

A New Method for Wide-Field Near-IR Imaging with the Hubble Space Telescope

Ivelina G. Momcheva,^{1,2} Pieter G. van Dokkum,³ Arjen van der Wel,⁴ Gabriel B. Brammer,¹ John Mackenty,¹ Erica J. Nelson,³ Joel Leja,³ Adam Muzzin,⁵ Marijn Franx⁶

ABSTRACT

We present a new technique for wide and shallow observations using the near-infrared channel of Wide Field Camera 3 (WFC3) on the *Hubble Space Telescope* (HST). Wide-field near-IR surveys with HST are generally inefficient, as guide star acquisitions make it impractical to observe more than one pointing per orbit. This limitation can be circumvented by guiding with gyros alone, which is possible as long as the telescope has three functional gyros. The method presented here allows us to observe mosaics of eight independent WFC3-IR pointings in a single orbit by utilizing the fact that HST drifts by only a very small amount in the 25 seconds between non-destructive reads of unguided exposures. By shifting the reads and treating them as independent exposures the full resolution of WFC3 can be restored. We use this “drift and shift” (DASH) method in the Cycle 23 COSMOS-DASH program, which will obtain 456 WFC3 H_{160} pointings in 57 orbits, covering an area of 0.6 degree² in the COSMOS field down to $H_{160} = 25$. When completed, the program will more than triple the area of extra-galactic survey fields covered by near-IR imaging at HST resolution. We demonstrate the viability of the method with the first four orbits (32 pointings) of this program. We show that the resolution of the WFC3 camera is preserved, and that structural parameters of galaxies are consistent with those measured in guided observations.

Subject headings: cosmology: observations — galaxies: evolution — instrumentation: miscellaneous — techniques: image processing — telescopes

1. Introduction

Over its lifetime the *Hubble Space Telescope* (HST) has imaged many “blank” fields at many wavelengths, to obtain a census of the Universe over most of its history. The survey strategy of the extra-galactic community has been to image a few individual HST pointings to great depth (examples are the Hubble Deep Fields, the Ultra

Deep Field, and the Frontier Fields; [Williams et al. 1996](#); [Ellis et al. 2013](#); [Illingworth et al. 2013](#); [Lotz et al. 2014](#)), and larger areas to shallower depth (the GEMS survey, COSMOS, the GOODS North and South fields, CANDELS, etc.; [Rix et al. 2004](#); [Scoville et al. 2007](#); [Giavalisco et al. 2004](#); [Grogin et al. 2011](#); [Koekemoer et al. 2011](#)). This “wedding cake” strategy of tiered surveys is driven by the form of the luminosity function of most astronomical objects. The number density of faint objects is almost always larger than that of bright objects, which means that representative samples of faint objects can be obtained in deep, pencil-beam surveys and representative samples of bright objects in shallow, wide-area surveys.

Ground-based surveys have extended the shallow, wide-area domain to degree-scales and beyond: SDSS covers $\sim 1/3$ of the entire sky in the optical, and a plethora of optical and near-IR sur-

¹Space Telescope Science Institute, Baltimore, MD 21218, USA

²imomcheva@stsci.edu

³Department of Astronomy, Yale University, 260 Whitney Avenue, New Haven, CT, USA 06511

⁴Max Planck Institute for Astronomy, D-69117, Heidelberg, Germany

⁵Institute of Astronomy, University of Cambridge, Cambridge, UK

⁶Leiden Observatory, Leiden University, Leiden, The Netherlands

veys are covering areas up to thousands of square degrees (e.g., the 155 degree² CFHTLenS survey, the ESO Kilo Degree Survey, and the 5000 degree² Dark Energy Survey in the optical, and the tiered UKIDSS near-IR surveys: the 12 degree² VIDEO survey, and the deep 1 degree² UltraVISTA survey in the near-IR; Erben et al. 2013; Arnaboldi et al. 2007; de Jong et al. 2013; Abbott et al. 2005; Jarvis et al. 2013; McCracken et al. 2012). These surveys measure the high mass end of the galaxy mass function at $0 < z < 4$ and also address questions such as the number density of bright Lyman break galaxies out to $z \sim 8$, the clustering of galaxies, the evolution of galaxy groups and clusters, the properties and demographics of AGNs, and the prevalence of short-lived events such as mergers.

Many of these science questions would benefit greatly from imaging at HST resolution. However, the widest/shallowest tiers of the wedding cake are devoid of HST imaging, particularly in the near-IR. The largest area imaged with HST in the optical is the 2 degree² COSMOS field, which was carried out with ACS in the I_{814} filter in Cycles 12 and 13 (Scoville et al. 2007) at a cost of 640 orbits. By contrast, the widest area imaged in the near-IR is an order of magnitude smaller. The five fields of the 900-orbit CANDELS survey, which used the J_{125} and H_{160} filters of the WFC3 camera in a Multi-Cycle Treasury program in Cycles 18, 19, and 20, add up to 0.25 degree².

The reason for the lack of very wide HST near-IR surveys seems obvious: the price of HST’s excellent resolution and small pixels is a small field of view, and a single WFC3 pointing covers a mere 4.6 arcmin². However, this explanation is not sufficient, as ground based surveys routinely cover thousands of times their detector area. As an example, the UKIDSS Large Area Survey (LAS) covers 4000 degree² in Y , J , H , and K using $\sim 20,000$ pointings with the 0.21 degree² WFCAM instrument. An HST/WFC3 survey with a similar strategy (that is, a similar number of pointings) would cover 25 degree². Furthermore, as the sensitivity of WFC3 is similar to that of a 30 m telescope on the ground, even short exposures reach depths comparable to the deepest ground-based surveys in existence: a 300 s exposure in H_{160} gives a 5σ point source sensitivity of $AB = 25.4$, comparable to the deepest ground-based surveys in existence.

The real limitation when designing wide-field programs is that HST has a natural lower limit to the exposure time per pointing, which stems from the time that is required for a guide star acquisition and other overheads associated with moving the telescope. The UKIDSS Large Area Survey has an exposure time of 40 s per band, and this would be hopelessly inefficient with HST: even if many guide star acquisitions were allowed in a single orbit (the limit is two), only four would fit and nearly the entire ~ 50 minutes of orbital visibility would be taken up with overhead. As a result, the natural lower limit to the exposure time is a single orbit,⁷ and the natural upper limit to the area of an HST survey is the number of orbits of the program multiplied by the detector area of the instrument.

2. “Drift and Shift”: A New Method for Wide-Field WFC3 Surveys

2.1. Overview

There is a way to circumvent the limitations imposed by the guide star acquisitions. If no new guide star is acquired between pointings the overheads decrease dramatically, and as we show below it is possible to fit eight distinct pointings, each with an exposure time of ~ 300 s, in a single orbit.

Operating without a guide star does not change the telescope control: ever since the last servicing mission the pointing of HST has been controlled by three gyros, and this is the case irrespective of whether a guide star is acquired or not. In a standard guided exposure the three gyros receive continuous corrections from the Fine Guidance Sensors (FGS), and turning off guiding merely stops the stream of corrections from the FGS. The effect of the lack of corrections is that the telescope begins to drift by an expected $0''.001 - 0''.002$ per second. The drift can be larger during orbits when the telescope experiences unusually strong atmospheric drag.⁸

The result of operating in a gyro-only mode is

⁷Note that it is possible to split the orbit into shorter observations with different filters; this reduces the per-filter observing time but does not increase the area of the survey.

⁸This is an expectation; there are no systematic measurements of this effect because all observations (with the exception of spatial scans for bright objects) are done under FGS control.




Fig. 1.— Illustration of the “drift and shift” (DASH) method of restoring unguided HST images. The top panel and inset show the standard data product (the FLT file) of an unguided, gyro-controlled exposure. The objects are smeared due to the lack of fine guidance sensor corrections. The middle panels show the twelve individual samples that comprise the exposure, created from the non-destructive reads. The smearing is small in each individual sample. The bottom panels show the reconstructed image after shifting the twelve samples to a common frame and adding them.

therefore that exposures longer than a minute are smeared. For CCD detectors, such images are scientifically unusable. When a guide star acquisition fails, or a guide star is lost during an observation, the visit is flagged and typically such observations are repeated. However, a property of the WFC3 detector (and IR detectors in general) is that an exposure is composed of multiple non-destructive, zero-overhead reads. The exposure time of an individual read can be set,⁹ and for times up to 25 s the drift *in between reads* is $\lesssim 0''.05$, or less than half a $0''.13$ pixel. The dataset obtained in the interval between two reads is simply the difference between read i and read $i - 1$. Therefore, an unguided, gyro-controlled, 300 s exposure with 25 s reads effectively consists of 12 independent, dithered exposures that can be shifted and combined into a full resolution image with hot pixels and cosmic rays removed. A schematic of this procedure is shown in Figure 1.

After the 300 s exposure the telescope can then be offset to a new position without the need to acquire a guide star at that position; as soon as the telescope move is completed the next integration can begin. As we detail below, we can observe eight distinct positions in a single orbit with this method.

2.2. Implementation: Structure of a Single Orbit

The optimal way to implement the “drift and shift” strategy is dictated by the amount of data that can be stored in memory during the exposure. We first consider the minimum time between reads during an exposure.¹⁰ With 10 s intervals between reads (SPARS10 mode) the buffer fills so quickly that memory dumps have to be conducted during the exposure, drastically reducing the observing efficiency and negating the benefits of the method. Fortunately, 25 s intervals (SPARS25 mode) are possible, and the typical drift in that time is still significantly smaller than a pixel. For our purposes, longer intervals do not provide significant benefits in terms of the total on-target exposure

time, reduce the number of independent samples that are available for cosmic ray rejection, and begin to show appreciable drift within each sample. However, for other scientific applications such longer intervals may be more advantageous.

Next we determine how many independent pointings we can fit in a single orbit, by varying the number of samples taken during a single exposure (i.e., the per-pointing exposure time) and the number of pointings. This calculation depends on the field that is observed, as that determines the length of the visibility window. We only consider the COSMOS field (at $RA=2^h15^m00^s$ and $Dec=10^\circ00'00''$), which is the target area of our GO-14144 program (see §3) and is typical for most of the sky.¹¹ The maximum number of pointings is eight: when a larger number is attempted (with a smaller number of samples taken during each exposure), buffer dumps cause the ninth pointing to spill over into a second orbit. The structure of the orbit is summarized in Table 2.1 and shown graphically in the left panel of Figure 2. The pointing pattern for this particular orbit (Visit 1 of GO-14144) is shown at right in Figure 2: the total covered area is $\sim 8\times$ greater than that of a single WFC3 pointing.

The exposure time per pointing is approximately the exposure time per sample multiplied by the number of samples. The number of samples is 11 (NSAMP=11) for the first four pointings and 12 (NSAMP=12) for the last four, however the first sample is read out at the very beginning of the exposure (at 2.9 seconds). The realized per-pointing exposure times are 253 s and 278 s respectively. The total number of independent 25-second observations obtained in this single orbit is $4 \times 10 + 4 \times 11 = 84$. The observing efficiency, expressed as the on-target realized exposure time divided by the total orbital visibility, is 66 %, similar to the typical efficiency of standard observing modes.

⁹To certain specified values; e.g., 10 s, 25 s, 50 s, 100 s, and 200 s.

¹⁰We used the Astronomer’s Proposal Tool Phase II software (<http://www.stsci.edu/hst/proposing/apt>) for this analysis, and implemented this strategy in our GO-14144 program described in §3.

¹¹It is possible to fit more pointings in continuous viewing zone observations, at the expense of an increased background in some of the pointings and possibly larger drift rates.

Orbit 1

Server Version: 20150609




Fig. 2.— Left: graphical representation of a single orbit (Visit 1 of GO-14144). Only the first pointing is guided (**PCS Mode=FINE**). Blue, dotted bars indicate science exposures. Black bars are buffer dumps. The total exposure time is 2124s, corresponding to 66 % of the total orbital visibility. Right: pattern of the eight WFC3 pointings observed in this orbit. The green lines illustrate the shifts relative to the first pointing in the orbit. When the observations are carried out we step from one pointing to the next, without returning to the starting position.

TABLE 1
STRUCTURE OF A SINGLE GO-14144 ORBIT

Step	Event	Changes to keywords	Duration	Exposure time
1	guide star acquisition	...	333 s	
2	guided exposure, position 1	PCS Mode=FINE SAMP-SEQ=SPARS25 NSAMP=11	295 s	253 s
3	stop FGS corrections	...	21 s	
4	offset to position 2	...	52 s	
5	unguided exposure, position 2	PCS Mode=GYRO	295 s	253 s
6	offset to position 3	...	50 s	
7	unguided exposure, position 3	...	295 s	253 s
8	offset to position 4	...	52 s	
9	unguided exposure, position 4	...	295 s	253 s
10	offset to position 5	...	50 s	
11	unguided exposure, position 5	NSAMP=12	320 s	278 s
12	offset to position 6	...	52 s	
13	unguided exposure, position 6	...	320 s	278 s
14	offset to position 7	...	50 s	
15	unguided exposure, position 7	...	320 s	278 s
16	offset to position 8	...	52 s	
17	unguided exposure, position 8	...	320 s	278 s
Unused orbital visibility:			38 s	
Total duration and exposure time:			3209 s	2124 s




Fig. 3.— Layout of the 456 pointings of the COSMOS-DASH program. The background image is a composite of the UltraVISTA deep H band imaging and the CANDELS H_{160} WFC3 imaging (the “hole” in the mosaic above the center of the middle stripe; [Grogin et al. 2011](#); [Koekemoer et al. 2011](#)). Small gaps between pointings are deliberate: they coincide with bright sources that would cause persistence problems.

3. A Wide-Field Survey of the COSMOS Field – COSMOS-DASH

We have undertaken a survey with the drift-and-shift technique during Cycle 23, adding a wide-field tier to the “wedding cake” of near-IR imaging surveys with HST. Covering 0.6 degree^2 , COSMOS-DASH (for “Drift And Shift”) more than triples the extragalactic survey area that HST has observed in the near-IR. Here we describe the observing strategy of this program (GO-14114); in § 4 we analyze the first four visits and demonstrate that the drift-and-shift technique produces the expected results.

The 57-orbit COSMOS-DASH program is targeting the COSMOS field (Scoville et al. 2007), as this is the only field with optical ACS imaging (in the I_{814} filter) over a sufficiently large contiguous area. The longest wavelength filter, H_{160} , is used to maximize the color baseline at HST resolution. We do not cover the entire 2 degree^2 that has ACS imaging but only the UltraVISTA deep stripes (McCracken et al. 2012); these regions have extremely deep complementary ground-based Y , J , H , and K imaging as well as deep Spitzer IRAC imaging from the SMUVS Exploration Science program (Spitzer GO-11016, PI: K. Caputi).

The layout of the $57 \times 8 = 456$ pointings is shown in Figure 3. The data are taken at two orientations to facilitate scheduling. The large gap in the mosaic in the central stripe is the area of the COSMOS field that was observed in J_{125} and H_{160} in the CANDELS Multi-Cycle program (Grogin et al. 2011; Koekemoer et al. 2011). We ensured that some pointings partially overlap with CANDELS in order to test our ability to recover the photometry and structural parameters of galaxies (see § 4). Small gaps in the mosaic are deliberate, and coincide with very bright stars or galaxies. Those objects would cause severe persistence, with the potential to affect the other pointings in an orbit as well as consecutive orbits. The total area covered by GO-14114 is 0.59 degree^2 , which can be compared to the existing total five-field CANDELS area of 0.24 degree^2 . GO-14114 therefore increases the survey area that has WFC3 imaging in addition to deep ground- and Spitzer data by a factor of 3.5.

The depth that will be achieved depends on the realized drift rate. If the drift during a point-

ing is $\lesssim 1$ pixel the limiting magnitudes will be identical to those of regular 253s and 278s exposures. With the average zodiacal background in the COSMOS field the 5σ point source limit will then be $H_{160} = 25.1$ (on the AB system). If the drift rate is several pixels the point source depth will be slightly lower, as the shifts place different independent pixels onto the same output pixel; a conservative expectation is $H_{160} \approx 24.9$ in those circumstances. As we show in the next Section the drifts in the first four orbits of our program are of order 1 pixel per pointing.

Figure 4 places COSMOS-DASH in the context of other WFC3 H_{F160W} observations, in terms of their depth and covered area. We probe a regime that was previously unexplored with space-based IR imaging, despite the fact that the orbit-total of GO-14114 is an order of magnitude smaller in terms of time allocation compared to, for example, COSMOS and CANDELS.




Fig. 4.— Area vs. depth for WFC3/IR F160W imaging. The technique described in this paper as implemented in GO-14114 (red point) allows us to add a wide shallow tier to infrared surveys with HST. Furthermore, the orbit total of GO-14114 (57) is an order of magnitude smaller than that of the other surveys shown in this Figure.

4. Early Data

We observed four orbits of GO-14114 on UT 2015 October 15 and 16. The decision to obtain these data was a trade-off: these early data allow

us to assess the viability of the methodology, but they come at a price. The COSMOS field has a strongly variable near-IR background throughout the year, as there is a relatively large amount of zodiacal dust in its direction. The background in October is so high that the depth of the data is compromised; as we detail below these four orbits are ~ 0.6 magnitudes shallower than the rest of the survey will be, when taken at lower background levels. The four orbits (32 pointings) cover the top area of the middle ULTRAVISTA deep stripe, North of the CANDELS field.

4.1. Reduction and Analysis

We downloaded the raw and calibrated images from the Mikulski Archive for Space Telescopes (MAST¹²). The images were processed on the fly with the best available calibration by the `calwfc3` pipeline. In our reductions we make use of both the flat-fielded final pipeline outputs (FLT) and the calibrated intermediate MultiAccum exposures (IMA). Each orbit consists of one guided and seven unguided pointings. These two types of pointings are processed through similar steps, but we note the differences where necessary.

We first carry out a basic reduction of the FLT, similar to that described in Skelton et al. (2014). We mask all sources in the image and subtract a second order polynomial fit to the background. We then align the image to an external reference world coordinate system using the `tweakreg` task. We use the I_{F814W} images¹³ and catalogs provided by the COSMOS collaboration as a reference (Koekemoer et al. 2007; Massey et al. 2010). As in Skelton et al. (2014), we use all sources in the image for alignment, not only stars. Even though the FLT images are slightly smeared, we prefer to use them for this rough alignment because the sources have higher signal to noise. A fine alignment step is performed later in the reduction. The aligned FLTs are drizzled together to create a preliminary mosaic for the orbit. We run `SExtractor` (Bertin & Arnouts 1996) on this mosaic to create a segmentation map which is used as a mask of sources in later steps of the reduction.

The `tweakreg` task provides the offsets of each

of the 8 exposures from the commanded pointing position. Generally, the relative pointing accuracy of *HST* is very high, and the RMS precision of offsets within an orbit is ~ 2 to 5 milli-arcsec with fine lock on two guide stars and when the offsets are small (Gonzaga et al. 2012). In the absence of guide stars, we find that the unguided exposures show larger relative pointing errors, as large as $8''$. The left panel of Figure 5 shows the offsets from the commanded position as a function of time since the start of the orbit. The offsets increase monotonically with time as can be expected from the build-up of uncompensated forces on the spacecraft during the orbit. Furthermore, sudden jumps in pointing accuracy of $\sim 4''$ can be seen for the last position in COSMOS-15, 16 and 17. In orbit COSMOS-18 the jump happens after the fourth position. The cause for these larger offsets is not yet understood. The offsets are small relative to the instrument FOV (10% of the WFC3 FOV); however users should make sure that there is sufficient overlap between pointings to compensate for this effect if a contiguous mosaic is needed.

Next, we proceed to reconstruct unsmearred images from the non-destructive reads. These reads are preserved in the IMA images: these are multi-extension files which contain all the individual reads from the original exposure. Each read has been bias-, overscan- and dark-current-subtracted as well as flat-fielded. The up-the-ramp fits have also been carried out, flagging cosmic rays. The IMA images are in units of electrons per second. For each read i , starting with the second, we subtract the preceding read, $i - 1$. These difference images are treated as the science arrays in our analysis. We construct an error array based on the read noise, Poisson noise and flat field uncertainties. The data quality (DQ) array is taken directly from the IMA. These three arrays (science, error, DQ) are paired with the header from the aligned FLT image and saved as a new file. In this manner, each pointing is split into 10 or 11 new 25-second images, depending on the number of reads. Below we refer to the set of 10 or 11 such images produced from a single original image as a “sequence”. Guided and unguided exposures are treated in the same manner.

We mask all sources (using the masks created from the FLT files) and perform a median background subtraction for each of these new images.

¹²<http://archive.stsci.edu>

¹³http://irsa.ipac.caltech.edu/data/COSMOS/images/acs_mosaic_2.0/tiles/




Fig. 5.— *Left*: Offsets relative to the commanded position as a function of time for the eight exposures in each of the four observed orbits. The increase with time is expected, however the source of the large offsets at the end of the orbit (and halfway through COSMOS-18) is not yet understood. *Right*: Drift rate during the pointing in pixels per 25 seconds as a function of time. Following the first guided pointing (no drift), the drift rate increases with time.

As noted above, the zodiacal background in the data is very high, at $\approx 2.2 \text{ e}^-/\text{s}$. Typical values of the background in the CANDELS observations of the COSMOS field range from $0.6 \text{ e}^-/\text{s}$ to $0.8 \text{ e}^-/\text{s}$. For background levels below $0.9 \text{ e}^-/\text{s}$, a 250 second exposure would be read-noise limited ($\text{RN} \sim 15 \text{ e}^-$), however, the current observations are background-limited.

The treatment of cosmic rays deserves careful attention. Cosmic rays are flagged by the up-the-ramp fits of the `calwfc3` pipeline. However, there are two issues with this procedure. First, due to the telescope drifts, objects move across the detector during the exposure. This change in flux from one read to the next causes the up-the-ramp fits to flag the varying pixels as cosmic rays. This issue primarily affects point sources, where the size of the source is commensurate with the drift during the exposure. To remedy this, we reset all cosmic ray flags within the boundaries of objects (as identified by the segmentation map). This has the obvious negative effect of resetting the flags for cosmic rays that fall on top of objects. The second issue is that cosmic rays which occur between the zeroth read (at 2.9 seconds) and the first read (at 27.9 seconds) cannot be identified by

the ramp fits. While the first issue affects only unguided pointings, the second affects all pointings. In order to identify those cosmic rays, and also to flag any cosmic rays which were erroneously reset in the previous step, we use a combination of L.A. Cosmic (van Dokkum 2001) (run on the FLT) and `astrodrizzle` (run when combining the individual reads; Gonzaga et al. 2012). The 4096 flag is added to the DQ array for all pixels identified by both methods. We visually verified that our adopted procedure correctly identifies nearly all cosmic rays and leaves the central pixels of compact objects intact.

The images in all non-guided sequences are aligned to the COSMOS I_{F814W} mosaic. The alignment is performed with the `tweakreg` task in `astrodrizzle` (Gonzaga et al. 2012). Again, we use all sources within the image to calculate the offsets. Here we only solve for x and y shifts and do not fit for rotation or change in scale. The typical rms in the differences of matched positions is ~ 0.4 native pixels, identical to that achieved when aligning the much deeper CANDELS COSMOS images (Skelton et al. 2014). For each sequence we calculate the offset between the first and the last read to determine the mean rate of




Fig. 7.— The COSMOS-DASH H_{F160W} mosaic. The insets in the top right corner show a portion of the mosaic in pointing 8 of COSMOS-15 before and after the alignment of the individual reads. The drift is evident in compact objects but overall still very small. The inset in the lower right part of the image shows a cutout of the mosaic compared to the same portion of the CANDELS COSMOS mosaic (Grogin et al. 2011; Koekemoer et al. 2011; Skelton et al. 2014).




Fig. 8.— SExtractor’s FLUX_RADIUS in pixels vs. MAG_AUTO H_{F160W} (black points) for the COSMOS-DASH mosaic (left) and ratio in the fluxes measured in two different apertures (2’’ and 0.5’’) vs. MAG_AUTO H_{F160W} (right). Stars that are used to construct the PSF are marked with red points (see text). For comparison we also show the distribution of points in the CANDELS COSMOS mosaic of Skelton et al. (2014, gray points), accounting for the different pixel scales. The stellar sequence in the COSMOS-DASH mosaic is in the same location and as tight as the guided CANDELS COSMOS mosaic.




Fig. 6.— Distribution of measured drift rates for the GO-14114 observations (solid red histogram) compared to those for 87 archival WFC3 exposures where guiding failed (dotted black histogram). The drift rates are shown in units of pixels per 25 seconds, which measures the drift between reads in our program. The expected drift based on engineering predictions is up to 0.002 per second (gray region) and our observations fall well within that. The pointing shown in Figure 1 is indicated with an arrow.

drift. In the right panel of Figure 5 we show the drift rate as a function of time for each of the four orbits. The drift rate increases with time, as expected, but stays within 0.002 per second, or 0.42 pixels per 25 seconds. In Figure 6 we show the distribution of drift rates for all unguided sequences (solid red line) compared to archival WFC3 IR observations that were taken in gyro-only control (dotted black line). The mean drift rate in our observations is 0.001 per second or 0.2 pixels per 25 seconds, smaller than that of the archival observations. It should be noted that in the archival data the switch to gyro-control was not deliberate, but the result of the failure to acquire (or the loss of) a guide star. Our data suggest that when the switch to gyro control is done deliberately, the drift rates are very small and within the expected range.

Final mosaics are produced with *astrodrizzle* using all sequence images. We use exposure time weighting, a square kernel and $pixfrac = 0.8$. We drizzle the images to a pixel scale of 0.1’’. The final mosaic of the current observations is 9100×10200 pixels, centered at RA=10 : 00 : 25.4, DEC=+2 : 34 : 51.2. In Figure 7 we show the full mosaic. We also show a zoomed-in cut-out (top right) to demonstrate the difference between the

original drifted image and the final processed image. The example shown is from the eighth pointing of COSMOS-15, i.e., one of the pointings with the largest drift rate (Figure 5). Another set of cut-outs (bottom right) compare our mosaic to that from the CANDELS observations. The mosaic is of high quality, with no obvious problems or defects: it looks like a shallow version of the CANDELS imaging, as intended. In the following subsections we characterize the data quality.

4.2. Point Spread Function

In order to assess the quality of the reconstructed data, we first turn to the point spread function (PSF). A key question is whether the fact that the exposures were unguided has led to a net broadening of the PSF compared to guided exposures.

We construct the PSF directly from the H_{F160W} mosaic using stars in the image. We run **SExtractor** on the final mosaic to detect objects and perform photometry. Fluxes are measured in a series of apertures with different radii. In the left panel of Figure 8 we show the **SExtractor** FLUX_RADIUS (the radius of the circle centered on the barycenter that encloses about half of the total flux) vs. MAG_AUTO (a total magnitude measurement). Point sources follow a tight sequence with small sizes at all magnitudes and match well the stellar sequence in the CANDELS COSMOS catalog of Skelton et al. (2014, gray points). Points that scatter above the stellar sequence at bright magnitudes are saturated stars. In the right panel of Figure 8 we show the ratio of the flux in a large aperture ($2''.0$) to that in a small aperture ($0''.5$) as a function of magnitude which shows a similar tight sequence, which also matches the CANDELS COSMOS catalog of Skelton et al. (2014). The median FWHM of the stars in these sequences is $0''.21$ (2.1 pixels) compared to $0''.19$ (3.2 pixels at $0''.06/\text{pix}$) from the CANDELS COSMOS mosaic (Skelton et al. 2014), which is fully guided, indicating that we are recovering the original resolution. The sequences in both panels of Figure 8 are useful diagnostics of the image quality such that large spreads or offsets in the stellar sequences can indicate problems with alignment or clipping of the cores of stars.

For a quantitative comparison, we follow the PSF construction method of Skelton et al. (2014).

Stars are selected based the stellar sequence in the right panel of Figure 8 such that the ratio between the fluxes is $1.1 < f(2''.0)/f(0''.5) < 1.2$ (red points). We visually inspect all selected stars and exclude 21 objects which are either too close to the edge of the mosaic or have close neighbors. Despite our effort, some stars still have suppressed weight at the center of their weight maps. We exclude 39 such stars, leaving a final sample of 106 objects with $H_{F160W} < 22$. Stars are distributed evenly across the field and, since only 4 of the 32 pointings are guided, the final PSF will be dominated by the unguided exposures. We mask neighboring objects within a postage stamp cut-out of 84 pixels around each star. The postage stamps are recentered, normalized and then averaged to determine the PSF and the PSF weight map (Figure 9).




Fig. 9.— Point spread function (PSF, left) constructed from stars within the COSMOS-DASH mosaic and the corresponding combined weight map (right). Both the PSF and the weight map have been normalized by the peak value. The selection of stars is shown in Figure 8. 106 stars with $H_{F160W} < 22$ were used to create the PSF.

The curve of growth, which is the fraction of enclosed light as a function of aperture size (normalized at $2''.0$), is shown in the top panel of Figure 10 (solid line). For comparison, we also show the H_{F160W} PSF derived from the CANDELS COSMOS observations (Skelton et al. 2014). The two growth curves show excellent overall consistency. A quantitative measure of the consistency is the ratio between the two curves, which is shown in the bottom panel of Figure 10. The COSMOS-DASH PSF has $\sim 15\%$ less energy in the central pixel (for a pixel scale of $0''.1$ per pixel), however within the typical aperture used for photometry, $0''.35$ (solid vertical line), and out to $2''.0$ the dif-




Fig. 10.— Growth curve of the point spread function. *Top:* The fraction of light enclosed as a function of radius relative to the total light within $r = 2''$ (solid curve). For comparison we also show the growth curve for the COSMOS CANDELS field constructed by Skelton et al. (2014). *Bottom:* Ratio between the COSMOS-DASH and the Skelton et al. (2014) growth curves as a function of radius. At the typical aperture used for photometry ($0''.35$, vertical line) the differences between the two PSFs are $\sim 1\%$.

ferences between the two PSFs are $\lesssim 1\%$. This demonstrates that by aligning the individual reads we have recovered the resolution of guided *HST* images.

4.3. Noise Characteristics

Here we analyze the noise characteristics of the reconstructed images. As a comparison sample we obtained from MAST all other observations of extragalactic fields carried out with the SPARS25 sampling mode and matching the exposure times of our observations. In total we find a total of 39 public datasets, containing 107 exposures (FLT) that match these criteria. We process these datasets in a manner identical to the COSMOS-DASH observations: we split them into individual reads and drizzle those to a final distortion corrected image (no alignment is necessary). In this analysis we use the native pixel scale of $0''.12$ per pixel.

For each dataset we measure the pixel-to-pixel noise in the original FLT) and in the final drizzled image, masking sources and clipping the image edges. In the top left panel of Figure 11 we compare the noise in the FLT) to the drizzled images. As expected the scatter is lower in the drizzled images: drizzling artificially suppresses the noise in the images as a result of the rebinning of flux (Fruchter, Sosey, et al. 2009, §3.3.1). The COSMOS-DASH observations (red symbols) have higher noise overall as a result of the elevated background levels (see §4.1). The four *guided* pointings (red open symbols) follow the trend of the comparison sample (broken line): they behave in the same way as other guided data, except with higher noise due to the increased background. The difference is broadly consistent with the difference in the background level: the zodiacal background is a factor of ~ 3 higher than for typical exposures, which translates into a factor of $\sqrt{3}$ higher noise.

However, the drizzled *unguided* pointings (filled red symbols) exhibit higher noise than the drizzled guided images. On average, the noise in the unguided pointings is 15% higher than in the unguided ones (and up to 28% higher). This corresponds to a loss in depth of 0.15 (and up to 0.27) magnitudes. The reason for this behavior is that, as a result of the shifts within each pointing, a larger number of independent pixels contribute to a single drizzled output pixel.

This loss of depth due to the drift-and-shift method depends on the spatial scale. It is rare that one is interested in individual pixels; typical aperture photometry is performed over scales of several pixels, to match the size of the PSF or to match the size of spatially-extended objects. As an example, the aperture used in the Skelton et al. (2014) catalogs has a diameter of $0.7''$. To assess the noise increase on larger scales we rebinned the FLT and DRZ images to coarser grids and re-measured the pixel-to-pixel variation. In the 2×2 binned images (Figure 11, second panel), the noise in the unguided exposures is on average 5% higher than that in the unguided exposures. In the 3×3 and 5×5 binned images the noise is only $\sim 3\%$ higher. We conclude that, in “real world” applications, the loss in depth due to our method is only $\sim 0.03 - 0.07$ magnitudes.

Finally, we stress that *for these particular orbits* the depth is limited by the high zodiacal background, which results from our decision to schedule them early in the Cycle. A full analysis of the depth of the COSMOS-DASH data, including a completeness analysis, will be performed when all data are taken in Fall 2016.

4.4. Galaxy Structural Parameters

A small section of the mosaic of 32 pointings overlaps with existing H_{160} data from the CANDELS survey. We measure structural parameters of objects in this overlap region from our mosaic, and compare them to measurements of the same objects in CANDELS by van der Wel et al. (2012). The same methodology is used here, and we refer to van der Wel et al. (2014) for the details. Briefly, GALFIT (Peng et al. 2002) is used to fit Sersic profiles (Sersic 1968) to the images, using the GALAPAGOS wrapper (Barden et al. 2012). The local background is not a free parameter in the fit, but determined in an initial step.

For the comparison we selected objects brighter than $H_{F160W} = 22$ and excluded objects with `use_phot = 0` which removes edge objects and stars (Skelton et al. 2014). We further excluded galaxies with uncertain GALFIT fits ($f \geq 2$) in either catalog, corresponding to bad fits and no fits. The two catalogs are crossmatched based on position with a tolerance of $0''.5$. No attempt is made to match the segmentation maps. The final sample consists of 48 objects.

Figure 12 compares the structural parameters measured in COSMOS-DASH to those measured in CANDELS. Overall there is excellent agreement. The total magnitudes are offset by 0.05, with very small scatter. The median offset in axis ratios is 0.004, again with almost no scatter ($\sigma \sim 0.02$). There is a small systematic offset in sizes, which correlates with a small systematic offset in Sersic indices. This offset is negligible for galaxies with low Sersic indices and/or small sizes. The median size difference for all objects is 0.04 dex or 10%, with objects in COSMOS-DASH larger than in CANDELS, and the scatter is 0.05 dex. It is not clear whether this small difference is caused by a small remaining error in our PSF, issues with the background subtraction, or other effects. There is also the (perhaps unlikely) possibility that CANDELS slightly underestimates the sizes. Further analysis of these effects is limited by the high background in these four orbits and the small number of stars that are suitable for PSF reconstruction; we expect to explore this further when the program is completed. Based on the data we have in hand, we conclude that sizes can be estimated with an accuracy of $\sim 10\%$.

5. Conclusions

In this paper we show that HST can obtain wide-field near-IR data in a relatively small number of orbits. The “drift and shift” (DASH) technique that makes this possible is well understood, and has now been demonstrated to produce science-grade data at the full resolution of the WFC3 camera. The low near-IR background from space means that such wide-field surveys reach depths that are competitive with the deepest ground-based surveys, despite the short ~ 300 s per-pointing integration times.

Our medium-sized, 57-orbit Cycle 23 program covers an area more typically associated with a Treasury survey. Were our technique applied in an actual Treasury program, one could cover an area of five square degrees in approximately 500 orbits. This opens up regimes previously regarded as out of Hubble’s reach, both for “blank” extragalactic surveys and for specific targets such as M31 and the Magellanic Clouds.

We thank Merle Reinhard for advice on the




Fig. 11.— Comparison between the pixel-to-pixel scatter in the noise of the FLT images, σ_{FLT} , and the noise of the drizzled individual reads, σ_{DRZ} . The left-most panel shows the noise at the native pixels of the image, while the remaining panels show how the noise changes when we rebin the images in coarser grids (2×2 , 3×3 and 5×5). The guided pointings (red open circles) have noise characteristics in agreement with analogous archival observations (gray symbols, dotted line fit), while the drizzled un-guided pointings have higher noise.




Fig. 12.— Comparison between CANDELS and COSMOS-DASH structural parameters, as determined with GALFIT. Overall there is good agreement. The sizes and Sersic indices show a small systematic effect, with the sizes in COSMOS-DASH larger than those in CANDELS by 10% on average (see text).

observation planning and Katherine Whitaker for help with the data reduction. This paper is based on observations made 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, Inc., under NASA contract NAS 5-26555. These observations are associated with program GO-14114. Support for GO-14114 is gratefully acknowledged.

REFERENCES

- Abbott, T., et al. 2005, arXiv:astro-ph/0510346
- Arnaboldi, M., Neeser, M. J., Parker, L. C., et al. 2007, *The Messenger*, 127, 28
- Barden, M., Häußler, B., Peng, C. Y., McIntosh, D. H., & Guo, Y. 2012, *MNRAS*, 422, 449
- Bertin, E., & Arnouts, S. 1996, *A&AS*, 117, 393
- van Dokkum, P. G. 2001, *PASP*, 113, 1420
- Ellis, R. S., McLure, R. J., Dunlop, J. S., et al. 2013, *ApJ*, 763, L7
- Erben, T., Hildebrandt, H., Miller, L., et al. 2013, *MNRAS*, 433, 2545
- Fruchter, A. and Sosey, M. et al. 2009, “The MultiDrizzle Handbook”, version 3.0, (Baltimore, STScI)
- Giavalisco, M., Ferguson, H. C., Koekemoer, A. M., et al. 2004, *ApJ*, 600, L93

- Gonzaga, S., Hack W., Fruchter, A., and Mack, J. et al. 2012, "The Drizzlepac Handbook", version 1.0, (Baltimore, STScI)
- Grogin, N. A., Kocevski, D. D., Faber, S. M., et al. 2011, ApJS, 197, 35
- Illingworth, G. D., Magee, D., Oesch, P. A., et al. 2013, ApJS, 209, 6
- Jarvis, M. J., Bonfield, D. G., Bruce, V. A., et al. 2013, MNRAS, 428, 1281
- de Jong, J. T. A., Verdoes Kleijn, G. A., Kuijken, K. H., & Valentijn, E. A. 2013, Experimental Astronomy, 35, 25
- Koekemoer, A. M., Aussel, H., Calzetti, D., et al. 2007, ApJS, 172, 196
- Koekemoer, A. M., Faber, S. M., Ferguson, H. C., et al. 2011, ApJS, 197, 36
- Lotz, J., Mountain, M., Grogin, N. A., et al. 2014, American Astronomical Society Meeting Abstracts #223, 223, #254.01
- Massey, R., Stoughton, C., Leauthaud, A., et al. 2010, MNRAS, 401, 371
- McCracken, H. J., Milvang-Jensen, B., Dunlop, J., et al. 2012, A&A, 544, A156
- Peng, C. Y., Ho, L. C., Impey, C. D., & Rix, H.-W. 2002, AJ, 124, 266
- Rix, H.-W., Barden, M., Beckwith, S. V. W., et al. 2004, ApJS, 152, 163
- Scoville, N., Aussel, H., Brusa, M., Capak, P., Carollo, C. M., Elvis, M., Giavalisco, M., Guzzo, L., Hasinger, G., Impey, C., Kneib, J.-P., LeFevre, O., Lilly, S. J., Mobasher, B., Renzini, A., Rich, R. M., Sanders, D. B., Schinnerer, E., Schminovich, D., Shopbell, P., Taniguchi, Y., & Tyson, N. D. 2007, ApJS, 172, 1
- Sersic, J. L. 1968, Cordoba, Argentina: Observatorio Astronomico, 1968,
- Skelton, R. E., Whitaker, K. E., Momcheva, I. G., et al. 2014, ApJS, 214, 24
- van der Wel, A., Bell, E. F., Häussler, B., et al. 2012, ApJS, 203, 24
- van der Wel, A., Franx, M., van Dokkum, P. G., Skelton, R. E., Momcheva, I. G., et al., 2014, ApJ, 788, 28
- Williams, R. E., Blacker, B., Dickinson, M., et al. 1996, AJ, 112, 1335