Control of Antenna Focus – Detailed Requirements

This topic has been discussed several times and was written up in May 2009 in the attached note ("Control of Antenna Focus and Related Topics"). That note was based on what it says in the antenna specifications about the performance of the subreflector positioning mechanism. Following that I checked what the performance of the drives actually is with the relevant people in the antenna IPT. The significant new information is that 1) the resolution of the subreflector drives is sufficiently fine that we can adjust the subreflector continuously without having to make a correction for the difference between the desired position and the actual position due to limited resolution, and 2) the mechanisms are specified to perform large numbers of operations without wearing out, so that reasonably frequent updates to the position should not be a problem. After further discussion, and with guidance from CIPT, the following implementation is proposed. The intention is that this will give us the freedom to do either continuous updating of the subreflector position or updating at chosen moments only, while handling the resulting pointing and path length changes in a consistent manner.

Introduction

The key components of the software involved here are: the focus model, the pointing model and the delay model. For each of quantity, focus, pointing and delay, these models describe the behaviour of the antenna in a parameterized form. Various pieces of information need to be supplied – most obviously the desired pointing position and which receiver band is in use[1], but also temperature information – so that the models can be used to calculate the necessary offsets to the subreflector position, the pointing and the delay. The details are given in the sections below.

On the question of temperatures we have to take account of the fact that the optimum focus position (and the path length) will depend on the temperature of the antenna structure and temperature gradients in it, and not just on the ambient air temperature. The design should therefore allow for a reasonable (10?) number of temperatures to be read in (presumably these are available in the ACU) and used in the models. If temperature differentials turn out to be what is important, that will be covered by using positive and negative coefficients for the relevant temperatures. There is an important caveat here. The temperature sensors and readout electronics are not 100% reliable, so any time we use such data in this sort of application we must have sanity checks in place – e.g. comparisons between several sensors to see whether the differences in readings are physically sensible – and a mechanism for then dealing with this. We can't for example just set unreasonable values to zero and carry on using them in the model: we have to do something like interpolate between the other values. This will take a good deal of work so we should not plan to implement the temperature corrections for some time yet.

The suggestion here is that we implement the temperature corrections in two steps: we first put in a correction term for the ambient temperature, which we know is present on the antennas that we have accepted so far, and provide the tools for further corrections based on the differences between the other readings and ambient. That way we can leave those additional factors turned off until we find evidence that they are worth adding and when we introduce them they should not disturb the existing model significantly.

---

[1] One important question here is whether it is sufficient to define just the operating band or whether the actual observing frequency is required. One might expect some small changes in all these coefficients as a function of frequency. The reason for this is the illumination will change with frequency and if the dish has a complex pattern of deformations the best focus will then change. I believe however that these changes will be small and in most cases the effects will normally be taken out by our standard observing procedures – e.g. if we are doing high precision work we will in any case need to do focus checks reasonably often and these should be done at the observing frequency. There is in any case little prospect of us being able to characterize each antenna at multiple frequencies in every band. I propose therefore that the observing frequency should not be included as a parameter in the model, only the band.

Focus Calculation

We assume that the model for the position of the optimum antenna focus (X_A, Y_A, Z_A) depends only on the commanded[2] elevation (El) and temperatures (T_a, T_1, T_2 …), where T_a is the ambient temperature and the others are from sensors on the antenna. For generality we can adopt the form:

$$X_A = X_R + X_C \times [\cos(El) - \cos(El_R)] + X_S \times [\sin(El) - \sin(El_R)] + X_{Ta} \times (T_{amb} - T_R) + X_{T1} \times (T_1 - T_{amb}) + X_{T2} \times (T_2 - T_{amb}) \ldots$$

$$Y_A = Y_R + Y_C \times [\cos(El) - \cos(El_R)] + Y_S \times [\sin(El) - \sin(El_R)] + Y_{Ta} \times (T_{amb} - T_R) + Y_{T1} \times (T_1 - T_{amb}) + Y_{T2} \times (T_2 - T_{amb}) \ldots$$

$$Z_A = Z_R + Z_C \times [\cos(El) - \cos(El_R)] + Z_S \times [\sin(El) - \sin(El_R)] + Z_{Ta} \times (T_{amb} - T_R) + Z_{T1} \times (T_1 - T_{amb}) + Z_{T2} \times (T_2 - T_{amb}) \ldots$$

Here El_R is the reference elevation, which we need to choose and then keep fixed. El_R = 50 degrees is recommended. X_R, Y_R and Z_R define the reference focus for a given antenna at the reference elevation. This will be determined during the initial characterization of an antenna but should then remain fixed unless some major piece of work is carried out, such as replacing one of the legs of the quadrapod, which might be expected to change the nominal focus by a large amount. Using a reference elevation is not strictly necessary but by choosing a point in the middle of the range of elevation we limit the size of the resulting pointing corrections, which is clearly a good thing to do (see later for details).

Similarly T_R is a reference temperature. Assuming that the temperatures are in Kelvin, I suggest we set this to 273 K since this is a good estimate for the mean temperature at the AOS. Notice how the other temperatures are referenced to the ambient temperature and not to T_R. This produces the behaviour suggested in the introduction, where all except the first term are assumed to be associated with differences with respect to ambient rather than absolute temperatures. Obviously this will only work correctly if all temperatures are measured in the same units, i.e. Kelvin, and are reasonably well calibrated.

The other terms account for the gravitational deformations of the dish and the subreflector support structure. For a symmetrical structure we expect X_C, X_S, Y_S and Z_C all to be zero, but it is possible that we will find it necessary to have non-zero terms for some of these, so it is simpler to include them at this stage. We expect these parameters to be very different for the different antenna designs and slightly different from one antenna to another of the same type.

Similarly, the majority of the temperature terms will presumably be zero but it is better to put them in now rather than add them later. We do know that there is a significant variation of the Z focus with temperature so we can introduce the term Z_{Ta} right away. This definitely needs to be determined separately for each antenna.

So far the question of dependence on band was left out. There are two ways in which this could be dealt with: a) having a different model with all the terms above for each band, and b) having additional terms (X_B, Y_B, Z_B) which are the focus offsets between cartridges. The offsets will be small and there is no reason to think that they should vary with elevation or temperature. For this reason and because of how we want to treat the tilting of the subreflector, which comes next, I propose that we adopt option b). The offsets will be determined experimentally and they represent the focus offsets for a given receiver band with respect to the reference position focus (X_R, Y_R, Z_R). They are needed to compensate for the distortions in the wavefronts produced by the cartridges and the differences in the heights of the beam waists above the focal plane[3]. We will refine these values as we characterize the system and should expect that we may need to re-determine them if we replace the front-end. The instruction from the higher level software to the focus module should therefore include the Band number. (I assume that

[2] It seems clear that this should be the commanded elevation rather than the actual elevation, so that the subreflector will start to move at the beginning of a slew rather than lagging behind the telescope.

[3] Again some frequency dependence within a band can be expected here but, for the reasons given in footnote 1, it is proposed not to try to account for that here.

there is already a system of band names in use. It would be wise to allow for a few additional bands being added in the future.)

In addition to these, we need to allow "user offsets" to the focus. These presumably take the simple form (X_{Off}, Y_{Off}, Z_{Off}) and are added to the values above before they are sent to the sub-reflector. The user offsets are what are manipulated by the system when doing a "focus" measurement, i.e. moving the focus to determining the optimum position, to which it is then set. There is also the possibility, especially on the total power antennas, that we may wish to make observations where we move the subreflector back and forth in Z on alternating scans as a way of removing residual baseline ripple. I strongly recommend that the Mount Panel be updated to display separately the total focus position (i.e. the values send to and/or returned from the ACU) and the user offsets.

Note from Robert Lucas here: there is a 'delta' field for focus coordinates in the MountPanel and in the ICD, but I'm afraid it has a special meaning in terms of metrology (that is user offsets above the internal ACU Focus model, which I think we chose not to use). What we might want to do is to display our user offsets instead of these. Comment from REH. Yes. Also I am not sure that all the antenna designs have implemented those ACU offsets.

Subreflector Tilts

In general we will also need to set the "X-tip" and "Y-tilt" according to which band is being used. These are rotations of the subreflector about the X and Y axes respectively. According to the official ALMA document on coordinate systems[4] we are supposed to use α and β for these. (There is a Z-rotation capability, but nobody has been able to think of any uses for it, so I will ignore that here.) Note that the intention is that these rotations take place around axes passing through the focus of the secondary mirror and that is what I have assumed here. We need to check that this is in fact the case. As discussed in the older note, the tilts are used to realign the illumination pattern which would otherwise be skewed by the fact that the feeds are off axis. According to this model then, the frequency dependence is covered by having five numbers per band X_B, Y_B, Z_B, α_B and β_B which are loaded when the prime observing band is changed. In the case of (X_B, Y_B, Z_B) these are added to the other offsets. By contrast α_B and β_B will be absolute values (i.e. we assume that the antennas have been well enough constructed that setting them to zero puts the subreflector perpendicular to the axis to within the required accuracy, which is only a tenth of a degree or so). The values of the tip and tilt for each band are calculated directly from the positions of the feeds and are fixed[5] – i.e. they are the same on all the 12m antennas – a different set of numbers will be needed for the 7m's, assuming that tilting is implemented on them. The second attachment here ("Pointing Offsets and De-focus") gives the values of the tilts to be applied, as well as initial estimates of the pointing offsets.

Again we do need to have user offsets to the tilts, α_{Off} and β_{Off}. These will probably only be used for commissioning purposes – e.g. to check how much the tilt effects the pointing and focus and that we do get the expected improvement in sensitivity when we use the correct tip and tilt. Ideally the mount panel would display the total tip and tilt and the user tip and tilt separately but, given the limited usage we will make of these, that may not be worthwhile. Obviously we do need to be able to "set" and "get" them from the command line.

The values found by the steps above should be summed, i.e.

$$X_T = X_A + X_B + X_{Off}\,,\ \ Y_T = \ldots \quad \text{and} \quad \alpha_T = \alpha_B + \alpha_{Off}\,,\ \beta_T = \ldots$$

and the five values (X_T, Y_T, Z_T, α_T, β_T) are what should be sent to the hexapod drive.

[4] ALMA-80.05.00.00-009-A-SPE

[5] In the analysis of the Melco antenna design, a fit for the best value of X-tip as a function of elevation is included. We need to check whether this additional adjustment provides any significant advantage. Here I have assumed that this is not done. If we wanted to add this we would need to add an item U_A with a dependence on elevation of the same form as X_A, etc.

Updating

The attached memo discusses the question of how frequently the focus should be updated. The conclusion is that when tracking a source it only strictly needs to be done every few minutes but, if we move to a calibrator a several degrees away in elevation, an update would need to be done then. Now that we know that the drives have both higher resolution and can make many more cycles than is required in the specifications, it makes more sense for us to use continuous updating. We should not however overdo it. Updating on every 48ms tick would give something like 2×10^{10} operations in a 30 year life, which is definitely too high for any mechanism. A rate of once every few seconds is easily fast enough for all practical purposes. We do also want to be able to turn off the updating at least for experimental purposes and quite probably for when we are doing fast-switching – e.g. we may wish to not do an update when we move to the phase calibrator but to turn on the updating again when we come back to the astronomical source.

It seem to me that the simplest method is to have a parameter, P, for the period between updates, presumable with units of seconds (although in practice the update would take place at the next 48 msec tick after the period has elapsed) and with P = 0 meaning no updating. I suggest a default value of 5 seconds. The response to the arrival of a non-zero value of P should be to force an update on the next tick. This means that the higher level software can control the behaviour easily when we do switching: sending P = 0 at with the command to move to the reference and P = 5 with the command to move to the source will prevent the subreflector moving in response to the switch but ensure that the position is updated each time we come back to the source. Note that it is essential that the update uses the elevation of the new commanded pointing position, not the old one. Sending P = 5 with each command would update the focus for both the source and the reference if we find that is preferable. (Note that the motion of the drives should be fast enough to get to the appropriate focus for the other band while the antenna is moving.) We will decide which approach to adopt on the basis of experiments. *(I appreciate that there are several alternative designs possible here and computing should decide which is the most secure and natural.)*

In principle the tip and tilt only need to be updated when we change the band that is being used for the astronomical source. In particular, when we are doing fast-switching using a different band for the phase calibration, the tip and tilt should <u>not</u> be updated on each switch: the moves required are quite large so the subreflector would not be able to get to the required position in time. Because the subreflector mechanism is a hexapod all six commanded positions of the actuators change every time we change one of the input values – X, Y, Z, α and β. It is probably best therefore for us to send all five commands on every update, but to make sure that the higher level software is always sending to the focus module the band number, B, that corresponds to the astronomical source and not changing it to the reference band, when we go to the reference. (This is in contrast to what has to be done with the pointing – see below.)

Pointing Model

We are planning to change to a pointing model that resembles the above. This is not the place to discuss the details, but my suggestion is that we have a basic model that represents the antenna, call it Q_A consisting of the usual terms – zero points, collimation, alignment and various functions of azimuth and elevation – plus a second set of terms, call them Q_B, which are associated with the band. In addition there will be user offsets Q_{Off} which enable us to move the beam around and keep track of the drifts that occur in the pointing over and above those included in the model. Once again I think it is important to have these user offsets displayed on the Mount Panel. *(Note that at present we effectively have these user offsets through the auxiliary pointing model. The proposal here is to drop that. One of the weaknesses of the present arrangement is that it does not display these although they are stored in the ASDM.)*

The feature that we need to concentrate on here is that all the movements of the subreflector that we have discussed above effect the position of the beam on the sky so we have to add a further set of terms to allow for this. Let's call them Q_F. These will need to be calculated by the focus model and applied on top of the pointing model.

The vectors Q above have different lengths: Q_A has 18 terms in our present model. I am not sure whether we plan to limit ourselves to just this number or keep space for a few more in case we find we need them. Q_B has just two (large) terms, collimation and elevation offset, corresponding to CA and IE in Q_A. If the receiver and cartridge structures are stiff enough these should just be fixed numbers, but there could be flexure and, since this has never been measured or even calculated, we may need to allow additional terms for this[6]. The user and focus vectors should have only two terms each, again collimation and elevation offset, so we can write $Q_F = (C_F, E_F)$. It is worth noting that, unlike the focus, we do not include temperature terms in the pointing model but instead assume that, if there is a temperature dependence, the metrology system in the antenna takes care of it. It is possible that we may find that explicit temperature dependence outside that covered by the metrology model but I think that it is again reasonable to assume that this will be removed by our normal procedures of doing pointing checks reasonably frequency when we are making critical observations.

Note that, since we are going to introduce a non-linear calculation of the pointing from these terms, it is important that any large terms, specifically Q_A and Q_B, must be summed before the non-linear part of the calculation is done. The important term is in fact the collimation one. The contribution to the collimation offset from the focus will normally be relatively small, since it is associated with the X-motion, so that can be added as a simple pointing offset afterward the other large terms. The other point to note here is that, unlike the focus module, the pointing module does have to use the right band number when we switch to a phase reference source. If we don't to this we will miss the reference source completely! By contrast, the fact that we have not applied the right focus offsets and tilts when we go to the reference (which will normally be at a low frequency) will only lower the gain slightly and this will have a negligible effect on accuracy with which we can measuring the phase.

Calculation of Effect of Focus and Tilt on Pointing

For the calculation of the two terms in Q_F we assume that the focus offsets are small enough for a linear model to apply and that the antennas are well enough built that there are no cross terms – i.e. that the collimation shift depends linearly on the X focus and not on Y or Z, and the elevation shift depends linearly on Y and not on X or Z. I believe that our experience so far suggests that this is good enough but I am not sure that we have done any accurate experiments to confirm it. Theory says that if the subreflector is a long way from its optimum position then the relationship between pointing offset and focus offsets becomes non-linear, but this is not an issue because we don't make critical measurements with the system a long way out of focus.

With these approximations the expressions are very simple:

$$C_F = C_X \times (X_T - X_R) + C_V \times \beta_T$$

$$E_F = E_Y \times (Y_T - Y_R) - E_U \times \alpha_T$$

In both cases the coefficients can be calculated from theory – see the second of the attached notes. For the translation motions of the subreflector we have $C_X = E_Y \approx 34$ arc-seconds per mm. A weak dependence the illumination and therefore on band and frequency is expected. The value of 34.1"/mm has been measured by experiment, mostly at band 3, but we should probably check it other bands. In principle it needs to be rather accurate because the excursion in Y is quite large over the full range of elevations – several millimetres – which implies that we need better than 1% accuracy to get arc-second accuracy. In practice any small deviations will get taken out by slight modifications to terms with the same elevation dependence that are in

---

[6] The plausible terms are a collimation term and an elevation offset that depend on sine and/or cos of elevation. Note that not all of these occur in the present pointing model, so carrying a set of delta values for the existing 18 term model does not in fact cover this need. Instead I think it would probably be better to add explicit band-dependent terms to the pointing model like C (collimation) = $C_R + C_C \times [\cos(El) - \cos(El_R)] + S_C \times [\sin(El) - \sin(El_R)]$ and the equivalent in elevation. We have no evidence that these terms are necessary so I do not think that they should be added yet.

the pointing model. I therefore think that we should be able to use a fixed value for all antennas (except the 7m's) for this.

So long as the rotational motion takes place about the focal point of the secondary, the coefficients for the tilt C_V and E_U should both be given by the ratio of the distance between the foci (6.177m) and the Cassegrain focal length (96m), i.e. 0.0643. This is purely geometric – we do not expect any band or frequency dependence here. In practice the centre of rotation may not be exactly correct, which will change this value somewhat. We need to check this for the different antenna designs (the 7m will of course be quite different). Again great accuracy is not needed because small errors will be taken out elsewhere, in this case by the pointing offsets associated with the individual bands.

Note that, in the expressions above for the pointing offsets due to the focus adjustments, the reference position has been subtracted. The reason for that is simply to keep the corrections as small as possible, which will reduce errors due to uncertainty in the coefficients, and to arrange things so that the nominal pointing model corresponds to an antenna with the secondary at the reference position, which seems appropriate.

These two focus offsets (C_F, E_F) should be calculated from the current focus position and applied on top of the pointing model each time the pointing is updated.

Path Length Model

As discussed in the following note, there is a relatively large path length change associated with the Z-focus changes and a much smaller one associated with X and Y moves. These should be linear for modest movements and so the form to be used is

$$P_F = P_X \times (X_A - X_R) + P_Y \times (Y_A - Y_R) + P_Z \times (Z_A - Z_R) ,$$

where (X_A, Y_A , Z_A) are the focus positions, and the reference positions are subtracted for the same reasons as before. *(Point for debate: it could be argued that these should be the reported positions and not the commanded positions, but I think this is an unnecessary complication because of the time that it takes the subreflector to come into position. It should be adequate to assume that the subreflector mechanism does what it is asked to.)*

We expect that the Z-coefficient will be close to P_Z = 1.74 (assuming that path and motion are in the same units – presumably meters). We will be able to check this experimentally. Rather high accuracy is in principle needed to get the specified micron-level accuracy in the path. The main variation of Z is, however, proportional to sine of elevation, so if we make an error in the value of P_Z the form is exactly the same as an error in the altitude of the pad and so it should be taken out in fitting the baseline parameters.

The other coefficients, P_X and P_Y, are likely to be of the order of 0.05 and they will be band-dependent: depending on which side of the axis a receiver is on, moving the subreflector laterally will add phase in one case and take it away from the other. This means is that we need to have band-dependent parameters here. Clearly the presence of these terms complicates the behaviour during band-switching phase reference observations and this needs more investigation both on the theoretical and experimental side.

Interactions between the Models

The outline given above was chosen, at Robert Lucas's suggestion, so as to avoid interactions between the different models, i.e. we don't want a change in the focus model to cause a change the pointing model. Key to this is the way in which the measurements of the current positions are stored. The proposal is to do the following:

1) The ABM receives the required astronomical position and calculates the Az,El required assuming perfect antenna and no atmosphere.
2) Add the user offsets Q_{Off}. *(I think this is the correct place to do it so that we are applying the offsets in the right frame. So long as they are small it would not matter if they are done in a different order.)*

3) Applies refraction correction to El.

4) Apply pointing model. This now includes the basic antenna model Q_A, the receiver terms Q_B. As noted these need to be combined to do the non-linear calculation correctly.

5) Apply pointing corrections implied by the focus positions Q_F.

6) Send the resulting positions as commands to the ACU.

7) The ACU applies the metrology corrections and then closes the servo.

8) The ACU reads back the encoder positions.

9) The ACU subtracts the metrology corrections *(Is this right?)* and sends the "actual" position back to the ABM.

10) The ABM subtracts focus offsets Q_F from both the actual and the commanded positions.

11) The ABM does not reverse the pointing model terms Q_A and Q_B or the user offsets Q_{Off} .

12) The ABM subtracts the refraction from both the actual and the commanded positions.

13) Both these positions and the user offsets are stored in the pointing table of the ASDM.

If one looks closely at this one sees that the outcome is that items that are assumed to be corrected internally in the antenna system – the metrology, the focus correction terms and the refraction – do not show up in the values that are stored: they go in as part of the control of the antenna but they are taken out again. Because of line 11) however, the effects of the pointing model, including the band-dependent terms and the user offsets, are left visible. This means that when we take the "actual" values measured for the apparent position of a source and reduce them using, e.g., T-point then the result will be the full pointing model. This is much more convenient than if they had been taken out and it has the important consequence that we will, if necessary, be able to do this analysis without knowing what pointing model was being used at the time.

It is however important for our ability to process data from different bands correctly that we also store the band-dependent offsets in the ASDM. That way we should be able to separate the antenna pointing model proper from the offsets associated with the receiver bands.

Summary of Information that Needs to be Stored

(I do not have a good understanding of this – needs to be clarified by CIPT.)

1) The ASDM needs to contain (once) the parameters in use in the models identified above, i.e. the pointing model, the focus model (including the band-dependent terms, etc.) and the path model. I understand that there is a table set of tables for these.

2) I assume that the ASDM also needs to contain the values that are derived from the model, i.e. the actual values sent to the ACU for the five degrees of freedom of the secondary, and also the values read back. These numbers have the same status as the commanded and what I am calling "actual" pointing positions that are, as I understand, already being stored with the individual data samples. I do not know what flexibility there is in the time interval. Clearly it is only necessary to store the command values each time they are updated. Storing the values that were read back at those same times is clearly not the right thing to do, since the mechanism will not have moved at that point. I suppose that the logical thing is to store "actual" values that are read back when the integration is going on. It appears to me that we do not need to store the output of the path length calculation since all the values that go into that are already included.

3) I do not know what should go into the Telescope Monitor and Configuration Database but it is clear that we want to be in a position to track the historical evolution as function of time of the parameters that we are using in these models in each antenna.

Richard Hills with the important ideas coming from Robert Lucas and Jeff Kern 19th Jan 2010

Control of Antenna Focus and Related Topics

This is a discussion of the approach to be adopted in the control of the secondary mirror and, in particular, how to adjust it to allow for gravitational and thermal deformations in the dish and the apex support structure.

The main issue is whether or not the subreflector should always be placed in the position that gives the optimum antenna performance, which is assumed here to be the maximum on-axis gain. The obvious answer is that it should be, but the complication is that any change in the subreflector position will change the signal path and hence the phase, so one may prefer not to make subreflector movements either during integrations or when moving between a source and a calibrator. This conflict needs resolution.

Specifications

The antenna specifications call for a range of movement of +/- 5mm in X and Y and +/-10 mm in Z. The maximum step size (effectively the resolution) is 20 microns in all three coordinates but the "differential accuracy" is set at 10 microns in X and Y and 5 microns in Z. The required speeds for the motion are quite high: 0.5 mm/sec in X and Y and 2 mm/sec in Z. This means that the motion from one end of the range to the other takes 20 seconds in X and Y and only 10 seconds in Z. By contrast the antenna takes 30 seconds to cover the full range of elevation at the nominal maximum speed of 3 degrees/second. Nothing is said in the antenna specification about acceleration and deceleration of the subreflector movements, or about settling times. We should check on what those are in the real implementations. On the basis of the numbers that are given, however, it is clear that the intention was to make it safe to assume that the subreflector will always get to the correct position as fast as the antenna can. So long as we are dealing with the elevation-dependent gravitational errors is concerned, and assuming that the command to move the subreflector to a new position is sent at the same time as the command to move the antenna, the subreflector should in fact always get there first. A possible problem is that the requirement only calls for 20 million steps of adjustment in the lifetime of the device. I estimate that we might well do that in <u>one</u> year of normal use.

Tolerances – 1) amplitude effects

In order to work out how frequently we need to update the position of the subreflector, we should first ask how accurately we need to control it. We look first at the effect on the amplitude of the signal. There is an ALMA memo – number 479, by Bryan Butler – that discusses the loss in signal due to focus errors. I have run some Zemax and GRASP simulations and get similar (although not quite identical) results. The conclusions in the memo are that for a loss of 1% the axial focus should be held correct to 0.09 wavelengths and the lateral focus to 0.45 wavelengths. Now since we are trying to get 1% overall calibration accuracy we obviously cannot allow the errors due to each axis of the subreflector positioning to be as large as 1%. For example if the gain changes by 1% just due to the change in focus when we move from a source to a calibrator we are obviously in trouble. The 1% goal does however only apply at the longer wavelengths – frequencies of up to 370GHz. The following table gives the loss for some important frequencies for a particular set of assumed errors.

		X	Y	Z	
	Move for 1%	0.45	0.45	0.09	waves
	Assumed errors	120	120	40	microns
Freq	Wavelength	Loss			Sum
GHz	Microns	percent			percent
300	1000	0.07	0.07	0.20	0.34
370	811	0.11	0.11	0.30	0.52
650	462	0.33	0.33	0.93	1.59
920	326	0.67	0.67	1.86	3.20

Since the adjustment in X and Y are less critical than in Z, I have set the values so that most of the loss is associated with the error in Z. It can be seen that these values give a total loss of just over half a percent at 370 GHz and over 3% at 920 GHz. It could well be argued that these losses are still too high but I suggest we take them as working numbers for now.

The general formula for estimating the loss due to an error in positioning the subreflector is

$$\text{Loss (\%)} = (F/300)^2 \times [\,(\Delta X/450)^2 + (\Delta Y/450)^2 + (\Delta Z/90)^2\,] \qquad \text{equation (1)}$$

where F is the frequency in GHz and ΔX, ΔY and ΔZ are the errors in microns.

Now the movements required to correct for gravitational deflections are of the form:

X = X_0

Y = Y_0 + Y_c × cos(Elevation)

Z = Z_0 + Z_s × sin(Elevation)

where the coefficents are Y_s ≈ –3.5 mm and –7 mm for the Vertex and Melco antennas respectively and Z_s ≈ +2 mm in both cases. From the prototype measurements, we expect the AEM design to have smaller values.

By differentiating the above expressions one can immediately loss resulting from a change in elevation if we do not adjust the position of the secondary. We see that the Z focus is changing most rapidly at low elevations while the Y focus changes most quickly near the zenith. If we restrict ourselves to the elevation range 20 to 80 degrees we see that in both cases the tolerance limits are already reached with changes of elevation of about 1 degree. This is however not quite the right thing to do since the tolerances were set assuming that errors were present in all three axes simultaneously. Here is a plot of the actual situation:

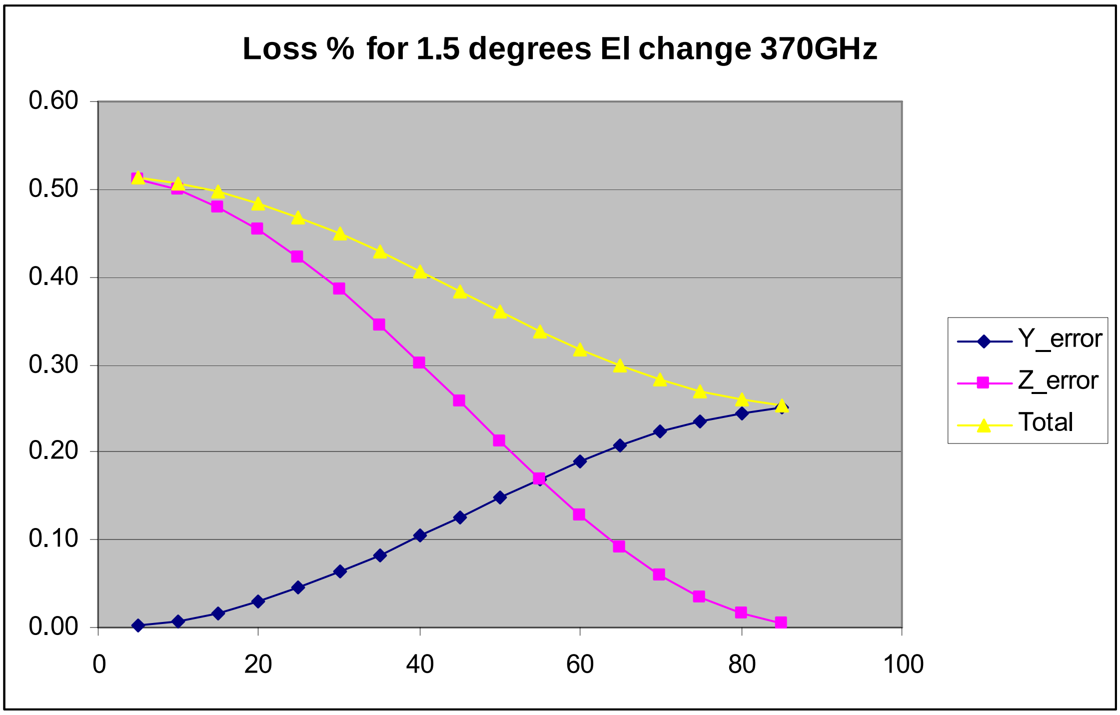

This uses the Melco numbers but the situation, except at high elevations, is not much different with the Vertex ones. If we regard this loss as an error on the calibration the values in this plot are about as much as we can tolerate – remember that we also have to allow for thermal effects on the focus as well as these gravitational ones. This is for 370 GHz and for a 1.5 degree change in elevation – the losses scale with the square of frequency and with the square of the change in elevation.

Since most <u>amplitude</u> calibration sources will be more that 1.5 degrees in elevation away from our astronomical target, the most obvious approach is that we should always update the focus position when we move the antenna to do an amplitude calibration. As an alternative however we could keep the focus position set for the astronomical source, calculate the loss resulting from the error in focus on the calibrator and apply a correction for this as part of the calibration.

So if we wish to limit the number of times the subreflector is a possible strategy is as follows: the focus would be set to the correct position at the start of an observing block and the control system would record the antenna elevation at which this update was made. It would then keep watch on the commanded elevation and update the subreflector position whenever the difference exceeds a certain threshold. A suitable expression for the threshold on the change of elevation would be (0.7 + 300 / F) degrees, where F is the observing frequency in GHz. (I have fiddled the numbers here to give the 1.5 degrees found above at 370 GHz and to provide somewhat smaller errors at lower frequencies and greater at high frequencies.)

Note that, if we are tracking a single source, the 1.5 degrees of change in elevation quoted above corresponds to times of from ~6 to perhaps 15 minutes. These times are comparable to or longer than the time between calibrations, so in most cases it would be the move to the calibrator that would trigger a re-focus.

In addition to the gravitational deflections we have to deal with thermal deformations. We think that (at least on the Vertex dish) these will be of order 30 or 40 microns of change in Z focus for each degree of change in the ambient air temperature. According to the numbers above this means that for high frequencies we will need to update the focus if the temperature changes by more than about 1K. Clearly it would be possible to incorporate this into the strategy suggested above by recording the temperatures (probably one recorded on the dish rather than the air temperature) and testing the change in this relative to a second threshold. It would however be more direct and complete to calculate the required secondary position, allowing for both temperature and gravity, compare it to the current position, estimate the losses according to equation (1) above and issue the command if the loss is greater than some value, which could for example be (0.3 + F/1000)% with F again be the frequency in GHz.

There is no simple theory to predict the changes in the X and Y focus as a function of temperature (in fact the FEA does not predict the Z shift correctly either!). Since the structure is symmetrical we do not expect any change with the ambient temperature, but we may find significant shifts due to asymmetric heating and if so we can try to correct for them by using the measured differences between the temperature sensor readings. Clearly this is a likely future development which should be allowed for in the software design.

Tolerances – 2) phase effects

Now let us look at the path length errors that result from these focus changes. It is clear that the most important effect is due to the axial movement of the subreflector, i.e. the Z focus. On axis the change of path is of course twice the movement (because of the reflection) but when one integrates over the range of incident angles (taking account of the aperture weighting) one finds that the mean path change is 1.74 times the Z movement. To first order there is no effect due to X and Y movements for a completely symmetrical system, but because our feeds are off-axis there will in fact be a small change in path. I estimate this at only about 0.05 (TBC) times the X or Y movement, which is small but unfortunately not negligible.

It is important to appreciate that there are two components to the corrections that we have to apply to the position of the subreflector: 1) the movement of the subreflector with respect to the some reference frame (say the elevation axis) and 2) the change in the position of the prime focus – i.e. the location of the focus of the best-fitting paraboloid. This implies that the path length will change as a function of elevation whether or not we actively change the focus. In the case of the Vertex design, the Finite Element Analysis (FEA) shows that turning on gravity in the zenith position causes the subreflector to move down by about 1.25 mm while the prime focus moves up by about 0.65 mm – together these give the prediction correction in Z-focus of ~1.9mm. The changes in path length, however, include the movements of the primary and the receiver as well as the subreflector. It turns out that there is a good deal of compensation here: the numbers I have from the Vertex FEA indicate that, if we do not move the subreflector to the optimum Z focus, the change in path due resulting from turning on gravity in the zenith pointing case is about -0.3 mm, whereas if we do move the subreflector it is +3.0 mm.

As far as I know there is no separate item in the phase error budget for the subreflector positioning. The relevant numbers in the systems specifications on the total delay associated with the antenna movements are as follows:

> Systematic, for (az,el) change of 2.0deg, < 8 fsec
> Random, for (az,el) change of 2.0deg: < 15 fsec
> Systematic, az ± 180 º rms < 100 fsec, el ± 40 º rms < 50 fsec
> Random, for full range of permitted az,el : < 32 fsec

Light travels 0.3 microns in 1 fsec so these numbers correspond to path length changes of only 2.4(!), 4.5, 30, 15 and 9.6 microns respectively.

In the antenna specifications we find:

> The repeatable residual delay for an antenna shall not change by more than 20 microns when the antenna moves between any two points 2 degrees apart in the sky. (Note however that this residual delay is defined as the difference between a "nominal" antenna and the actual one. If we accept that the different designs can be referenced to different "nominal" antennas, the real delay changes could be larger than this even though the antennas are deemed to have met the spec.)

and

> The non-repeatable residual delay under Primary Operating Conditions (section 4.4.3) must be less than 15 micrometers RSS when tracking an astronomical source at sidereal rate.

It is not at all clear to me that these numbers are consistent, but let us see what values we can expect from just the subreflector positioning, which is of course part of the antenna.

First of all the Z motion has a specified accuracy of 5 microns. This corresponds to a path length error of 8.7 microns. It is not clear whether this is intended as a peak or an rms value.

The step size of 20 microns corresponds to 35 microns of path. If we regard this as an error and if there are many steps during and observations, so that any systematic offset is smoothed out, then the effect on the phase coherence is just an rms error of 10 microns. Unfortunately we will not be tracking the calibrators for long, so we must allow for the full effect of the resolution on the accuracy of the phase measured on the calibrator. Assuming that the resolution of the Z motion is really limited in this way, then, if we decide update the subreflector position continuously, it is clear that we cannot simply go to the nearest step and ignore the error. Instead we will have to calculate the phase error due to the difference between the actual position and where we want it to be and apply that as a correction.

Taking the values quoted above from the Vertex FEA we see that the change in going from say elevation 30 to elevation 32 degrees if we do not update the Z-position is -9 microns, compared to +90 microns if we do update it. (The change is largest at low elevation.) The effects due to the Y movements are likely to be smaller, but this needs to be checked. As expected the conclusion from this is that the safest thing to do is NOT to change the focus position when going between the source and the phase calibrator. Assuming that we will normally use a calibrator for phase that is no more than 2 degrees away in the sky and that the FEA is correct, this will limit the error to ~9 microns. Note, however, that this does not actually meet the 2.4 microns required in the system spec. In principle therefore we should still be applying a correction to the phase for these effects, which since they are gravitational are in principle completely repeatable. We can certainly do this by using the FEA for the three different dish designs and calculating the correction. Whether we will be able to measure the changes to confirm the model is much less clear – a term that varies as sin(El) will be very hard to separate from many other effects such as baseline errors, LO changes and atmospheric refraction. Conversely, if there is a significant effect that differs from the prediction, then that will lead to errors in the baseline determination, etc.

The thermal effects on path are likely to be quite large but of course this is one of the reasons why we plan to do frequent observations of phase calibrators: so long as we do this fast enough the thermal effects should get removed rather well. Thermal drifts will however complicate the process of baseline determination. To achieve high accuracy it will probably be necessary to do this at night.

Tolerances – 3) beam shape effects

When we are observing extended sources, we do of course need to know the beam patterns of the antennas accurately. The fact that when we are out of focus the peak gain goes down means of course that other parts of the beam pattern come up. I have not done anything quantitative on this but clearly this is another argument that points in the direction of updating the secondary position continuously, or at least at frequent intervals and especially when we are observing at high frequencies.

Tilting of the subreflector

In all our antenna designs the subreflector has 5 degrees of freedom: i.e. three displacements and two tilts. The intention of the specification is that the tilts take place around the location of the near focus of the hyperboloidal secondary mirror. This de-couples the effects of the displacements and the tilts. In the following it is assumed that this has been achieved exactly[7]. The requirement to tilt the subreflector was not in the original antenna specification but was added later when it was realized that, because the receivers are off-axis there is a significant advantage in sensitivity to be gained by tilting it (ALMA memo 545 – see also the change request at http://edm.alma.cl/tiny/9m64d.html for further discussion). In fact the optimum tilt is such that that axis of the subreflector lies half-way between the axis of the primary and the line from the feed to the prime focus. The plan is therefore to command the tilts of the subreflector according to the band being used for the astronomical observations. The tilts required are in the range 0.5 to 1 degree. Note that tilting the subreflector will produce a significant shift in the pointing, but so long as the tilt is about the correct position it should not be necessary to alter the focus. When we are doing fast-switching using a different band as the phase calibrator, we will NOT change the tilt. This would in any case be undesirable because it is likely to introduce phase changes, but in fact the drives would not be able to make these large changes in tilt quickly enough. The fact that there is a loss of a few, or at most ten, percent in sensitivity on the phase calibrator is not very critical.

Pointing Offsets

The X and Y subreflector motions and the tilts all produce changes in the pointing offsets. The attached note contains some discussion and numbers. These offsets are quite large – for example a move from 70 to 72 degrees elevation would produce a difference in the pointing of ~5 arc seconds depending on whether or not the subreflector position in Y is updated. It is clear that whatever strategy we adopt for updating the subreflector needs to include the corrections to the pointing as well as to the phase and amplitude.

Proposed Strategies

A) The Conservative case

In the following I assume that we are doing a “standard” interferometric observation with an amplitude calibration source being observed rarely and phase calibrators being observed frequently.

Normally the sequence will start with an amplitude calibrator and the subreflector should be set to the correct position for that elevation.

The next object will normally be a phase calibrator. If this is far enough away from this to require a change in position an update should be done when the antenna there.

[7] If the point of rotation is not at the focus, then a tilt also produces a lateral displacement. As far as I can see errors of a few millimeters here will not have any significant effect on performance so long as the tilts are used in the way described here – i.e. fixed positions for different bands.

The subreflector should then not be moved for the subsequent observations of the astronomical source and calibrator until it is found that the change in focus due to gravitational and thermal effects exceeds the threshold. At that point a final phase calibration should be made.

The subreflector position should then be updated and the sequence restarted, i.e. calibrator, source, calibrator…etc., ending on the phase calibrator again.

Note that in this case we will be able to check from the change in phase, whether the subreflector has actually moved by the correct amount.

The problem with this is that, according to the estimates above, the accuracy of the resolution of the positioning is not good enough to make the resulting errors negligible. This means that we will have to keep track of the phase errors produced and apply a correction based on the difference between the nominal position and the actual one. We will also need to keep updating the pointing corrections to account for the effects of the subreflector positioning. This is all getting rather complicated so we should at least look at the alternative.

B) The Optimistic case

Here we simply adjust the subreflector all the time, either by sending frequent commands, say once a second and certainly every time we move to a calibrator, or simply by letting the ACU in the antenna take care of it. (I think this is called "autonomous mode".) We would then apply a correction to the phase for the resulting path difference relative to some nominal position for the subreflector. This correction would be calculated from the position actually read back – i.e. taking account of the limited resolution of the hexapod. If we just used a fixed position, say (0,0,0), for this reference, then the resulting phase will still contain an elevation dependent term due to the gravity as explained above. To the extent that we understand the gravity deformations we can in fact work out what correction is needed to give us a constant phase independent of elevation and apply that. Similarly if we apply refocusing for thermal changes we can allow for the movements we have made by correcting the phase. In principle at least, we would not need to do anything about pointing corrections in this case, since the subreflector position would always vary in a simple way with elevation and the offsets would be absorbed into the pointing model.

Discussion

As far as I can see, since the limitations in both cases are set by out lack of knowledge of the true deflections, as opposed to what the models say, the two methods should give the same answers. The second one is apparently rather simpler to implement and slightly more efficient. Note however that we might find that we are wearing out the drives if we tell them to move on every cycle of the fast switching.

The other practical difference is that in the second option we rely completely on the subreflector position drive to do exactly what it is asked to do. By contrast in the first we tell it to stay fixed while we make the critical move from source to reference.

In the end we may want to try both these, so the immediate question is, "Which should we try to implement first?"

Zero-spacing (Single Dish) Observations

Here the considerations are somewhat different. Phase is not important so one would again think that continuous updating is the way to go. On the other hand for some cases the baseline ripple due to standing waves between the feed and the subreflector may be important, despite the fact that we have done our best to suppress them in the subreflector design. This means that the strategy of only updating when necessary is probably best here. In this case one would complete a cycle of On and Off observations before doing the update.

Richard Hills with key input from Alison Peck 19th May 2009 (Updated 20th May taking account of comments received from Bojan Nikolic, Ruediger Kneissl, Stuartt Corder and Peter Napier.)

Pointing Offsets and De-focus

Because the receiver cartridges are not on the nominal axis of the antenna, we expect there to be pointing offsets. This picture shows the locations of the cartridges and the windows:

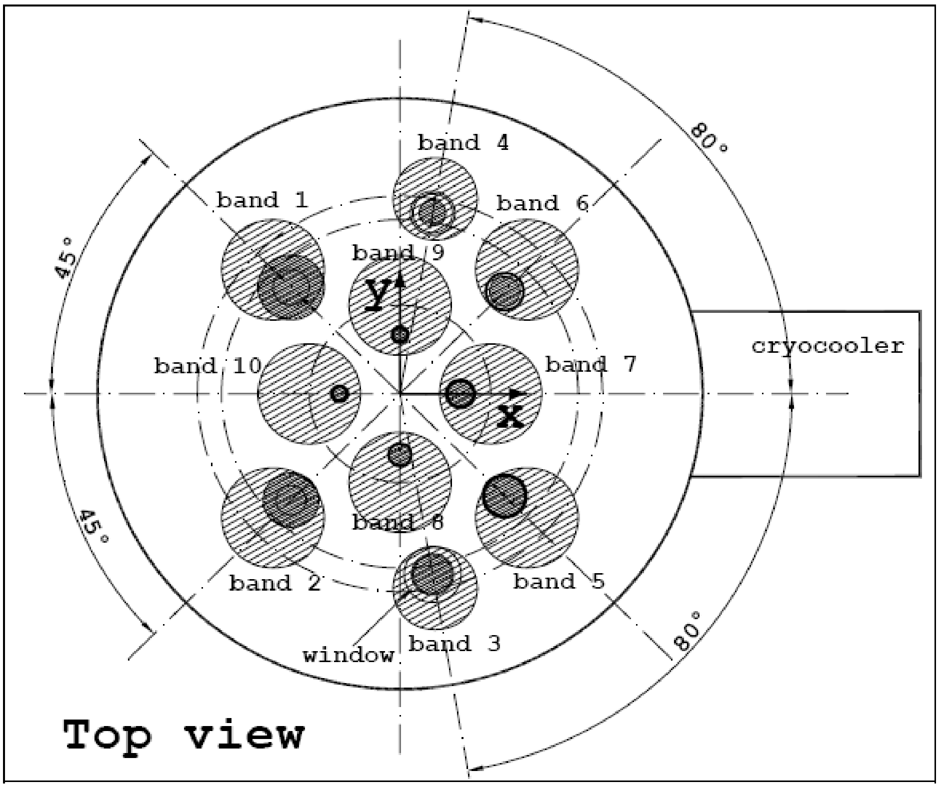

Note that because of the warm optics, the positions of the beams on bands 3 and 4 are not as far off axis as the positions of the windows.

Here are the expected pointing offsets with respect to the nominal boresight. (Signs to be confirmed!)

	Millimeters		Arcsecs		Subreflector (degrees)	
Band	Beam X	Beam Y	Off_Az	Off_El	Tilt_X	Tilt_Y
1	-180.3	180.3	-387.4	387.4	-0.861	0.861
2	-180.3	-180.3	-387.4	-387.4	-0.861	-0.861
3	32.7	-185.1	70.2	-397.8	0.156	-0.884
4	33.7	191.0	72.4	410.5	0.161	0.912
5	173.2	-173.2	372.2	-372.2	0.827	-0.827
6	173.2	173.2	372.2	372.2	0.827	0.827
7	100.0	0.0	214.9	0.0	0.477	0.000
8	0.0	-103.3	0.0	-221.9	0.000	-0.493
9	0.0	100.0	0.0	214.9	0.000	0.477
10	-100.0	0.0	-214.9	0.0	-0.477	0.000

The offsets are calculated simply by taking the beam positions in the focal plane, dividing by the focal length (96m at the secondary focus), taking the arcsine and converting into seconds of arc.

Also shown in the table are the tilts of the subreflector that should be implemented to optimize the illumination pattern on the primary. These are obtained by dividing the beam position offsets by the distance to the secondary mirror (nominally 6m at the rim), taking the arcsine and dividing by 2. These have not presently been implemented. When they are the pointing offsets will be reduced by a factor of 2. *(Note added Jan 2010: With the model now proposed, no such reduction in the pointing offsets should be applied. Instead the effects of the tilts will be corrected separately and automatically.)*

Offsets associated with Secondary Mirror Movements

The coordinates of the secondary mirror motion are (Xs, Ys, Zs) with Zs pointing outwards along the telescope axis and Ys pointing downwards at the groud when the dish is pointing at the horizon. When the secondary is moved in Zs (axial focus) there should be no effect on pointing. When it moves in Ys there will be a change in elevation and when it moves in Xs there will be a change in the Azimuth offset. You can think of this as being due to the movement of the virtual image of the feed horn formed by the secondary mirror near the prime focus.

For a telescope with a large f-ratio the relationship would be just angle = dX / Fp where Fp is the primary focal length, 4.8m in our case. This would give a movement of ~43 arc second per millimeter.

Because our dishes have a small f-ratio (0.4) the angle is reduced by the amount of the “Beam Deviation Factor”. This depends slightly on the illumination taper but should be of order 0.8. There is a second correction for the fact that we are using a Cassegrain telescope which depends on the magnification – the ratio of the secondary to primary focal lengths, which is 20 for our dishes. The result of all this is that expect a movement of about 33 arc seconds per millimeter of lateral movement of the subreflector.

The subreflector can also be tilted and the point about which the tilt is defined to take place is chosen to the location of the prime focus. These tilts will also cause pointing offsets. The relevant number is the ratio of the distance between the primary and secondary foci (conventionally called 2c), which is 6.177m, to the secondary focal length, which is 96m. The result is a pointing offset of 0.0643 arcseconds per arcsecond of tilt of the subreflector.

We will not normally introduce tilts of the subreflector other than those needed to optimize the illumination mentioned above. We will nevertheless have to take account of the effects of tilts when we are doing fast switching because we will not readjust the tilt when we go to the reference. When we are observing the source with Band X and the reference with Band Y we will set the tilt at the values for Band X. Then when we go to the reference we have to allow for this difference in tilt.

We should confirm these figures and come up with an improved version that includes signs.

REH 14th April 09