Open Research Online
The Open University's repository of research publications and other research outputs

An improved model of charge transfer inefficiency and correction algorithm for the *Hubble Space Telescope*

Journal Item

oro.open.ac.uk

An improved model of charge transfer inefficiency and correction algorithm for the *Hubble Space Telescope*

Richard Massey,[1]⋆ Tim Schrabback,[2] Oliver Cordes,[2] Ole Marggraf,[2] Holger Israel,[1] Lance Miller,[3] David Hall,[4] Mark Cropper,[5] Thibaut Prod'homme[6] and Sami Matias Niemi[5]

[1]*Institute for Computational Cosmology, Durham University, South Road, Durham DH1 3LE, UK*
[2]*Argelander-Institut für Astronomie, Auf dem Hügel 71, D-53121 Bonn, Germany*
[3]*Department of Physics, University of Oxford, The Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, UK*
[4]*Centre for Electronic Imaging, Department of Physical Sciences, The Open University, Walton Hall, Milton Keynes MK7 6AA, UK*
[5]*Mullard Space Science Laboratory, University College London, Holmbury St Mary, Dorking, Surrey RH5 6NT, UK*
[6]*European Space Agency, ESTEC, Postbus 299, NL-2200 AG Noordwijk, the Netherlands*

ABSTRACT

Charge-coupled device (CCD) detectors, widely used to obtain digital imaging, can be damaged by high energy radiation. Degraded images appear blurred, because of an effect known as Charge Transfer Inefficiency (CTI), which trails bright objects as the image is read out. It is often possible to correct most of the trailing during post-processing, by moving flux back to where it belongs. We compare several popular algorithms for this: quantifying the effect of their physical assumptions and tradeoffs between speed and accuracy. We combine their best elements to construct a more accurate model of damaged CCDs in the *Hubble Space Telescope*'s *Advanced Camera for Surveys/Wide Field Channel*, and update it using data up to early 2013. Our algorithm now corrects 98 per cent of CTI trailing in science exposures, a substantial improvement over previous work. Further progress will be fundamentally limited by the presence of read noise. Read noise is added after charge transfer so does not get trailed – but it is incorrectly *untrailed* during post-processing.

Key words: instrumentation: detectors – methods: data analysis – space vehicles: instruments.

1 INTRODUCTION

The harsh radiation environment above the Earth's atmosphere, in particle accelerators, or in medical contexts, gradually degrades Charge-coupled device (CCD) imaging detectors. CCD detectors work by converting photons to electrons inside a silicon lattice, then collecting the electrons in electrostatic potential wells that form each pixel. At the end of an exposure, the photoelectrons are transferred through a chain of pixels to the edge of the device, where they are amplified for external counting. However, radiation damage to the silicon lattice creates charge traps that capture electrons for a short time. When electrons are delayed during their transfer to the amplifier, the effect is known as 'charge transfer inefficiency' (CTI). Those electrons emerge several pixels later, as spurious trails behind every source (Holland et al. 1990; Janesick 2001). Different 'species' of traps (different configurations of the damaged silicon lattice) capture electrons for different lengths of time, resulting in a variety or complex superposition of trail profiles.

CTI trailing is particularly troublesome because the amount of flux trailed is a non-linear function of the flux, size and shape of a source, and the previous illumination history. The effect on images is therefore *not* a simple shape transformation such as might be described by a convolution. To first order, this can be roughly understood by considering that there are only a finite number of charge traps, so the *fraction* of electrons trailed from faint sources will be greater than that from bright sources. The process of adding CTI trails can be reproduced in software that models the flow of electrons past the charge traps.

Since trailing is (almost) the last process to happen during data acquisition, it can be corrected during the first stages of data analysis. Bristow (2003) pioneered an iterative method to reverse the effect of CCD readout in images from the *Hubble Space Telescope* (*HST*). With this technique, CTI trailing can be removed by repeatedly running the software to *add* CTI trailing. Massey et al. (2010b) used this technique

---

⋆ E-mail: r.j.massey@durham.ac.uk

to correct 90 per cent of the trailing in *HST* Advanced Camera for Surveys/Wide Field Channel (ACS/WFC) data from 2006, and the approach was incorporated into the STSCI CALACS software by Anderson & Bedin (2010). As the radiation damage accumulated, the trailing increased and became easier to measure. Massey (2010) updated the model and achieved a 95 per cent correction in data from 2010.

Although most of the trailing can now be removed, the residual will still limit high precision science with *HST*. There are many examples where measurements of photometry, astrometry and morphology are required to better than 5 per cent (e.g. Capak et al. 2007; Anderson & van der Marel 2010; Ragozzine et al. 2012; Suzuki et al. 2012; Ghez et al. 2013; Kitching et al. 2013) – and *HST*'s detectors continue to degrade. The same effect will also limit the European Space Agency (ESA)'s future missions *Gaia* (Perryman et al. 2001) and *Euclid* (Laureijs et al. 2013). For example, *Euclid* will require 99 per cent correction (Cropper et al. 2013; Massey et al. 2013). In this paper, we attempt to build a new model for *HST* ACS/WFC that meets this level of accuracy.

This paper is organized as follows. In Section 2, we study the details of three popular algorithms to correct CTI, comparing their physical assumptions and their approximations made for speed. In Section 3, we measure relevant properties of the *HST* ACS/WFC detectors using real data, and quantitatively compare the algorithmic options using simulated data (where the truth is known). In Section 4, we identify suitable algorithmic choices to achieve high accuracy in a minimum of time, and test our improved correction on real data in Section 5. In Section 6, we summarize lessons learned, and discuss possible avenues for future work.

2 DETAILED CODE COMPARISON

CTI correction splits naturally into two tasks. It is first necessary to be able to add trails by mimicking the hardware process that happens during readout (Section 2.1). Attempting to then undo the trailing is a separate process (Section 2.2). We shall compare approaches taken in the literature, attempting to understand their unique advantages and disadvantages – and aiming to combine their best practices.

2.1 Adding CTI trailing

In recent years, a series of ever-more sophisticated algorithms have been developed to model charge trapping in CCD readout. Three codes are now in widespread use, each adopting slightly different model assumptions. The Massey et al. (2010b) and Anderson & Bedin (2010) codes were developed for *HST* and generally aim for accuracy at the expense of computation speed. They share a common heritage, but considerable development has been put in since they forked. The Short et al. (2010, 2013) code was developed independently for *Gaia*, and is built for speed.

All three algorithms share the same basic framework for adding CTI trails, which is illustrated for parallel CTI[1] in Fig. 1 . The different choices made for various parameters are listed in Table 1, but the independent codes share some remarkable similarities. For example, while in hardware, electrons are transferred down through traps to a readout node, in software all the algorithms scan traps up through the electrons. This change of 'inertial reference frame' saves computational overheads by keeping the main image array static.

The core components of each method are two functions: one controlling which traps capture an electron (box 3), and one controlling when the electrons are released (box 5). Fundamental differences also arise in the accounting practices adopted to monitor which traps are full at any time, and in approximations used to speed up the inner loop (boxes 3–5). We shall discuss each of these steps in more detail.

2.1.1 Initializing charge traps (box 2)

To model parallel CTI, the charge traps are generally created empty. After a long, static integration on sky, most traps will have had ample opportunity to release electrons. In a few regions of the CCD, a steady state may have been reached in which some traps are full. However, the first operation during a readout algorithm is for traps to capture electrons. This sets occupancy to the desired levels before readout continues, rendering as immaterial their occupancy during integration (if capture is instantaneous, which is assumed for volume-driven models – see below).

One notable exception to this is the *Gaia* observations modelled by Short et al. (2013). These observations are performed in 'time-delay integration' mode, in which the shutter is left open and electrons are continuously read out at the same rate as the telescope is slewed. When initializing traps, it is therefore necessary to consider nearby sources that may have recently transited the focal plane and pre-filled traps (Prod'homme et al. 2012) (this is exacerbated if capture is not instantaneous, as assumed in density-driven models – see below).

2.1.2 Capture of electrons into charge traps (box 3)

To determine which traps can capture an electron, we start by considering the location of the traps and the electrons.

---

[1] To add serial CTI, Rhodes et al. (2010) rotated the image by 90° and rerunning the same code with a different trapping model. This is not perfectly accurate, because parallel CTI codes generally do not model the illumination history from the (long ago) previous exposure. The history relevant for serial CTI is the trap occupancy from readout of the row below, which passed through the serial register *immediately* before. Serial CTI in *HST*/ACS is very small, and we do not consider it in this paper.
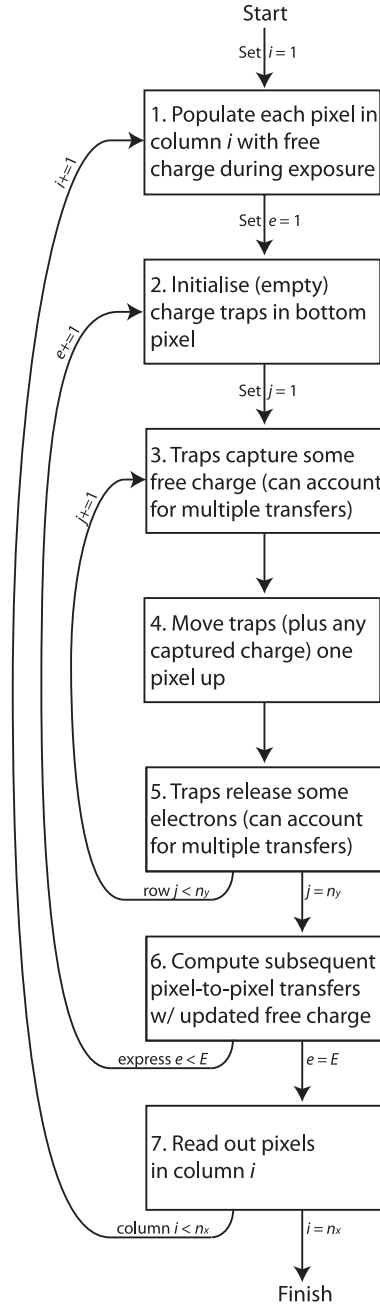
Start

Set $i = 1$

1. Populate each pixel in column i with free charge during exposure

Set $e = 1$

2. Initialise (empty) charge traps in bottom pixel

Set $j = 1$

3. Traps capture some free charge (can account for multiple transfers)

4. Move traps (plus any captured charge) one pixel up

5. Traps release some electrons (can account for multiple transfers)

row $j < n_y$ $j = n_y$

$j += 1$

6. Compute subsequent pixel-to-pixel transfers w/ updated free charge

express $e < E$ $e = E$

$e += 1$

7. Read out pixels in column i

column $i < n_x$ $i = n_x$

$i += 1$

Finish

Figure 1. A generic forward algorithm to add CTI trailing to an image, mimicking the effects of radiation-damaged hardware. The number of pixel columns and rows is $n_x \times n_y$, while the parameter E governs a tradeoff between speed and accuracy that is discussed in Section 2.1.5. All CTI correction algorithms currently in operation on astronomical telescopes follow the same framework.

Charge traps of all species are scattered throughout a CCD's 3D silicon lattice. In-orbit measurements show that each species of traps in ACS/WFC have uniform density ρ_t (Massey et al. 2010b, section 2.5), and are not significantly clustered (Ogaz et al. 2013 show column-to-column shot noise in ρ_t that appears Poissonian).

The distribution of electrons $\rho_e(\boldsymbol{x})$ is more complex. For storage and transportation, electrons are confined to a 'buried channel' (BC) region of the silicon lattice by an electric field, which can be shaped by doping the silicon with other elements during manufacture. Information about the precise doping structures is industrially sensitive, but Seabroke et al. (2011) and Clarke et al. (2012a,b) use ATLAS semiconductor device simulation software (Silvaco Inc. 2010) to model the geometry and volume of a cloud of n_e electrons in generic devices (top panel in Fig. 2). The volume $V_e(n_e)$ within any electron density threshold typically grows as $V_e \propto n_e^\beta + c$ (Hall et al. 2012), and converges to the geometric volume of (the BC in) a pixel V_{pix} as $n_e \to w$, the full well depth (bottom-left panel in Fig. 2). The constant c prevents the volume of a cloud of electrons tending to unphysically small values for low n_e.

During each clock cycle (pixel-to-pixel transfer) of CCD readout, the electrons are held stationary for a time t_{dwell}. During this time, the probability $p_{\mathrm{capture}}(\boldsymbol{x})$ that a trap at position $\boldsymbol{x}$ captures an electron, depends upon $\rho_e(\boldsymbol{x})$ and the capture cross-section of the trap σ_t (Hall 1952;

Table 1. Summary of the parameter choices made for various forward algorithms that add CTI trailing to an image. The core framework shared by all algorithms is to: capture electrons in traps, move the electrons (or traps), then release some electrons. Different approximations have been made in these three steps, and in accounting practices to monitor which traps are occupied. Note that, where numbers are >1 or <n_y (the number of rows of pixels in a CCD), they are user configurable and have been varied to fit the CCD in question.

Algorithm	Trap initialization	Trap capture	Trap monitoring $n_{\text{species}} \times n_{\text{levels}}$	Transfer many at once ($P = n_y/E$)	Transfer thro' gaps (n_{phases})	Trap release	Serial transfer
Bristow (2003)	all empty	β, d	$3 \times$ discrete[a]	1	3	τ	✓
Massey et al. (2010b)	all empty	β, d	$3 \times$ discrete[b]	1	–	τ	–
Rhodes et al. (2010)	all empty	β, d	$5 \times 10\,000$	1	3	τ	✓
Anderson & Bedin (2010)	all empty	lookup	$1 \times \rho_t V_{\text{pix}} n_y$	n_y	–	lookup	–
Massey (2010)	all empty	β	$3 \times 10\,000$	$n_y/3$	–	τ	–
Hall et al. (2012)	all empty	full sim	$5 \times$ discrete[a]	1	full sim	τ	–
Short et al. (2013)	nearby objects	α, β	4×1	n_y	–	τ	–
This work	all empty	β	$3 \times \rho_t V_{\text{pix}} n_y$	$n_y/5$	–	τ	–

[a]A population of traps is scattered throughout the model CCD.
[b]The population of traps is oversampled, each trap holding a fraction of an electron.
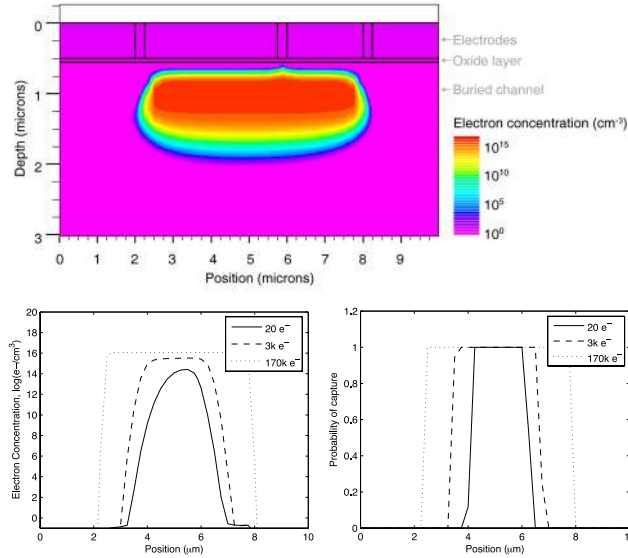

Figure 2. Top: the density of electrons within one pixel of a generic CCD. This is a two-dimensional slice through the full three-dimensional ATLAS model. The charge cloud in this particular example contains 172 000 electrons, and the detector model includes a four-phase, uneven electrode structure for readout in the left–right direction on the page. Bottom-left: a one-dimensional slice through the electron density. Bottom-right: the probability that an electron will be captured by a trap at that location, computed by combining the electron density with the capture cross-section for a specific trap species (here Si-E) and a set dwell-time (here 1 ms), through the use of Shockley–Read–Hall theory.

Shockley & Read 1952). The capture cross-section varies between trap species and also depends on operating conditions like the temperature T. Integrating over all the traps (all the volume) within a pixel, the total number of electrons that will be captured into each species of charge trap is

$$n_{\text{c}} = \int_0^{V_{\text{pix}}} \mathrm{d}V \; \rho_t^{\text{empty}}(\boldsymbol{x}) \; p_{\text{capture}}(\rho_e(\boldsymbol{x}; n_e), \sigma_t(t_{\text{dwell}}, T)), \tag{1}$$

where only the initially empty traps (or fractions of traps) are considered in the trap density $\rho_t^{\text{empty}}(\boldsymbol{x})$.

ATLAS models suggest that the probability of capture during a dwell time ~ 1 ms is typically close to a step function in the electron density. That is, inside some density threshold $p_{\text{capture}} \approx 1$ and all the traps capture an electron, but outside that threshold, $p_{\text{capture}} \approx 0$ (bottom-right panel in Fig. 2). Empirical tests with in-orbit data confirm that this step function is a suitable approximation for ACS/WFC (Anderson & Bedin 2010, section 3.2). This simplifies equation (1): $p_{\text{capture}}(\rho_e, \sigma_t)$ becomes a delta function inside V_e, which is realized as a volume when the integration is done. Since ρ_t is constant, it can also be moved outside the integral. This leads to a 'volume-driven' model, in which the number of electrons captured during each dwell time,

$$n_{\text{c}}(n_e) = \rho_t^{\text{empty}} \; V_e(n_e), \tag{2}$$

depends solely on the effective volume $V_e(n_e)$ of the packet. A packet of only few electrons is confined to a small volume of silicon, and thus 'sees' to only a fraction of the charge traps. As the number of electrons n_e grows or shrinks, so does the number of exposed traps.

If a cloud of electrons must traverse y pixels to reach the readout register, the volume of silicon (and the number of traps) to which they are exposed increases by this factor. Incorporating that factor of y (but we shall return a more sophisticated treatment of this issue in Section 2.1.5), the Massey (2010) algorithm empirically fitted a form

$$n_{\rm c}(n_e) = \min\left\{\rho_{\rm t}^{\rm empty} y\, V_{\rm pix}\left(\frac{n_e}{w}\right)^\beta,\; n_e y\right\},$$
(3)

where the min{} explicitly ensures the number of electrons stays positive, and $\beta \approx 0.5$. Like the Hall et al. (2012) fitting function, the code also allows $V_e(n_e)$ to include a constant. In early analysis, Massey et al. (2010b) (their equation 4) found a slightly negative best-fitting constant, and interpreted it as tentative evidence for a supplementary buried channel (SBC). It remains unknown whether the ACS/WFC CCDs have an SBC (Anderson, private communication): one was designed in the silicon doping structure, but it may never have properly formed, or evaporated away after manufacture. The constant was subsequently set to zero in Massey (2010), because the data analysed contained a relatively bright sky background so was insensitive to it.

The Anderson & Bedin (2010) algorithm is also volume driven, but uses an empirical look-up table for $n_{\rm c}(n_e)$ instead of an analytic function. They analysed data with lower sky backgrounds, and found a relative *decrease* in the number of traps exposed to small ($DN_2 = 10$) electron clouds (Anderson & Bedin 2010, fig. 6) that supports the existence of an SBC.

Not assuming that $p_{\rm capture}$ is a step function led Short et al. (2013) instead to a partially 'density-driven' algorithm. This is particularly important because the *Gaia* CCDs will often need to transfer very small numbers of electrons between pixels (because in TDI observations, the exposure time for the first pixel-to-pixel transfers is almost nil). In a density-driven model, a packet of a few electrons is assumed to occupy a relatively larger[2] volume ($\beta \sim 0.3$) and 'see' more traps. However, with a low electron density inside that larger volume, only a fraction of the traps capture an electron. The fraction is governed by a parameter α, which varies between trap species and depends on $\sigma_t(t_{\rm dwell}, T)$. Prod'homme et al. (2012) and Short et al. (2013) used this approach to reproduce lab test data in a much wider range of operating conditions and lower signals than is typically encountered by ACS/WFC.

The number of captures gains an extra term

$$n_{\rm c}(n_e) = \frac{2\rho_{\rm t} y V_{\rm pix}(n_e/w)^\beta}{2\rho_{\rm t} y V_{\rm pix}(n_e^{\beta-1}/w^\beta)+1}\left(1-\exp\left\{-\alpha\, n_e^{1-\beta}\right\}\right)$$
(4)

(cf. Short et al. 2013, equation 22, assuming that the traps do indeed start empty). Note that the factor of 2 (and the counterintuitive appearance of y in the denominator) arises from a peculiar approximation of $\min\{A, B\} \approx AB/(A + B)$ in Short et al. (2013) equation 5. There is no reason for this, and is more physically intuitive rewritten as

$$n_{\rm c}(n_e) = \min\left\{\rho_{\rm t} y V_{\rm pix}\left(\frac{n_e}{w}\right)^\beta,\; n_e y\right\}\left(1-\exp\left\{-\alpha\, n_e^{1-\beta}\right\}\right),$$

(5)

where we see the first term identically recovers equation (3).

2.1.3 Monitoring trap occupancy

Monitoring which traps are full at any time during the many pixel-to-pixel transfers requires an accounting scheme. Early codes (e.g. Bristow 2003) modelled populations of discrete traps, each able to capture a single electron, and kept a ledger of their occupancy. This would be ideal if the model traps could be placed at the true trap positions (which could in principle be measured by in-orbit trap-pumping, Mostek et al. 2010). However, the position of traps in *HST* CCDs are unknown. To implement this approach, model trap positions would therefore have to be scattered at random in 3D. Every random realization of trap positions (in both hardware and software) adds shot noise to the trails. The random locations of traps in the hardware is unavoidable but, to remove the software half of this noise, Rhodes et al. (2010) instead modelled a fixed grid of $n_{\rm levels}$ fractional traps. Each fractional trap is able to hold up to $\rho_{\rm t} V_{\rm pix}/n_{\rm levels}$ fractions of an electron, and samples a small volume within a pixel. Because V_e grows monotonically with n_e (see Fig. 2), these small volumes can represent successive regions of a pixel into which a cloud of electrons will next grow, and the 3D integral in equation (1) can be converted to a 1D integral over the ordered set of traps.

This approach has been followed ever since, but the adopted density of fractional traps varies considerably between modern algorithms. At the most dense extreme, the Massey (2010) algorithm monitors an array of $3 \times 10\,000$ traps in every pixel. The $n_{\rm species} = 3$ trap species observed in *HST* each delay electrons by a different characteristic time. After a coarse convergence test, $n_{\rm levels} = 10\,000$ equal-sized traps in every pixel were found necessary to sample the volume within each pixel. Such an enormous array is slow to manipulate (execution time-scales roughly as the total number of traps to the power 0.55), so the other codes each use an approximation to shrink one of its dimensions.

The Short et al. (2013) algorithm monitors $n_{\rm species} = 4$ species of traps that contribute to parallel CTI (Prod'homme et al. 2010) – but the entire volume of the CCD pixel is bundled together. Only a single number is monitored for each trap species in each pixel: the total fraction

---

[2] The volume of electrons $V_e \propto (n_e/w)^\beta$ increases as β decreases because $n_e/w < 1$.

of currently occupied traps. This less detailed accounting saves a factor of 10 000 traps and provides a significant speed boost, but may not be sufficiently accurate for our purposes.

The Anderson & Bedin (2010) algorithm monitors traps at many n_{levels} within the CCD – but the different species are combined into one 'composite' type. The composite traps represent a mixture of the three physical species of traps, so they have complex behaviour (not exponential release), and their behaviour may also vary at different heights within the CCD.

Anderson & Bedin (2010) make two especially useful points about the trap accounting. First, although intermediate calculations deal with fractions of electrons, the number of electrons in each pixel after readout (and the file formats that store these data) must be integers. Achieving greater precision is therefore unnecessary. Staying within a precision of one electron after n_y transfers requires only $n_{levels} \geq \rho_t V_{pix} n_y$ traps, each of which interact with one marginal electron. Secondly, the traps need not be all the same size. Where marginal electrons occupy a different volume, it could be most efficient to use large/small traps holding many/few electrons. Indeed, in accounting terms, the traps could be spaced linearly in terms of volume, marginal-free electrons, or marginal exposed traps.

2.1.4 Pixel-to-pixel parallel transfers (box 4)

During every cycle of pixel-to-pixel parallel transfers, electrons still in the image array are held stationary for a dwell time t_{dwell}, while electrons in the serial register are transferred to the readout electronics. At the end of the cycle, electrons in the image array are then moved rapidly to an adjacent pixel. They pass fleetingly through the volume of silicon between pixels (i.e. at position $<2\,\mu m$ or $>8\,\mu m$ in the top panel of Fig. 2). The time spent in this volume is orders of magnitude shorter than t_{dwell}, but the electrons are exposed to additional traps.

Accounting practices for traps between pixels varies throughout the literature. Bristow (2003) and Rhodes et al. (2010) included the ability to step electrons through this region, and explicitly encounter new traps. Massey (2010) and Anderson & Bedin (2010) assumed trapping is instantaneous, so account for these traps by enlarging V_{pix} to include the region (only the degenerate combination $\rho_t V_{pix}$ can be measured). Short et al. (2013) assuming trapping is not instantaneous, and the transfer time is so short that no electrons will be captured. In practice, this decision will only matter at the image level if the clock sequence changes or, if trapping is not instantaneous, by influencing the initialization of traps (traps under integration phases become filled during exposure, or the first operation of Fig. 1, while traps under barrier phases remain empty; Hall et al. in preparation).

2.1.5 Accounting for multiple transfers at once

Early codes (e.g. Bristow 2003; Massey et al. 2010b; Rhodes et al. 2010) computed the effect of every pixel-to-pixel transfer individually. In a CCD with $n_x \times n_y$ pixels, electrons in the row of pixels closest to the serial register take only one parallel transfer to be read out, but electrons in the farthest row much undergo n_y parallel transfers. Full readout requires $n_x n_y (n_y + 1)/2$ parallel transfers, which is more than 10^9 for ACS/WFC ($n_y = 2048$ in ACS/WFC and *Euclid* or 4500 in *Gaia*). This is fast in hardware, but performing this many iterations of the inner loop (boxes 3–5) in Fig. 1 takes a long time to reproduce in software.

To speed up their software, Short et al. (2013) and Anderson & Bedin (2010) independently introduced the same approximation. If $\rho_e \gg \rho_t$, the size of the electron cloud does not change dramatically from its first transfer to its last. If the density of charge traps in each pixel is (modelled as) constant, the population of exposed charge traps is also the same at each transfer. Every transfer is thus identical.

When the above assumptions hold, we need to compute only the first pixel-to-pixel transfer (which shifts electrons by one pixel), and multiply its effect by the number of transfers. In every column of the ACS/WFC CCDs, the number of transfers faced by electrons in pixel rows 1 to 2048 is simply the vector

$$y_1 = y = [1, 2, 3, \ldots 511, 512, 513, \ldots 2046, 2047, 2048], \tag{6}$$

i.e. electrons in the first row of pixels take one step straight into the serial register, electrons in the row above take two steps . . . and the farthest electrons take 2048 steps. Multiplying the effect of one transfer (with suitable error catching) is algorithmically equivalent to multiplying the number of electrons captured in one step, and this is the origin of the factor y in equations (3), (4) and (5). We shall call this the 'express' $E = 1$ approximation, because the effect of each pixel-to-pixel transfer has been computed only once. Reading out the entire CCD requires only $n_x n_y$ iterations of the inner loop in Fig. 1.

Anderson & Bedin (2010) suggested a generalization for better accuracy, which Massey (2010) implemented. In reality, not all pixel-to-pixel transfers are identical. The assumption $\rho_e \gg \rho_t$ breaks down if ρ_t is large (e.g. in a severely damaged CCD) or ρ_e is small (e.g. in dark frames, short exposures or at wavelengths where the sky is very faint – see Section 3.2.2). In this situation, once a cloud of electrons has undergone all transfers but one, the size of the charge packet may have changed significantly. If the charge packet is smaller (or larger) during the final transfer, it will be exposed to fewer (or more) traps than during the first.

We can account for this evolution by applying the first transfer only a limited number of times, then recalculating the effect of subsequent transfers. As a concrete example, consider a case in which we apply the first transfer up to only 512 times. We dub this, the $E = 2048/512 = 4$ approximation, because the effect of some transfers will need to be recomputed four times. The effect of the first 512 transfers can be computed by performing one transfer as before, but capping $y_{4,1} = \max\{y, 512\}$. The effect of the next 512 transfers is then recalculated,

using the updated size of each charge packet. The procedure is then repeated for a third and a fourth time. In each case, the number of captured electrons is multiplied by one row of

$$
y_4 = \begin{bmatrix} 1 & 2 & 3 & \dots & 511 & 512 & 512 & 512 & \dots & 512 & 512 & 512 & 512 & \dots & 512 & 512 & 512 & 512 & \dots & 512 & 512 & 512 \\ 0 & 0 & 0 & \dots & 0 & 0 & 1 & 2 & \dots & 511 & 512 & 512 & 512 & \dots & 512 & 512 & 512 & 512 & \dots & 512 & 512 & 512 \\ 0 & 0 & 0 & \dots & 0 & 0 & 0 & 0 & \dots & 0 & 0 & 1 & 2 & \dots & 511 & 512 & 512 & 512 & \dots & 512 & 512 & 512 \\ 0 & 0 & 0 & \dots & 0 & 0 & 0 & 0 & \dots & 0 & 0 & 0 & 0 & \dots & 0 & 0 & 1 & 2 & \dots & 510 & 511 & 512 \end{bmatrix}.
$$

For the first 512 pixels, only one transfer is ever calculated; for pixels 513–1024 there are two calculations; for pixels 1025–1536 there are three calculations; and for pixels 1537–2048, the maximum number of transfers explicitly calculated is four. It can be checked that each pixel is acted upon by the correct number of transfers because the columns of y_4 add up to the same, monotonically increasing values as y_1.

The above approach can be generalized to an approximation with any value E between 1 and 2048 (note that we also refer later to $P \equiv n_y/E$). In the limit of $E = 2048$ (or $P = 1$), this recovers the slow algorithms of Bristow (2003) and Massey et al. (2010b) in which the effect of every transfer is computed afresh. For this, y_{2048} is a square, upper-triangular matrix containing only ones and zeroes. Since later transfers are not applied to many pixels (the zero entries in y_E), they do not even need to be computed for all pixels. Counting the non-zero entries in y_E, an efficient implementation will require $n_x n_y(E + 1)/2$ iterations of the inner loop in Fig. 1.

2.1.6 Release of electrons from charge traps (box 5)

After a short delay inside a charge trap (the length of which depends on the trap species), captured electrons are eventually released into whichever charge packet is then nearby. The Massey (2010) and Short et al. (2013) algorithms model the probabilistic release of electrons as an exponential decay. Each species of charge trap (at least three species in *HST*, and four in *Gaia*) has a characteristic release profile with a different half-life. Releasing a constant fraction of electrons at each time step is computationally easy. More importantly, the half-lives changed in a predictable way when *HST*'s operating temperature was lowered in 2006 (for more details see fig. 4 of Massey et al. 2010b), and behave as expected within a wide range of operating conditions for *Gaia* (Prod'homme et al. 2012; Short et al. 2013).

The Anderson & Bedin (2010) algorithm instead releases traps according to a function described by a look-up table, which depends upon the time since the electrons were captured (and thereby reproduces a three-exponential decay, Massey 2010). Only one computational species of trap is required to produce any trail profile. There is extra overhead to keep track of the time but is still faster than computing the trails for four separate trap species.

Anderson & Bedin (2010) also found evidence for two different sets of trail profiles in *HST* data: electrons near the bottom of a pixel (smaller packets of electrons than those considered by Massey 2010 or in this work) saw only fast-release traps, and electrons near the top saw additional slow-release traps. They suggested (see their section 6.2) that this discrepancy might disappear when $E > 1$ algorithms are implemented. However, if the effect remains, an elegant solution in the regime when the number of exposed traps is greater than the number of electrons might be a density-driven charge capture model in which the fast traps are filled first, and the slow ones are filled more slowly.

Various algorithmic tricks can be used to increase speed, for example by considering (for electron release) only traps that have previously been exposed to electrons. Monitoring a high water mark of electrons against the fixed grid of traps, and considering only traps below this mark, offers a large speedup until a saturated pixel is encountered. It is also possible to lower the water mark as the electron content in the traps approaches zero, and thus regain speed.

2.1.7 Loop over columns (box 7)

This loop is trivially parallel, and may thus be ideal case for massive parallelization techniques that exploit either GPUs or MPI on multiple CPUs.

2.2 Removing CTI trailing

The above approach describes a forward operation to *add* CTI trailing. It is a software simulation that approximates the effect of CTI in radiation damaged hardware. What we really need is a software algorithm to *remove* CTI trailing by reversing the readout. Finding such an algorithm is a classic 'inverse problem', and two generic approaches exist.

2.2.1 Forward modelling

The data analysis pipeline for *Gaia* (e.g. Prod'homme et al. 2012; Seabroke et al. 2012) generates a model of every source and applies the Short et al. (2013) forward CTI transform, before matching the trailed model to the (trailed) data. This approach requires accurately simulated data, which must be passed through the entire data analysis pipeline. Since *Gaia* observes mainly stars, the simulations [of just a point spread function (PSF)] have the potential to be sufficiently accurate. Observations of galaxies and extended sources are harder to simulate
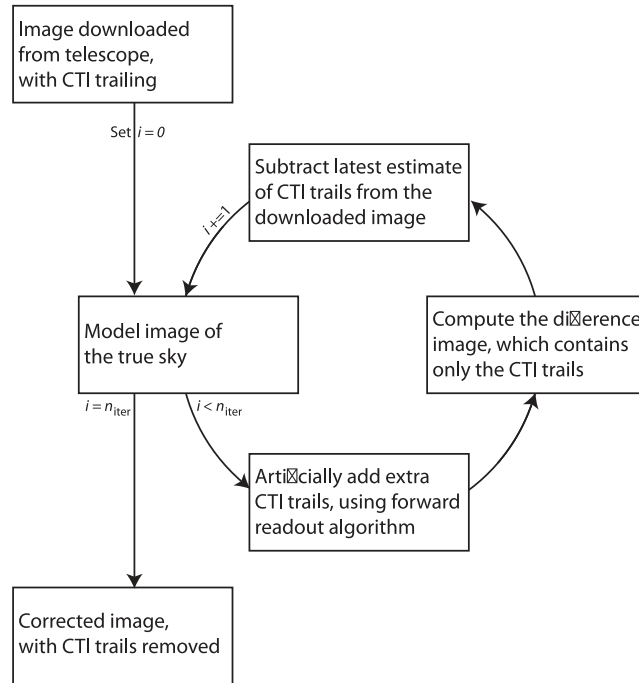
Figure 3. Inverse operation to remove and thus correct CTI trailing. This works by repeatedly applying the forward algorithm to add CTI trailing (Fig. 1, shown here as the bottom box in the loop). It converges towards an image that, when read out again, reproduces the (trailed) image downloaded from the spacecraft.

accurately. For limited applications, such as weak lensing with *Euclid*, it may be possible to simulate galaxies with sufficient detail for a similar forward-modelling approach (e.g. Miller et al. 2013; Zuntz et al. 2013).

One unique ability of a forward approach is to model cosmic rays that hit the CCD during readout. Electrons create by these cosmic ray hits undergo a smaller number of transfers than other electrons, become less trailed, and offer less shielding to downstream sources than would otherwise be expected. Correcting their trailing probably requires that exact sequence of events to be followed precisely.

2.2.2 Iterative inversion

For more general workhorse applications, a flexible and stand-alone image correction scheme is better. Trailing is (almost) the last process to happen during data acquisition, so it should ideally be corrected during the first stage of data analysis (after only bias subtraction and gain correction). The Massey et al. (2010b) and Anderson & Bedin (2010) pipelines developed for *HST*/ACS use an iterative method first suggested by Bristow 2002 and illustrated in Fig. 3 (for a more mathematical treatment, see table 1 of Massey et al. 2010b). This method iterates towards a (corrected) version of the image that, when trailed a final time, matches the (trailed) data. Convergence requires $n_{\rm iter} = 5$ (Anderson & Bedin 2010) or 3 (Massey 2010) iterations, which correspondingly lengthens runtime – but once it is run, a detailed image is available for all subsequent analysis and CTI can be forgotten.

2.2.3 Dealing with read noise

The forward-modelling approach requires models of the true sky, which can be very accurately trailed because that model can be created without noise. On the other hand, any iterative CTI correction will be fundamentally limited by the presence of read noise (Anderson & Bedin 2010; Cropper et al. 2013). Read noise is added after charge transfer so does not get trailed during readout, but it is incorrectly *untrailed* with the rest of the image during post-processing. As evidenced by convergence to the wrong answer in even the idealized test of Fig. 6, the spurious untrailing of read noise introduces errors in measurements of photometry, astrometry and morphology. Furthermore, spuriously untrailing the read noise increases its apparent rms, and thus increases the noise level noise in the final image (CTI correction is similar to an unblurring/sharpening operation, in that it increases the autocorrelation of pixels while introducing an anticorrelation of adjacent pixels).

Anderson & Bedin (2010) attempt to find the 'minimum' CTI correction consistent with knowledge of the image and the read noise. They apply a high-/low-pass filter to the image (see their section 5.1), which isolates a component that is considered to be read noise, and running the correction on only the smooth component. This certainly achieves a smoother final product, but in an unsatisfactory manner as e.g. the choice of filter threshold is ad hoc, and (high frequency) shot noise in the image *was* trailed.

We propose a different solution, based around the idea that the incorrect measurements after correction are purely caused by the presence of noise that is (anti)correlated between pixels. In similar situations, Kaiser (2000), Bernstein & Jarvis (2002) and Huff et al. (2011) have noted

that noise correlated more in the *y*-direction than the *x*-direction can lead to biases in centroids in the *y*-direction, plus a larger uncertainty in the centroid in the *y*-direction, which in turn creates biases in galaxy shapes. Furthermore, this correlated noise can also produce selection biases, whereby more objects aligned in the *y*-direction are detected above an S/N threshold. The three papers above developed a method to adjust the noise covariance matrix, by *adding extra random noise* that has then been correlated between adjacent pixels in exactly the right way to make the total noise white. The decision to add noise in software is somewhat analogous the decision to implement a pre-/post-flash illumination in hardware. For the cost of a small amount of extra noise, the data products become easier to handle for precision measurements of photometry, astrometry and morphology. We shall test whether this also improves the CTI correction.

To determine how much noise to add, we first identify a region of blank sky and measure the noise power spectrum $P_N(\boldsymbol{k})$ (in practice, we use a truly blank image containing only simulated sky background passed through the forward readout process, plus read noise). We then follow the method described in section 4.2 of Huff et al. (2011) to construct an additional amount of synthetic noise such that the covariance of the total noise is four-way symmetric. In addition, we also try adding synthetic noise that makes the total noise white (zero covariance between pixels). This uses a very similar method and just requires a larger rms. We shall test in Section 3 how much noise each version adds to a typical image.

3 CALIBRATION ON REAL DATA

We shall now construct an up-to-date model of CCD detector readout for the *HST* ACS. Using a combination of real, in-orbit data and simulations, we shall attempt to construct a code in which all approximations lead to sub-1 per cent inaccuracies. Within this limit, we shall attempt to maximize computation speed.

3.1 Measurements from real data

We measure CTI trails behind warm pixels in long *F*814*W* exposures taken as part of the CANDELS survey (GO-12444) in 2013 January. Warm pixels are useful because they should be perfect delta functions in the absence of CTI. The observed trails have the same profile shape as those seen in Anderson & Bedin (2010). After accounting for 'self-CTI' (multiple captures of an electron during readout, so that trail becomes longer than an exponential, in a way that depends on the background level and capture model), Massey (2010) demonstrated that this profile is well fitted by a model in which three trap species have exponential release profiles,

$$\tau = \{0.74, 7.7, 37\} \text{ pixels} \tag{7}$$

and relative densities

$$\frac{\rho_t}{\sum \rho_t} = \{0.17, 0.45, 0.38\}. \tag{8}$$

However, recent trails have increased in *amplitude* – reflecting an increased trap density. From the *HST* data archive, we identify long (1/3- or 1/4-orbit) exposures in *F*606*W* or *F*814*W* of uncrowded, extragalactic fields obtained between mid-2011 and early 2013. These images have sky backgrounds $\gtrsim$50 electrons. A fit to the inferred trap densities during this period is shown in Fig. 4. We have not updated our trap model for data obtained before Servicing Mission 4, because the lower trap density then makes the existing correction sufficient.

We introduce two improvements when fitting the trap densities. In Massey (2010), we measured the trail amplitudes throughout an image, then fitted the global trap density, weighting all the trails equally. In our improved procedure, we now recognize warm pixels far from the readout register are trailed most, so we now give measurements of their trails an appropriately higher weight during fitting than those near the readout register. We also recognize that in these pixels, which have undertaken many transfers through a harshly degraded CCD, not all of the transfers will have been identical (equivalent to the $E > 1$ readout algorithm described in Section 2.1.5). To account for this (and self-CTI), we pass a hot pixel through the full readout process (this can be done fast and analytically for a delta function), and iterate its flux until it contains n_{warm} after readout. Both changes result in an increase in reported trap densities: approximately 7 per cent from upweighting pixels farther from the readout register and a further 1 per cent by accounting for the changing charge cloud size, although these numbers depend mildly on the number of electrons in a pixel.

Fitting the amplitude of trails for warm-hot pixels now requires a well-filling model with $\beta = 0.478$ and full well depth $w = 84\,700$ electrons. Given this, the observed density of traps after Servicing Mission 4 is

$$\rho_t V_{\text{pix}} = (1.66 \pm 0.01) + (5.65 \pm 0.24) \times 10^{-4} (t - 3374) \, \text{pixel}^{-1}, \tag{9}$$

where *t* is the number of days since launch (see Fig. 4). The rate of degradation from radiation damage increased during the 2007–2011 Solar minimum, and is showing no sign of slowing.

3.2 Tests on simulated data

The challenge with real data is that we see the sky only after the effect of CTI, and we rarely know its truth. Here, we construct simulated data – which we degrade with a readout model, then correct using the same readout model. To test the internal stability of (the inversion of) the algorithm, we can explicitly compare corrected data to its appearance in the absence of CTI. Note, however, that any inconsistency will not account for error on the model, such as the density or positions of charge traps.
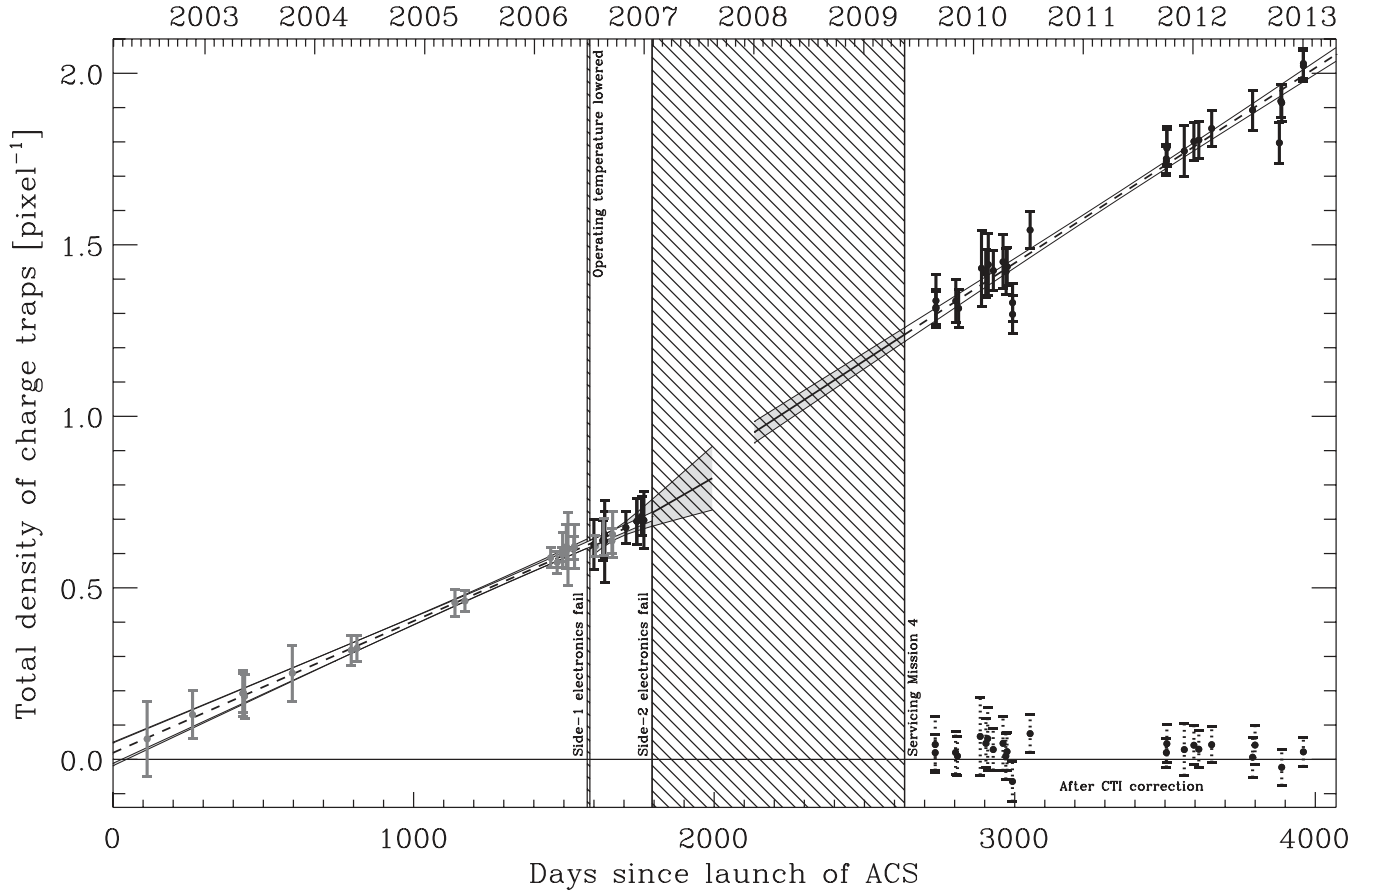
Figure 4. Density of charge traps in the *HST Advanced Camera for Surveys/Wide Field Channel*. Points with solid error bars show trap densities measured from the amount of trailing behind warm pixels (grey and black are different gain settings). Points with dashed error bars show measurements behind the same warm pixels after correction with $n_{iter} = 6$ and other algorithmic choices as described in the text. Hatched regions show significant periods when ACS was offline.

3.2.1 Tests on simulated science data

To investigate the effect of CTI on complex, extended sources, we simulate images of faint galaxies seen by ACS/WFC. We simulate several hundred thousand noise realizations of a circular, exponential, just-resolved (scale size 0.06 arcsec) galaxy, convolved with a TINYTIM (Krist 1995; Hook & Stoehr 2008; Krist, Hook & Stoehr 2011) model of the *HST* PSF at 814 nm, centred randomly within a pixel at position $y = 2048$ from the serial register. We place this galaxy on a $\sim$100 electron (and noisy) sky background, simulate parallel CCD readout, then add Gaussian read noise with 4.0 electron rms per pixel. We scale the flux of the galaxy and repeat this analysis for a galaxy detected by SOURCE EXTRACTOR (Bertin & Arnouts 1996) at a mean signal-to-noise ratio of S/N = 10, 50 and 100. In these cases, SOURCE EXTRACTOR measures a mean full width at half-maximum (FWHM) size of 2.48, 2.00 and 2.00 pixels, respectively.

We use SOURCE EXTRACTOR to measure the total flux (FLUX_AUTO), signal-to-noise ratio S/N (FLUX_AUTO/FLUXERR_AUTO) and size (FWHM_IMAGE) of the image $I(x, y)$, and RRG (Rhodes, Réfrégier & Groth 2000) to measure its y position and PSF-convolved ellipticity

$$e_1 = \frac{\int I(r, \theta) \, W(r) \, r^2 \cos(2\theta) \, r \mathrm{d}r \mathrm{d}\theta}{\int I(r, \theta) \, W(r) \, r^2 \, r \mathrm{d}r \mathrm{d}\theta}, \tag{10}$$

where the weight function $W(r) = \exp(-r^2/2\sigma^2)$ is centred such that the first moments are zero and σ, which is fitted to each noisy image, is $\sim$3.3, 4.4 and 4.6 pixels on average for the S/N = 10, 50 and 100 galaxies. We also record the numerator of e_1, which is dimensionful but linear in image flux and therefore avoids potential noise-induced shape measurement biases (Massey et al. 2007; Zhang & Komatsu 2011; Melchior & Viola 2012; Réfrégier et al. 2012).

We first measure all these quantities on a version of the simulated image without any CTI, and record their 'true' values. We then remeasure the quantities on versions of the simulated images in the presence of CTI, using our ACS/WFC readout model (assuming an observation date of 2013 January 1, 3959 d since the launch of ACS and requiring a total of 1.989 traps per pixel). The effect of forward CTI trailing is shown in Fig. 5, for various algorithmic choices.

Finally, we apply our CTI correction algorithm to the degraded simulated images, and remeasure the quantities. The net effect of CTI trailing after correction is shown in Fig. 6, for various algorithmic choices.
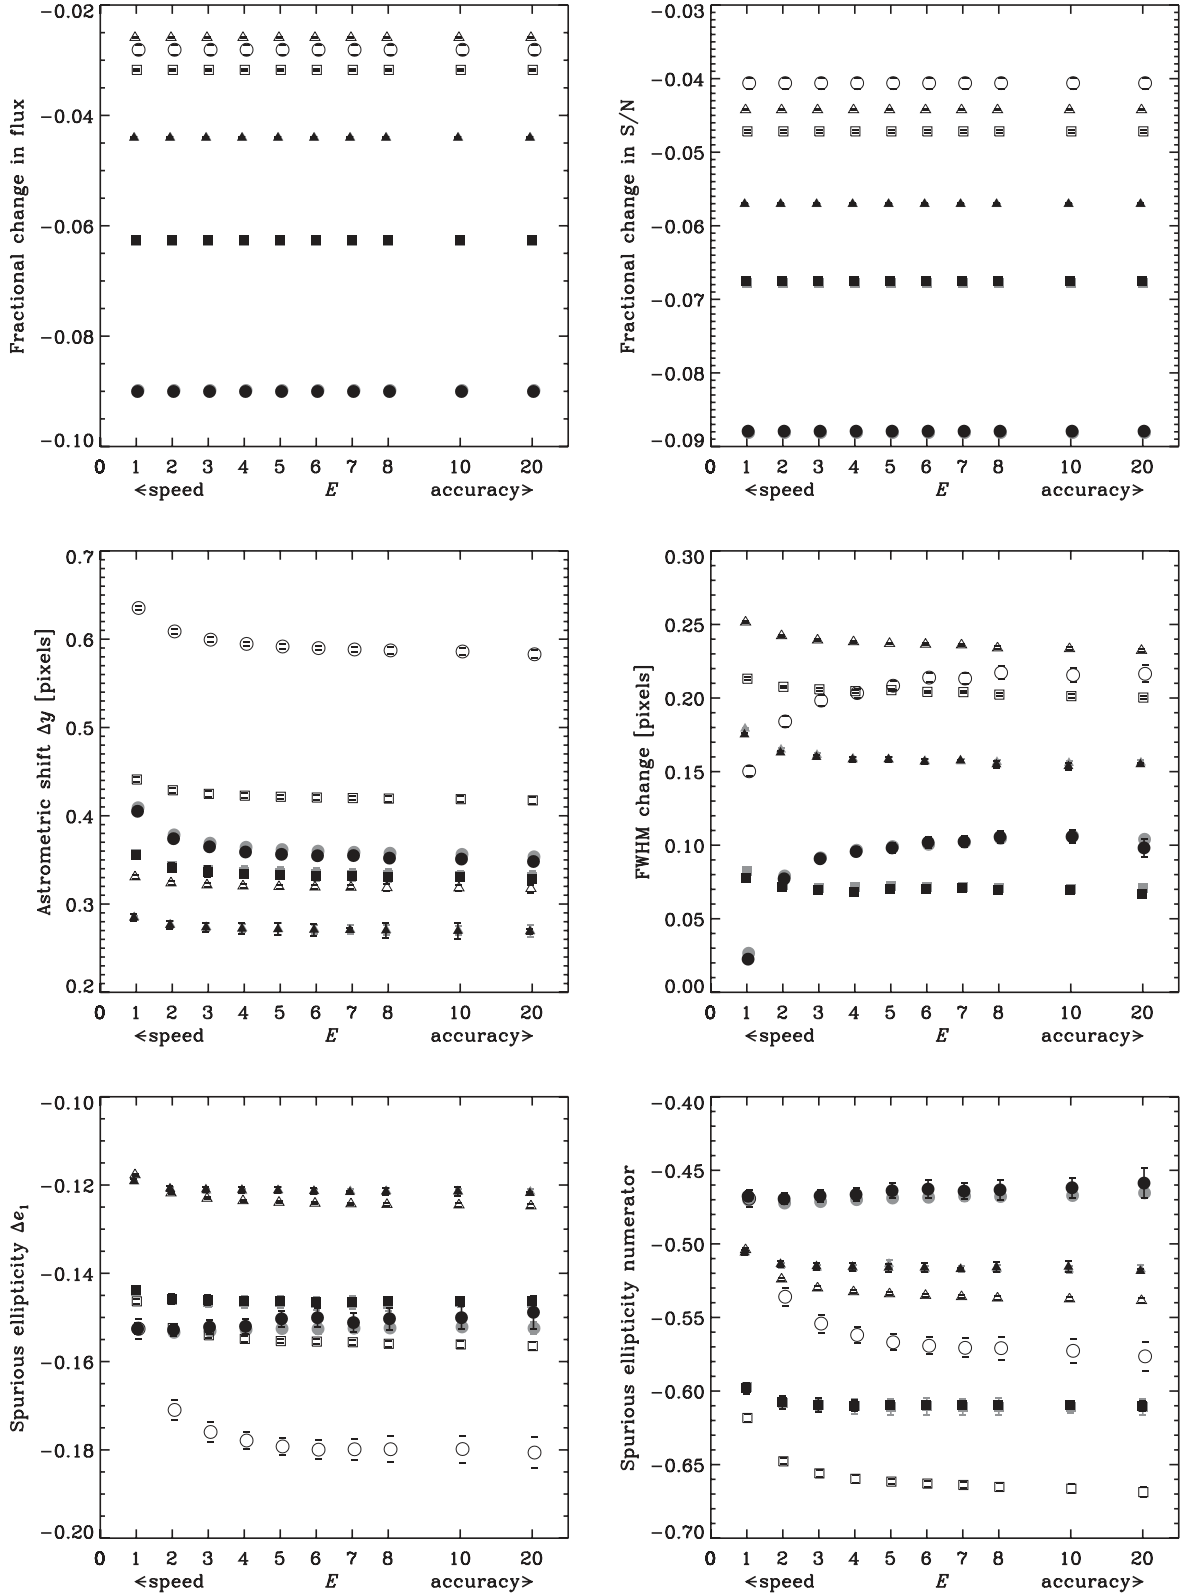
Figure 5. Convergence of our (forward) CCD readout model. The express parameter E governs a tradeoff between speed ($E = 1$, the result of a single pixel-to-pixel transfer is applied identically many times) and full realism ($E \to 2048$, the effect of each transfer is calculated afresh). Circles, squares and triangles, respectively show the effect on a galaxy with S/N = 10, 50, 100. Black, grey and white symbols show outputs of algorithms identical except that $n_{levels} = 10000$, $2048V_{pix}\max(\rho_t)$, 1: i.e. mimicking Massey 2010, Anderson & Bedin 2010 and Short et al. 2013. The different-sized errors mainly reflect the different number of galaxies simulated in each case.
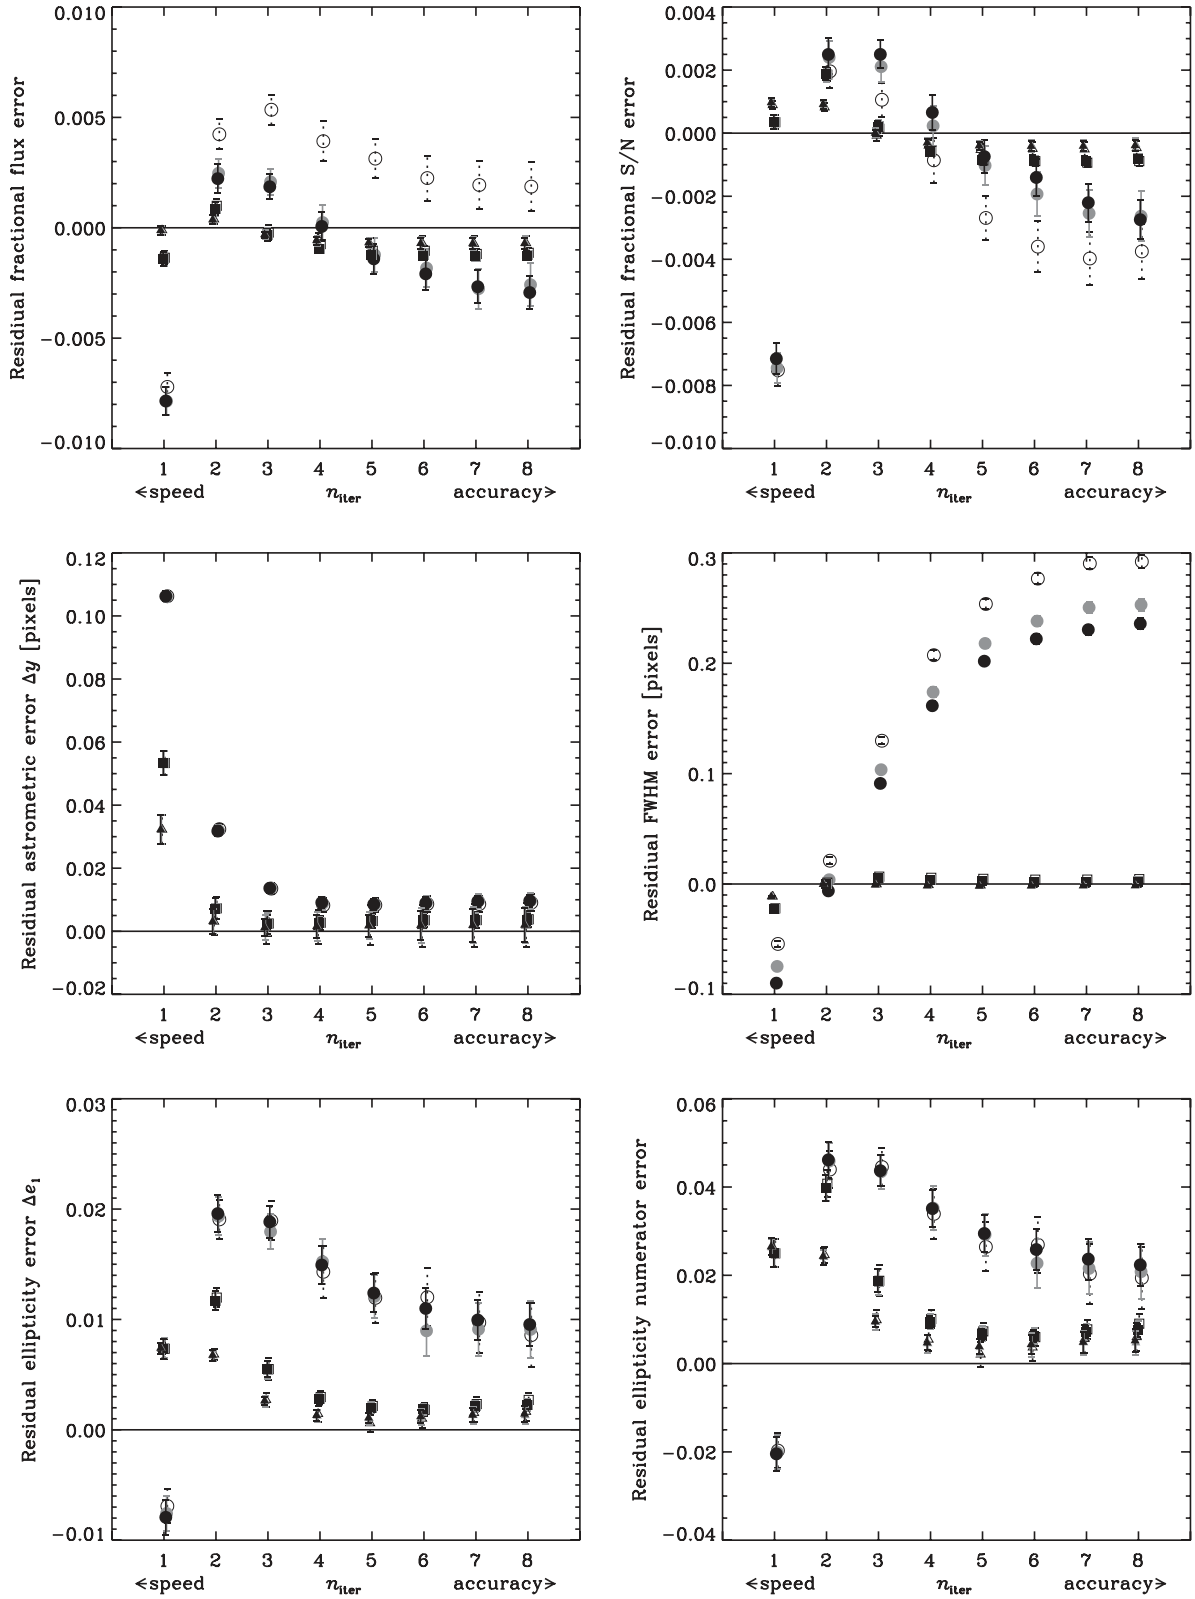
Figure 6. Convergence of our (inverse) CTI-removal algorithm, after some number n_{iter} of iterations. The same forward algorithm is used to add then remove CTI trailing. In the absence of read noise, correction is perfect (this case is not shown); with read noise, the iteration converges to the wrong image. Circles, squares and triangles, respectively, show the miscorrection of a galaxy with S/N = 10, 50, 100. Black, grey and white symbols show results after a raw CTI correction, noise whitening to make total noise four-way symmetric, and noise whitening to make the noise completely white.

3.2.2 Tests on simulated dark (low-signal) exposures

Accounting for trap occupancy in a finite number of n_{levels} (Section 2.1.3), and accounting for multiple transfers simultaneously via the 'express' E approximation (Section 2.1.5), relied on the condition that $\rho_e \gg \rho_t$. This assumption breaks down if sufficient radiation damage created a high trap density, or if the signal and background are low (e.g. in dark exposures or at wavelengths with faint emission and sensitivity such as with WFC3/UVIS data). In terms of CTI correction, the low-signal regime is defined in terms of the *total* number of photoelectrons, because charge transfer does not discriminate between source counts from astronomical objects or the sky background. High sky backgrounds pre-fill traps that do not then affect electrons, while a low sky background makes images more susceptible to CTI.

To investigate the behaviour of our algorithm in the low-signal regime, we construct test data consisting of delta function peaks of varying heights (similar to the hot pixels used in Section 3.1), 2048 pixels from the serial register and on zero background. We then apply an artificial readout, through a high density of traps that all have $\tau = 1$ pixel. Note that all the traps we use for this test are of the same species in order to avoid the complicating ambiguity of which to fill first. The upper panel of Fig. 7 shows an example of our test data on input (a ramp in the *x*-direction or a series of delta functions in the *y*-direction), and after simulated readout with three different choices of E.

Since $\rho_t \gg \rho_e$ for these tests, one might expect that some of the electrons will be captured and not released until after readout. Such electrons would effectively be lost from the image. However, the parabolic shape of the trailed regions in Fig. 7 reveals an interesting aspect of any volume-driven CTI model – even in the case of best accuracy ($E = n_y$ or $P = n_y/E = 1$), in which the effect of each pixel-to-pixel transfer is calculated afresh. During the first transfer, traps indeed capture all the electrons up to a certain height. Depending on the trap species' characteristic release time τ, the electrons are gradually released. The electrons fall to the bottom of the potential well where, because of the high trap density, they are immediately recaptured. For initially small charge clouds, the electrons never emerge from the traps back into the image (they are transported off the edge of the CCD). With larger charge clouds, however, the traps at the bottom of the potential eventually become saturated. Electrons released from traps at the top of the potential fall to the bottom and can no longer be recaptured. These electrons finally re-emerge back into the image.

To quantify the flux loss from an image, we sum over the CTI trail downstream of the original peak for 52 pixels, and compute the fractional charge loss from the image before and after readout, $\delta \equiv (\sum n_e^{\text{trailed}} - \sum n_e^{\text{true}})/\sum n_e^{\text{true}}$. We find that fast and intermediate trap species (τ from a few tenths to a few pixels) rapidly deplete the signal whenever the first trap captures a significant fraction of electrons. After 2048 transfers, the signal loss from the image is

$$\delta \approx -8.46 \times 10^{-4} \left(\frac{n_e}{w}\right)^{-0.647} \left(\rho_t V_{\text{pix}}\right)^{1.221} \tau^{0.114}, \tag{11}$$

assuming $\beta = 0.478$ and full well depth $w = 84\,700$ electrons (errors on the exponents are less than 0.03). Slow trap species deplete the signal even more severely. For example, traps with $\tau = 40$ pixels remove 20 per cent of a signal from an image in the same conditions (n_e and ρ_t) where traps with $\tau = 1$ pixel remove only 1 per cent.

Unfortunately, we find that additional flux can be artificially lost when approximations are made in the readout algorithm. In particular, the timing of the electrons' re-emergence into the image depends upon the values of n_{levels} and E. The n_{levels} parameter affects this because, although our algorithm is globally volume driven, it is set up to be density driven inside each level. Low values of n_{levels} increase the number of traps in the bottom level, delay trap saturation, and lengthen the time before any electrons re-emerge. However, this effect is minor: lowering n_{levels} from 10 000 to 2048 only affects trails with $n_e < 1$ electron by more than a few per cent.

Performance is more severely degraded by lowering E. As described above, some flux is lost from an image even in the slow-but-accurate extreme $E = n_y$ (see the right-hand side of the middle panel in Fig. 7). As E decreases, the effect of each transfer is updated only after $P \equiv n_y/E$ transfers, so the number of electrons captured in each but the first transfer is overestimated. Furthermore (but less importantly), because more electrons are allowed to be captured during each computed transfer, the 'bottom' level is effectively enlarged. Both effects act to delay the eventual re-emergence of trapped electrons back into the image. We find that there is a minimum E at which a cloud initially containing n_e electrons can be registered with a certain fraction of its flux for a given detector (defined by $\rho_t V_{\text{pix}}$ and τ). For $n_e = 100$, the most accurate output ($E = n_y$) is reproduced to within 10 per cent if $E > 5$, or within 1 per cent if $E > 20$.

Conveniently, the factor by which flux loss is artificially enhanced by the 'express' approximation is nearly constant with respect to n_e and ρ_t (i.e. it depends only on E). The bottom panel of Fig. 7 replots data from the middle panel, renormalized to the 'natural' performance with $E = n_y$. While 'natural' flux loss from the image cannot be avoided, the ordinate quantifies additional flux loss for $E < n_y$, which can be avoided at the expense of computational speed. The lined-up data points can be used to define a minimum suitable E in the low-signal regime. The multiplicative factors can be as large as 30 – but consider that this is 30 times a very small amount. Our overall goal of <1 per cent residual from CTI applies to the electrons in a typical science image, rather than to the very few electrons in these simulations.

4 IMPROVED CCD READOUT MODEL

We now have all the necessary data in hand to make suitable choices for the algorithmic options discussed in Section 2. Our choices are summarized in the bottom row of Table 1.

4.1 Adding CTI trailing

Here, we describe and justify our choices for the forward readout algorithm, in sub-sections mirroring those in Section 2.1.
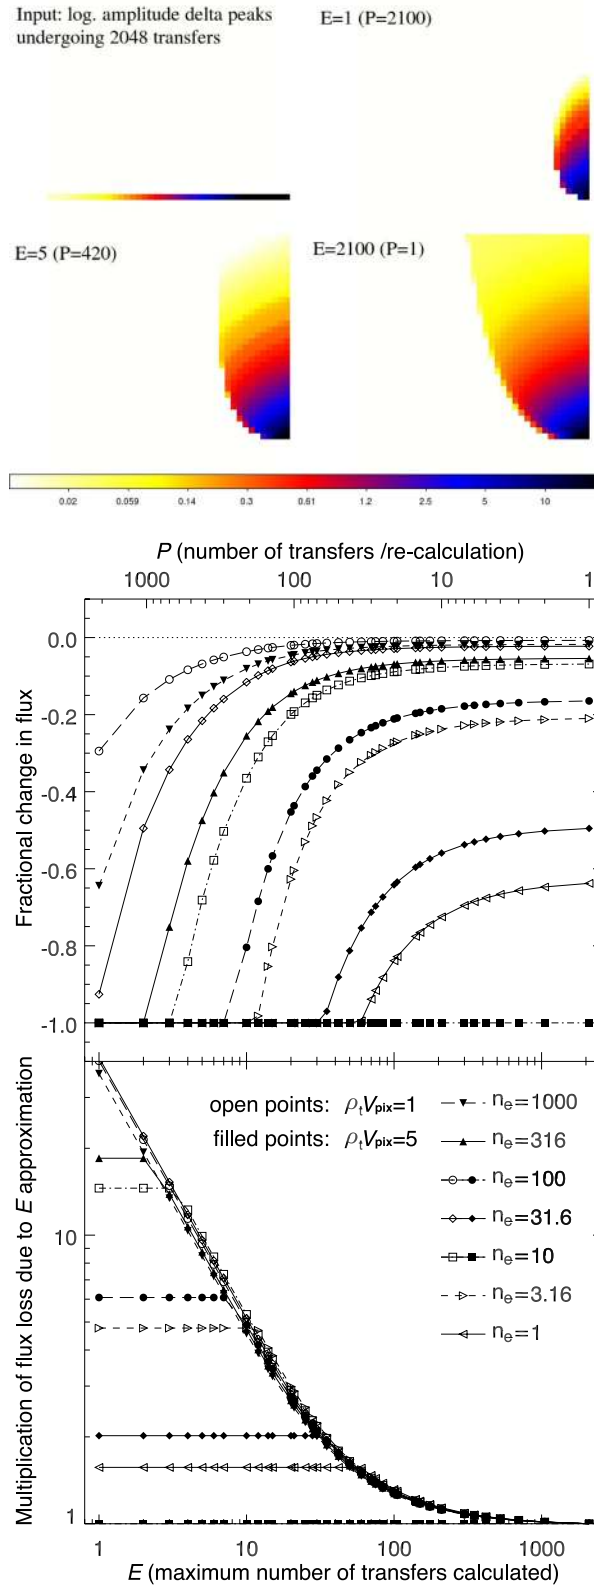
Figure 7. In the low signal-to-trap regime, approximations made to speed up processing have a greater impact on performance. Top panel: test data showing a ramp of hot pixels before and after CTI trailing (with readout towards the bottom of the page), for different choices of E. The colour scale is logarithmic. Middle panel: fractional charge loss from the image as a function of E. Open (filled) symbols denote different input signal levels and a trap density $\rho_t V_{\rm pix} = 1$ pixel^{-1} (5 pixel^{-1}), for $\tau = 2$ pixels. Bottom panel: excess charge loss due to the express approximation, compared to the 'natural' flux loss with slow-but-steady $E = n_y$.

4.1.1 Initializing charge traps

We initialize all traps as empty, assuming a long exposure during which electrons could escape.

4.1.2 Capture of electrons into charge traps

We use a volume-driven model of charge capture (equation 3) with $\beta = 0.478$ and $w = 84\,700$ electrons. We model the growth of total trap density via equation (9). Supported by Anderson & Bedin (2010)'s empirical findings that all ACS/WFC traps that could capture an electron do so on their first opportunity, we assume electrons are instantaneously captured by any exposed traps.

4.1.3 Monitoring trap occupancy

The value of n_{levels} required for suitable performance can also be learned from Fig. 5. The slow-but-sure limit of $n_{\text{levels}} = 10\,000$, in which the accounting for trap occupancy is finely grained, is shown by black points. We space our traps linearly in terms of the occupied pixel volume, as it is the captured traps that make a difference, rather than the free ones. Using $n_{\text{levels}} = 2048 \times \max(\rho_t V_{\text{pix}}) = 1860$ boosts speed by a factor of 3.7 while maintaining <1 per cent accuracy, as shown by grey points. Halving the number of traps further (as tried by Massey 2010, and reproduced by us, but not shown on the plot for the sake of clarity) does indeed introduce ~ 1 per cent errors.

Using $n_{\text{levels}} = 1$ (as advocated by Short et al. 2013, but implemented within our code) is not sufficiently accurate for this application. The white points in Fig. 5 show the consequences. The model's approximations break down when the size of the electron cloud changes between adjacent pixels. For example, the large electron cloud representing the central peak of a star image will fill many charge traps. After that charge cloud has moved along, some of the traps will immediately release some of their electrons. When the same traps see fewer electrons in the trailing wing of the star, the Massey et al. (2010b) and Anderson & Bedin (2010) algorithms will allow some of those electrons to be captured. However, the Short et al. (2013) algorithm will know only that many traps are still full, so capture no more free electrons.[3]

We use $n_{\text{species}} = 3$ and account independently for each species. Bundling all three trap species together (as advocated by Anderson & Bedin 2010) results in a speed boost by a factor of 1.6, at no loss of accuracy. This is not as great as it could be because of the overhead from monitoring the time since they were filled (and emptying them accordingly). There is also overhead in the complexity of this code, particularly when it comes to changing the operating temperature etc., so we choose to keep it simple. However, if the instrument parameters were completely fixed, the freedom to arbitrarily adjust the shape of trail profiles has the potential to increase accuracy via a small number of extra parameters, and might be useful in the future.

4.1.4 Pixel-to-pixel parallel transfers

If charge capture is indeed instantaneous, we can also model parallel transfer as instantaneous, and the region between pixels can be generically considered to enlarge V_{pix}. Exactly how much it enlarges V_{pix} may depend on the shape of the electric potential in the CCD during transfer phases. However, the only observable quantity is $\rho_t V_{\text{pix}}$, so we can merely fix V_{pix} and obtain an effective ρ_t.

Note that some codes (Bristow 2003; Rhodes et al. 2010) include an (fairly crude) implementation of parallel transfer through n_{phases} multiple phases. Before running the algorithm, the image array is artificially expanded from $n_x \times n_y$ to $n_x \times n_{\text{phases}} n_y$ (with each row of pixels interleaved with $n_{\text{phases}} - 1$ rows of empty pixels), and trap properties are redefined as $\rho_t V_{\text{pix}}/n_{\text{phases}}$ and τn_{phases}. We have implemented this approximation and begun testing its consequences. For a volume-driven model assuming instantaneous charge capture, we find no difference except that short traps effectively delay fewer electrons (and if the trail profiles are exponential, this decrement is degenerate with their effective density ρ_t; Israel et al., in preparation).

4.1.5 Accounting for multiple transfers at once

We first test convergence of the forward CTI algorithm that adds trailing. A crucial balance between speed and accuracy in the forward algorithm is set by the 'express' E parameter (see Section 2.1.5). The change in measurements in the presence of CTI, compared to an ideal image, is shown in Fig. 5. These converge as E increases (not to zero, but to whatever value a long-winded CTI algorithm would predict).

In typical science data, to achieve ~ 1 per cent precision in most observables within a minimum CPU runtime, we recommend using $E \geq 5$. In images with low ρ_e/ρ_t (especially low background), the artificial flux losses introduced by the express approximation can be mitigated by increasing E. In dark exposures, we recommend that E should be ≥ 20. This slows runtime on the dark exposures by a factor of 4. However, this may not significantly affect *total* runtime, because (a) the number of dark exposures is typically much smaller than the number of science exposures, and (b) the 'high watermark' method (Section 2.1.6) greatly speeds execution when n_e is low, counteracting this effect.

---

[3] This limitation is present only *within* individual objects. The full *Gaia* data processing pipeline allows bright 'up-stream' sources (or charge injection lines) to pre-fill and shield traps from 'down-stream' faint sources by treating them as a variable background during (box 2) trap initialization (Prod'homme et al. 2012; Seabroke et al. 2012).

4.1.6 Release of electrons from charge traps

Since the CTI trails in *HST* observations are still well fitted by a sum of exponentials, we keep the Massey (2010) trap model with $n_{\mathrm{species}} = 3$ species that have exponential release in characteristic times $\tau = \{0.74, 7.7, 37\}$ pixels and relative densities $\{0.17, 0.45, 0.38\}$. This parametrization remains sufficiently flexible to easily model performance over the range of past operational temperatures.

4.1.7 Loop over columns

We tried recasting the algorithm to implement this loop via NVidia CUDA on a GPU. In practice, optimizing the code for GPUs is limited by the large amount of data needed to be transferred to and from the internal memory in administrating the traps. Also, our tests showed that double precision is indispensable in the calculations, so this limits the GPU hardware (using NVidia's CUDA) to expensive *Tesla* GPUs. We therefore find that, at present, it is thus more economic to perform parallel processing on a large multicore CPU system (either with openMP or just running simultaneous jobs) than on a GPU system.

In contrast, note that experiments at STScI found single point precision sufficient, and GPUs there have resulted in more than 100× speedups (Anderson, private communication).

4.2 Removing CTI trailing

Here, we describe and justify our choices for an iterative method to invert CCD readout and remove CTI trails at the pixel level. Sub-sections mirror those in Section 2.2.

4.2.1 Number of iterations for inversion

The iterative method (Fig. 3) operates for some number n_{iter} of cycles. Our tests on simulated images (black points in Fig. 6) show that $n_{\mathrm{iter}} \geq 6$ iterations are required for all measurements to converge, especially with the very faint (S/N = 10) galaxies.

In some situations, fewer iterations may be sufficient, or even preferred. Measurements of brighter galaxies (S/N = 50, 100) and low-order (e.g. astrometry) measurements of very faint galaxies have converged by $n_{\mathrm{iter}} = 3$. This takes only half the computation time. Furthermore, each iteration amplifies the read noise. The contribution of read noise to the autocorrelation function of test image pixels at $y = 2000$ is ideally 4.0, but 5.03 after three iterations, and 5.15 after 6 iterations. It may therefore be prudent to perform no more iterations than necessary.

Note that even with very many iterations, the algorithm does not converge to the desired solution (shown in each panel as a horizontal line), despite exactly the same algorithm being used to add and remove CTI trails. If read noise is not added to the simulations, the convergence is perfect. Two issues related to read noise cause the incorrect convergence. First, the read noise is amplified during CTI correction and, at low S/N, SOURCE EXTRACTOR in not stable to background noise level. Secondly, the read noise is spuriously untrailed during CTI correction, when it was never trailed to begin with. This introduces biases in even the more robust RRG measurements of astrometry and shape, which we shall attempt to deal with in the next section.

4.2.2 Dealing with read noise

We have verified that all the measurements converge to the correct answer (i.e. that in the absence of CTI) if we add CTI and immediately correct it using the same readout model. However, if we add read noise between adding and removing CTI, the correction converges to the wrong answer. If we blindly ignore the read noise, astrometric shifts for the faintest galaxies are undercorrected from about 0.36 pixels to 0.01 pixels (but not zero), and their ellipticity is overcorrected from $\langle e_1 \rangle = -0.15$ to $+0.01$. This latter overcorrection is roughly proportional to the amount of read noise squared (Israel et al., in preparation).

Our measurements of galaxy position and shape were obtained with RRG, which is carefully designed to be able to average away the effect of background noise from a large sample of galaxies. These (and RRG measurements of flux, which we have not shown) are well behaved during the correction, and even if the noise whitening process increases the amount of noise (see Section 2.2.3). Indeed, these measurements are completely unaffected by the addition of this correlated noise.

Measurements with SOURCE EXTRACTOR (flux, S/N, FWHM in Fig. 6) of very faint galaxies (S/N = 10) are not robust to the extra noise. In this case, the noise whitening process has actually made the correction worse – although the instability of the measurements suggests that this conclusion is random and that the correction was probably equally as likely to have got better.

We conclude from this that adding noise to an image after CTI correction does not improve any measurements, even if the noise is correlated in such a way to make the total noise white.

5 PERFORMANCE TESTS ON REAL DATA

The final proof that our choices are appropriate lies in the correct removal of charge trailing from real imaging. We shall now correct a large back-catalogue of ACS/WFC data, then test the shapes of warm pixels (which should become delta-function spikes) and galaxies (which should be randomly oriented).
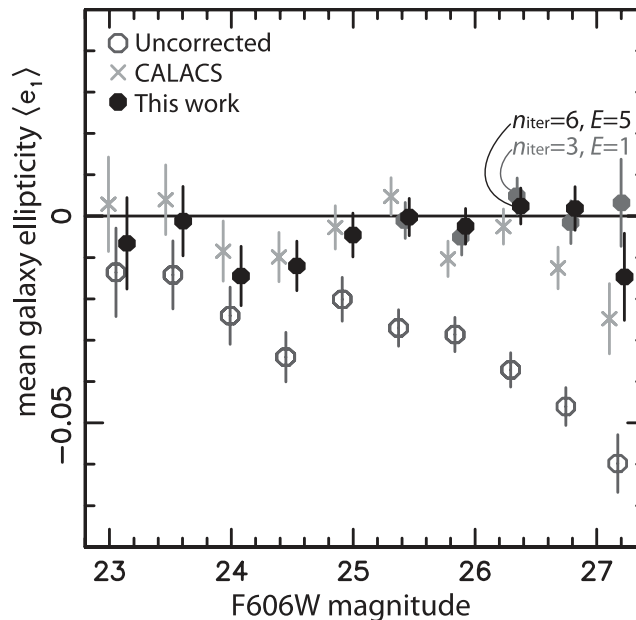
Figure 8. The measured shapes of faint galaxies in real *HST* ACS/WFC imaging from 2011 to 2012, as a function of their brightness in the *F*606W band (different points use the same magnitude bins, but have been offset for clarity). We show the mean ellipticity, averaged over a range of galaxy sizes and positions in the ACS field of view. Parallel CTI causes mean ellipticity $\langle e_1 \rangle < 0$ (net alignment along the *y*-axis). Spurious elongation in the raw data is successfully removed when our CTI correction algorithm is applied as the first step of data analysis (the choice of n_{iter} does not affect bright galaxies but its effect on noise is just discernible for the faintest). Note that the error bars are slightly underestimated because some galaxies are located in the overlapping regions of adjacent tiles, and contribute more than once.

5.1 Warm pixels

Recall from Section 3.1 that the density of charge traps is $\rho_t V_{pix} = (1.66 \pm 0.01) + (5.65 \pm 0.24) \times 10^{-4}(t - 3374)$ pixel^{-1}, where t is the number of days since launch. The performance of our CTI correction algorithm on warm pixels is shown by the points in Fig. 4 with dashed error bars. After correction with $n_{iter} = 6$ iterations, the residual trap density is well fitted by $\rho_t V_{pix} = (0.032 \pm 0.023) - (0.24 \pm 5.46) \times 10^{-5}(t - 3374)$ pixel^{-1}, where t is the number of days since launch. Comparing the constants in these two fits, we find that 98 per cent of the trail amplitude has been removed.

If we stop the iteration after $n_{iter} = 3$, the residual trap density (this is not shown in the figure) is fit by $\rho_t V_{pix} = (0.036 \pm 0.012) - (0.35 \pm 2.90) \times 10^{-5}(t - 3374)$ pixel^{-1}. This amplifies noise less (see Section 4.2.1), but produces almost as good a correction. In some scientific contexts where statistical errors are sufficiently large to dwarf the CTI residual, we therefore advocate the use of lower $n_{iter} = 3$.

As an incidental part of our testing, we performed a sensitivity analysis by rerunning the CTI correction for six of the data sets from late 2012 with 1 per cent fewer charge traps. This produced corrected images with a (1.10 ± 0.13) per cent smaller change in trail amplitude – consistent with the expectation that both the amount of trailing and the amount of correction depends linearly on the density of traps.

5.2 Extended sources

Since the warm pixels used to check the performance of our algorithm above are the same pixels that were used to develop and calibrate the model, being able to correct them is necessary but not sufficient. The performance of our CTI correction algorithm on independent images of real galaxies is demonstrated in Figs 8 and 9. We also compare this to the performance of on-the-fly Anderson & Bedin (2010) processing, using CALACS PixelCTE 2012 v3.2 (with reference table `w591643nj_cte.fits`). We used both packages to correct all the CANDELS *F*606*W* data taken in the COSMOS field between 2011 and 2012 (Grogin et al. 2011; Koekemoer et al. 2011), then measured galaxies' mean ellipticity $\langle e_1 \rangle$ (see equation 10). To correct for *HST*'s PSF, we followed the method described in Schrabback et al. (2010). Positive values of ellipticity indicate that the galaxies' major axes tend to be aligned with the detector's *x*-axis, and negative values indicate alignment with the *y*-axis.[4] If there is no preferred direction in the Universe, the mean ellipticity ought to be consistent with zero. This particular test is useful because it is self-contained (cf. tests on galaxy positions or fluxes require comparison to 'pristine' external imaging). It is also possible to perform this measurement very accurately, using software that has been developed to measure galaxy shapes at high precision for weak gravitational lensing analysis (e.g. Hoekstra & Jain 2008; Massey, Kitching & Richard 2010a).

---

[4] For some science analyses, it may be possible to correct any residual at the catalogue level by simulating the data, the radiation damage, and the imperfect CTI removal (e.g. Cawley et al. 2002; Leauthaud et al. 2007; Chiaberge et al. 2009). The correction will be smaller from the residual than from a raw measurement. However, accuracy relies on realism throughout the simulation, and every observable must be calibrated separately. It is therefore still preferable to meet requirements on CTI correction at the pixel level.
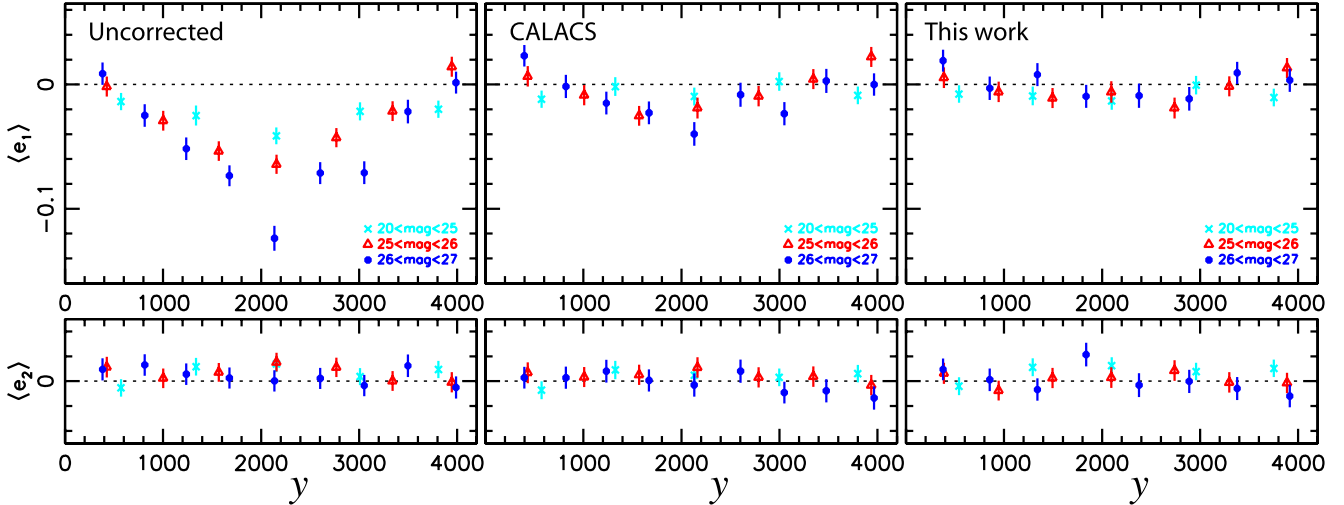
Figure 9. The measured shapes of faint galaxies in real *HST* ACS/WFC imaging from 2011 to 2012, as a function of their y position on the detector. The left-hand column shows raw data, and the other columns show analyses after correction with the public versions of different CTI correction packages (we use $n_{\rm iter} = 3$, which limits noise amplification and slightly improves performance for the faintest galaxies). The top (bottom) panels show the elongation of galaxies along (at 45° to) the y-axis. In this camera, readout nodes are at the bottom ($y = 0$) and top ($y = 4096$) of the detector, so the CTI is worst in the middle. Note that the error bars are slightly underestimated because some galaxies are located in the overlapping regions of adjacent tiles, and contribute more than once.

Parallel CTI is dominant in the raw images, producing a spurious negative ellipticity of ~0.12 for the faintest galaxies that are farthest from the readout register (averaging over all positions on the detector in Fig. 8 reduces this by half). Our measurement accuracy will be limited by statistical noise owing to the finite number of galaxies in even this very large data set. After correction, the measured ellipticities are consistent with zero, demonstrating that the CTI correction is successful.

The small residual variations from zero are probably due to imperfections in our model of the PSF. The PSF imprints itself on to the ellipticity of faint galaxies, and can only be removed to the accuracy of the PSF model. We are now actively rereducing ACS images of dense stellar fields, to remove CTI trailing and improve our PSF model.

6 CONCLUSIONS

The performance of the CCD detectors in *HST* ACS/WFC has steadily degraded since they were placed in orbit in 2002. They now exhibit strong trailing due to CTI, including a spurious elongation of faint galaxies. Such trailing now limits the application of *HST* for many high precision analyses, including supernova photometry, gravitational lensing and proper motions. If uncorrected, it may similarly limit science exploitation of future space missions including *Gaia* and *Euclid*.

6.1 Accurate CTI correction for HST

We have developed an algorithm to reproduce CTI trailing in *HST* ACS/WFC (Fig. 1). This can also be used iteratively to remove CTI trailing (Fig. 3). Compared to the Massey (2010) algorithm, which achieves ~95 per cent correction on recent data in various scientific contexts, our new algorithm is 8 times faster and 2.5 times more accurate (successfully removing ~98 per cent of the spurious trailing). These advances are due to a combination of algorithmic improvements and recalibration upon in-orbit data up to early 2013.

In particular, it is now important that an $E > 1$ CTI correction algorithm should be used for high precision applications. In the words of Anderson & Bedin (2010), the CTI trailing in ACS is rapidly becoming 'pathological', and more than a perturbation about the raw image. An implementation of our code in c can be downloaded from http://www.astro.uni-bonn.de/download/software/cte-tool/.

6.2 Meta-model from detailed literature survey

As part of the process of optimizing our CTI correction, we have compared (in detail) several successful algorithms from the literature that have been developed to model CCD readout. Our main finding is the remarkable similarities between even completely independent codes. Indeed, the codes can be summarized as sub-sets of an overarching model that takes several input parameters, described in Table 1.

In the trade-off between accuracy and speed of execution, the various codes each adopt a slightly different balance. They already share several computer–science tricks to speed up runtime, and we have incorporated others. We have suggested a balance in these trades appropriate for *HST* ACS/WFC.

In terms of physics, the codes' two main differences are: the procedure to determine which charge traps can (or are likely to) capture an electron, and the monitoring of which traps are full. Theory (Hall 1952; Shockley & Read 1952) indicates that the probability for capture, $p_{\rm capture}(\boldsymbol{x}, \sigma_t; n_e)$, should depend upon the traps' 3D position within a pixel $\boldsymbol{x}$ and capture cross-section σ_t. The position can be reduced to a

1D variable, by carefully ordering the sub-pixel volume. No working CTI correction algorithm has yet incorporated a full trapping model. Seabroke et al. (2011) retain 3D structure, and the integrals over the pixel volume take too long to be practical; Massey (2010) and Anderson & Bedin (2010) collapse p into a delta function; and Short et al. (2013) do not keep track of any dependence on $\boldsymbol{x}$.

6.3 Possible directions for future work

6.3.1 Specific recommendations for HST

An algorithm with variable $E > 1$ should be incorporated into the standard STSCI data analysis pipeline. To achieve 1 per cent CTI residuals, it will also be necessary to start correcting serial CTI.

The flexibility of Anderson & Bedin (2010)'s non-parametric form for $n_c(n_e)$ bodes well for the translation of ACS/WFC code to WFC3/UVIS data (Anderson et al., in preparation), in which the sky background can be very low. This has the ability to reproduce Short et al. (2013)'s $n_c(n_e)$ at low n_e (for a fixed clock speed and operating temperature), where experience from *Gaia* suggests a density-driven approach might be more suitable. A density-driven model might also provide an elegant way to implement Anderson & Bedin (2010)'s observation that fast-release traps may be filled before slow-release traps, in the regime where $n_e < n_c$.

6.3.2 Recommendations for future missions

Gaia should consider using $n_{\mathrm{levels}} > 1$ to monitor trap occupancy. Even a modest increase to $n_{\mathrm{levels}} \sim 50$ could improve accuracy, if the new trap levels are concentrated at low n_e where they will have the most effect. If the corresponding increase in run time is prohibitive, a similar saving could be achieved by accounting for only one species of trap instead of four, but with a complex release profile to represent the mixture of species.

Euclid requires 99 per cent correction of CTI in the shapes of the faintest galaxies at the end of the mission (Cropper et al. 2013; Massey et al. 2013), which is better than we have achieved for *HST* in this paper. Laureijs et al. (2013) had suggested noise whitening as a promising avenue to improve performance. We find that the basic implementation of noise whitening does not have the desired effect. However, steps have been taken to mitigate the problem via specifically optimized hardware, including slower readout to lower read noise (and avoid resonance between clock speed and trap release that affects observables of interest Rhodes et al. 2010) plus multilevel clocking (Murray et al. 2013a).

Noise added by CTI correction can be further reduced if the 3D locations of individual charge traps in damaged detectors are known. An STSCI program has been initiated to better measure these using long and short combinations of dark exposures taken during 'internal' *HST* time (http://www.stsci.edu/hst/wfc3/tools/cte_tools). A potentially even greater benefit – in both noise reduction and the calibration of software correction to unprecedented accuracy – may be obtained from in-orbit trap pumping measurements (Janesick 2001; Murray et al. 2013b). Where it is still possible to adapt readout electronics, we recommend that they should be enhanced to enable in-orbit trap pumping. Since CTI may otherwise seriously limit long-term science exploitation, modest investment in such hardware (and software to build models from trap pumping data) is likely to reap significant return.

We have not considered the effect of cosmic rays that hit during readout. Their associated cloud of electrons will not be subject to as many pixel-to-pixel transfers as others, so will not be as trailed. Identifying these cosmic rays will be especially important for missions with continuous or slow readout. Anderson & Bedin (2010) suggest an iterative procedure to identify and avoid undercorrecting them.

6.3.3 Recommendations for future software

We have explored most of the physics choices and the tradeoffs between speed and accuracy that are summarized in Table 1. However, a starting point for future software development (especially for new detectors) should probably be a meta-code that can easily take any or all of the options from Table 1 as input parameters. Code optimization starting from that point may further improve accuracy, over a wider range of operating conditions.

A useful addition to readout algorithms would be a (more sophisticated) treatment of multiphase clocking. Multilevel clocking (Murray et al. 2013a) could then be modelled by adjusting the direction travelled by released electrons. Multiple phases could probably be implemented in n_{phases} static events, but the dwell times should be allowed to differ between phases, and the changing electrostatic potentials may need different full well depth and filling parameters. If the charge clouds physically overlap in adjacent phases, it may even be necessary to use full 3D modelling like Hall et al. (2012) – or even allow charge to flow through the silicon at different speeds and densities at different positions. However, this would become very slow.

The most important improvement for iterative CTI correction will be a method to circumvent the limitations of read noise. Read noise is added after charge transfer, so is never trailed. Algorithms to undo the trailing may perfectly correct the underlying image, but 'overcorrect' the read noise. Anderson & Bedin (2010) proposed a tuneable high-/low-pass filter to identify some component of the read noise; we proposed a scheme to add correlated noise in such a way that the total noise becomes white. Our method works in other contexts (Huff et al. 2011) but did not have the desired effect here. Estimators for galaxy shapes, sizes, etc. that are carefully constructed to average out background noise,

are unaffected by adding more noise to a final image (by construction). However, even such carefully obtained measurements yield the wrong values if read noise is added between the addition and the removal of CTI trailing.

We remain unsure as to the exact mechanism through which read noise prevents perfect measurements. *Something* fundamental about the image is being irrevocably lost by the overcorrection of read noise – but it is not merely 'amplification' of read noise (Anderson & Bedin 2010) or even 'anticorrelation' of read noise (upon which assumption was our method based). We speculate that it may be due to a more subtle asymmetry in properties of the noise after correction, even though the read noise was Gaussian when it was added. Because $dn_c/dn_e < 1$, the number of traps to which each marginal electron gets exposed is smaller than the last. We therefore note that

(i) Positive fluctuations in read noise peaks are trailed (and untrailed) slightly less than negative fluctuations.
(ii) Read noise (or photon shot noise) on top of an object is untrailed less than those in the wings of an object or an area of blank sky.

There are also two asymmetries from one side of an object to the other.

(i) Read noise on the leading side of a source is untrailed more than read noise on the (shadowed) lee side.
(ii) Read noise on the side of an object closest to the readout amplifier is (very slightly) untrailed less than read noise far from the amplifier.

Determining which of these (or other) effects drives the current limitation is now beyond the scope of this paper. However, we recommend that it will be worth investigating the asymmetries of post-correction read noise in future work.

REFERENCES

Anderson J., Bedin L., 2010, PASP, 122, 1035
Bernstein G., Jarvis M., 2002, AJ, 123, 583
Bertin E., Arnouts S., 1996, A&AS, 117, 393
Bristow P., 2002, Instrument Science Report CE-STIS-2002-01
Bristow P., 2003, Instrument Science Report CE-STIS-2003-001
Capak A. et al., 2008, ApJS, 172, 1
Cawley L., Goudfrooij P., Whitmore B., Stiavelli M., , 2002, Instrument Science Report WFC3-2001-005
Chiaberge M., Lim P. L., Kozhurina-Platais V., Sirianni M., Mack J. the STScI CTE Working Group, 2009, Instrument Science Report ACS-2009-01
Clarke A., Hall D., Holland A., Burt D., 2012a, J. Instrum. 7, C01058
Clarke A., Hall D., Murray N., Holland A., Burt D., 2012b, Proc. SPIE 8453, 84531I
Cropper M. et al., 2013, MNRAS, 431, 3103
Ghez A. et al., 2008, ApJ, 689, 1044
Grogin N. et al., 2011, ApJS, 197, 35
Hall R., 1952, Phys. Rev., 87, 387
Hall D., Holland A., Murray N., Gow J., Clarke A., 2012, Proc. SPIE, 8453, 15
Hoekstra H., Jain B., 2008, Annu. Rev. Nucl. Part. Syst., 58, 99
Holland A., Abbey A., Lumb D., McCarthy K., 1990, Proc. SPIE, 1344, 378
Hook R., Stoehr F., 2008, Instrument Science Report WFC3 2008-014
Huff E., Hirata C., Mandelbaum R., Schlegel D., Seljak U., Lupton R. H., 2011, MNRAS, preprint (arXiv:1111.6958)
Janesick J., 2001, Proc. Vol. PM83, Scientific Charge Coupled Devices. SPIE, Bellingham, p. 920
Kaiser N., 2000, ApJ, 537, 555

Kitching T. et al., 2013, ApJS, 205, 12

Koekemoer A. et al., 2011, ApJS, 197, 36

Krist J., 1995, in Shaw R. A., Payne H. E., Hayes J. J. E., eds, ASP Conf. Ser. Vol. 77, Astronomical Data Analysis Software and Systems IV. Astron. Soc. Pac, San Francisco, p. 349

Krist J., Hook R., Stoehr F., 2011, Proc. SPIE 8127, 81270J

Laureijs R. et al., 2011, preprint (arXiv:1110.3193)

Leauthaud A. et al., 2007, ApJS, 172, 219

Massey R., 2010, MNRAS, 409, L109

Massey R., Rowe B., Réfregier A., Bacon D., Bergé J., 2007, MNRAS, 380, 229

Massey R., Kitching T., Richard J., 2010a, Rep. Prog. Phys. 73, 086901

Massey R., Stoughton C., Leauthaud A., Rhodes J., Koekemoer A., Ellis R., Shaghoulian E., 2010b, MNRAS, 401, 371

Massey R. et al., 2013, MNRAS, 429, 661

Melchior P., Viola M., 2012, MNRAS, 424, 2757

Miller L. et al., 2013, MNRAS, 429, 2858

Mostek N., Bebek C., Karcher A., Kolbe W., Roe N., Thacker J., 2010, Proc. SPIE 7742, 774216

Murray N., Burt J., Holland A., Stefanov K., Gow J., MacCormick C., Cropper M., 2013a, Proc. SPIE, 8860, 88600K

Murray N., Holland A., Burt D., Hall D., 2013b, Proc. SPIE 8860, 37

Ogaz S., Anderson J., Maybhate A., Smith L., 2013, Instrument Science Report ACS-2013-02

Perryman M. A. C. et al., 2001, A&A, 369, 339

Prod'homme T., Weiler M., Brown S., Short A., Brown A., 2010, SPIE, 7742, 29

Prod'homme T., Holl B., Lindegren L., Brown A., 2012, MNRAS, 419, 2995

Ragozzine B., Clowe D., Markevitch M., Gonzalez A., Bradač M., 2012, ApJ, 744, 94

Réfrégier A., Kacprzak T., Amara A., Bridle S., Rowe B., 2012, MNRAS, 425, 1951

Rhodes J., Réfrégier A., Groth, 2000, ApJ, 536, 79

Rhodes J., Leauthaud A., Stoughton C., Massey R., Dawson K., Kolbe W., Roe N., 2010, PASP, 122, 439

Schrabback T. et al., 2010, A&A, 516, 63

Seabroke G., Prod'Homme T., Hopkinson G., Burt D., Robbins M., Holland A., 2011, EAS, 45, 433

Seabroke G., Cropper M., Huckle H., Smith M., Benson K., Dolding C., Baker S., 2012, Gaia technical note GAIA-C6-TN-MSSL-GMS-003

Shockley W., Read W., 1952, Phys. Rev., 87, 835

Short A., Prod'homme T., Weiler M., Brown S., Brown A., 2010, in Holland A. D., Dorn D. A., eds, Proc. SPIE, 7742, 774212

Short A., Crowley C., de Buijne J., Prod'homme T., 2013, MNRAS, 430, 3078

Silvaco Inc., 2010, ATLAS User's Manual, Device simulation software, April 20

Suzuki N. et al., 2012, ApJ, 746, 85

Zhang J., Komatsu E., 2011, MNRAS, 414, 1047

Zuntz J., Kacprzak T., Voigt L., Hirsch M., Rowe B., Bridle S., 2013, preprint (arXiv:1302:0183)

This paper has been typeset from a TEX/LATEX file prepared by the author.