February 12, 2001,Version 3

Some Centroiding & A/D Results

Rob Olling

USNO/USRA, Washington, DC

ABSTRACT

The effects of point-spread function (PSF) mismatch has been investigated. Centroiding errors can be divided in two categories: 1) due to pixel integration in the central region where PSF gradients are large, and 2) at larger distances where the PSF gradients are smaller. A metric has been developed to quantify the effects of #2. Simplified model PSFs have been used to study the effects of pixel-integration.

Using an "analytical perturbation" to a symmetric Gaussian PSF, it is found that deviations at the percent level in its central region cause well-reproducible, systematic, centroiding errors of order several milli-pixel. This is caused by the integration over the width of the pixel, and is true even if the true PSF is symmetric. These results are consistent with recent studies reported by Anderson & King (PASP, 2000, 112, 1360). As a result, the RMS centroiding errors for undersampled PSFs are of order 1 mas, or about twice larger than FAME's design goal of $1/350^{th}$ of a pixel. Thus, it is imperative to accurately determine the PSF, whatever the color of the star. Astrometric results are independent of the color of the stars as long as there are sufficient stars to provide detailed measurements of the PSF as a function of the myriad (time-varying?) parameters. This may be a problem for blue stars. The pixel-phase bias is a strong function of PSF width. For the non-Gaussian PSFs investigated in this memo, the pixel-phase bias is all but gone for FWHMaxima larger than about 1.7 pixels.

The large-scale centroiding error can be estimated to have magnitude of $\delta x_0 = (PSF_{true}(x) - PSF_{used}(x)) \times x$, where PSF($x$) is the area-normalized PSF, and x the distance from the true centroid. The FAME requirement will be achieved when $x \times \delta PSF(x) \leq 1/350$. Thus, at large distances, the template PSF must match the actual PSF to better than 1 part per 350, or better.

An very nasty problem has been discovered in the errors introduced by discrete nature of the linear A/D conversion scheme. I estimate that A/D conversion contributes a value of $gain/18$ to the reduced χ^2_ν values, or about 8 for a V=9.001

star. Such large χ^2_ν would "normally" indicate that the centroiding is bad, with 100% likelihood. As far as I am aware, there are no statistical tools available to handle such situations adequately. Several remedies exist for this problem: 1) using a linear 14 or 15 bit A/D converter, 2) employing a >10-bit square-root A/D converter, 3) observe all the bright stars ($V \in (9, 13]$) in 2D, 12-bit linear mode, 4) invent a new? statistical tool to handle digitization drop-outs.

1. Introduction

Monet & Zacharias report a strong systematic behavior of the centroiding precision as a function of pixel-phase (ϕ_x). I find very similar behavior employing the code that I used to determine the photometric precision. This memo presents an analytical and numerical investigation into the reasons for this pixel-phase behavior. I also present some estimates for the number of required observations to reach FAME's astrometric goals.

2. The Works

To first approximation, the centroid x_0 of a function PSF is given by the value of its first moment:

$$x_0 \;=\; \frac{\int_{-\infty}^{\infty} PSF(x')x'dx'}{\int_{-\infty}^{\infty} PSF(x')dx'} \tag{1}$$

The difference between the centroid of this function and another, PSF', is given by:

$$\delta x_0 \int PSF(x')dx' \;=\; \int [PSF(x') - PSF'(x')]x'dx' \;=\; \int \delta PSF(x')x'dx \tag{2}$$

where we dropped the integration limits and assumed that both functions have the same total number of photons. As the PSF looks roughly like a sinc-function, it has a approximately Gaussian core and some weaker wiggles to the side, and possibly some asymmetries.

There are two regimes where it is easy to make large contributions to δx_0: 1^{th}) at large distances, small $\delta PSF(x')$ values result in large centroiding errors, 2^{nd}) large $\delta PSF(x')$ values at small distances also have large effects. Few photons land at large distances, so that the determination of the PSF is difficult for large x. Because of the under-sampling, the gradients in the PSF are large in the central pixel and a half: a slight mismatch between true and assumed centroid of the PSF results in large $\delta PSF(x)$ values, and hence substantial centroiding errors.

It is thus imperative to minimize the effects of cases #1) and #2). The central PSF gradients can be minimized by using a somewhat wider PSF (see below). The effects at large

distance must be controlled by the design of the instrument, by making the wings of the PSF stable and predictable. Here is why. Let us consider the case that the PSF is perfectly well established, except for a position x pixels away from the centroid. Assuming $S(x)$ photons in pixel number x and a total of S_{tot} photons, eqn. 2 leads to the following condition:

$$\begin{aligned} \delta x_0 &\sim \sum_x \delta x_0(x) \\ &= \sum_x x \frac{\delta S(x)}{S_{tot}} \sim \sum_x x \frac{\sqrt{S(x)}}{S_{tot}} \end{aligned} \tag{3}$$

Employing equation 3, I find that it produces numbers that are very close to the numerical simulations reported below. By plotting $\delta x_0(x)$ as a function of pixel position, it is evident that the pixels with the largest gradients contribute most to the centroiding error. However, their relative contribution decreases as the width of the PSF increases. This confirms the hand-waving argument based on eqn. 1 above. Further, experimentation shows that an additional Gaussian "perturbation" several FWHM from the centroid adds substantially to the centroiding error. However, a wider PSF reduces this effect. For example, when placing 15% of the total counts in a two peaks ± 3 pixels from the centroid, the 2^{nd} peaks contribute twice as much to δx_0 as does the central 85% of the flux[1].

Equation 3 could be used as a metric to define the quality of the PSF. It is well suited to quantify "distant" pixels, but since it does not take biases due to pixel-phase effects into account, equation 3 only works for PSFs with FWHM$\gtrsim$ 1.5 pixels (see below).

2.1. Point Spread Functions

My approach in rather naiver than that of Monet & Zacharias in that I do not try to mimic the true FAME PSF in too much detail. Instead I use three different models: 1) A purely Gaussian PSF, 2) a Gaussian PSF that is integrated over the width of a pixel, and 3) a low-order Gauss-Hermite PSF. It can be shown that the second PSF corresponds to the convolution of a Gaussian and a top-hat function. The disadvantage of this approach is that it does not approach reality. On the other hand, the simplicity of the above PSFs make them amenable to some analytical rigor. The area-normalized PSFs employed are parameterized as follows:

$$PSF_{1,G} = G(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}((x-x_0)/\sigma)^2} \tag{4}$$

---

[1]If the secondary peaks are at constant offset *in pixels*, their contribution decreases with PSF width. If their distance scale with PSF width, their contribution to δx_0 is independent of FWHM.

$$PSF_{2,TH*G}(x) = \int_{x-\frac{1}{2}}^{x+\frac{1}{2}} G(x')TH(x')dx' \quad (5)$$

$$= \frac{1}{2}\left(\mathrm{errf}(\frac{1}{\sigma\sqrt{2}}(\mathrm{x}+\mathrm{x}_0+\tfrac{1}{2})) - \mathrm{errf}(\frac{1}{\sigma\sqrt{2}}(\mathrm{x}+\mathrm{x}_0-\tfrac{1}{2}))\right) \quad (6)$$

$$TH(x) = \mathrm{Heaviside}(x+\tfrac{1}{2}) - \mathrm{Heaviside}(x-\tfrac{1}{2}) \quad (7)$$

$$PSF_{3,GH} = \frac{2}{\sigma\sqrt{\pi}(2\sqrt{2}+h_4\sqrt{3})}e^{-\frac{1}{2}y^2}\sum_i h_i H_i(x) \quad (8)$$

$$H_3(x) = \frac{y}{\sqrt{3}}(2y^2-3) \quad (9)$$

$$H_4(x) = \frac{1}{24}(3+4y^2(y^2-3)) \quad (10)$$

where σ is the dispersion[2] of the Gaussian and erf(x) is the error function[3], and $y = (x - x_0)/\sigma$. The top-hat function is defined as: TH(x) $= 0, \frac{1}{2}, 1$ for $|x| > \frac{1}{2}, |x| = \frac{1}{2}$ and $|x| < \frac{1}{2}$, respectively. Obviously, neither $PSF_{2,TH*G}$ nor $PSF_{3,GH}$ can be equal to some Gaussian function ($\tilde{G}$) with some other values for the peak and dispersion. Figure 1 shows that the maximum difference between $\tilde{G}(x)$ and $PSF_{2,TH*G}(x)$ is of order 1%. In the left-hand panel, the FWHM of the Gaussian is 1.05 pixels, whereas the pixel-integrated PSF is about 20% wider. For the right-hand panel, I used FWHM$_G$=1.42 pixels and FWHM$_{2.TH*G}$=1.56 pixels. The widths of the functions displayed in these figures were chosen to more-or-less match the range of PSF widths delivered by FAME. Only a small range of pixel values is displayed to emphasize the residuals. Note that these curves are not pixelized.

2.2. Centroiding

The centroiding procedure works as follows: for each pixel-phase[4] tried (30 in total), 50 random[5] stellar images[6] are generated. Each model is then digitized using a 12-bit digitization stage. Three gain settings were used: g1=650k/4096=158.7 for V $\in$ [9,10), g2=g1/2.512^2=25.1 for V $\in$ [10,11) and g3=g2/2.512^2=3.98 for V $\in$ [11,18]. As it turns

---

[2]For a Gaussian function, the full width at half maximum is given by: $FWHM_G \sim 2.3548\,\sigma$

[3]$\mathrm{errf}(x) = 2/\sqrt{\pi}\int_0^x \exp -t^2 dt$

[4]Following Anderson & King (2000), the pixel-phase is defined as: $\phi_x \equiv x - \mathrm{int}(\mathrm{x}+0.5)$. As a result, ϕ_x is zero when the PSF falls exactly on the center of a pixel. Negative ϕ_x means that the PSF is centered to the left of the pixel center. The PSF is centered exactly between two pixels for $\phi_x = \pm 0$.

[5]Poissonian shot noise and 7 e$^-$read noise added

[6]The peak of the Gaussian is given by: $P = 650k \times (FWHM/1.414) \times 2.512^{(9-V)}e^-$. This sums up to 950k electrons for a V=9 star.

out, digitization does not greatly influence the astrometric results. However, the statistical properties (reduced χ^2) are strongly influenced by digitization (see § 3 below). This fact will be of crucial importance for any decision tree that is to discriminate between single and multiple systems.

The astrometric "measurement" process consists of fitting a Gaussian function to the noisy, digitized profiles (using the Levenberg-Marquardt Method). Then, at each magnitude, the RMS centroiding error is defined as the second moment of the distribution of true minus fitted centroids, for all 30x50=1500 models. Note that the standard non-linear least squares techniques employed assumes that the errors are Gaussian. Thus, the individual errors returned by the fitting algorithm may be biased as the true errors are a combination of Poissonian shot-noise and Gaussian read-noise. These differences are subtle ($\lesssim 20\%$ for the current experiment) but important for any decision tree that will discriminate between single and multiple systems.

2.3. Centroiding Results

The astrometric results for the four PSFs described in §§ 2.1 are summarized in figure 2. In each panel, four curves are presented. The drawn lines are the results of numerical simulations described above, The dotted lines are photon-statistics extrapolations from V=9. As reported in my photometry memo, the astrometric precisions in the photon-noise limit are astonishingly small if the template PSF that is used to fit the data is exactly equal to true PSF that was used to generate the data (black lines). Also, the Gaussian-Gaussian case follows the photon-statistics predictions well. However, this is not the case when a Gaussian template function is used to fit the non-Gaussian $PSF_{2,TP*G}$ (red lines): stars brighter than V=12 show essentially the same centroiding precision. Note that this leveling off is significantly stronger for narrow PSFs (left panel) than for better sampled cases (right panel)[7]. Consequences, causes and cures for this bias will be discussed below.

When fitting a Gaussian PSF with a Gaussian template, the limiting centroiding precision is 0.106 (0.136) mas for a FWHM of 1.055 (1.414) pixels (for a 9^{th} mag star): when the template and PSF match, very good precision can be achieved, even when only one pixel spans the width of the PSF. In fact, figure 2 indicates that at V=9, the attainable precision exceeds the FAME requirement of $1/350^{th}$ of a pixel (0.59 mas) by a factor of about five, and that photon-statistics pushes the astrometric precision above the design goal around V=12.

---

[7] For the $\mathrm{PSF}_{2,TH*G}$ case I find: $\delta x_{0,bias} \sim -71.31 - 91.95 \times FWHM_G + 29.61 \times FWHM_G^2$ or $\mid \delta x_{0,MAX} \mid \sim$ 12, 7, 2.5, 0.7, 0.05 milli-pixel for FWHM_G=0.94, 1.05, 1.23, 1.41, 1.64 pixels, respectively.

2.4. The Bias

Figure 2 also shows that the centroiding precisions for the $\mathrm{PSF}_{2,TH*G}$ case are factors 9.7 and 1.27 worse than for the case that the true PSF is Gaussian for the FWHM presented. Given the relation between the centroid bias and the PSF width (see footnote 7), I find that the FAME requirement of 0.59 mas centroiding can be achieved if the FWHM is larger or equal than 1.23 pixels[8], if the true PSF is given by $\mathrm{PSF}_{2,TH*G}$.

Above we have seen that a 1% mismatch between the actual and template PSF results in a centroid bias that is 9.7 times larger than what can be achieved when PSF and template match. A simple-minded linear scaling implies that photon-statistics centroiding precision can be reached if PSF and template match to within 0.1%. Extrapolating linearly, FAME statistics would be achieved in case the template-mismatch would be less than 0.5%.

2.4.1. Other Consequences Of The Bias

In figure 4, I present plots of the centroiding precision as a function of pixel-phase in the top panels. The thin black and red lines show the results of Gaussian fits to $\mathrm{PSF}_{1,G}$ and $\mathrm{PSF}_{2,TH*G}$, respectively. While the $\mathrm{PSF}_{1,G}$ errors hover around zero, the $\mathrm{PSF}_{2,TH*G}$ results show a strong pixel-phase dependent behavior: the "bias." As a result, the rms centroiding error is very large when averaged over all ϕ_x, and is in fact dominated by the bias. At fainter magnitudes, the magnitude of the bias is overwhelmed by the variance introduced by photon-statistics. This bias does not result from shot noise because the no-noise results (continuous thick red line) follows the with-noise results (jagged thin red line). The bias is also not caused by the digitization of the signal since the no-noise case show the same bias whether it is digitized (left-hand panels) or not (not shown).

We conclude that the *pixel-phase bias is caused by PSF-template mismatch.* In order to achieve the FAME centroiding requirement it is of crucial importance that FAME's actual PSF can be measured to 1 part in 500 or better. If the correct PSF is used to measure the stellar "images," no bias is present. As we have seen, this is so for the Gaussian-Gaussian case. The same results are obtained when $\mathrm{PSF}_{2,TH*G}$ stellar images are fitted with a $\mathrm{PSF}_{2,TH*G}$ point-spread function (cf. eqn. 6). These results are not shown. Makarov reports (private communications) similar findings for a PSF that resembles the worst Lockheed PSF at $\lambda = 400$ nm.

An alternative to fitting the data with a derived true PSF could be to fit Gaussians

[8]The precision will decrease again if FWHM is made too large, as $\delta x_0 \propto FWHM^{3/2}$ in the photon-statistics limit: for this example, $\mathrm{FWHM}_G \lesssim 5.8$ pixels.

to the observed stellar images and derive the bias empirically and iteratively correct for it. Since the centroiding bias is such a strong function of pixel-phase, it should be possible to determine it. A complication is that there are two pixel-phases that have exactly the same centroiding bias. An arbitrary non-Gaussian PSF may even be multi-valued.

At first it seems that this multi-valuedness poses a problem in determining the true pixel-phase. However, when "centroiding" we not only estimate the position, but also the peak and width of the stellar image. The fits for these parameters are displayed in the middle two panels of figure 4. From these panels it is clear that the bias-ϕ_x degeneracy will be broken if peak and/or width measurements are included.

2.4.2. Gauss-Hermite Polynomials

I have experimented with other non-Gaussian PSFs from the Gauss-Hermite (GH) family (cf. eqn. 8), which differ in normalization from those described by Makarov (FTM2001-2) but conform the the practice of workers in the stellar kinematics community. A nice illustration of how GH point-spread function look like can bee found in Binney & Merrifield (Galactic Astronomy, 1998, p 702). In essence, positive (negative) h_3 parameters produce an asymmetry on the right (left) side of the centroid. Positive h_4's result in PSFs that are more "wingy," while negative h_4's produce PSFs that are less wingy than Gaussians. The pixel-values of GH PSFs with non-zero $h_{i=3,4}$ terms differ from Gaussians at the $\sim h_i$ level.

Experimentation with these GH functions shows that, for small h_i, the pixel-phase bias can be written as a low-order Fourier series: $O-C \sim a_0(h_3, FWHM)+a_c(FWHM)h_3\cos(\phi_x\pi)-a_s(FWHM)h_4\sin(2\phi_x\pi)$ with $a_0 \sim 0.0$, $a_c \sim 1.8$ and $a_s \sim 1.5$ for FWHM$_G$=1.055. For under-sampled images, the systematic bias due to the even h_4 term is zero when averaged over pixel-phase, with as a net effect an increase of the scatter. Due to its $\cos(\phi)$ dependence, the bias due to a positive (negative) h_3 term will not average to zero, but will be systematically positive (negative). Note that the coefficients describing $O-C$ all depend on the width of the function, and that the constant term has an additional h_3 dependence (all in a non-transcendental manner).

The FWHM-dependence of the centroiding bias is illustrated in figure 3. Here I summarize the results for Gaussian fits to a $PSF_{1,G}$, a $PSF_{2,TH*G}$ and a series of GH PSF's. This figure clearly indicates that the bias decreases with PSF width. Note that none of the tested PSFs have second humps, so that this experiment only describes the "close-in" centroiding effects.

Thus, the behavior of the bias for a GH PSF is very different from that of the pixel-convolved Gaussian used above. In the latter case, the non-Gaussianness of PSF$_{2,TH*G}$

decreases with increasing FWHM because the pixel-integration effect is less important. This also the case for the symmetric (h_4) part of the GH PSF, but not for the asymmetric term: the h_3 bias just increases as the PSF grows fatter. **These results also imply that, since the bias depends on the shape and width of the PSF, it will be as hard to determine the bias as it is to determine the PSF.**

For FWHM$_G$=1.055 (1.414) pixels, in order to just meet the FAME requirement, $\mid h_3 \mid$ has to be smaller than 0.004 (0.016) and $\mid h_4 \mid \leq 0.002(0.008)$, *if a Gaussian is used to fit those GH PSFs.* That is to say, PSF knowledge of order several tenths of percent is required, as was the case for the PSF$_{2,TH*G}$ discussed above. Again, if a GH template is used to fit GH stellar images, all biases and systematic effects disappear.

2.5. Bias Explained

The PSFs presented in figure 1 hold the clue for the true cause of the bias. Consider the left-hand panel of this figure. The cyan curve is the difference between the true PSF$_{2,TH*G}$ and a Gaussian that fits this PSF$_2$ best. This $O-C$ curve is the basis for explaining the behavior of the pixel-phase systematics. To see this, consider the pixel-phase $\phi_x = 0$ situation. In this case, the central pixel (-$\frac{1}{2}$, $\frac{1}{2}$) "covers" the $O-C$ curve symmetrically. The first right-side pixel ($\frac{1}{2}$, $\frac{3}{2}$) contains as much flux as the first left-side pixel ($\frac{3}{2}$, $-\frac{1}{2}$), et cetera. Thus, in this case, the centroid would be identical zero. Similar arguments for the case that the star is located exactly between pixels also lead to zero bias for $\phi = \pm\frac{1}{2}$.

Taking $\phi_x = 0.2$, the left-hand plot of figure 1 shows that the pixels centered -1.8, -0.8 and +0.2 have $O-C \sim 0$, while the +1.2 pixel has positive flux of ~ 0.005. The result is a positive centroiding bias, as observed in left-hand panel of figure 2. For the wider PSF in the right-hand panel, the asymmetries are much smaller, and when one slides a pixel array across the $O-C$ curve, no astrometric biases are evident: consistent with figure 2.

3. Nasty Digitization Effects

Since the PSF$_{2,TH*G}$ and the GH PSFs are definitely non-Gaussian, it is expected that a Gaussian fit to non-Gaussian PSFs should result in a poor fit, statistically speaking. As a function of pixel-phase, magnitude and PSF model, I calculated the reduced χ^2 values:

$$\chi^2_\nu \;=\; \frac{1}{\nu}\sum_{i=1}^{i=13}\left((O_i - C_i)/\sqrt{(C_i + r^2)}\right)^2 \qquad (11)$$

with O_i and C_i the observed and computed counts, r the read-noise, and ν the number of degrees of freedom ($\nu = 13 - 4 = 9$). The results are plotted in the bottom panels of

figure 2 for the $\mathrm{PSF}_{2,TH*G}$ (thin red curves with error bars) and PSF_G (thin black lines with error bars) cases. In these panels we see that indeed the reduced χ^2 values differ from the expected value of unity, and that χ^2_ν is largest in the region with the largest bias: $\chi^2_\nu(\phi_x) \sim 60$. Thus, template-PSF mismatch can be identified in regions where the reduced χ^2 values are large. On the other hand, χ^2_ν can reach 15 for the case that a Gaussian PSF is fitted with a Gaussian template. That is to say, a reduced χ^2 value of 15 is *not* an indicator of template-PSF mismatch.

The arguments presented above indicate that even if we somehow know the true PSF, it is still possible to obtain a reduced χ^2 of about 15. In "normal" circumstances, such values would indicate that it is highly unlikely ($Probability < 1.310^{-38}$) that the stellar image resembles the template PSF, possibly due to the presence of a cosmic-ray, an electron trap or a visual binary. Above we saw that this need not be the case. These results pose severe problems for the interpretation of FAME astrometry since "normal" statistical tools can not be used to compare data and fits.

Which χ^2_ν values are "good" is determined by three factors: 1) the shape of the true PSF, 2) the brightness of the star and 3) the pixel-phase. This effect is illustrated in figure 5 (the FWHM=1.414 equivalent of figure 4), where the shape of the Gaussian-Gaussian $\chi^2_\nu(\phi_x)$ curve has changed substantially, while the maximum amplitude is still very large (10; $P \sim 210^{-15}$). Experimentation with GH PSFs yields similar results. Thus, in order to employ reduced χ^2 in the normal fashion, a criterion has to be developed that tells us whether a given χ^2_ν is "good" or "bad." Again, such criterion will depend on m_V, PSF and ϕ_x, not only on a comparison between data and model (and errors). To my knowledge, this state of affairs in not handled properly by commonly used statistical tools.

3.1. χ^2 Explained

In fact, it is quite odd that the Gaussian-Gaussian fitting yields such horrendous statistics. It turns out that this is due to the way the pixel data is digitized. The reduced χ^2 values for no-noise Gaussian-Gaussian cases (thick black lines) are: 15, 6, 1, 0 for 12, 13, 14, 15 bit linear digitization, respectively. Since there is no template-mismatch nor any noise, in the *G-G* case, the reduced χ^2_ν *should be zero*. The observed large positive values must be attributed to digitization effects. For the no-noise digitized $\mathrm{PSF}_{2,TH*G}$ case, χ^2_ν decreases when more bits are used (50 at 12, 40 at $\geq$ 13 bits). Thus, largest contribution to reduced χ^2 values for $PSF_{2,TH*G}$ is due to template-mismatch, not digitization.

On the up-side, a 14 or 15-bit linear A/D converter, would allow us to use standard statistical tools to analyze the FAME data. On the down-side, 15 bits/pixel require at most 25% larger data rate and requires an A/D converter that may or may not be available. If

Table 1: Digitization effects illustrated for a Gaussian function with FWHM=1.414 pixels and a peak of 650k electrons. The Gaussian is centered at $\phi_x = 0.4$. Position, flux and $\sqrt{\text{flux}}$ are listed in the first three columns. The digitized quantities in the subsequent columns. First for 12-bit linear digitization with $gain_L$ 158.7 and followed by 12-bit sqrt digitization with $gain_S$ 0.1968. The equivalent linear digitization step for the sqrt-scheme digitization scheme is tabulated in the last column.

x	F_o	$\sqrt{F_o}$	ADU_L	$F_{d,L}$	$N_{d,L}$	$\chi^2_{d,L}$	ADU_S	$F_{d,S}$	$N_{d,S}$	$\chi^2_{d,S}$	$gain_S^L$
-6.0	0.000	0.000	0	0.000	7.0	0.00	0	0.000	7.0	0.0	0.04
-5.0	0.000	0.000	0	0.000	7.0	0.00	0	0.000	7.0	0.0	0.04
-4.0	1.4E-6	0.001	0	0.000	7.0	0.00	0	0.000	7.0	0.0	0.04
-3.0	0.0692	0.263	0	0.000	7.0	0.00	1	0.0387	7.0	0.0	0.11
-2.0	218.1	14.77	1	158.7	14.4	1.88	75	217.9	16.3	0.0	5.84
-1.0	42.72k	206.7	269	42.69k	206.7	0.00	1050	42.71k	206.8	0.0	81.37
0.0	520.5k	721.5	3280	520.5k	721.5	0.00	3665	520.4k	721.4	0.0	283.9
1.0	394.2k	627.8	2484	394.2k	627.9	0.00	3190	394.3k	627.9	0.0	247.1
2.0	18.57k	136.3	117	18.57k	136.4	0.00	692	18.55k	136.4	0.0	53.64
3.0	54.37	7.374	0	0.000	7.0	6.70	37	53.04	10.1	0.0	2.90
4.0	0.0099	0.099	0	0.000	7.0	0.00	1	0.0387	7.0	0.0	0.11
5.0	1.1E-7	3.E-4	0	0.000	7.0	0.00	0	0.000	7.0	0.0	0.04
6.0	0.000	0.000	0	0.000	7.0	0.00	0	0.000	7.0	0.0	0.04

at this point it is still possible to consider an alternative A/D converter, it is worthwhile to consider a non-linear A/D conversion scheme. In the right-hand panel of figure 4, I present the pixel-phase results for the case of 12-bit *square-root* digitization. In such a sqrt-digitization scheme, the square-root of the data value rather than the data value itself is digitized. Such a scheme allows for a full recovery of the astrometric precision (not shown). Furthermore, the reduced χ^2 values are virtually identical to the case of $>$ 16-bit linear digitization. The reason for this is as follows. In any digitization scheme, pixels with less electrons than half the gain will be rounded to zero. A pixel with readout noise r (=7 e^{-1}) and observed counts *just* below $gain/2$ will have "O-C" of $\frac{1}{2}gain$ and will thus contribute to reduced χ^2 a value of:

$$\frac{1}{\nu}\left(\frac{(O-C)}{\sqrt{N^2+r^2}}\right)^2 \sim \frac{gain^2}{4\nu(N^2+r^2)} \tag{12}$$

where the noise value N could equal $gain/2$ or 0.0, depending on whether is is calculated from the computed model or the data values, respectively. If the read-noise is small and model noise values are used, eqn. 12 reduces to: $\chi^2_\nu = gain/(2\nu)$. Which, and how many pixels satisfy this "half-way" criterion depends on the detailed shape of the PSF and the pixel-phase. Since it is the digitization that matters, narrow and broad PSFs will have similar ranges of digitization-induced reduced χ^2 values.

Consider the 12-bit linear digitization of a V=0 star with FWHM=1.414, P=650k e^-, and a digitization step of $gain_L = 650k/4096 = 158.7$ e^-. The noise values for "computed

noise" and "observed noise" are: $N_C = \sqrt{gain_L/2 + 7^2} = 14.4$ and $N_O = 7$. Thus, the respective reduced χ^2 values are: $\chi^2_{\nu;C} = 1/9 \times (158.7/2/14.4)^2 = 3.3$ and $\chi^2_{\nu;O} = 1/9 \times (158.7/2/\ 7.0)^2 = 14.3$. In my numerical simulations I estimate the noise based on the observed count rates, not on computed model values, and I thus expect to see reduced χ^2 values of order 15. This expectation is borne out by the results presented in figure 5. In table 1, I present, as an example, the relevant quantities as a function of pixel position. Columns 2 and 3 are the computed flux (F_c) and the square-root thereof. Columns 4-7 tabulate the 12-bit linear digitized values ($ADU_L = \mathrm{ROUND}(F_c/gain_L)$), the estimated original flux ($F_{d,L} = ADU_L * gain_L$), the estimated noise from the digitized counts ($N_{d,L} = \sqrt{F_{d,L} + 7^2}$), and $\chi^2_{\nu;d,L}$ ($= 1/\nu \times ((F_c - F_{d,L})/N_{d,L})^2$). The equivalent quantities for square-root digitization are tabulated in the subsequent columns.

Inspecting table 1, it is clear why square-root digitization does not produce large reduced χ^2 values. With a $gain_S$ of 0.1968, the low count-rate regime is digitized with steps of $gain_S^2 \sim 0.039$, so that the maximum contribution per pixel to χ^2_ν is very small indeed (0.02/9). The astrometric precision is actually somewhat better when using 12-bit sqrt digitization because the *linear* flux step that is made at *all* sqrt digitization levels is always about 37% of the noise (cf. the last and second-before last columns of table 1). For linear digitization, the noise is comparably well digitized at flux levels $\gtrsim (3 \times gain_L)^2 \sim 220k$, for the bright stars. Only two points listed in table 1 satisfy this criterion.

Finally, since sqrt-digitization more faithfully follows the data at low flux levels, more pixels will have non-zero ADU values than in the linear case, so that more data points are available so that better fits can be obtained. Because the sqrt-digitization scheme uses steps that are related to the noise, it does a better job at characterizing the data *and* the noise. As a result, the digitization errors are small so that standard statistical tools can be used to analyze sqrt-digitized FAME data.

Because the digitization error is of order $gain/2$, linear digitization works best for faint stars that employ lower gain settings. For FAME, the second and third gain settings are 2.512^2 and 2.512^4 times smaller than the first ($g_2 \sim 25, g_3 \sim 4$ and $P_{2,MAX} \sim 100k$ and $P_{3,MAX} \sim 16.4k$). This leads to χ^2_ν biases of 12/9=1.33 and 2/9=0.222 per half-way pixel, respectively. This is consistent with my numerical simulations. In 2D FAME postage stamps, peak fluxes are of order 100k electrons so that $g_2 = 25$. Thus, the digitization effects on the image statistics will be 12/9 rather than 79/9 χ^2_ν units per half-way pixel. On the other hand, in a 2D image, more pixels will have a chance to collect $\sim g_2/2$ electrons.

3.2. Solutions to the digitization problem

Several solutions are suggested for the digitization-induced χ^2_ν problem:

1. Use a 14 or 15-bit linear A/D converter
   Requires a new A/D; 25% more data will be generated
2. Use a 10 or 12-bit sqrt A/D converter **(preferred from data-analysis perspective)**
   Requires a new A/D
3. Observe all 2M stars brighter than V=13 in 2D mode.
   This is disastrous for the data rate
4. Observe stars brighter than V=13 in 2D mode, use 12-bit linear digitization. To limit the data rate, perform marginalization of the digitized 2D image and use a 12-bit sqrt compression of the marginalized data.
   (best results for smallest change wrt current design)
   This is a complex scheme
5. Use a 15-bit linear followed by a 12-bit sqrt compression
   Requires: new A/D, and compression software
6. Develop a fitting scheme that utilizes the digitization drop-outs below the $gain/2$ level
   (least desirable)
   Requires specialized software

The first five fixes are rather unattractive as they require a substantial change of the current design. However, doing nothing and hoping that a ground-based software fix will be found might be worse: this issue needs immediate attention.

In the maximum-likelihood parameter estimation method, the errors are determined from the model rather than the observed values. In this case the digitization errors are mitigated, but are still significant for the bright stars. One could construe the following fitting scheme: 1) a model is estimated *without* the 1^{th} "zero-count" pixel on either side of the maximum, (this will change the width of the PSF and move the true centroid) 2) the "zero-count" pixels are checked for digitization-dropout effects. This is a dufficult process as it is not obvious how the "zero-count" pixels can be determined and how exactly the PSF width and location will be affected by the digitization dropouts. As explaned above and illustrated in figs. 4 and 5, *these effects depend sensitively on the PSF.*

4. PSF Construction

In the Tycho-2 catalog, there are about 8,000 to 12,000 stars per d(B-V)=0.1 with V=[8.5,11.5]. It is those stars that will be used to determine the empirical PSFs. Fainter

stars can be used as well, but they increase the "information" only by a factor of three or so (since the number of photons received per magnitude bin is approximately constant). So let us assume that we 20,000 stars –half of which are binaries– per color bin. Each star will yield 5 pixels of useful information, so we get 20k/2*5=50k data points. Each CCD will observe these stars 1000/24 times, so we have 2 million observations per star-group per CCD.

Now suppose that we have to live with the worst case, that a PSF needs to be determined for each of the 2048 rows per CCD, that yields 1k observations per row. Also suppose that the PSF has to be determined 10 times during the 2.5 years (= every 3 months = twice between eclipse seasons). The result is 100 "observations" per row. Since the apparent PSFs will vary with in-scan smearing, we need possibly 5-10 in-scan smeared PSFs, leading to 10-20 observations per PSF.

One could use these 10-20 observations to fit a GH function to the data to determine 10-20 constants per PSF, or 1 observation per GH term: a rather underwhelming number. Alternatively, one might split these 10-20 observations between between the 5 pixels that make up the PSF, or 2-4 independent samples per pixel per PSF. We are in trouble if the PSF needs to be determined as a function of sub-pixel location. A wider PSF would help in that it reduces the required number of GH-terms/sub-pixel-locations.

This problem can be solved in several ways: 1) use the fainter stars as well to yield 3 times as much "information," 2) live with an intrinsically worse astrometric accuracy for these rare (blue) stars, 3) make sure that the PSF is so well behaved that redder stars can be used as a template for the rare blue stars 3a) have better optics, 3b) de-focus to $FWHM \gtrsim 1.4$ pixels, 3c) have more in-scan smearing, 3d) cut out "blue" light. These options are listed more-or-less in order of desirability. As discussed earlier and shown in figure 3, we could live with a bit of additional "smearing" since even with FWHM=4.9, we have δx_0=0.459 mas ($1/449^{th}$ pixel) at V=9 and photon statistics beyond to δx_0=8.4 mas at V=15.

5. Conclusions

Centroiding errors arise from two regimes: 1) pixels close to the maximum where the PSF gradient is large, and 2) distant pixels that sample the wings/humps of the PSF. The first source of error can be significantly reduced by employing a wider, better sampled, PSF. The contribution due to distant pixels can be reduced by "better optics, " of fewer "bad photons," or both.

Linear digitization on a 12-bit scale of the FAME data produces unanticipated, very

large reduced χ^2 values. This problem can be solved in several ways. A ground-based statistical correction in software is anticipated to be very complicated. If at all possible, an in-flight software and/or hardware solution is preferred.

This investigation has benefitted substantially from conversations between the author and Valeri Makarov, Arsen Hajian, Marc Murison. George Kaplan, Norbert Zacharias and Dave Monet.

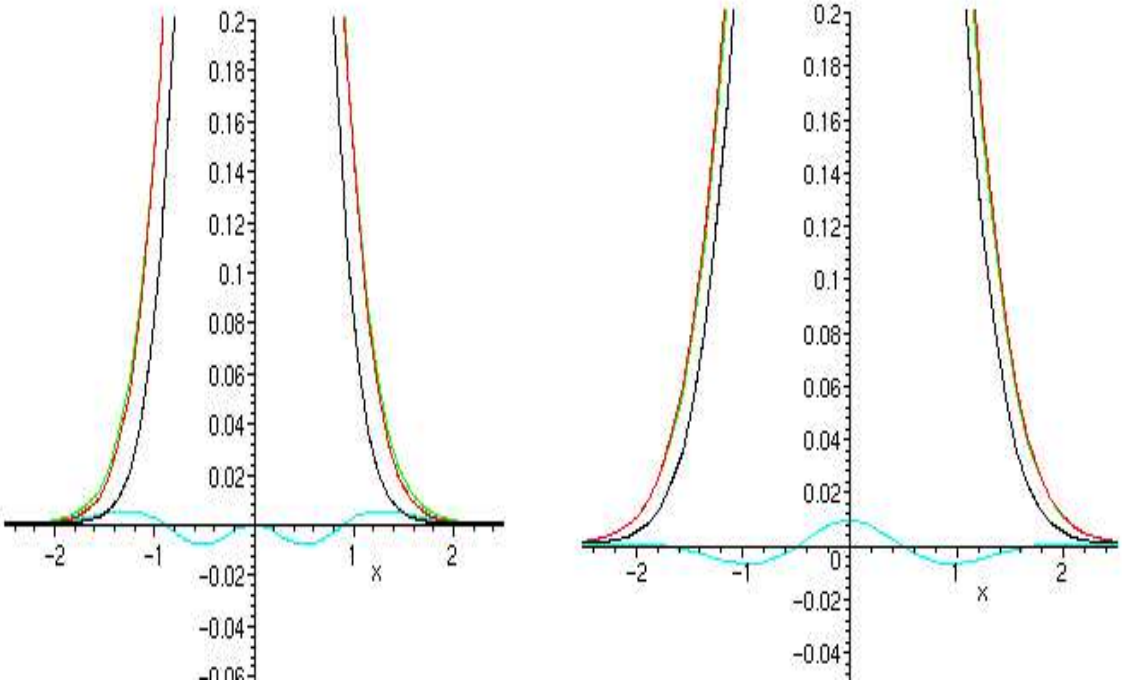

Fig. 1.— The template PSF (a true Gaussian [$G(x)$, black]), the true $PSF_{2,TH*G}$ (red) and the Gaussian ($\tilde{G}(x)$) that matches the true PSF as close as possible [$\tilde{G}(x)$, green]. I also display the difference between $\tilde{G}(x)$ and $PSF_{2,TH*G}$ (cyan). The left-hand panel is for $FWHM_{\tilde{G}} \sim FWHM_G \times 1.21 \sim 1.05 \times 1.21 \sim 1.27$ pixels, the right-hand panel for $FWHM_{\tilde{G}} \sim FWHM_G \times 1.1 \sim 1.42 \times 1.1 \sim 1.56$ pixels. The residuals in both panels are of order 1%, independent of the width of the PSF (I've tested up to $FWHM_G$=4 pixels). That is to say, there is **NO** Gaussian that will perfectly well fit $PSF_{2,TH*G}$.

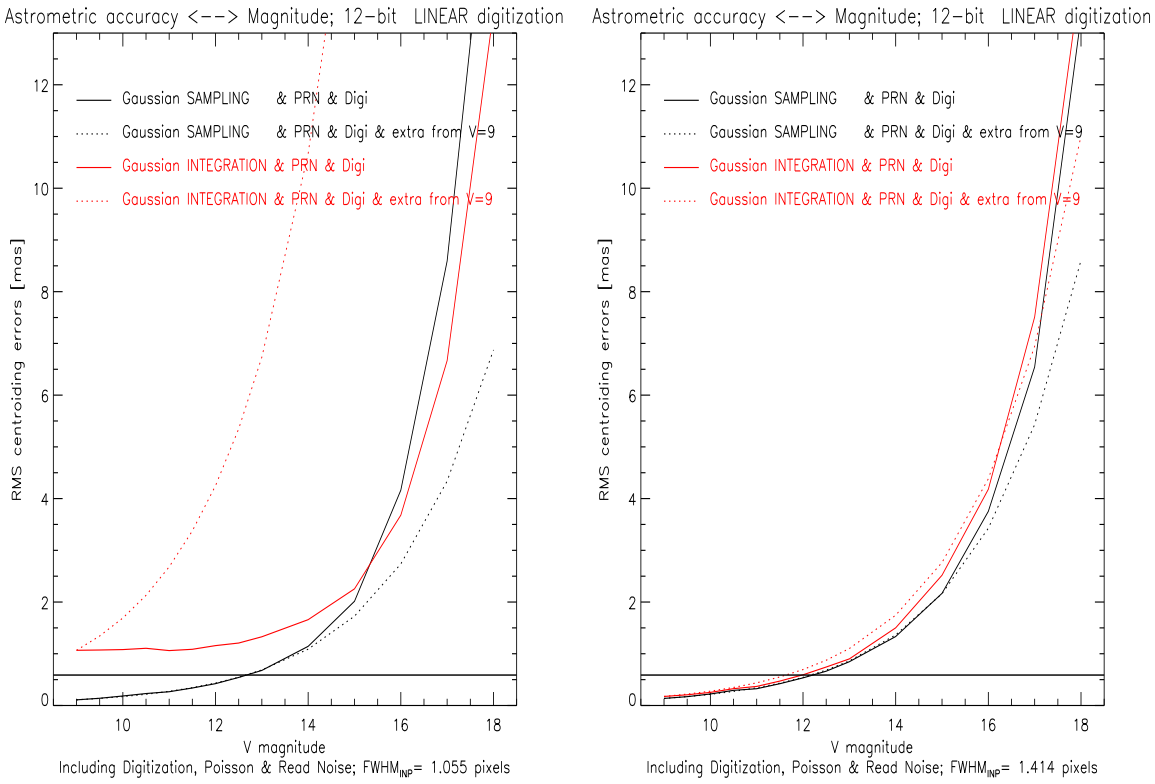

Fig. 2.— Astrometric precisions as a function of apparent magnitude for true PSF of FWHM 1.05 & 1.41 pixels (left & right, respectively). The thick horizontal line is the FAME requirement for astrometry: $1/350^{th}$ of a pixel ~ 0.59 mas.

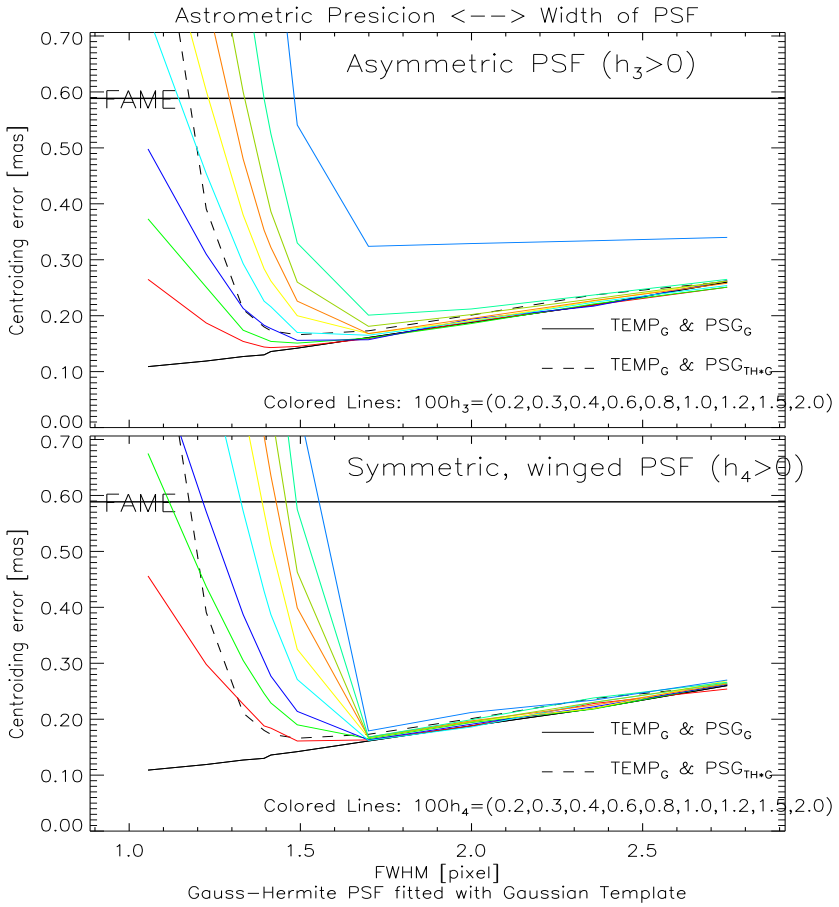

Fig. 3.— Astrometric precisions as a function of the width of the PSF. In all cases, a Gaussian was fitted to the true PSF. For the latter I present: a true Gaussian (thick black line), a pixel-integrated Gaussian (thick dashed line) and a series of GH PSFs (colored lines). The top panel is for $h_4 = 0$ and non-zero h_3 (increasing from left to right). The bottom panel for $h_3 = 0$ and non-zero h_4.

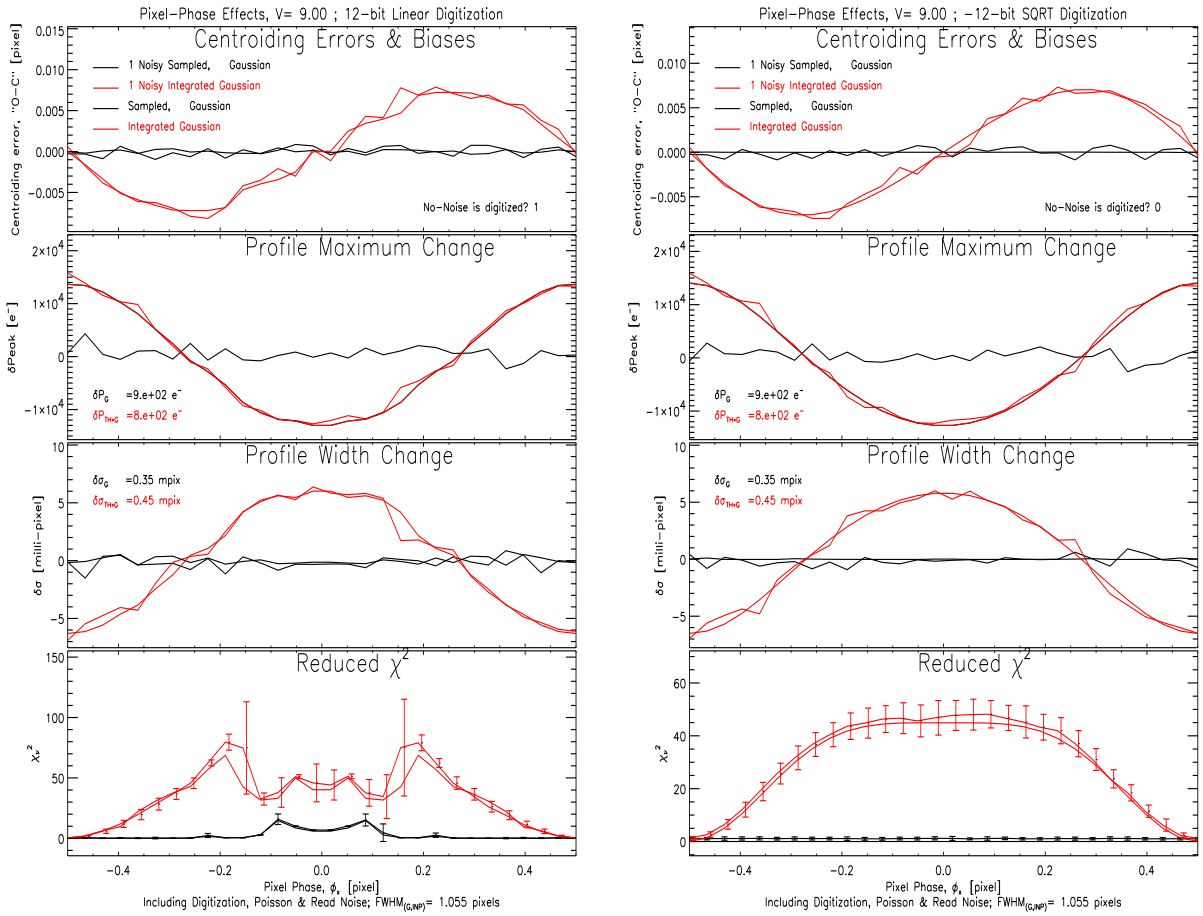

Fig. 4.— Several astrometric parameters as a function of pixel-phase ϕ_x. From top to bottom the panels display: the fitted centroid, the variations in peak flux, the profile width and the reduced χ^2. The black lines are for $PSF_1 = G(x)$, the red lines for $PSF_{2,TH*G}$. The PSF_1 results show the smallest systematics. In the top three panels, the thin lines represent one of the 50 random realizations, whereas the thick lines are the results for the no-noise profiles. For the χ^2_ν plot the thin lines represent the mean of 50 random models. The expected fitting errors are also displayed in the profile-peak and profile-width panels. For the left-hand set of panels, all data was digitized using a 12-bit, *linear* digitization scheme before fitting. The right-hand set of panels was digitized using a 12-bit, *square-root* algorithm digitization (except for the no-noise case [thick lines]). Comparing the χ^2_ν panels, it is clear that the linear digitization scheme introduces very large (15) reduced χ^2 values, even in the case that both fitting template and the PSF are Gaussians (thick black lines). Square-root digitization behaves much better.

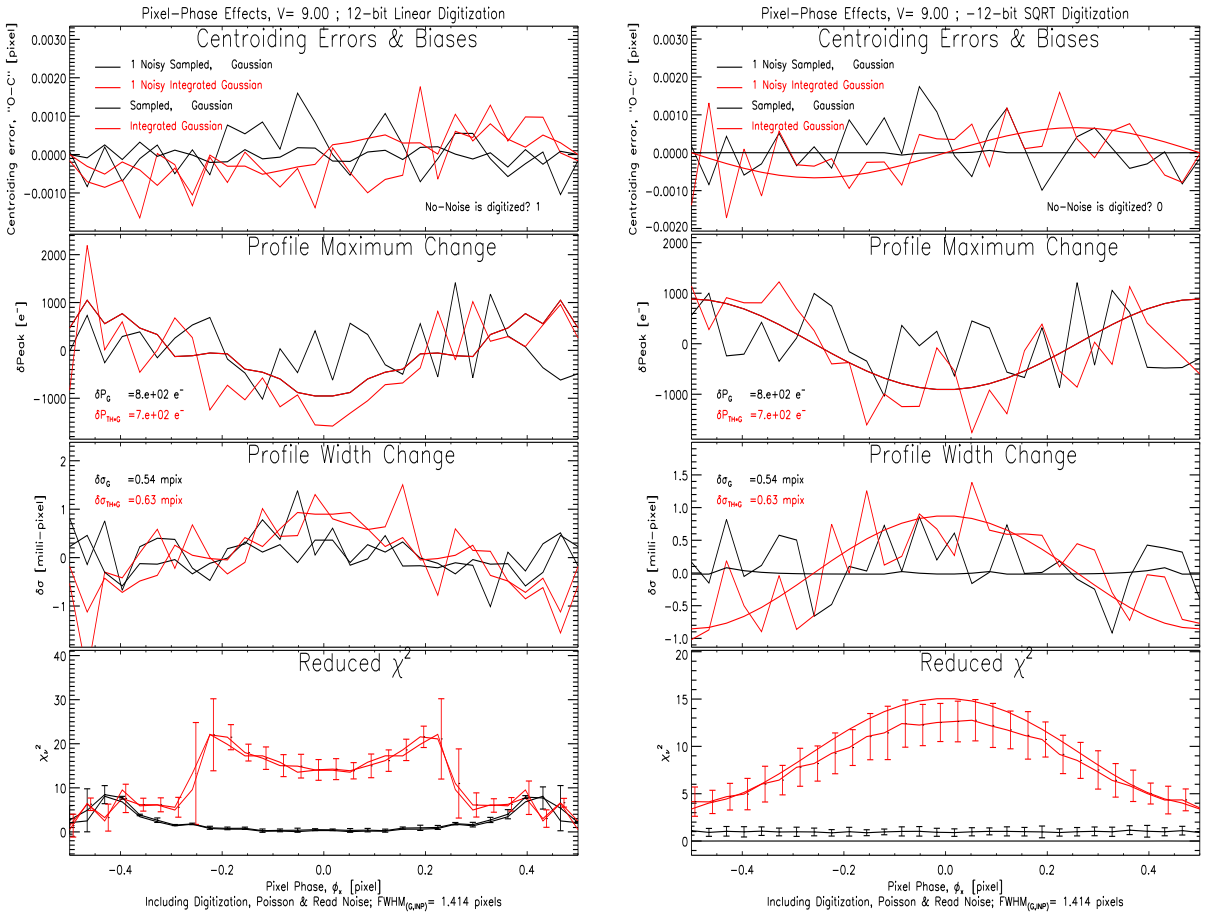

Fig. 5.— Same as for figure 4 but for FWHM$_G$=1.414 pixels. Note that the pixel-phase effects on the "astrometric" parameters are much smaller. The lower panels show that the reduced χ^2 values are about twice weaker. However, the G-G case (thick black lines) produces still unacceptably large χ^2.