

ASTROMETRY AND ORBITS OF NIX, KERBEROS, AND HYDRA

MARC W. BUIE¹, WILLIAM M. GRUNDY², AND DAVID J. THOLEN³

¹ Southwest Research Institute, 1050 Walnut Street, Suite 300, Boulder, CO 80302, USA; buie@boulder.swri.edu

² Lowell Observatory, 1400 West Mars Hill Road, Flagstaff, AZ 86001, USA; grundy@lowell.edu

³ Institute for Astronomy, University of Hawaii, 2680 Woodlawn Drive, Honolulu, HI 96822, USA; tholen@ifa.hawaii.edu

Received 2013 March 3; accepted 2013 September 24; published 2013 November 8

ABSTRACT

We present new *Hubble Space Telescope* observations of three of Pluto’s outer moons, Nix, Kerberos, and Hydra. This work revises previously published astrometry of Nix and Hydra from 2002 to 2003. New data from a four-month span during 2007 include observations designed to better measure the positions of Nix and Hydra. A third data set from 2010 also includes data on Nix and Hydra as well as some pre-discovery observations of Kerberos. The data were fitted using numerical point-spread function (PSF) fitting techniques to get accurate positions but also to remove the extended wings of the Pluto and Charon PSFs when working on these faint satellites. The resulting astrometric data were fitted with two-body Keplerian orbits that are useful for short-term predictions of the future positions of these satellites for stellar occultation and for guiding encounter planning for the upcoming *New Horizons* flyby of the Pluto system. The mutual inclinations of the satellites are all within 0.2° of the plane of Charon’s orbit. The periods for all continue to show that their orbits are near but distinct from integer period ratios relative to Charon. Based on our results, the period ratios are Hydra:Charon = 5.98094 ± 0.00001 , Kerberos:Charon = 5.0392 ± 0.0003 , and Nix:Charon = 3.89135 ± 0.00001 . Based on period ratios alone, there is a trend of increased distance from an integer period ratio with decreasing distance from Charon. Our analysis shows that orbital uncertainties for Nix and Hydra are now low enough to permit useful stellar occultation predictions and for *New Horizons* encounter planning. In 2015 July, our orbits predict a position error of 60 km for Nix and 38 km for Hydra, well below other limiting errors that affect targeting. The orbit for Kerberos, however, still needs a lot of work as its uncertainty in 2015 is quite large at 22,000 km based on these data.

Key words: astrometry – ephemerides – Kuiper belt objects: individual (Pluto, Charon, Nix, Hydra, Kerberos)

Online-only material: color figures, supplemental data

1. INTRODUCTION

In our recent paper, we presented an analysis of the motion of Pluto’s largest satellite, Charon, in which we claimed that the orbit is circular (Buie et al. 2012). Many of those observations also include useful astrometric information on the smaller outer satellites. In this companion paper, we present astrometric results and simple orbit fitting for Nix, Hydra, and Kerberos.

Showalter & Hamilton (2011) recently announced the discovery of a fourth satellite, originally designated as S/2011 P1, alternatively as P4, and now named Kerberos, orbiting Pluto in images taken with the Wide Field Camera 3 (WFC3) on board the *Hubble Space Telescope* (HST) during Cycle 18 in 2011. Even more recently, an even fainter, fifth satellite named Styx, also known as S/2012 P1, alternatively as P5, was discovered (Showalter et al. 2012). Inspection of our pre-discovery Cycle 17 images of the Pluto system taken with the same instrument for purposes of refining the orbits of Nix and Hydra revealed Kerberos at a low signal-to-noise ratio (S/N) in the 2010 June 25 data (Weaver et al. 2011). We have now carefully examined all the long-exposure images from all 12 visits during Cycle 17 and have successfully identified Kerberos in all but one visit. Searches for pre-discovery observations of Styx were unsuccessful, as it is too faint to be detected in our data.

The discovery of these small, outer satellites is especially timely in light of NASA’s *New Horizons* spacecraft’s impending encounter with the Pluto system in 2015. In order to be able to image the new satellites, accurate ephemerides are needed. This paper presents extracted astrometry with measurement uncertainties of Nix, Kerberos, and Hydra. We also show simple two-body unperturbed orbit fits to these observations. These

orbits are not expected to be useful for predicting long-term motions but provide a practical tool for error analysis of the data and the resulting orbit quality at the time of the *New Horizons* encounter. We will show that the orbits of Nix and Hydra are already good enough to support the encounter imaging but much more work is still needed for Kerberos and especially for Styx.

The orbit fitting work presented here has limited value but has the virtue of being extremely simple both conceptually and computationally. In reality, the mean combined gravity field of Pluto and Charon is non-spherical and acts like a single body with a large J_2 gravitational moment. A simple two-body orbit will have limited temporal validity. The actual gravity field will cause the orbit pericenters to precess and nodes will regress (Murray & Dermott 1999). These first order perturbations to the two-body elements are easiest to recognize and fit if working in a coordinate frame aligned to the Pluto–Charon plane unlike the system we have chosen for historical purposes. In our analysis, we neglect these effects largely because this is a data analysis paper and providing a more accurate dynamical description of the system is beyond the scope of this work. Even so, our derived orbital elements are still useful over short timescales, certainly on the scale of the errors inherent in our astrometric data. In our longer-term fits, the values for the semimajor axes and periods should be accurate but all of the other elements are likely to be seriously biased and should be used with caution.

2. OBSERVATIONS

2.1. Cycle 11 ACS/HRC Observations

These data have been documented, analyzed, and discussed in Buie et al. (2006, 2010a, 2010b, 2012) and Tholen et al.

(2008). Due to the relatively short exposure times, only Nix and Hydra were seen in these data. The previous measurements are still valid and we have not improved on the extracted positions. However, our improved understanding of the orbit of Charon permits us to more accurately place these satellites relative to the Pluto–Charon barycenter.

To recover fainter satellites in our data, we shifted and stacked multiple images, requiring a reference position to define the necessary shifts. In our Cycle 11 astrometry of Nix and Hydra, we used the measured position of the much brighter Charon as this reference position (Buie et al. 2006; Tholen et al. 2008). The process of stacking images can introduce a temporal bias from the assumed or chosen time of the data point. In the case of these data, the error could be as much as an hour since the sets of images took one to two visits to collect. In Buie et al. (2006) and also in Tholen et al. (2008), the time tags for all the astrometry were taken to be the mean of all the individual image mid-times. This would be correct if the mean registration point were at the middle of the timespan. The method worked reasonably well but some stacks had a slight amount of trailing when looking at either Nix or Hydra due to the differential motion between the outer satellite and Charon. The measurements published in 2008 improved the results by using the ephemeris motion of either Nix or Hydra (as appropriate) to stack the images on the faintest pixel and with no interpolation. The key to understanding the time of observation is the coordinate system for the output image of the stack since all measurements were made in this frame and became the astrometry to be fitted. Each frame going into the stack has its own time-tagged coordinate system but only one can survive on output. In this case, the first frame in the stack was used to define this coordinate system for output. Thus the time-tag for each Nix and Hydra point must be the mid-time of the first exposure. The data presented for Nix and Hydra in this paper have been revised to this corrected interpretation of the time of observation. The positions reported are unchanged from the 2008 publication except for the time.

2.2. Cycle 15 WFPC2/PC Observations

These observations are described in detail in Buie et al. (2012). Nix and Hydra are detected on most visits while Kerberos and Styx are too faint to be seen. The details of the observations are unimportant in this work except to note that a variety of filters were used, leading to a variable S/N on the observations. Also, a modest dithering pattern was used to combat the undersampled nature of the data.

2.2.1. Extraction of Results

For all images, we had already generated model image fits for Pluto and Charon based on a synthetic point-spread function (PSF). This part of the fitting process was described in Buie et al. (2012). Unlike the fitting process for Pluto and Charon, successful fitting of the Nix and Hydra images also required fitting for background sources. The fits for the stellar background and for Nix and Hydra are iteratively processed. First, the stellar background model is fitted, and after subtracting the stars, Nix and Hydra were fitted. Then, from the images with Nix and Hydra subtracted, the stars are refit, and so on until a converged solution was determined. For this entire process, the Pluto and Charon model images were already subtracted.

As with the Pluto astrometry, the location of the peak pixel for Charon is the primary reference point in each image. All satellite measurements are ultimately referenced to this point

prior to any orbit fitting. Just like the Pluto–Charon fitting step, this initial Charon position is used to define the precise and unchanging location of the fitting region in each image. Sky-plane coordinates for the other objects were computed based on the best ephemeris available as of 2007 July and their image locations were computed on each frame relative to the measured position of Charon.

Additionally, a single star was chosen for each visit as an anchor star. This star was required to be well-separated from all other sources in the image and be as bright as possible without being saturated. All of the field stars were measured relative to this anchor star.

Prior to fitting, all images were processed in the same manner. Flattened images from the STScI pipeline were the starting point. Next, the sky signal was subtracted. The sky-subtracted image was multiplied by the pixel-area map (f1k1552bu_r9f.fits obtained from the STScI archive). This step corrects the image for the photometric effects of the varying pixel size across the field of view. This is a small correction for WFPC2 but follows the same procedure we used for the Cycle 11 Advanced Camera for Surveys (ACS) fitting process where the correction was much more severe. The final step is to subtract a model image of any source that should not be present for the fitting about to take place. For example, when fitting Nix, the images for all field stars, Pluto, Charon, Nix, and Kerberos are subtracted. Another array was computed that records the pixel-by-pixel uncertainty in the image and is computed using quadrature summation of the photon noise based on the signal recorded in the original flattened image with the read-noise, taken to be $5 e^-$, while the gain was taken to be $13.99 e^-/\text{DN}$ (McMaster & Biretta 2008). The final uncertainty image was further limited to ensure no uncertainty was less than 1.5 DN. This latter correction was made late during the data analysis to prevent unrealistically high weights for sky pixels. This issue was most important for the final uncertainties on the fitted quantities.

A final element of the fitting process is a mask image whose function is to inform the fitting program of which pixels were to be fitted. The mask image records separately those pixels flagged as cosmic-ray strikes (CRS) or saturated pixels ($\text{DN} > 3000$). This mask image is initially built with automated routines. The mask images in the area around all Pluto system objects were carefully examined and modified to ensure that the masks were correct. It was a common occurrence for the pixels at the core of Nix and Hydra to be flagged as a CRS by the automated routines. A similar review was carried out for any field stars that fell at least once within the Pluto system fitting regions sometime during a visit.

The final set of arrays provided for fitting is thus a corrected image array, sigma array, and mask array given a set of predetermined optical distortion corrections.

2.3. Distortion Corrections

Image distortions are small for WFPC2/PC but must still be removed from the positional measurements. During the fitting process, a fit for a single source in a single image uses raw image positions that are in the distorted reference frame. For any fit involving multiple images within a visit, the fitting coordinates are carried in an undistorted tangent plane. Whenever an undistorted coordinate needed to be used in an image, it had to be converted to the distorted absolute pixel coordinates. The distortion correction was computed using the `wfpc2_metric` program contained within the IDL Astronomy User's Library

Table 1
Cycle 15 Ephemeris

ID	JDT (days)	X_{Nix} (arcsec)	Y_{Nix} (arcsec)	X_{Hydra} (arcsec)	Y_{Hydra} (arcsec)	OTAT (deg)
u9qa0105m	2454183.5285417	1.135	0.741	0.689	-2.762	-40.36
u9qa010am	2454183.5405440	1.138	0.736	0.684	-2.761	-40.36
u9qa0203m	2454190.5954051	0.654	-1.995	-1.539	-1.219	-51.55
u9qa0206m	2454190.6028472	0.651	-1.995	-1.540	-1.215	-51.55
u9qa040dm	2454204.3264583	-0.266	2.132	-0.214	2.752	-47.71
u9qa0506m	2454211.1695139	1.554	-0.564	1.798	0.660	-48.83
u9qa0604m	2454218.2266551	-0.307	-1.733	1.719	-2.199	-50.07
u9qa0905m	2454239.2778472	1.115	-1.995	-1.178	2.567	-55.95
u9qa100dm	2454246.3570139	-1.347	-0.513	1.093	1.724	-59.73
u9qa1103m	2454252.9926273	-0.639	2.018	2.077	-1.143	-65.57
u9qa1105m	2454252.9979861	-0.637	2.019	2.076	-1.145	-65.57
u9qa1302m	2454247.9292361	-1.621	0.420	1.443	1.221	-60.86
u9qa130bm	2454247.9495718	-1.623	0.432	1.448	1.215	-60.86
u9qa1407m	2454250.0049306	-1.417	1.454	1.929	0.358	-62.55
u9qa140cm	2454250.0160417	-1.415	1.458	1.931	0.353	-62.54
u9qa1701m	2454239.3397106	1.094	-2.011	-1.164	2.574	-55.98
u9qa1703m	2454239.3508218	1.091	-2.014	-1.161	2.575	-55.98
u9qa1901m	2454225.4836343	-1.319	1.494	-0.455	-2.372	-51.60
u9qa1904m	2454225.4919676	-1.317	1.496	-0.458	-2.370	-51.60
u9qa1905m	2454225.4940394	-1.317	1.497	-0.458	-2.370	-51.60

(Supplemental data for this table are available in the online journal.)

and is based on the distortion coefficients published in McMaster & Biretta (2008).

2.4. Anchor Star Fitting

The anchor star fitting process is essentially the same as that for fitting Pluto and Charon except in this case there is a single unresolved and isolated source fitted on each image. The half-width of the fitting box was set to 20 pixels. Thus, the fitting process returns a location in the distorted frame and flux for the star. The observations are all tracked at the Pluto rate and include on-orbit parallax corrections for Pluto as well. Thus, all of the stars are trailed to some degree. This trailing is ignored during the fitting process and must degrade the fit at some level. However, the star subtractions look quite good cosmetically and seem to work well enough for the overall fitting process.

2.5. Nix and Hydra Fitting

The S/N for Nix and Hydra is low enough and the stellar background is so crowded that we could not fit against individual images. Instead, the fitting process is simultaneously constrained by all images within a single visit. If Nix and Hydra were fixed in position relative to Charon during a visit, we could simply stack all the images and measure a position from the stack. Unfortunately, both objects move enough that the resulting image stack would be slightly smeared and would degrade the fitted results. The per-visit smear of Nix ranged from 0.09 to 1.40 pixels if uncorrected. The per-visit smear for Hydra ranged from 0.13 to 1.12 pixels. To preserve the concept of a single positional measurement for multiple images, we encoded the position as an offset relative to an ephemeris. This assumption requires that the offset be constant over the time interval of the fitted data—a condition that is easily met. We used the best-fit orbit from Tholen et al. (2008) and the computed values used are tabulated for each frame in Table 1. This table contains the original frame identifier assigned by STScI followed by the mid-time of the exposure. The position lists for Nix and Hydra

are J2000 offsets from the center of Pluto in arcseconds. Lastly, the orientation of each image is indicated by listing the value of the ORIENTAT keyword from the image header. All calculations used $0.0455547 \text{ arcsec pixel}^{-1}$ from McMaster & Biretta (2008) for the plate scale. The ephemeris position is converted to a Charon-relative offset and is added to the measured location of Charon to get the position in the image. If the orbit were already perfect, the fitted offset would be zero for all visits. The orbit chosen does not affect the final result as long as there is no differential motion between the true orbit and the chosen orbit. This difference was clearly not zero, but it was significantly smaller than the smear and less than the size of a pixel. To get the final astrometry for orbit fitting, we take the fitted offset and add it to the original ephemeris used for the visit.

The images that were fitted had all known and fitted sources subtracted. This of course includes Pluto and Charon, but also includes removal of all field stars. This latter step of removing field stars was especially important along with the bad pixel masks. This process made sure that all useful information was used to constrain the fits while avoiding contamination. Note that the uncertainty for each pixel depends on its original flux, not its subtracted flux. Therefore, any region near or on the Nix or Hydra image that had a star subtracted would naturally be deweighted due to the larger noise in those pixels.

2.6. Field Star Fitting

Identifying and fitting the field stars was a very similar process to that for fitting Nix and Hydra. To get starting locations for all the stars, the images in the visit were stacked relative to the anchor star. This image was scanned for point sources and their location was recorded as an offset relative to the anchor. These sources were also searched to find those that are close enough together that their PSFs overlapped. Each clump of stars (one or more sources) was independently fitted for their positions relative to the anchor and their flux. The fitting was done with all images providing simultaneous constraints. There was some




Figure 1. Example images from Cycle 15 data—visit 11 shown. A—original image. B—model image with everything but Nix and Hydra included. C—data minus model (A–B). D—navigation image with objects labeled. E—Pluto thumbnail images; left to right: data, model, sky-scaled residual image, and full-range residual image; top to bottom are the different images within the visit. F—same as E but for Charon. G—same as E but for Nix. H—same as E but for Hydra. See the text for full details.

(A color version of this figure is available in the online journal.)

manual intervention required to define new sources and manage the clumps. This intervention was applied only for sources that were near enough to objects in the Pluto system to affect their fits. As was done with Nix and Hydra, confusing sources were subtracted before fitting. In this case, all objects in the Pluto system were subtracted. This step needed to be iterated with the Nix and Hydra fitting to get cleanly separated fits for field stars and satellites. This process converged quickly in all cases.

2.7. Error Analysis

The uncertainties on positions and fluxes were computed from the variation of $\chi^2 = \sum (f_i ((D_i - M_i) / \sigma_i)^2)$, where D_i is the post-subtraction image pixel value, M_i is the model image value, σ_i is the noise of the pixel, and f_i is the flag value, set to 1 for pixels to use and 0 for pixels to ignore. All parameters were assumed to have uncorrelated noise. With the use of properly weighted fits, an increase in χ^2 by 1 due to the change of a single parameter corresponds to a 1σ change in that parameter (Press et al. 1992). When computing the final errors on the differential astrometry for Nix and Hydra, the errors of the fitted parameters have to be added in quadrature with the positional uncertainty of the reference.

2.8. Fitting Results

An example of the graphical products produced during the fitting process is displayed in Figure 1. The images A–D show a

sub-image from the data centered on the midpoint between Pluto and Charon. Image C is the input to the fitting process for Nix and Hydra. Image D is a diagram showing the objects plotted in the previous images but with object identifiers displayed. P, C, N, and H stand for the objects in the Pluto system. The color coding of the circles matches the colors in E–H. The other (purple) circles labeled with two-digit codes are the field stars that were identified and fitted. Images E–H present the data and model for each fitted object in more detail as an array of thumbnail images zoomed in on the object. The Pluto results are shown in E, where the data are in the left-most column. The images with a black spot in the center were overexposed and those pixels have been replaced with zero to make them easier to see. In all of the thumbnail images, the circle (red in the online journal) shows the fitted position. Rows 3, 4, 10, and 11 are from short-exposure data where Pluto was not saturated. There is no difference in the astrometry between the saturated and non-saturated data. The second column shows the model image. The third column shows the residual image after subtracting the model from the data. All three of these columns are scaled the same. The fourth column is the same as the third column except the scaling is from the minimum to the maximum in the residual image. This columnar organization of the data is used for all four objects. The same type of information is shown for Charon in F. Note that in the bottom three rows there is a field star impinging on the image of Charon. The fitted star properties are included in the model and the residuals are shown for that as

Table 2
Cycle 15 Nix Observations

ID	JDT (days)	Filter	Exp (sec)	dx (pixel)	σ_{dx} (pixel)	dy (pixel)	σ_{dy} (pixel)	sky (DN)	jit (pixel)	χ_v^2
u9qa0101m	2454183.5327083	F555W	100.0	0.064	0.160	-0.276	0.160	0.307	0.46	1.58
u9qa0201m	2454190.6157330	F555W	100.0	0.464	0.060	0.154	0.060	0.391	0.46	1.25
u9qa0301m	2454197.3861806	F555W	100.0	0.112	0.080	0.639	0.080	0.286	0.49	1.63
u9qa0401m	2454204.3132639	F555W	100.0	0.258	0.160	0.082	0.160	0.309	0.60	2.78
u9qa0501m	2454211.1717516	F555W	100.0	0.204	0.080	-0.032	0.080	0.298	0.41	1.30
u9qa0601m	2454218.2347917	F555W	100.0	0.237	0.040	0.397	0.100	0.307	0.50	1.32
u9qa0701m	2454225.4299306	F555W	100.0	0.015	0.160	0.359	0.040	0.325	0.51	1.45
u9qa1901m	2454225.4971759	F814W	140.0	0.062	0.080	0.333	0.080	0.527	0.51	1.34
u9qa0801m	2454232.2229861	F555W	100.0	0.282	0.040	-0.170	0.080	0.312	0.57	1.75
u9qa1801m	2454232.2897106	F675W	50.0	0.355	0.320	0.032	0.320	0.205	0.49	1.39
u9qa0901m	2454239.2820139	F555W	100.0	0.586	0.040	0.112	0.040	0.318	0.54	1.30
u9qa1701m	2454239.3508218	F439W	350.0	0.342	0.480	0.641	0.480	0.123	0.51	0.66
u9qa1001m	2454246.3438194	F555W	100.0	0.137	0.080	0.603	0.080	0.384	0.62	2.08
u9qa1201m	2454247.1411960	F555W	100.0	0.065	0.080	0.650	0.080	0.285	0.61	2.02
u9qa1301m	2454247.9403472	F555W	100.0	0.043	0.160	0.511	0.080	0.293	0.52	2.26
u9qa1401m	2454250.0049306	F555W	100.0	-0.097	0.120	0.448	0.160	0.329	0.52	1.31
u9qa1101m	2454253.0021528	F555W	100.0	0.014	0.080	0.149	0.080	0.420	0.58	1.31

(Supplemental data for this table are available in the online journal.)

Table 3
Cycle 15 Hydra Observations

ID	JDT (days)	Filter	Exp (sec)	dx (pixel)	σ_{dx} (pixel)	dy (pixel)	σ_{dy} (pixel)	sky (DN)	jit (pixel)	χ_v^2
u9qa0101m	2454183.5327083	F555W	100.0	0.561	0.020	0.422	0.080	0.307	0.46	1.31
u9qa0201m	2454190.6157330	F555W	100.0	0.130	0.060	0.909	0.060	0.391	0.46	1.44
u9qa0301m	2454197.3861806	F555W	100.0	-0.117	0.020	0.531	0.030	0.286	0.49	1.76
u9qa0401m	2454204.3132639	F555W	100.0	0.186	0.020	-0.264	0.080	0.309	0.60	1.23
u9qa0501m	2454211.1717516	F555W	100.0	0.446	0.030	-0.468	0.080	0.298	0.41	1.75
u9qa0601m	2454218.2347917	F555W	100.0	0.601	0.020	-0.011	0.080	0.307	0.50	1.19
u9qa0701m	2454225.4299306	F555W	100.0	0.171	0.050	0.743	0.080	0.325	0.51	1.30
u9qa1901m	2454225.4971759	F814W	140.0	0.158	0.060	0.643	0.080	0.527	0.51	1.43
u9qa0801m	2454232.2229861	F555W	100.0	-0.025	0.030	0.678	0.050	0.312	0.57	1.55
u9qa1801m	2454232.2897106	F675W	50.0	-0.193	0.160	0.858	0.160	0.205	0.49	1.28
u9qa0901m	2454239.2820139	F555W	100.0	-0.005	0.080	-0.044	0.080	0.318	0.54	1.44
u9qa1701m	2454239.3508218	F439W	350.0	0.510	0.200	0.032	0.320	0.123	0.51	0.66
u9qa1001m	2454246.3438194	F555W	100.0	0.649	0.030	-0.610	0.080	0.384	0.62	1.24
u9qa1201m	2454247.1411960	F555W	100.0	0.574	0.040	-0.514	0.040	0.285	0.61	1.28
u9qa1301m	2454247.9403472	F555W	100.0	0.590	0.040	-0.578	0.040	0.293	0.52	1.36
u9qa1401m	2454250.0049306	F555W	100.0	0.571	0.020	-0.453	0.080	0.329	0.52	1.34
u9qa1101m	2454253.0021528	F555W	100.0	0.725	0.020	-0.257	0.080	0.420	0.58	1.27

(Supplemental data for this table are available in the online journal.)

well. Finally, we see the results for Nix in sub-image G (orange circles) and Hydra in sub-image H (green circles). The small circles show the fitted location as before with Pluto and Charon. The large circles in Column 3 show the ephemeris position that was fitted against. If the prior orbit had been perfect, these two circles would be concentric. The offset shows the error in the prior orbit. In this particular example, the predicted location for Nix was extremely good while the Hydra position was clearly in error. In general, the orbit for Hydra was more in error across all the visits. This example is typical of most cases for Hydra. Not all of the Nix data were as close as this, showing that the orbit should indeed improve by including the new data. The empty thumbnails in rows 3, 4, 10, and 11 are from the short-exposure navigation images. These empty cells are shown to make it easier to compare results from the same images. The

first row is the same data image for all four objects and this pattern repeats for all rows.

The results from the fitting process are contained in Table 2 for Nix and Table 3 for Hydra. For each image, the identifier and exposure mid-time of the first frame in the visit are listed, followed by the filter and exposure time. The final fitted values (dx , dy) for the offset from ephemeris are shown with their uncertainties. Lastly, the computed sky signal, jitter, and the goodness of fit per degree of freedom in the image fit are tabulated. In most cases, χ_v^2 is slightly larger than 1, indicating the uncertainties were underestimated. This discrepancy is not large and we could have applied an a posteriori correction to force each visit to have $\chi_v^2 = 1$ but this was left to be adjusted in a later step. The results are tabulated in order of increasing time. The offsets for Nix range between -0.28 and 0.65 pixels

(−13 to 30 mas). Most values are positive but the offsets clearly change with time. The offsets for Hydra also vary and range between −0.61 and 0.91 pixels (−28 to 42 mas). In the case of Hydra, the offsets are somewhat more evenly distributed between positive and negative values.

3. CYCLE 17 WFC3/UVIS1 DATA

Newer images of the Pluto system were obtained with the UVIS detector of WFC3 during 12 visits by *HST* distributed between 2010 March 14 and 2010 September 7. These data were fully described in Buie et al. (2012). Each visit consisted of 10 exposures, some purposely kept short to prevent saturation on Pluto, and some long to obtain adequate S/Ns on Nix and Hydra for purposes of extracting photometric measurements. All observations were made with the F350LP filter to maximize throughput.

3.1. Extraction of Results

The Cycle 17 data taken with WFC3/UVIS1 were much simpler to process than the Cycle 15 data. The primary reason for this is a significantly higher per-frame S/N on Nix and Hydra for the long-exposure data. The penalty for this was much heavier saturation on Pluto and Charon and somewhat greater dependence on subtracting the wings of the Pluto and Charon PSFs from the region around the satellites. The higher S/N was fortuitous in light of the new discovery of the fourth satellite around Pluto (Showalter & Hamilton 2011). Kerberos is present and visible in most of the long-exposure data. When Kerberos is not visible, it is always as a consequence of obscuration from a field star, diffraction spike from Pluto or Charon, or internal reflection ghosts from Pluto. For these data, all objects are fitted for position and flux in each individual image.

3.2. Image Fitting Setup

The fitting process for the Cycle 17 data is very similar to the Pluto and Charon fitting described in Buie et al. (2012). The main difference is that Pluto and Charon model images plus any satellite not being fitted were subtracted before fitting the target satellite. Similar to the Charon fitting, the faint satellites are inserted into the model image as a 0.5 pixel radius top hat convoluted with the numerical PSF. No correction was made for the pixel-area map since the dithering pattern is very small and the pixel-area variation across our sub-array is negligible. Once again, we use TinyTim to provide the spatially variable PSF (specific to each filter). For these data, the jitter component was not used at all. The half-width of the fitting box was set to 6 pixels for Nix, Kerberos, and Hydra.

3.3. Pixel Masking

The pixel masks for these data were handled similarly to those used for the Cycle 15 data. The biggest difference is that there was no attempt to automatically detect bad pixels other than saturated pixels. The saturation threshold was set to 30,000 e^- (same units as the pipeline processed images). Manual processing was used to mark CRSs as well as field stars. Pixels affected by field stars are simply excluded from the fit rather than being fitted and subtracted. There was an additional manual step where extra pixels were flagged near the automatic saturation-flagged pixels. This was done because of the systematic residual pattern seen around saturation. There were pixels immediately above and below the bleeding trails that were clearly too bright. There

were also pixels to the left and right of the core of saturated sources that were always too bright in the residual image as well. There is also a small detached knot of brightness to the upper right of the saturated core that appeared in proportion to the brightness of the object. This latter artifact is treated as if it were a saturation artifact when it is more likely to be an internal reflection. In either case, TinyTim cannot model its presence and it was safest to block it from affecting the fit.

3.4. Object Fitting

This subtracted image is used for the fitting of Nix, Kerberos, and Hydra. Some images could not be fitted for the small satellites due to unsubtractable image artifacts (typically field stars or image ghosts). During the fitting process, a single TinyTim PSF was used. This PSF was computed for the rough position (to the nearest 50 pixel boundary) and the filter used for the observation. The PSF image accounts for the focal plane distortion at this location on the image. We also used a model value for the telescope focus based on the web-based tool provided by STScI (<http://focustool.stsci.edu/cgi-bin/control.py>). The position and flux of the object is encoded in the model image by placing a circle into the image. The area of the circle defines the flux while the position determines which pixels it overlaps. The size of the circle chosen is the smallest value that can still accurately encode the position without aliasing. This choice was verified by extensive synthetic image testing. The fractional area overlap between pixels is computed and this determines the fraction of the target flux for each pixel. The final model image is the result of convolving the PSF with the emplaced flux by the circular aperture. The fitting proceeds by adjusting the position and flux until a minimum χ^2 value is found. The uncertainties are computed the same way as for the Cycle 15 data. More details on the fitting and error analysis are described in Buie et al. (2012).

The fitted positions are in distorted image coordinates and must be converted to sky-plane offsets to obtain the desired astrometry. The WFC3 FITS image headers contain the mapping function from pixels to the sky that includes the pointing of the telescope and the distortion of the camera. This transformation was applied to all the positions to get raw sky-plane coordinates. The absolute registration provided by the guide stars is not accurate enough for this orbit determination work. Thus, the position of the system barycenter was computed using the orbit from Buie et al. (2012) and the Charon/Pluto mass ratio from Tholen et al. (2008). All of the positions are then referenced to this barycenter position to derive the differential astrometry. This process removes all effects of the optical distortion of the camera as well as removing the spacecraft roll angle. All positions are in the J2000 reference frame when the transformation is complete.

3.5. Fitting Results

An example of the data and the image fitting results taken from visit 17 is displayed in Figure 2. The image labeled A is the original image with the bright object in the center being a blend of Pluto and Charon. Across the rest of the image are field stars and CRSs. Those pixels that have been flagged as CRSs or stars have been replaced with sky values for display purposes. A few small regions interior to the Pluto PSF appear like dropouts because they were replaced with sky values. In the fitting process, these pixels are ignored, no matter why they were flagged. Only those pixels near enough to the objects of interest were flagged since it is a largely interactive process. Image B




Figure 2. Example image from Cycle 17 data—visit 17 shown. A—original image. B—model image with Pluto, Charon, and bad pixels subtracted out or removed. The inset at the upper left corner shows the Pluto–Charon fitting region residuals with different image scaling settings. C—data minus model where all fitted objects are removed. D—model image containing Pluto and Charon, scaled the same as in A. The outer satellite positions are labeled here. E—model image containing Nix, Hydra, and Kerberos, scaled the same as in A. F—image stack where all images from the visit are registered on the Pluto system then averaged to remove features that are not fixed with respect to the system. See the text for full details.

(A color version of this figure is available in the online journal.)

takes the data from A one step further by removing the model image for Pluto and Charon. The region fitted is marked by the (green) box. Also, any pixel that was automatically flagged as saturated is indicated in red. The inset in the upper left is the fitted region. The thumbnail on the left has the same stretch as the main image but with no overlain symbols. On the right is the same residual image scaled to the minimum and maximum in the thumbnail. In the main image, the Pluto system objects are marked (but not labeled). The outer satellites are marked with (red) diamonds. Image C is the same as B except that the model images for the outer satellites have been subtracted as well. Also, the pixels flagged as bad are marked in blue. The model image for Pluto and Charon is shown in D, with the outer satellite position labeled with the object name. The scaling in this image is clearly evident in this image. The model image for the outer satellites is shown in E—Hydra is to the left, Nix is to the upper right, and the faintest of the three is Kerberos. The scaling here is also the same as the previous images. The final image, F, shows a stack of all the images in the visit. The images were registered on the position of Charon and a robust average was performed to remove the stellar background and the transient CRS blemishes. Only coherent structures fixed with respect to the Pluto system survive in this stack. This stacked image is shown for illustration purposes only and was not used in the fitting process. The three outer satellites do appear in this stack, though attempts to locate Styx were unsuccessful. Other residual structures are clearly evident as well. The diffraction spikes from Pluto and Charon

still remain, though at a very low level in the region where the model was subtracted. Another more troublesome feature appears to the upper right of Pluto’s position. There are a series of rings, small and bright near Pluto and fainter and larger further away. At least three distinct rings are seen here and must be the result of an out-of-focus ghost image of Pluto that is not incorporated into the TinyTim PSF model. Some of the potential observations of Kerberos were lost in visits where Kerberos fell on the brightest ghost image. In most cases, Nix and Hydra could be measured even when on the ghost image feature. Although the subtraction of the Pluto and Charon model image is pretty good, it is clear that significant residual structure remains. It is the noise and structure left behind that limits satellite searches and observations in the inner system.

A more detailed look at the data and fitting results on the outer satellites is shown in Figure 3. The difference in rows B and C clearly shows the improvement from subtracting the Pluto and Charon PSFs. This example is typical of the rest of the visits. The post-fit residual patterns clearly show non-random patterns. In this case, the model PSF must be broader than the observed PSF, almost certainly a consequence of our top-hat convolution needed to get accurate positions in these undersampled images. In an earlier version of the reduction, we were using a preliminary version of TinyTim that had quite a few optical prescription errors. In that version, the residuals looked more random as if the PSF size were a better fit. However, there were other artifacts, notably the diffraction spikes, that were clearly worse. In any case, the astrometry from




Figure 3. Example fitting results from Cycle 17 data—visit 17 shown. Row A is the original data, scaled from the image minimum to maximum. Row B is the original data, scaled from -3σ to $+5\sigma$ on the sky level. Row C is the same as B, except that the model Pluto and Charon are now removed. This table contains the data that are fitted for satellite positions and fluxes. Row D is the model image showing the actual location and sampling and is scaled the same as row C. Row E shows the residuals after subtracting the model from the data. The numbers at the bottom identify which image within the visit is displayed and the object name is displayed across the top. See the text for full details.

Table 4
Cycle 17 Observations

ID	JDT (days)	Object	Exp (s)	x (pixel)	σ_x (pixel)	y (pixel)	σ_y (pixel)	sky (e^-)	σ_{sky} (e^-)	Focus (μm)	χ^2
ib4w01ilq	2455310.83593	Hydra	197.0	302.369	0.040	200.458	0.040	44.759	0.019	-0.019	2.47
ib4w02fqq	2455338.18217	Nix	197.0	224.449	0.020	259.745	0.080	55.026	0.024	-0.003	2.74
ib4w03fcq	2455341.44726	Kerberos	197.0	219.668	0.240	296.547	0.320	52.242	0.023	-0.024	1.41
ib4w04j5q	2455349.44461	Nix	197.0	293.467	0.080	241.070	0.100	54.263	0.024	-0.034	2.10
ib4w04j6q	2455349.44790	Kerberos	197.0	238.999	0.400	205.593	0.640	55.980	0.026	-0.035	1.46
ib4w05yoq	2455358.81805	Hydra	183.0	284.101	0.080	311.541	0.040	52.291	0.022	-0.011	2.64
ib4w05z0q	2455358.83653	Kerberos	183.0	318.407	0.320	247.529	0.240	50.160	0.020	-0.048	1.90
ib4w06yxq	2455372.80248	Hydra	180.0	239.204	0.020	224.411	0.030	41.882	0.017	-0.018	3.16
ib4w06yxq	2455372.80248	Kerberos	180.0	205.674	0.120	246.515	0.160	41.882	0.017	-0.018	1.76
ib4w08gyq	2455383.58704	Hydra	190.0	270.305	0.020	320.872	0.080	32.192	0.014	0.010	4.82
ib4w08haq	2455383.60568	Hydra	190.0	270.363	0.020	322.456	0.020	32.039	0.014	-0.037	4.61
ib4w08hbq	2455383.60889	Nix	190.0	250.272	0.050	202.593	0.050	32.440	0.015	-0.041	2.68
ib4w09w1q	2455392.37163	Hydra	197.0	205.661	0.040	289.687	0.030	42.980	0.017	0.021	2.58
ib4w10gsq	2455402.22510	Kerberos	197.0	256.501	0.160	189.401	0.160	45.155	0.018	0.030	1.81
ib4w10gtq	2455402.22839	Hydra	197.0	248.849	0.030	186.149	0.030	45.583	0.018	0.022	2.05
ib4w12dyq	2455446.09636	Nix	197.0	226.345	0.020	297.508	0.020	41.366	0.016	0.052	1.65
ib4w12eaq	2455446.11517	Hydra	197.0	294.577	0.020	200.519	0.020	40.319	0.016	0.011	3.00
ib4w12ebq	2455446.11845	Hydra	197.0	297.000	0.080	203.147	0.050	40.314	0.016	0.008	1.89

(Supplemental data for this table are available in the online journal.)

this earlier reduction is statistically indistinguishable from the astrometry we present here. On that basis, we believe that the astrometry from this fitting is unaffected by the sub-optimal image residuals. It is clear that Kerberos is at the very faint end of the detectable range for this exposure time and explains why we were unable to recover Styx in these data.

The final results from the fitting process are contained in Table 4. Each image is listed with its exposure mid-time and can be shown up to three times depending on how many of the outer satellites could be fitted on the image. The object fitted and the exposure time are listed, followed by the location of the object in distorted image coordinates along with its uncertainty. Next, the sky value and its uncertainty is shown along with the focus position used. The last column shows the goodness of fit per degree of freedom. Once again, the simple noise model we used gets us close to $\chi^2_v = 1$, but the errors are still slightly underestimated. As with the Cycle 15 results, we will address this with the final astrometry. The final uncertainties for Hydra are generally around 2 mas with the worst getting as high as 6 mas. The uncertainties for Nix are generally around 3 mas

with a few as bad as 5 mas. The results for Kerberos are noisier and range from 3 to 4 mas for the best up to 38 mas for the worst with a lot of frame-to-frame variability.

4. FINAL ASTROMETRIC RESULTS

The raw results in the previous sections must be converted to angular offsets for the final astrometry. In all cases, the positions of the satellites are referred to the system barycenter defined by the orbit for Charon from Buie et al. (2012) and the Charon/Pluto mass ratio of 0.1166 taken from Tholen et al. (2008). The final astrometry is collected in Table 5. The data are listed in increasing order of time, showing the time of observation, the HST observing cycle for the observation, the instrument and mode used, the filter(s) used, and the object name. The barycentric offsets are in the J2000 reference frame. The uncertainty listed is a combination ($\sigma = \sqrt{\sigma_x^2 + \sigma_y^2}$) of the original raw uncertainties. It was not feasible to retain the original uncertainties since they are broken into components that do not align with right ascension and declination. In effect,

Table 5
Final Satellite Astrometry Relative to System Barycenter

JDT	Cyc	Inst	Filter	Object	$\Delta\alpha$	$\Delta\delta$	σ
2452436.8217600	11	ACS/HRC	F435W&F555W	Nix	-0.9729	1.8693	0.0298
2452436.8217600	11	ACS/HRC	F435W&F555W	Hydra	-1.5968	0.3505	0.0162
2452440.0270500	11	ACS/HRC	F435W&F555W	Nix	-0.1080	2.1565	0.0298
2452440.0270500	11	ACS/HRC	F435W&F555W	Hydra	-1.5808	1.7931	0.0162
2454183.5195139	15	WFPC2/PC	F555W	Nix	1.1920	0.7631	0.0156
2454183.5195139	15	WFPC2/PC	F555W	Hydra	0.7130	-2.7339	0.0087
2454183.5215972	15	WFPC2/PC	F555W	Nix	1.1926	0.7621	0.0156
2454183.5215972	15	WFPC2/PC	F555W	Hydra	0.7123	-2.7339	0.0087
2454190.5938194	15	WFPC2/PC	F555W	Nix	0.7025	-2.0324	0.0058
2454190.5938194	15	WFPC2/PC	F555W	Hydra	-1.5074	-1.2240	0.0065
2454190.6007639	15	WFPC2/PC	F555W	Nix	0.7000	-2.0330	0.0058
2454190.6007639	15	WFPC2/PC	F555W	Hydra	-1.5090	-1.2211	0.0065
2455310.8171204	17	WFC3/UVIS1	F350LP	Nix	0.0203	-1.9291	0.0048
2455310.8171204	17	WFC3/UVIS1	F350LP	Kerberos	-0.5431	-2.0208	0.0087
2455310.8171204	17	WFC3/UVIS1	F350LP	Hydra	-0.6421	2.6685	0.0063
2455310.8204097	17	WFC3/UVIS1	F350LP	Nix	0.0103	-1.9318	0.0048
2455310.8204097	17	WFC3/UVIS1	F350LP	Kerberos	-0.5701	-2.0218	0.0262
2455310.8204097	17	WFC3/UVIS1	F350LP	Hydra	-0.6420	2.6686	0.0063

(Supplemental data for this table are available in the online journal.)

we converted an ellipsoidal error boundary into a larger circular region that just encloses the major axis of the ellipse. This was a conservative choice that absorbed some of the errors in the uncertainty caused by an inaccurate noise model for the image data that led to $\chi_v^2 > 1$.

There is one additional subtlety in the Cycle 15 data reported here. Even though the result of fitting the images provided one offset per visit, the measurements are cast as one measurement per image. Clearly there is a degree of correlation between the images in a visit, but the correlation is not 100% because the offset needs to be related to the position of Charon in each image and this position is not correlated. To provide just a single measurement per visit would therefore underestimate the noise and could lead to a small systematic error for each visit. Splitting out the measurement for each image properly captures the noise in measuring Charon on each image, but will over-constrain the solution if the uncertainties are not adjusted. Since the aggregate fit comes from combining measurements, we adjusted the uncertainties upward by \sqrt{N} , where N is the number of images fitted in each visit. This adjustment works to provide a balanced constraint on the orbits between the three independent data sets.

5. ORBIT FITTING

The fitting process for the two-body orbits was the same as what we used in Buie et al. (2012). In summary, we use a nonlinear least-squares minimization method to find the best fitting orbit using the assigned weights. This task was handled by the built-in IDL routine, AMOEBA. An ideal fit would have $\chi_v^2 = 1$ if the model were a perfect representation of the actual motion. We do expect the motion to be perturbed and might naturally expect $\chi_v^2 > 1$, but for this crude analysis we will assume the perturbations can be neglected for short-arc observations. Making this assumption means we cannot extract any dynamically significant insight from our orbit fitting results. Instead, we use these fits as a guide for general characterization that can be helpful to future work. Our chosen fitting routine does not evaluate numerical derivatives and as a result does not directly provide the covariance matrix from the returned best-

fit parameters. To find the uncertainties on the resulting orbital elements, we cloned a copy of the observations with random noise added that is consistent with each point's uncertainty. For each cloned copy of the data, we refit an orbit. We used twenty trials, and the standard deviation of the results becomes the uncertainties for the elements. We have performed extensive tests of this procedure with synthetic data and find that it is very robust and is very effective at determining uncertainties on the fitted parameters. This same approach can be used to estimate the positional uncertainty of the satellites at some future date. We stress that these fitting results are clearly inferior to a proper dynamical model and should be used with caution.

As has already been noted by Lee & Peale (2006), the orbits of the outer satellites do not follow strictly Keplerian two-body motion. A fully perturbed solution is required for the best fit and to derive orbits that are fully consistent in the masses of the objects. However, if the system mass is left to vary for each satellite, the resulting two-body orbit provides a very good approximation to their position as a function of time. We take advantage of that property to derive these two-body orbits for the purposes of predicting stellar occultations and planning for the upcoming *New Horizons* encounter. In the case of encounter planning, the positional uncertainty is much more important than the actual orbit from this analysis. Our uncertainties set an upper bound to the actual error. Thus, these orbits give us useful near-term predictive ability while the more complete and difficult perturbed solutions are developed. As we found with the Charon orbit fitting (Buie et al. 2012), fitting different subsets of data is useful to see trends in the data and their noise behavior. In the following, we fit different groups of data and collect the results at the end of the section.

5.1. Fitting of Cycle 11 Data

The fitting of these observations with two-body orbits was already presented in Buie et al. (2006) using many the same techniques we used in this analysis. The main difference from our older work is that the Pluto–Charon barycenter position is now known much better (Buie et al. 2012) and we fixed the assigned time for each observation. These data

Table 6
Orbit Fitting Results

Object	Cycle	P (days)	a (km)	e	i (deg)	L (radians)	Ω (radians)	ω (radians)	$ \overline{O - C} _\alpha$ (mas)	$ \overline{O - C} _\delta$ (mas)	χ_v^2
Nix	11	24.8571(41)	48620(220)	0.0032(32)	96.1(5)	2.153(11)	3.894(7)	3.8(1.4)	12.7	9.0	1.02
Nix	15	24.86070(86)	48814(16)	0.00288(29)	96.21(5)	2.250(14)	3.8931(8)	0.65(8)	4.9	4.9	1.00
Nix	17	24.86046(24)	48744(52)	0.0012(5)	96.24(9)	2.313(7)	3.8958(17)	5.7(9)	4.9	4.4	2.53
Nix	11–15	24.85514(28)	48811(22)	0.00300(28)	96.219(54)	2.1585(48)	3.8930(8)	0.61(8)	6.0	4.8	0.85
Nix	11–17	24.854943(72)	48841(23)	0.00230(22)	96.287(33)	2.1567(14)	3.8938(8)	0.47(8)	5.9	5.4	1.47
Kerberos	17	32.1866(22)	57890(120)	0.0045(11)	96.37(21)	1.413(39)	3.8894(36)	0.31(24)	12.4	9.4	1.00
Hydra	11	38.2065(28)	64720(150)	0.0068(20)	96.36(26)	5.6354(27)	3.8940(24)	3.50(36)	6.8	6.6	1.29
Hydra	15	38.19652(52)	64736(18)	0.00739(16)	96.395(16)	5.5874(36)	3.88762(35)	5.890(21)	2.8	5.5	1.01
Hydra	17	38.19736(24)	64695(28)	0.00523(31)	95.995(47)	5.5699(30)	3.8885(6)	0.584(43)	2.6	3.7	1.00
Hydra	11–15	38.20324(28)	64727(15)	0.00734(16)	96.410(23)	5.6341(20)	3.8874(4)	5.858(17)	4.7	7.7	1.56
Hydra	11–17	38.20169(9)	64736(10)	0.00658(16)	96.331(21)	5.6241(7)	3.88785(27)	6.006(17)	5.1	11.9	5.22

Note. All elements are in a J2000 coordinate system from two-body fits relative to the Pluto–Charon barycenter as inferred from Buie et al. 2012.

Table 7
Position Uncertainty for *New Horizons* Flyby

Object	Cycle	Uncertainty (km)
Nix	11	5600
Nix	15	860
Nix	11–15	270
Nix	17	112
Nix	11–17	70
Kerberos	17	22000
Hydra	11	2500
Hydra	15	590
Hydra	11–15	160
Hydra	17	340
Hydra	11–17	40

Note. Uncertainties computed on 2013 Jul 15 relative to the Pluto–Charon barycenter as inferred from Buie et al. 2012.

did not have an uncertainty derived at the time of our original measurements. Given the very low signal, our normal techniques of measuring the uncertainty failed and instead we assigned a global uncertainty based on the scatter from the fit. Based on this analysis, we re-adjusted the uncertainties with a single scaling factor so that $\chi_v^2 = 1$ from this fit. The Hydra uncertainties grew from 9 mas in the old analysis up to 16 mas in this work, while Nix went from 15 mas to 30 mas.

The fitting results are summarized in Table 6. Additionally, the uncertainty in their position at the time of the *New Horizons* encounter is shown in Table 7. As expected, the orbits for Nix and Hydra based on these data have the largest uncertainties and do the worst job of predicting their position at the time of the *New Horizons* encounter. The latter is a direct consequence of the short observational arc (1 yr) and long extrapolation (13 yr) to the encounter.

5.2. Fitting of Cycle 15 Data

The fitting of the WFPC2 data was very straightforward. Of the three data sets, this one has the shortest observational arc, spanning just four months. As with the Cycle 11 data, we found a need for a global inflation of the uncertainties based on the value of χ_v^2 from the fit. Given the various adjustments already made, a small refinement seemed justified.

As expected, the orbital elements in Table 6 are quite uncertain. The predicted error at the Pluto encounter is also

large but not as bad as from Cycle 11 data alone since the extrapolation interval is much shorter.

5.3. Fitting of Cycle 17 Data

These data were fitted the same as the others except no global adjustments were made to the uncertainties. The value of χ_v^2 looks fine for Hydra and Kerberos observations but is larger for the Nix data. There should be no difference in the handling of measurement errors between the three objects and we did not feel justified in adjusting the Nix data separately from the others.

For a single epoch data set, these observations do a good job in determining the orbit. The encounter prediction is correspondingly smaller, largely due to the shorter extrapolation interval. The orbit for Kerberos comes from this data set alone and the uncertainties in the orbital elements and encounter position are all quite large. In these data, Kerberos is quite faint and the raw uncertainties on the positions are larger than the other satellites. However, with additional observations we can expect the orbit to improve much as was seen for Nix and Hydra. As long as *HST* observations continue, Kerberos (and also Styx) should be easy to predict for *New Horizons* during the flyby.

5.4. Combined Cycle Fitting

After completing the fitting for individual data sets, we combined ever larger groups of data and fitted orbits. All of the orbit fitting results are summarized in Table 6. Clearly, the uncertainties decrease as the observational time base increases. The best results come from combining all the available data. The results in Table 7 show that the positional uncertainties for Nix and Hydra are well under 100 km as of the *New Horizons* encounter in 2015. For these objects, the dominant targeting uncertainty is now the uncertainty in the exact time of closest approach. Kerberos still needs work to get to this level.

It is interesting, perhaps even puzzling, to see that χ_v^2 for the full data set on Nix is lower than that for Cycle 17 alone. Also, χ_v^2 for Hydra in the combined data set is significantly larger than unity. Perturbations on the satellites would definitely show up as an increase in χ_v^2 . These results suggest that Hydra is more perturbed than Nix. We can only reiterate our caution that a complete understanding of the motions of this satellite system will require a fully perturbed dynamical model fit.

The post-fit residuals for the full data set on Nix are shown in Figure 4 plotted against orbital longitude. Figures 5 and 6 show residual information for Kerberos and Hydra, respectively. Orbital longitude refers to the angle along the orbit measured




Figure 4. Post-fit residuals for Nix astrometry. This plot shows the scatter in the residuals from the fit of all our data on Nix. The residuals do not show any strong patterns versus orbital longitude. The green circles show the Cycle 11 residuals. The blue squares show the Cycle 15 residuals. The black triangles show the Cycle 17 residuals.

(A color version of this figure is available in the online journal.)




Figure 5. Post-fit residuals for Kerberos astrometry. This plot shows the scatter in the residuals from the fit of our limited data. The residuals do not show any strong patterns versus orbital longitude. These data still have large gaps in orbital coverage. All data are from Cycle 17.

from inferior conjunction. The residual pattern in right ascension does not show any obvious non-random patterns. There may be a slight non-random signature in declination but it is very weak and depends on a small number of measurements. There is also no apparent pattern with respect to the data set.

The residuals for the Kerberos data are understandably larger due to being fainter. As with Nix, there are no discernible patterns with orbital longitude, but the data do not uniformly fill all orbital longitudes. The Cycle 17 observing cadence clearly

missed large regions of longitude. At the time the observations were designed, Kerberos was not known and the observing pattern optimization only concerned sampling the orbits of Charon, Nix, and Hydra.

The residuals for Hydra show the most structure of the three, as expected from its larger value for χ^2_v in the orbit fit. Once again, there is no large pattern seen versus orbital longitude. However, these residuals show a clear trend or at least groupings by data set. The grouping is most evident in the declination




Figure 6. Post-fit residuals for Hydra astrometry. This plot shows the scatter in the residuals from the fit of all our data on Nix. The green circles show the Cycle 11 residuals. The blue squares show the Cycle 15 residuals. The black triangles show the Cycle 17 residuals. The residuals do not show any strong patterns versus orbital longitude. However, there are systematic patterns in the residual versus time. All of the residuals clustered around +50 mas in declination are from Cycle 11. This pattern suggests that these data contain useful perturbation signatures.

(A color version of this figure is available in the online journal.)

residuals. It is always possible that some unrecognized problem remains in the extracted results but the most likely source of difficulty is related to subtracting Pluto and Charon. In this case, it seems reasonable to expect Hydra, being the most distant satellite, would be least affected and Nix would be most affected. On this basis, we suggest that the residual patterns are a signature of perturbations and orbital evolution that a two-body orbit fit cannot match. Confirmation of the question of perturbations most likely lies with completing a fully perturbed solution with mass determination, perhaps requiring data over an even longer time base. Getting a complete census of the gravitationally important objects would also help greatly. It seems that we may be nearing completion on the inventory of satellites around Pluto, but that sentiment has proven to be incorrect multiple times already.

6. CONCLUSIONS

We present astrometric results from three epochs of data taken with the *Hubble Space Telescope*. The orbits of Nix, Hydra, and Kerberos are very close to co-planar and are all close to mean-motion resonances (MMRs) with Charon. Based on our results, the period ratios are Hydra:Charon = 5.98094 ± 0.00001 , Kerberos:Charon = 5.0392 ± 0.0003 , and Nix:Charon = 3.89135 ± 0.00001 . On the basis of strict ratios, Hydra appears to be closest to its MMR location. In fact, there appears to be a trend of getting further from MMR as the object gets closer to Charon. Also, not all of the objects are on the same side of the MMR. Nix and Hydra are interior to the MMR, while Kerberos is exterior. This type of orbit analysis cannot substitute for a fully perturbed solution, but the trends are intriguing nonetheless. We can look forward to additional astrometry from *HST* and *New Horizons* that will more fully reveal the current dynamical state of this interesting system of objects.

This paper is based on observations with the NASA/ESA *Hubble Space Telescope* obtained at the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Incorporated, under NASA contract NAS5-26555. Support for this work was provided by NASA through grant numbers HST-GO-10786, HST-AR-10940, and HST-GO-11556 from STScI. This project also was made possible, in part, by the software development efforts of Doug Loucks, Peter Collins, and Amara Graps. We are also grateful for help from STScI personnel Tony Roman and Bill Januszewski for their attention to scheduling the observations, with understanding PSFs from Linda Dressler and Remi Soummer, and Colin Cox for his development of the focus predictor tool.

REFERENCES

- Buie, M. W., Grundy, W. M., Young, E. F., Young, L. A., & Stern, S. A. 2006, *AJ*, **132**, 290
- Buie, M. W., Grundy, W. M., Young, E. F., Young, L. A., & Stern, S. A. 2010a, *AJ*, **139**, 1117
- Buie, M. W., Grundy, W. M., Young, E. F., Young, L. A., & Stern, S. A. 2010b, *AJ*, **139**, 1128
- Buie, M. W., Tholen, D. J., & Grundy, W. M. 2012, *AJ*, **144**, 15
- Lee, M. H., & Peale, S. J. 2006, *Icar*, **184**, 573
- McMaster, M., & Biretta, J. 2008, *WFPC2 Instrument Handbook*, Version 10.0 (Baltimore, MD: STScI)
- Murray, C. D., & Dermott, S. F. 1999, *Solar System Dynamics* (New York: Cambridge Univ. Press), 152
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 1992, *Numerical Recipes in C: The Art of Scientific Computing* (New York: Cambridge Univ. Press)
- Showalter, M., & Hamilton, D. P. 2011, *IAU CBET*, **2769**, 1
- Showalter, M. R., Weaver, H. A., Stern, S. A., et al. 2012, *IAUC*, **9253**, 1
- Tholen, D. J., Buie, M. W., Grundy, W. M., & Elliot, G. T. 2008, *AJ*, **135**, 777
- Weaver, H., Hamilton, D. P., Showalter, M. R., et al. 2011, *IAU CBET*, **2769**, 1