how accurately the telescope optics are aligned within the tube, whether the tube is at right-angles to the declination axis and any east-west displacement of the crosswires, the center of the CMOS/CCD, or whatever other aiming-point is being used. The declination axis is supposed to be at right-angles to the polar axis; NP describes any deviation from this condition. The terms MA and ME describe how far the polar axis is from the true pole, up-down in the case of ME and left-right for MA. (Altazimuth telescopes have a similar set of terms; the zero points are in azimuth and elevation, the collimation error is left-right rather than east-west, the nonperpendicularity is between the azimuth and elevation axes, and the mount misalignment terms describe north-south and east-west tilts in the azimuth axis.)

In addition to these purely geometric terms, three types of flexure are often found:

Table 2.
Three different forms of flexure; ϕ is the site latitude.

term	*description*	Δh	$\Delta\delta$
TF	tube flexure	TF cos ϕ sin h sec δ	TF (cos ϕ cos h sin δ - sin ϕ cos δ)
FO	fork flexure		FO cos h
DAF	δ axis flexure	–DAF(cos ϕ cos h + sin ϕ tan δ)	

TF describes a droop in the telescope which gets worse the lower you go. FO is fork flexure. It is always seen in fork equatorials, and sometimes in yoke mounts and horseshoe mounts. It can be a big effect; the Lick 120-inch telescope, for example, has an FO value of about 4 arcminutes. DAF is flop in a cantilevered declination axis, for example a German equatorial or a cross-axis mount. Most small telescopes need either FO or DAF and there may be signs of TF.

Harmonic terms are often seen as well, caused by miscentering and eccentricity in the various axes and drive wheels.

In the above tables, each of the coefficients IH...DAF is a small angle, usually expressed in arcseconds. The formulas for dHA and dDec give the corrections to be added to the telescope readouts. The set of coefficient values and the formulas constitute the telescope's pointing model, the different terms adding up to yield an overall correction in each axis. Applying a model consisting of the six geometrical terms plus, say, fork flexure, can have an astonishing effect on the pointing accuracy of a telescope. Uncorrected, it may be necessary to locate the target by using a finder, hard to do if the object is faint. With the corrections applied, it is usual for the target to appear centered in even a high-power eyepiece — every time, all over the sky.

We now know how to point a telescope. We know what calculations to perform on the catalog position of the star to predict its position on a given night; we know what needs to be done to take account of the location of the observatory and the distortions in the atmosphere; we know how to apply telescope pointing corrections and have a good idea of what form of model to use. But we're not quite there yet. How do we know what values to use for the coefficients IH, ID and so on? Enter TPoint.

Bridging the gap with TPoint™

TPoint is an interactive software tool that unscrambles observations of star positions into a pointing model for the telescope concerned. It is used by professional observatories worldwide (Keck, GBT, Gemini, ALMA, AAT, ARC, WIYN, WHT, UKIRT, JBO, IRTF, NSST, ESO, CTIO, SOAR, MMT, Magellan, LBT...) on telescopes of many different designs — equatorial and altazimuth, optical/IR and radio. TPoint does three things:

- It accepts lists of pointing observations specifying (i) where the star really was and (ii) where the telescope readouts said the star was.
- It fits a user-specified pointing model (the desired list of coefficient names) to the observations, so that the coefficient values give the best possible match between the star positions and the corrected telescope readouts.
- It displays in a variety of graphical formats the remaining pointing errors (the residuals). If the plots suggest that systematic errors remain, the operator can include additional terms in the model and try again.

Two TPoint implementations are available, with the same software at their core, but offering different styles of use. Both are available from Software Bisque. They are descendants of programs developed for the Anglo-Australian Telescope in the 1970s, subsequently used by many major observatories and running on VAX/VMS and Unix machines. One of these implementations, TPoint Professional (a.k.a, "ProTPoint"), has a similar look-and-feel to the original command-line-operated version but runs on PCs under Windows, macOS, and Linux. The other TPoint version has a greatly enhanced interface which is fully integrated with TheSky™ software's telescope-control facilities and can also run under Windows, macOS, and Linux. Both versions of TPoint use the same modeling techniques and give identical results.

This is how you use TPoint:

1. Point the telescope at a selection of stars all over the sky (TPoint comes with several catalogs containing suitable stars) and carefully center the image in a reticle eyepiece or on a selected sensor pixel. Log the star's catalog position, the telescope RA/Dec readouts, and the sidereal time.
2. Run the TPoint package: read in the observations and establish a pointing model (either by trial and error or by letting TPoint do it automatically for you).
3. Apply the pointing model to the telescope.

The only tricky part is Step 3. If the telescope has a computer control system that already supports TPoint corrections, such as TheSky from Software Bisque, then applying the model simply means entering the new set of numbers. If not, it will be necessary to provide whatever tools are required to apply the corrections. This is a problem which the individual telescope-maker must address: it might involve writing new control software, or, for a telescope with no computer, using a spreadsheet package to generate an all-sky look-up table.

It should always be borne in mind that an excellent result reported by TPoint is of little value if a comparable result is not delivered during normal operation. Accurate implementation of the model, with careful monitoring and adjustment where necessary, is obviously vital. But to provide arcsecond pointing on a large modern telescope means proper management not only of star position data but also the position in the focal plane into which the star image is to be sent. This requires artful presentation, powerful on-line calibration tools and meticulous setting-up each night.

What does it all mean?

The two tables present some of the TPoint correction descriptions of real physical effects: if NP is 50 arcseconds, TPoint is estimating that the two axes of the mounting are out of square by this amount. A question that is often asked is whether there is any need for the pointing model to reflect mechanical reality in this way — why grapple with HA/Dec nonperpendicularity and declination-axis flop when a bunch of spherical harmonics or polynomials would do just as well, maybe better? There are at least two reasons why TPoint's mechanically-based approach is preferable:

- A mechanical model should be a snugger fit, involving fewer terms and requiring fewer observations to pin it down. It is also likely to be better-behaved; polynomials in particular have a habit of doing well in between the observations but going berserk outside the region covered.
- You can learn useful things from a mechanical model. For example, you can establish how far out the polar axis is, even at sites where Polaris is invisible, and you might pick up signs of flop in a mirror cell or some other problem which has a mechanical remedy.

Once you've got as far as you can with the mechanical approach, TPoint still has the full armory of harmonics and polynomials at its disposal for a final mopping-up operation.

Whether using mechanical models or empirical ones, it is important to distinguish between repeatable effects and random noise. Given a set of observations, it is tempting to go on adding terms to the model until the RMS figure is as small as you can get it. But how do you know you've been describing real properties of the telescope and haven't simply been chasing noise? The answer is to study the repeatability of the model from one test to the next. If the terms that you expect to remain constant do so, then the model has some predictive value and the terms may mean something. If the terms change radically from test to test, they are worse than useless and should be omitted from the model. However, some terms may be expected to change. If the mounting axes do not have absolute encoders, so that you have to "synch" on a star to get started, then the values of the IH and ID terms will naturally vary. Similarly, if the telescope has been recollimated, or the sensor moved, the terms CH and ID will not be constant. It is particularly important to be aware of which terms should and should not persist from one test to the next when searching for new terms. TPoint includes facilities for combining data from many tests, so that systematic errors can emerge from the noise; before the data can be combined, it is important to remove from each data set the effects of the terms that will have varied.

Long-term monitoring of the pointing terms can be interesting. Figure 2 shows a plot of the polar-axis misalignment in elevation term ME for the Anglo-Australian Telescope over about a decade.

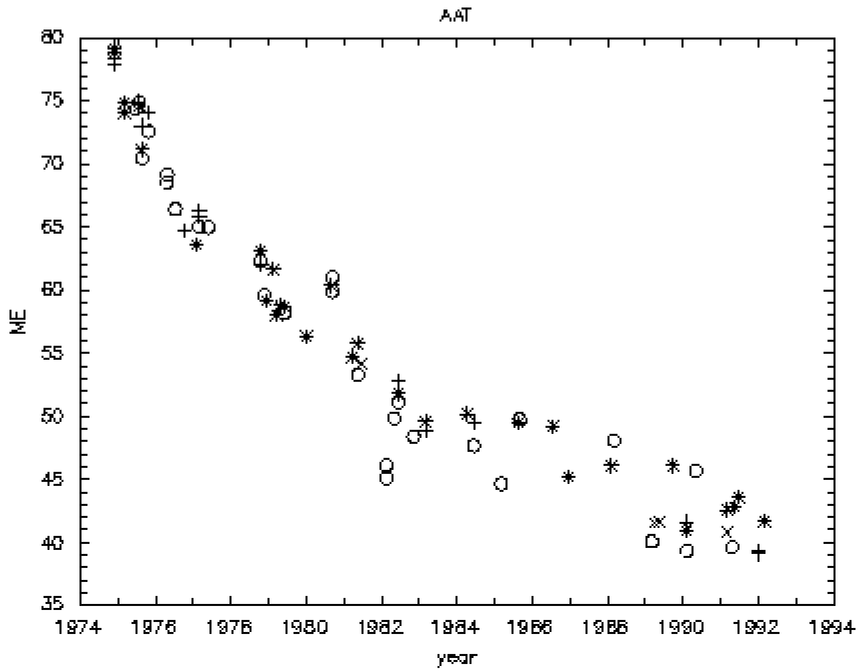

Figure 2.
The changing polar-axis elevation for the Anglo-Australian Telescope, revealing a downward drift of a few arcseconds per year. The effect may be the result of distortion in the concrete pier as the concrete ages, an explanation which is consistent with the slowdown in later years. Each marker comes from one pointing test; the different symbols correspond to the AAT's various interchangeable top-ends.

There is a constant drift downward, amounting to 50 arcseconds. The cause is unknown: curing in the concrete pier, settling of foundations or even changes in the surrounding water table have been suggested. More will be said about polar axis alignment later on.

Operational aspects

Even on the best telescopes, some start-of-night calibration is usual. This amounts to a mini-pointing test and a subsequent adjustment to just a few of the terms. On telescopes with absolute encoders, three stars will give a good estimate of CH and ID for the night (CA and IE on an altazimuth); five or six stars will allow IH (or IA) also to be adjusted; two or three more stars and the polar-axis alignment terms, ME and MA can be added to the list (AW and AN in the case of an altazimuth). Fewer stars will suffice if the terms are large and the accuracy objectives modest. A full-scale pointing test, which might involve 50-100 stars, can be time-consuming, and on computer-controlled telescopes is often done without operator intervention, using a CMOS/CCD or an autoguider to nudge the telescope into position before logging each star. Advanced amateur telescopes are perfectly capable of carrying out such robotic tests, but safety must be a primary consideration in such cases — it is wise to keep an eye on the telescope's gyrations, and vital to do so if there are people about.

A particularly useful feature of TPoint is its ability to set up the polar axis, without any form of polar trail tests or observations of Polaris. The procedure is straightforward. An ordinary pointing test is carried out, using as many stars as possible and covering the whole visible sky. TPoint is then used to fit a model which includes all the standard terms, including the polar-axis misalignment terms MA and ME. The polar axis can then be aligned in azimuth by rotating the mounting through an angle of MA/cos(lat) (with due regard to sign conventions) and in elevation by the difference between the actual and desired ME. The recommended ME value is that which corresponds to the refracted pole rather than the true pole (the latter corresponding to ME = 0), in order to minimize field rotation. (An interesting by-product of setting the polar axis to the refracted pole, the result which all the standard methods attempt to deliver in fact, is that the tracking rate near the zenith becomes the textbook 15 arcseconds per second. The refraction squashes the picture slightly, and if the polar axis were truly parallel to the Earth's axis the tracking rate would be a little less than 15 arcseconds per second; however, tilting the polar axis up to the refracted pole means that a slightly increased tracking rate would be needed; and it turns out the two effects cancel out.)

TPoint in action

As a demonstration of what TPoint can do, we will look at three very different telescopes: a large equatorial several decades old, an example of the new breed of altazimuth-mounted large telescopes, and finally an advanced amateur telescope.

I will start by looking at data from the Hale 200-inch telescope, Palomar Mountain — a file of 39 observations which includes stars all over the sky down to elevation 16°. Using only the simple zero-point corrections (IH and ID), the RMS error is 35 arcseconds: the TPoint "scatter" diagram (Figure 3) shows where the stars would have appeared when acquired blind.

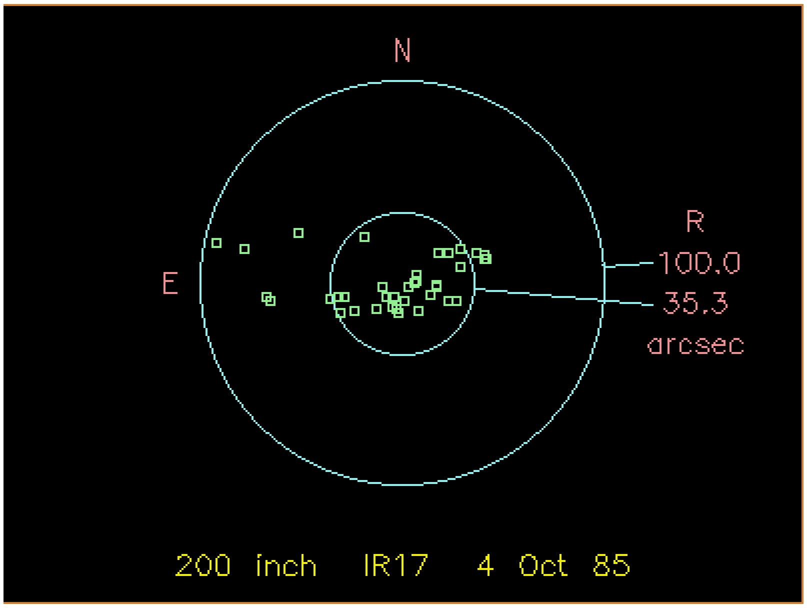

Figure 3.
Intrinsic pointing accuracy of the 200-inch Hale Telescope. The only corrections that have been applied are zero-points in hour angle and declination. This scatter plot shows where each star would have appeared by simply setting the telescope dials; the better the pointing, the tighter the grouping. The inner circle shows the pointing accuracy which about half of the stars in the sample will reach — in this case 35 arcseconds.

This result is already considerably better than the 1 arcminute often quoted and illustrates the importance of the basic positional-astronomy corrections — precession, nutation, aberration, refraction. We now ask TPoint to use the standard six-coefficient geometrical model. This produces an impressive 8 arcseconds RMS (Figure 4).

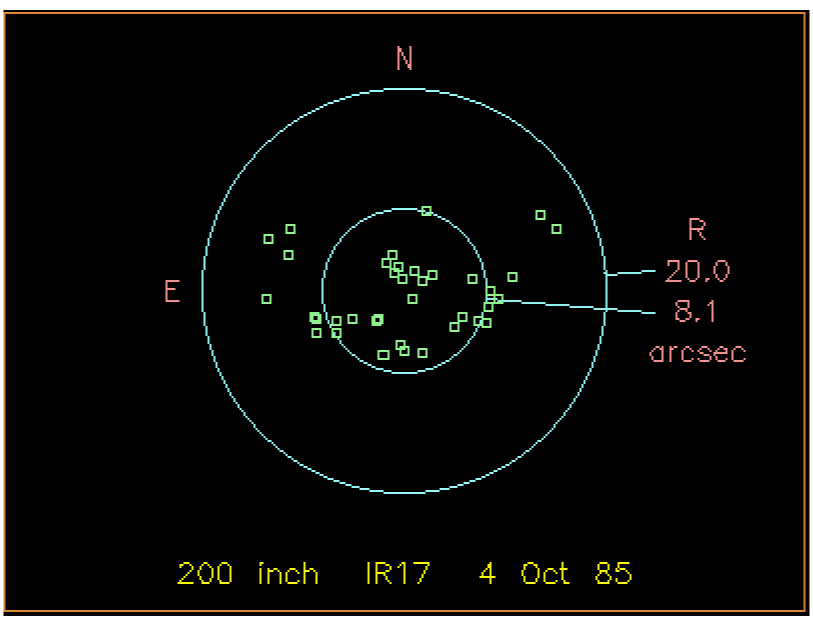

Figure 4.
By correcting for the other fundamental errors (NP, CH, ME and MA), the Two Hundred Inch Telescope achieves 8 arcsecond pointing.

At this stage we plot a variety of different graphs of the residuals in order to look for uncorrected effects (Figure 5).

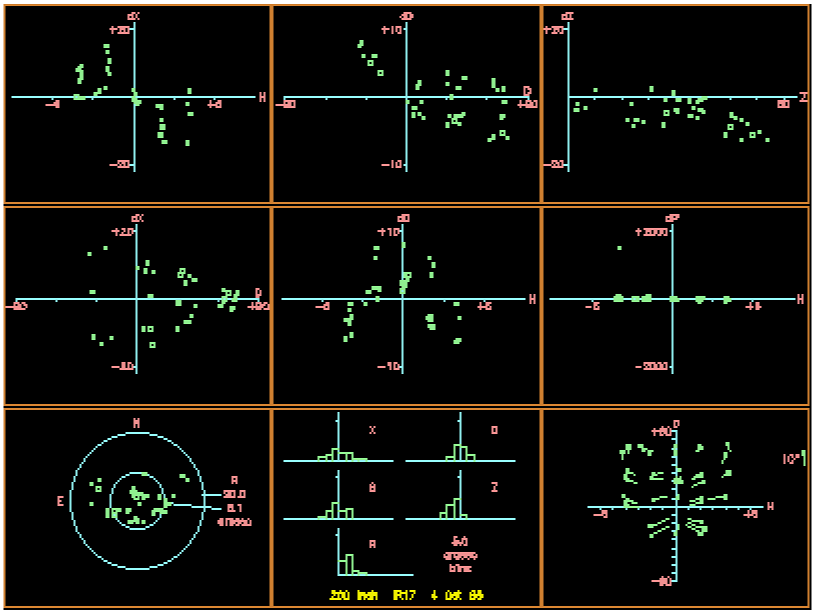

Figure 5.
A selection of TPoint plots of the 200-inch data fitted with the basic 6-term model. Runout is evident in both hour angle (top-left, east-west errors versus hour angle) and declination (top-center, declination errors versus declination). At this stage there also appears to be tube flexure (top-right, zenith distance errors versus zenith distance) and fork flexure (center, declination errors versus hour angle) but it turns out these go away when the runouts are corrected. The other plots are east-west errors against declination (center-left), h/d nonperpendicularity versus hour angle (center-right), the scatter diagram (bottom-left), the error distributions (bottom-center) and the map of error vectors on the sky (bottom-right).

There are signs of mis-centering in both hour angle and declination; adding to the model the HHSH and HDSD terms (the first-harmonic sine terms in each axis, to match the observed phase), we reach 3.3 arcseconds RMS (Figure 6).

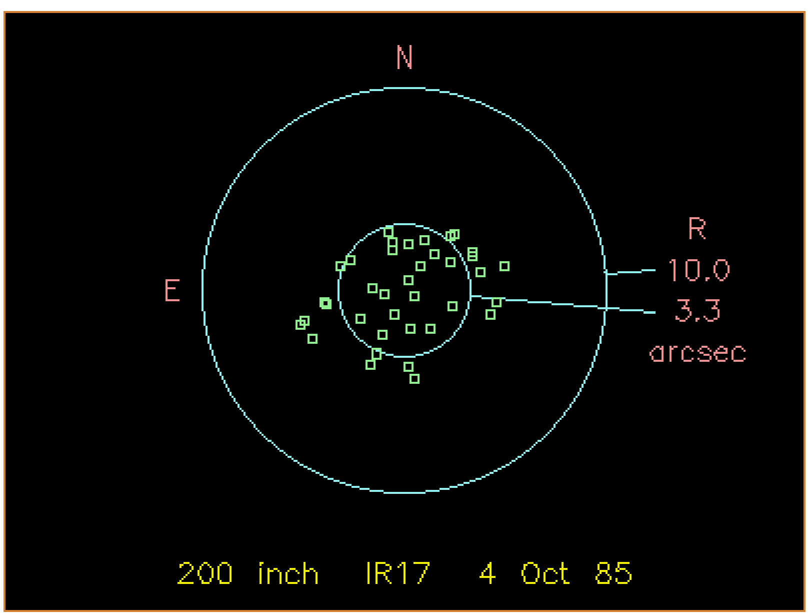

Figure 6.
The result of adding to the 200-inch model the terms HHSH and HDSD. The pointing accuracy is now 3.3 arcseconds RMS, an excellent result. The outer circle is about the size of Saturn's disk.

This is an extremely fine result by any standard, and an indication of what an amazing telescope the 200-inch is; in fact further work with TPoint suggests the ultimate performance may be below 2 arcseconds RMS. It is also a sobering thought that the original 200-inch control system design included an analog computer for applying TPoint-style pointing corrections, using cams and synchros instead of computers and encoders. The device was conceived by the astronomer Sinclair Smith, who tragically died before his work could come to fruition. The design was completed by Ed Poitras, but for various reasons the device was never built. Had it been, the 200-inch might have delivered 5-arcsecond pointing in the 1940s.

My second example is a modern altazimuth: the Multiple-Mirror Telescope in Arizona, which has recently undergone conversion to a 6.5-meter single-mirror configuration. The raw data, observations of 36 stars, were acquired using the central reference telescope. The RMS pointing accuracy after fitting the azimuth and elevation zero-points, IA and IE, is 7.7 arcseconds. Including CA, NPAE, AW and AN to complete the basic 6-term geometrical model for an altazimuth mount, there is only a marginal improvement, to about 7.4 arcseconds RMS, suggesting that the mount is set up very accurately. Plotting the vertical component of the pointing error against zenith distance (Figure 7) reveals significant tube flexure (or perhaps elevation runout).

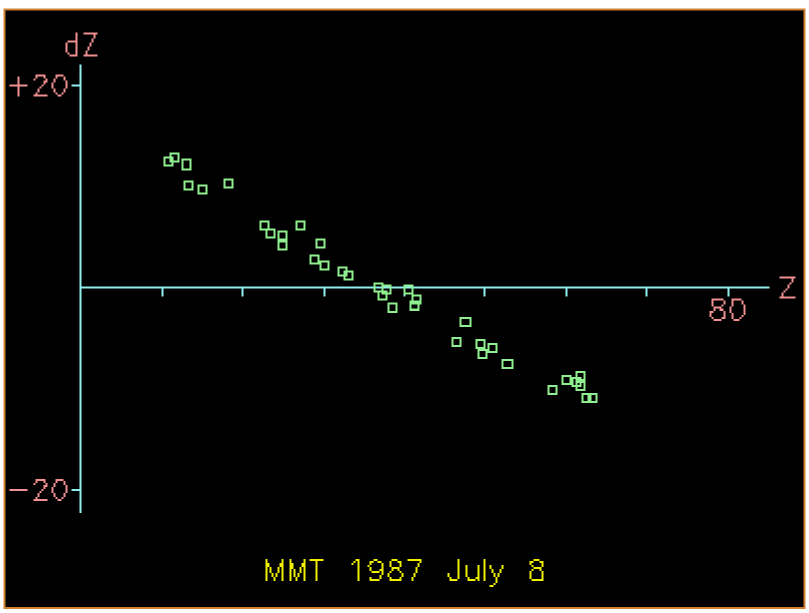

Figure 7.
The zenith-distance residuals in the MMT mount after the basic 6-term model has been applied: the plot suggests that there is significant tube flexure.

Once the TF term is added to the model, the RMS figure improves dramatically, to 1.9 arcseconds. There is also evidence of azimuth mis-centering, and adding the term HACA brings a further reduction, to 1.2 arcseconds RMS (Figure 8).

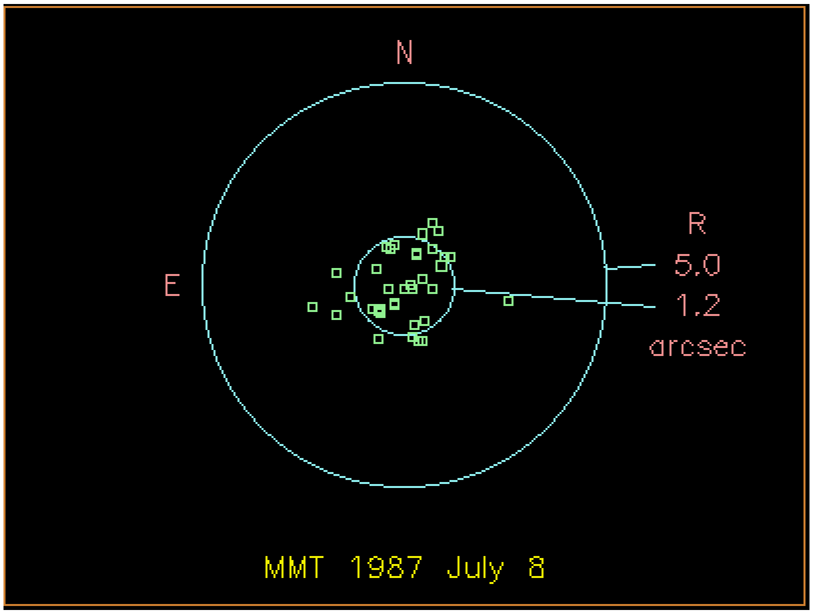

Figure 8.
Adding to the MMT model the terms TF and also HACA gives 1.2 arcseconds RMS. The ultimate pointing performance of the MMT mount is probably even better, considerably under 1 arcsecond.

Further work has suggested that the MMT mount could achieve a figure as low as 0.6 arcseconds, but a much denser test (or several test-runs combined) would be needed to confirm this result. Similarly spectacular results are currently emerging from the new ESO VLT (four 8-meter telescopes in Chile) and Gemini (8-meter telescopes in Hawaii and Chile), with dense tests of a hundred stars or more returning RMS results well under 1 arcsecond and in-service performance at almost that level.

Finally, as an example of what can be achieved in the amateur sphere, I will look at the 24-inch Cassegrain reflector of the Lone Star Observatory in Caney, Oklahoma, which was constructed by a group of 12 amateur astronomers from Dallas, Texas. Prior to TPoint analysis, the telescope delivered 10 arcminute pointing, making acquisition of faint objects quite difficult. Preliminary TPoint tests revealed about 0.2° of polar-axis misalignment, plus indications of mechanical slop and cyclic errors in hour angle. The test run shown here was carried out after (i) the polar-axis had been adjusted in accordance with TPoint's findings, (ii) loose components had been identified and tightened and (iii) the existing hour-angle cam drive had been replaced with a superior worm-based system. The run involved observations of 123 stars spaced roughly 10° apart and took just two hours. Earlier use of TPoint had already provided a sufficiently good model for the computer control system to place all the stars straight into the 800x illuminated-reticle eyepiece without any need for a finder. The basic six-coefficient geometrical model produced a promising 46 arcseconds RMS result. Cyclic errors, often present because of residual mis-centering of bearings etc., and fork flexure were then apparent. Five additional terms were added: HHSH (corrections to hour angle proportional to sin AH), HXCH (east-west corrections proportional to cos HA), HDCD (corrections to declination proportional to cos HA) and finally FO and HDSH2 (corrections to declination proportional to cos HA and sin 2xHA respectively). No significant tube flexure was detected. The final 11-term model delivered 22 arcseconds RMS (Figure 9), a fine result.

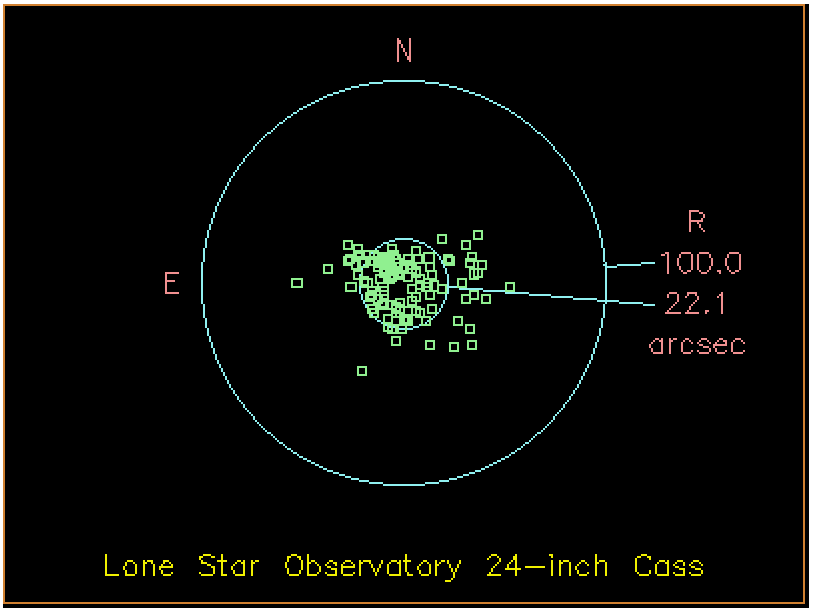

Figure 9.
The pointing performance of the 24-inch Cassegrain telescope of the Lone Star Observatory. The basic six-coefficient geometrical model delivered 46 arcsecond pointing. After correction for various flexures and misalignments (HHSH, HXCH, HDCD, FO and HDSH2), pointing accuracy of 22 arcseconds RMS was achieved on this amateur-built computer-controlled telescope. The inner circle of the plot is about the same size as Jupiter's disk.

To convey what this means, I cannot improve on what Barry Smith, Chairman of the Lone Star Observatory, wrote a few weeks later:

"Absolutely incredible pointing for the visual observer. I went to a host of Uppsala galaxies and it nailed every one. And they are damned easy to see when you know that they are virtually dead solid in the center of the field of view. Galaxy after galaxy with a rated magnitude of 15 to 16.5 could be seen. Granted, I didn't see a lot of detail in the mag 16.5 galaxy, but knowing where to look, and confirming that the ultra-faint spot in fact moved with the field, enabled me to add quite a few objects to my life list."