Operated for NASA by AURA

Instrument Science Report ACS 2007-08

Variation of the Distortion Solution of the WFC

Jay Anderson

jay@eeyore.rice.edu

July 2007

Abstract

During reduction of the Ultra Deep Field in 2004, it became apparent that the distortion solution constructed in 2002 did not allow images taken at different orientations to be mapped into each other with simple rectilinear transformations. In an effort to examine the time dependence of the WFC distortion, I cross-compare the many observations that have been taken of the outer calibration field of 47 Tuc. I find that the linear skew terms have changed monotonically since ACS was installed in 2002. The size of the change is appreciable: a star at the corner of the detector in late 2006 will be about 0.3 pixel off from where the original 2002-based distortion solution says it would be. These skew errors cancel out when relating data sets taken at the same epoch and same orientation, but they must be accounted for when dealing with observations taken at different roll angles. A simple correction to the 2002 solution is provided.

1. Background

Distortion is usually thought of as a non-linear mapping of the detector coordinates onto the sky, but there are also purely linear components to distortion. Linear distortion is present whenever the x and y axes either have different scales or are not perpendicular. Since the WFC is a square detector that projects onto a rhombus-like shape on the sky, it clearly suffers from a lot of linear distortion. Both the linear and non-linear distortions are treated by the distortion correction in Anderson (2002).

During the reduction of the UDF observations, it became apparent that the 2002-based distortion solution did not allow images taken at different orientations to be mapped into each other with a simple rectilinear transformation (rotation, expansion, and offset). Additional linear terms had to be added in order to make the stars and galaxies line up properly.

In an effort to determine whether this implied an error in the original distortion solution or simply a variation of the solution over time, I initiated a comprehensive study of the calibration field in 47 Tuc. This field is located $5'$ west of the cluster center, and has been visited many hundreds of times since the installation of ACS, with a variety of pointings and orientations. Such a varied data set is ideal for examination of the global terms in the distortion solution.

2. Solving for Distortion

The easiest way to solve for distortion (or to examine errors in a distortion solution) is to compare positions measured for stars in an observed frame against positions for the same stars in an absolute reference frame. This absolute frame does not have to be calibrated in terms of RA and dec, but it must be free of distortion.

In the absence of such a reference frame, we can still solve for the distortion, but it is harder. We must cross-compare many observations taken at different offsets and orientations and tease out the single solution that allows all frames to be transformed rectilinearly into each other. I will use this approach to examine the time-dependent distortion variations in the WFC. I will then use this solution to construct an accurate reference frame, so that in the future, distortion evaluation can be done the easy way.

3. Observations and reductions

The data for this study come from the many observations that have been taken of the outer calibration field in 47 Tuc. This field is located at (00:22:35,−72:04:00), about 5′ west of the center of the cluster. The field has been observed over a thousand times by the WFC at many different pointings and through many different filters. Since different filters can have slightly different distortion solutions, I will restrict this analysis to the F606W observations. There are 193 F606W observations to cross-compare, some short (∼30s) and some deep (∼350s).

For each exposure, the program `img2xym_WFC.09x10` was used to measure positions and fluxes for all the stars in the image. These positions were corrected for distortion using the best available solution (`wfc_gcf`). Both the measuring program and the distortion solution are documented in Anderson & King (2006). A collating routine was then used to generate a rough master frame for the entire field, and cross-identify the stars in the different exposures.

4. Transformations

The first step in solving for the distortion is to relate the observed frames to one another. In general, the telemetry data from the guide stars is not accurate enough to blindly map one image into another to better than 10 pixels (0.5″). The effect we are looking for is on the order of 0.1 pixel, so we must find better transformations empirically, based on the positions of stars in the two frames.

A general linear transformation from one frame (x, y) into another (u, v) has 6 parameters, and can be put into the following form:

$$\begin{pmatrix} u \\ v \end{pmatrix} = \begin{bmatrix} A & B \\ C & D \end{bmatrix} \begin{pmatrix} x - x_0 \\ y - y_0 \end{pmatrix} + \begin{pmatrix} u_0 \\ v_0 \end{pmatrix}. \tag{1}$$

The parameters A, B, C, and D are the linear terms, and x_0, y_0, u_0, and v_0 are the constant terms. It may look like there are 8 free parameters here, but we can choose the zeropoint in one frame arbitrarily, so only one of the two offsets is actually solved for. The easiest way to solve for the offsets is to adopt the centroid of the star list in the first frame as the offset for that frame, (x_0, y_0). Then, the centroid in the other frame will be the least-squares solution for the other offset, (u_0, v_0).

All that remains is to solve for the four linear terms, which can be done by least squares

(often iteratively rejecting extreme outliers). The least-squares solution for A is:

$$A = \frac{\Big[\sum(u-u_0)(x-x_0)\Big]\Big[\sum(y-y_0)^2\Big] - \Big[\sum(u-u_0)(y-y_0)\Big]\Big[\sum(x-x_0)(y-y_0)\Big]}{\Big[\sum(x-x_0)^2\Big]\Big[\sum(y-y_0)^2\Big] - \Big[\sum(x-x_0)(y-y_0)\Big]^2}, \qquad (2)$$

with similar equations for B, C, and D. With these equations, whenever we have a list of N stars with distortion-corrected positions in two frames $(x_n, y_n; u_n, v_n)$, we can construct a linear transformation to relate the positions in the two frames with each other.

The four linear terms can be decomposed into some more conventional parameters. The relative scale is given by $\mathrm{S} = \sqrt{AD-CB}$. The rotation angle $\theta = \arctan(B-C, A+D)$. The on-axis skew is $(A-D)/2$, and corresponds to a difference in scale between the x and y axes. Similarly, the off-axis skew is $(B+C)/2$, and corresponds to a non-perpendicularity between the x and y axes. If $A=D$ and $B=-C$, then there is no linear skew, and the transformation is rectilinear.

Note that as HST orbits the earth, and the earth orbits the sun, a second-order effect of velocity aberration introduces a small expansion or contraction to the nominal field-of-view for each exposure (Cox & Gilliland 2002). This effect is evaluated for each frame and stored in the `VAFACTOR` keyword for each image. This field expansion is a purely radial effect, so it will have no impact on the skew terms. When we compare two frames using the above transformations, the relative scaling cause by velocity aberration will be completely absorbed in the overall scale S in the transformation, leaving whatever other linear terms may be present unaffected. Therefore, in the following sections, we can concentrate on the linear skew terms of the solution, and defer discussion of any overall scale changes until §6.3.

5. Examining the residual distortion

Using the cross-referenced star lists and the above transformations, I cross-compared all 193 of the F606W exposures taken of the calibration field. This yielded 193×192 = 37,056 frame-to-frame comparisons. For each pair of images, I took all the stars that the two datasets had in common and found the transformation from the first frame into the second, yielding the 4 linear terms A, B, C, and D. Now, if the frames are all free of linear distortions, then these transformations will have $A = D$ and $B = -C$. To the extent that this is not the case, we have linear distortion present in one or both frames.

5.1. Initial comparison

The first step is to examine the transformations from a given exposure into each of the other exposures. The skew terms in each transformation have one contribution from the first data set and another from the comparison data set. From a single comparison, it is impossible to know how much distortion each of the two frames has since, if the frames are rotated, an on-axis skew in one frame can show up as an off-axis skew in the other frame. However, if we consider all these transformations together, the contribution of the comparison frames should average out to zero, and we will be left (on average) with the terms appropriate for the target data set.

Figure 1 shows the result of the many image-pair comparisons. Each dot in the plots on the left represents the skew observed when we compare each image against the others, as a function of observation date. There is a definite trend in the bottom plot, though it is very messy.

In the right panels, the comparisons for each exposure are averaged. The linear trend in the bottom plot is now very clear, and the top plot shows some variation. It is interesting to note that the skew terms for the 2002 observations appear to be closer to zero than for the later observations. This makes some sense, since the solution was constructed in 2002. Much of the remaining scatter comes from the fact that we are trying to "average out" the contribution of the comparison frames, each of which has its own distortion errors. Since our comparison observations are not taken uniformly with respect to orientation, everything does not always average out completely. Note that this trend is the same trend that was shown in Figure 6 of Anderson (2005).

5.2. Adjusting the linear terms in the solution

Figure 1 showed that the average skew terms present in the 47Tuc data set are different from the skew terms of the 2002 solution. Our first step in improving the solution, then is to adjust the skew terms by a constant amount. The adjustment took (xc, yc), the initial distortion-corrected position from `wfc_gcf`, and found a new distortion-corrected position (xc', yc'):

$$\begin{aligned} xc' &= xc + \beta\tilde{x} + \alpha\tilde{y} \\ yc' &= yc - \beta\tilde{y} + \alpha\tilde{x}. \end{aligned} \tag{3}$$

Here, $\tilde{x}$ and $\tilde{y}$ are the scaled positions ($\tilde{x} = [xc - 2048]/2048$) . The quantities α and β represent the adjustments to the off-axis and on-axis skew terms, respectively.

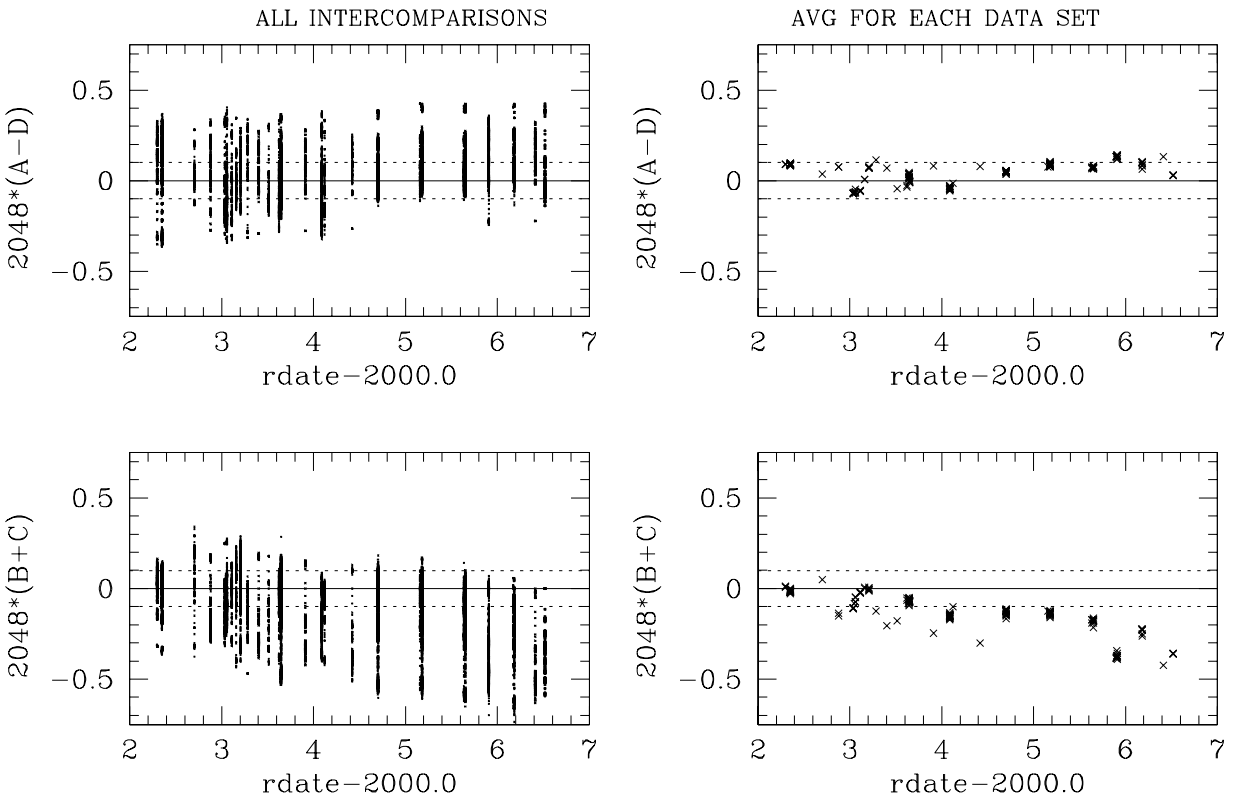

Fig. 1.— The trends in the two skew terms against time (using the original correction). Left: showing all cross-comparisons. Right: showing at the average skew term for each data set. The scaling by 2048 provides the size of the effect in pixels at the edge of the detector, where the effect is largest.

By trial and error, I found that $\alpha = 0.095$ and $\beta = 0.029$, did the best job tightening up the distribution and centering it about zero. Figure 2 shows the new average skew residuals for each exposure. Both skew terms now clearly exhibit a very linear trend with time.

The reason that a simple constant adjustment improves the relation so significantly is that each point represents the average comparsison between each exposure and all the others. Before the constant correction, the comparison exposures had systematic skew error, and there was therefore a larger error in each comparison. With the average systematic skew removed, the typical comparison is more accurate, thus the average of the many comparisons is more accurate.

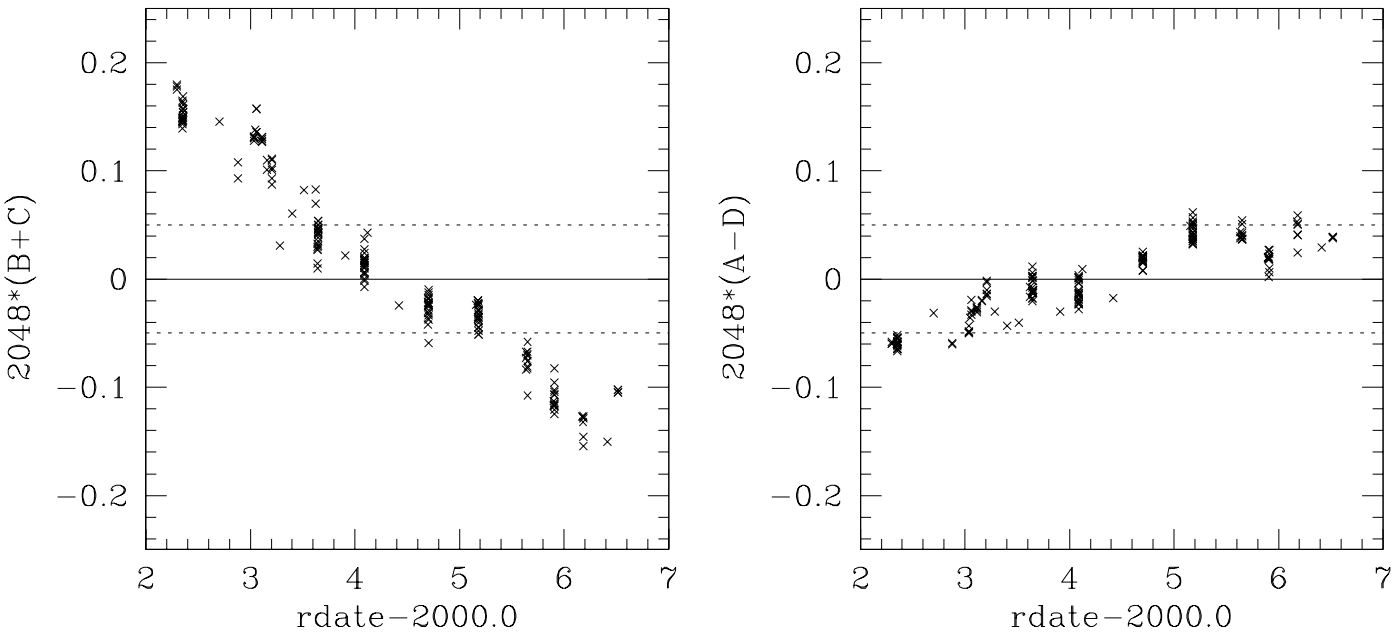

Fig. 2.— The trends in the two skew terms against time, using the original correction with a constant skew adjustment. Left: the off-axis term. Right: the on-axis term.

5.3. Making the solution time-dependent

The obvious next step was to make α and β functions of time:

$$\begin{aligned} \alpha(\texttt{rdate}) &= \alpha_0 + \alpha'(\texttt{rdate} - 2004.5)/2.5 \\ \beta(\texttt{rdate}) &= \beta_0 + \beta'(\texttt{rdate} - 2004.5)/2.5, \end{aligned} \tag{4}$$

where `rdate` is the date expressed in years. I found that by adopting α_0=0.095, β_0=0.029, α'=0.090, and β'=0.030, it both flattened out the distribution and tightened it up even more. Figure 3 shows this quite nicely.

After the linear correction, the residual skew terms still show some variation with time, though the amplitude of this variation is very small (0.02 pixel) compared with what was seen before (0.3 pixel). This remaining skew will not have a significant effect on most applications.

These time-dependent skew corrections have been independently verified in van der Marel *et al* (2007). Their Figure 2 shows that this result from 47 Tuc agrees with the skew term they measured using an astrometric reference field in the open cluster M35.

The above correction is provided in the Appendix in the form of a "wrapper" for the `wfc_gcf` routine.

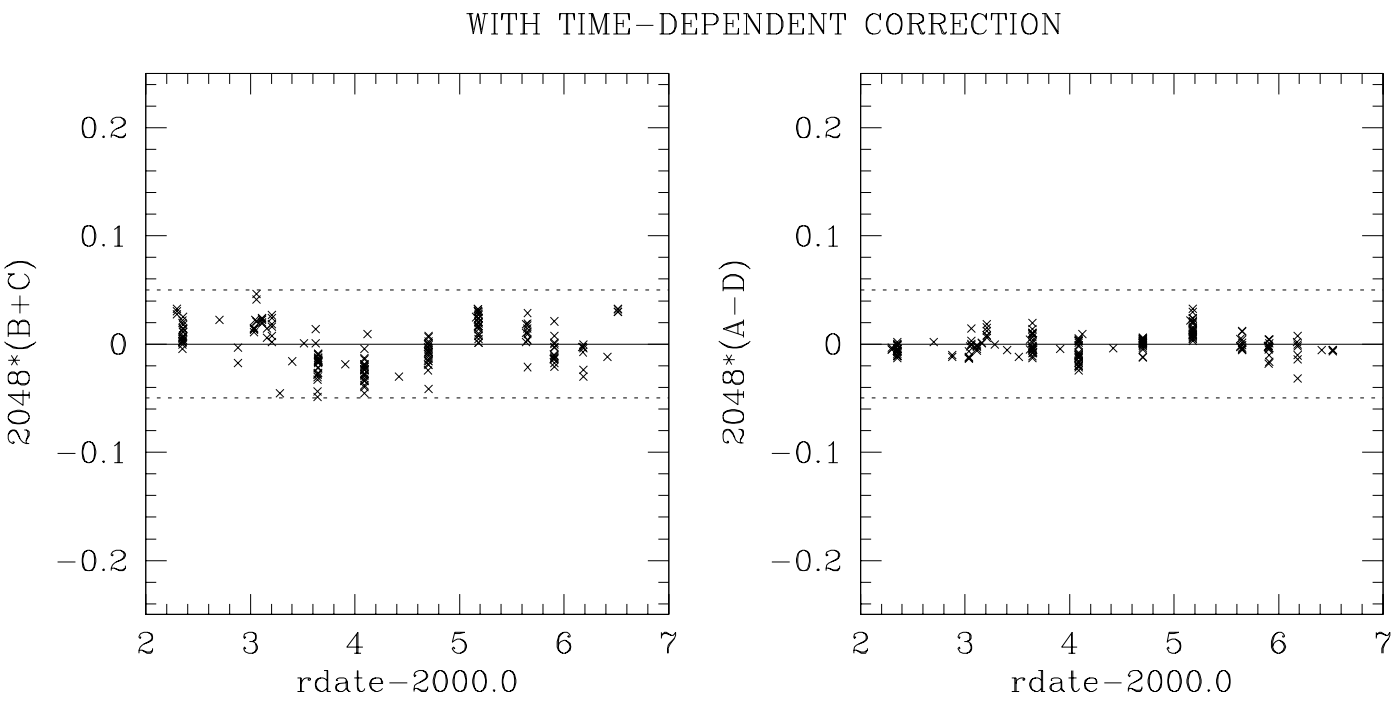

Fig. 3.— The remaining trends in the two skew terms against time, after adopting the time-dependent skew adjustment. Left: the off-axis term. Right: the on-axis term. The scaling by 2048 provides the size of the effect in pixels at the edge of the detector.

6. Making a rectified reference frame

Now that we have a distortion solution that works for any arbitrary observation date, it is finally possible to construct an accurate reference frame. To do this, I started with the rough master frame derived in § 3 and found the 4-parameter rectilinear transformation from each exposure into this frame. This allows us to rigidly transform each observation of each star into the master frame. Since each star is observed in 50 to 194 exposures, this gives us many estimates for the star's master-frame position. These estimates can be robustly averaged to improve the star's master-frame position. These improved master-frame positions, in turn, allow us to go back and improve the transformations from each exposure into the reference frame.

This procedure was iterated several times, using partial-step adjustments to facilitate convergence. After about 10 iterations, it had converged on a set of master-frame positions that did not change from iteration to iteration.

The final test is that when individual data sets are compared against this master frame, the skew terms should be zero. Figure 4 shows the remaining skew terms for each exposure's transformation into the master frame. The skew has been reduced to about 0.02

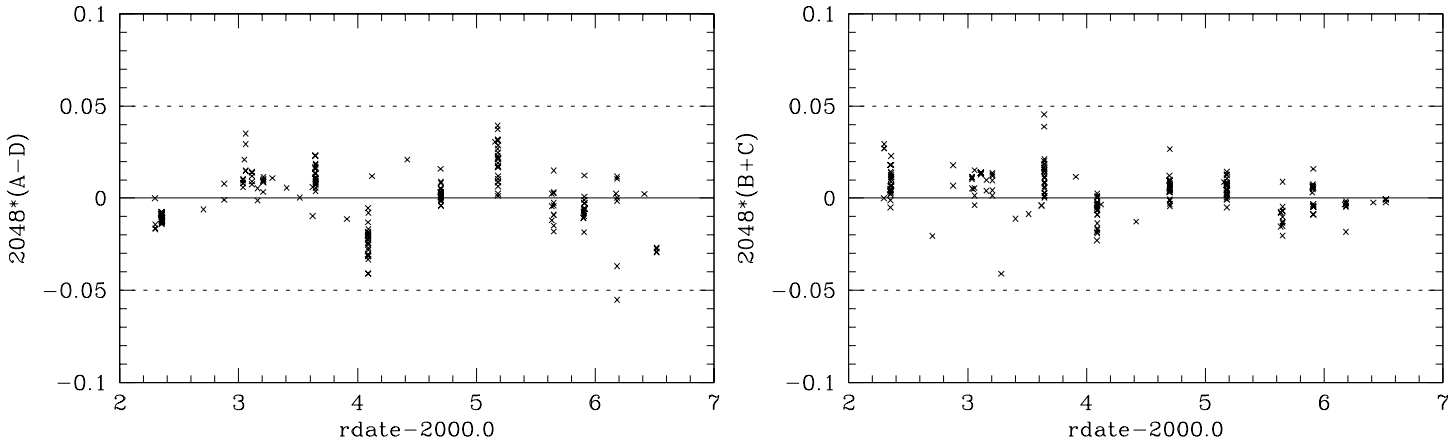

Fig. 4.— The skew-term residuals after transformation from the adjusted individual frames into the adjusted master frame.

pixel and shows no long-term trends with time.

6.1. The rigid reference frame

The reference frame was constructed with a nominal scale of 50 mas/pixel. It is nominally oriented with north up, and has a central RA/DEC of (00:22:35.0,−72:04:00) at frame coordinate (5000,5000).[1]

Positions for cluster member stars in the improved master frame can be found in the file `MEMBER.RIGID.XYM`. This star list can now be used as an absolute frame of reference, and it should be globally accurate to about 0.02 pixel (1 mas). (It should be noted, though, that the cluster stars are moving relative to each other at about 0.01 pixel/year.) I also created a stacked image of the field in the above coordinate system. All these files can be found at:

`http://www.stsci.edu/hst/acs/analysis/distortion`

It should be possible to incorporate the calibration-field observations from the other filters and make this astrometric reference list into a photometric reference list as well.

[1]A more accurate *absolute* astrometric analysis by van der Marel *et al* gives the actual scale as 49.9933 mas/pixel and finds true North to be 0.0006 degrees clockwise from the y-axis.

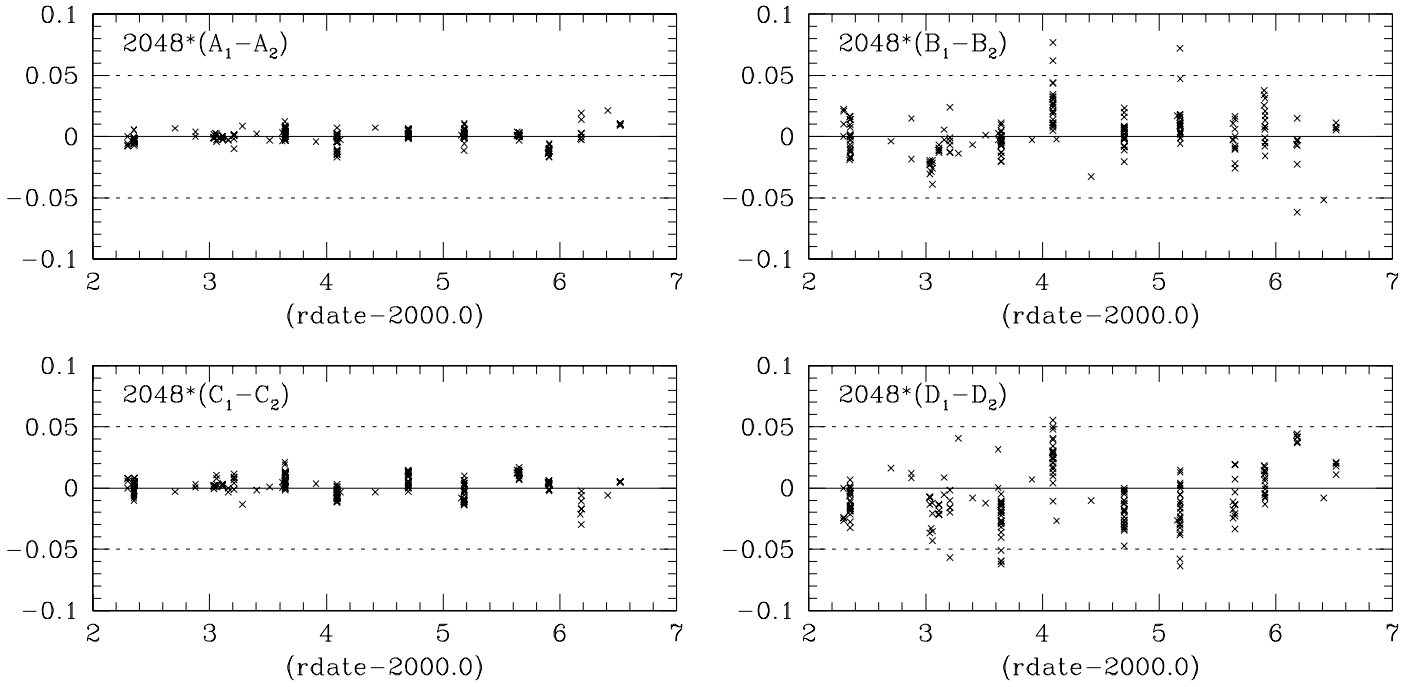

Fig. 5.— The difference in the four linear terms between the top chip (WFC1) and the bottom chip (WFC2). The scaling by 2048 provides the size of the effect in pixels at the edge of the detector.

6.2. Differences between the chips

One additional thing I looked at was whether the global skew-term adjustments are accurate for both chips, or if they just represent some average over both of them. To answer this question, I found the transformation from each chip into the master frame independently, using only the positions of stars in that chip and their positions in the master frame. We can then compare the linear terms fitted for the different chips.

Figure 5 shows the difference between the four linear terms for the two chips. Coefficients A and C, which relate to changes from the left of the chip to the right of the chip, are almost identical for WFC1 and WFC2. The other two coefficients, B and D, however, do seem to indicate some differences between the top and bottom chip, but there is no long-term trend. There also does not appear to be any indication of a shift between the two chips (not shown).

6.3. Additional findings

van der Marel *et al* (2007) perform additional comparisons between the exposures and the reference frame. In their Figure 2, they show that once velocity aberration is taken into account, the detector scale is stable to within 6 parts in 10^6, with no apparent long-term trends. This corresponds to ∼0.02 pixel across the entire field of view. Thus the secular changes in the skew linear terms are not accompanied by any overall scale change.

Since this study was done exclusively with F606W data, we have verified that the same linear skew adjustments can be used to correct data taken through F435W and F814W. Although it is unknown what is causing the skew changes, it is clearly independent of filter.

7. Conclusions

I have examined together the 193 F606W observations of the calibration field in 47 Tuc, and find that the linear skew terms of the distortion solution have been changing at the rate of about 0.04 pixel per year. I provide a simple correction for this, and show that the remaining solution is globally accurate to 0.02 pixel. There is no evidence for a change in the overall scale with time, or any dependence of skew changes on filter. The appendix provides a "wrapper" for the routine `wfc_gcf` that will include time variation in the distortion correction.

REFERENCES

Anderson, J. 2002, in the Proceedings of the 2002 HST Calibration Workshop, S. Arribas, A. Koekemoer, and B. Whitmore, eds.

Anderson, J. 2005, in the 2006 Proceedings of the 2005 Calibration Workshop, A. Koekemoer, P. Goudfrooij, and L. Dressel, eds.

Anderson, J. & King, I. R. 2006 ACS Instrument Science Report. ACS ISR 2006-01.

Cox, C. & Gilliland, R. 2002, in the Proceedings of the 2002 HST Calibration Workshop, S. Arribas, A. Koekemoer, and B. Whitmore, eds.

van der Marel, R. P., Anderson, J., Cox, C., Kozhurina-Platais, V. Lallo, M., and Nelan, E. ACS Instrument Science Report, ACS 2007-07.

APPENDIX

The following wrapper for the FORTRAN routine `wfc_gcf` will provide a time-dependent distortion solution. It should work for all 12 filters.

```
subroutine wfc_gcf_rdate(xi,yi,xo,yo,ki,fi,rdate)
implicit none
real*8 xi, yi !  raw position
real*8 xo, yo !  corrected position
integer ki    !  chip number:  2=bottom chip 1=top chip
integer fi    !  filter number:  1=F435W, etc
real rdate    !  in units of years:  2002.0 is 01jan2002

real*8 xtemp, ytemp
real alpha, beta
real fx, fy

alpha = 0.095 + 0.090*(rdate-2004.5)/2.5
beta = -0.029 - 0.030*(rdate-2004.5)/2.5

call wfc_gcf(xi,yi,xtemp,ytemp,ki,fi)

fx = (xtemp-2048)/2048
fy = (ytemp-2048)/2048

xo = xtemp + beta*fx + alpha*fy
yo = ytemp - beta*fy + alpha*fx

return
end
```