3-1-2010

Noise and Bias in Square-Root Compression Schemes

Gary M. Bernstein

University of Pennsylvania, garyb@physics.upenn.edu

Chris Bebek

Jason Rhodes

Chris Stoughton

R. Ali Vanderveld

See next page for additional authors

Follow this and additional works at: https://repository.upenn.edu/physics_papers

Part of the [Astrophysics and Astronomy Commons](#), and the [Physics Commons](#)

Recommended Citation

Bernstein, G. M., Bebek, C., Rhodes, J., Stoughton, C., Vanderveld, R. A., & Yeh, P. (2010). Noise and Bias in Square-Root Compression Schemes. Retrieved from https://repository.upenn.edu/physics_papers/295

Bernstein, C. M., Bebek, C., Rhodes, J., Stoughton, C., Vanderveld, R. A., & Yeh, P. (2010). Noise and bias in square-root compression schemes. *Publications of the Astronomical Society of the Pacific*, 122(889), 339-346. doi: 10.1086/651281

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/physics_papers/295
For more information, please contact repository@pobox.upenn.edu.

Noise and Bias in Square-Root Compression Schemes

Abstract

We investigate data compression schemes for proposed all-sky diffraction-limited visible/NIR sky surveys aimed at the dark-energy problem. We show that lossy square-root compression to 1 bit pixel⁻¹ of noise, followed by standard lossless compression algorithms, reduces the images to 2.5–4 bits pixel⁻¹, depending primarily upon the level of cosmic-ray contamination of the images. Compression to this level adds noise equivalent to $\leq 10\%$ penalty in observing time. We derive an analytic correction to flux biases inherent to the square-root compression scheme. Numerical tests on simple galaxy models confirm that galaxy fluxes and shapes are measured with systematic biases $\leq 10^{-4}$ induced by the compression scheme, well below the requirements of supernova and weak gravitational lensing dark-energy experiments. In a related investigation, Vanderveld and coworkers bound the shape biases using realistic simulated images of the high-Galactic-latitude sky. The square-root preprocessing step has advantages over simple (linear) decimation when there are many bright objects or cosmic rays in the field, or when the background level will vary.

Keywords

data analysis and techniques

Disciplines

Astrophysics and Astronomy | Physical Sciences and Mathematics | Physics

Comments

Bernstein, C. M., Bebek, C., Rhodes, J., Stoughton, C., Vanderveld, R. A., & Yeh, P. (2010). Noise and bias in square-root compression schemes. *Publications of the Astronomical Society of the Pacific*, 122(889), 339-346. doi: 10.1086/651281

Author(s)

Gary M. Bernstein, Chris Bebek, Jason Rhodes, Chris Stoughton, R. Ali Vanderveld, and Penshu Yeh

Noise and Bias In Square-Root Compression Schemes

GARY M. BERNSTEIN,¹ CHRIS BEBEK,² JASON RHODES,^{3,4} CHRIS STOUGHTON,⁵
 R. ALI VANDERVELD,^{3,4} AND PENSU YEH⁶

Received 2009 October 23; accepted 2010 January 13; published 2010 February 18

ABSTRACT. We investigate data compression schemes for proposed all-sky diffraction-limited visible/NIR sky surveys aimed at the dark-energy problem. We show that lossy square-root compression to 1 bit pixel^{−1} of noise, followed by standard lossless compression algorithms, reduces the images to 2.5–4 bits pixel^{−1}, depending primarily upon the level of cosmic-ray contamination of the images. Compression to this level adds noise equivalent to ≤10% penalty in observing time. We derive an analytic correction to flux biases inherent to the square-root compression scheme. Numerical tests on simple galaxy models confirm that galaxy fluxes and shapes are measured with systematic biases ≲10^{−4} induced by the compression scheme, well below the requirements of supernova and weak gravitational lensing dark-energy experiments. In a related investigation, Vanderveld and coworkers bound the shape biases using realistic simulated images of the high-Galactic-latitude sky. The square-root preprocessing step has advantages over simple (linear) decimation when there are many bright objects or cosmic rays in the field, or when the background level will vary.

Online material: color figures

1. INTRODUCTION

The unexpected discovery of the acceleration in the Hubble expansion of the universe has inspired numerous proposals for large-scale observational projects to seek further constraints on this “dark energy” problem (see Frieman et al. 2008 for a cof dark energy). Several such proposals involve space-based visible or NIR imaging surveys of a substantial fraction of the full celestial sphere at resolution near the diffraction limit of 1–2 m telescopes. Such surveys would also be extraordinarily valuable for general astrophysics, but the amount of data involved is large. The number of pixels in a survey of a fraction f_{sky} of the sky, with pixel scale p and N_{exp} exposures per sky location, will be

$$N_{\text{pix}} = 2.7 \times 10^{14} \left(\frac{p}{0.1''} \right)^{-2} \frac{f_{\text{sky}} N_{\text{exp}}}{5}. \quad (1)$$

A crate filled with disk drives is all that is needed these days for transmission of hundreds of terabytes of data between terrestrial sites. The transmission of these data from a spacecraft at the

Earth-Sun L2 Lagrange point 1.5 million km away is, however, a substantial expense. If the data require an average bits per pixel (bpp) to transmit, with a telemetry rate r , and a survey duration of T , then the number of hours of downlink per day required will be

$$\text{downlink time} = \text{bpp} \times 0.33 \frac{\text{hr}}{\text{day}} \frac{f_{\text{sky}} N_{\text{exp}} / T}{1.25/\text{yr}} \left(\frac{r}{150 \text{ Mbps}} \right)^{-1} \left(\frac{p}{0.1''} \right)^{-2}. \quad (2)$$

To avoid excessive resource use both on the spacecraft and ground stations, the data volume must be reduced well below the 16 bpp that is typical for digitization of astronomical array detectors. In this article we investigate means for compressing astronomical imaging data well below this, ideally bpp ≲ 4.

Significant compression of the detector output signals will require a lossy compression step to reduce the bit depth of the noise—sometimes called “preprocessing”—followed by a lossless algorithm to reduce the mean bit depth to the signal entropy level. Because most of the pixels in an extragalactic visible/NIR exposure will be close to the sky level, it is clear that the primary determinant of the final bpp value will be the number of noise bits retained by the lossy step in pixels near sky level—extensive demonstration of this for ground-based images is done by Pence et al. (2009). Optical astronomers are averse to lossy compression for fear of losing the information from their dearly acquired photons or fear of inducing biases in high-precision measurements. It should be pointed out, however, that lossy compression is already being performed on essentially all data when the analog detector outputs are

¹ Department of Physics and Astronomy, University of Pennsylvania, Philadelphia, PA 19104; garyb@physics.upenn.edu.

² Lawrence Berkeley National Laboratory, Berkeley, CA 94720.

³ Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109.

⁴ California Institute of Technology, Pasadena, CA 91125.

⁵ Fermi National Accelerator Laboratory, Batavia, IL 60150.

⁶ Goddard Space Flight Center, Greenbelt, MD 20771.

digitized. Analog-to-digital gains are usually set, however, so that the errors induced by digitization are below the read noise of the detector system. In radio astronomy it is well known that digitization to only 1 or 2 bits per sample can capture nearly all the signal information when the signal-to-noise ratio (S/N) per sample is low (Weinreb 1963; Thompson et al. 2001). The critical parameter for information loss is the number of bits of noise retained, i.e., the ratio $b = \sigma_N / \Delta$ of the signal noise to the digitization step size. We seek a lossy compression scheme which holds b fixed over the full dynamic range of the detector. When the signal noise is dominated by shot noise from the detected photoelectrons, it is straightforward to show that a *square-root compression scheme* holds b constant. The application of square-root compression to Poisson-dominated data has been discussed by Gowen & Smith (2003) and is part of the signal chain for the *James Webb Space Telescope* (Nieto-Santesteban et al. 1999) and *Supernova Acceleration Probe* (Lin & Marriner 2003) spacecraft designs. Further discussion is in Seaman et al. (2009). In § 2 we review the algorithms of such a scheme and § 3 shows that the square-root compression algorithm increases the noise by a fraction $1/24b^2$ for cases of interest.

Once a lossy compression algorithm is determined to reduce the data volume with acceptably small degradation of the noise level, it is critical to determine whether the codec process will induce any bias on the astronomical measurements from the data. Future dark-energy experiments will require very demanding, high-precision measurements. For use of Type Ia supernovae as distance indicators, systematic errors in flux measurements must be $\ll 1\%$. For determination of galaxy redshifts using the “photo-z” method, similarly small systematic errors must be present in galaxy photometry. We target codec-induced systematic errors in flux to be $< 10^{-3}$ so that they contribute negligibly to the photometric error budget.

Even more demanding are programs for weak gravitational lensing (WL) measurements, which depend upon determination of the shapes (ellipticities) of $\gtrsim 10^9$ galaxies on the sky. Huterer et al. (2006) and Amara & Réfrégier (2008) both conclude that measurements of galaxy ellipticities must have systematic errors of $\leq 10^{-3}$ in order to remain subdominant to statistical errors in the WL experiments. Because there are many possible sources of shape-measurement error—primarily uncorrected instrumental image distortions—we would like shape biases induced by the codec to be a minor contributor to the shape error budget, i.e., the shift δe_i in each component e_i of galaxy ellipticity should on average be $< 10^{-4} e_i$. The goal of this article is to examine the biases induced by square-root compression down to $b = 1$ bit of noise pixel $^{-1}$, and to show how these biases can be reduced to negligible levels. A companion article (Vanderveld et al. 2010, in preparation) tests for codec-induced shape bias on realistic simulations of space-based CCD imaging.

Many lossy image compression algorithms are based on decomposition of 2D images over a function library, followed by

truncation of the terms in the library—for example, wavelet transforms (White et al. 1992; Press 1992) or the JPEG algorithms. Even more radical compression is possible with “compressed sensing” techniques if the data are known to be sparse in some library (Candès et al. 2006). For WL applications, we are wary of lossy compression algorithms that examine groups of pixels and might hence induce correlated codec errors over multiple pixels. We therefore focus here on a lossy step that compresses data strictly on a pixel-by-pixel basis. If single-pixel codec errors are unbiased and completely uncorrelated between pixels, object shape measurements should be unbiased. We investigate the single-pixel codec in this work, leaving open the question of whether grouped-pixel algorithms can meet the shape bias requirements of weak lensing surveys while obtaining better compression than single-pixel codecs.

In § 4 we show that there is a significant bias in the mean values reconstructed by the square-root class of codecs. We also show, however, that a straightforward adjustment of the decoded values can remove this bias to remarkable accuracy, largely independent of the details of the probability distribution of the raw data. We therefore expect the codec to produce unbiased flux and shape measurements, which we test with simplified galaxy images in § 5. Vanderveld et al. (2010, in preparation) test for shape-measurement bias on data that are a realistic simulation of images expected from the *Joint Dark Energy Mission (JDEM)*⁷ CCD imager.

Finally in § 6 we compress realistic images with the lossy square-root algorithm followed by standard lossless algorithms, and measure the number of bits per pixel required to transmit the fully compressed images. In the conclusion we summarize our findings, namely that it is possible to compress *JDEM*-like images to 3–4 bits pixel $^{-1}$ with only 4% penalty in image noise, and to reconstruct them such that biases in flux and shape measurements induced by the codec are not detected, down to $\approx 10^{-4}$ levels.

2. COMPRESSION ALGORITHM

We consider most generally a compression algorithm that maps the digitized detector output value N into a compressed code i whenever $N_i - \Delta_i/2 < N \leq N_i + \Delta_i/2$. Here N_i and Δ_i are the center and width, respectively, of the data range that codes to compressed value i . Note that $N_{i+1} - N_i = (\Delta_i + \Delta_{i+1})/2$. We will assume that N is a debiased value that is linearly related to the estimated number of photoelectrons e by $N = e/g$, where g is the (inverse) gain of the system. We will assume initially that decompression consists of mapping $i \rightarrow N_i$, i.e., we restore code i to the mean of the input values that it codes. In a later section we will suggest an alteration to the decompressed values which reduces the bias in reconstructed values.

⁷ At <http://jdem.gsfc.nasa.gov>.

Any vector of increasing N_i (or equivalently a vector of $\Delta_i > 0$) defines a codec algorithm. We expect the bias and noise characteristics of the codec to depend primarily on the the number of bits of noise retained in the compression, which we define as

$$b = \sigma_N / \Delta, \quad (3)$$

where σ_N is the rms noise level of a pixel and Δ is the typical Δ_i for the codes i that span the noise distribution of the pixel. Many astronomical detectors have noise described by

$$\sigma_N^2 = [e + (RN)^2]/g^2 = (N + c)/g, \quad c \equiv (RN)^2/g, \quad (4)$$

with some read noise level, RN , added to the Poisson variance of the photoelectrons. In this case, if we wish to keep b independent of the signal level, we require

$$\Delta = \left[\frac{di}{dN} \right]^{-1} = \sigma_N / b \quad (5)$$

$$\Rightarrow \frac{di}{dN} = \frac{b}{\sqrt{(N + c)/g}} \quad (6)$$

$$\Rightarrow i = \sqrt{k(N + c)}, \quad k \equiv 4gb^2. \quad (7)$$

Since the compressed code i must be integral, we can more precisely express the nominal square-root compression algorithm as

$$i = [\sqrt{k(N + c)}] = [2b\sqrt{e + RN^2}], \quad (8)$$

where, in this equation only, $[x]$ is the nearest integer to x . Gowen & Smith (2003) discuss more general offsets under the square root, but we stick to these values since they hold σ_N/Δ fixed across the dynamic range.

2.1. Range-step Sequences

If we define

$$\Delta'_i = \Delta_i - \Delta_{i-1}, \quad (9)$$

then from equation (7) we find that the square-root codec has

$$\Delta_i \approx \frac{2\sqrt{k(N + c)}}{k} = \frac{2i}{k} \quad (10)$$

$$\Delta'_i \approx \frac{2}{k} = \frac{1}{2b^2g}, \quad (11)$$

i.e., the width Δ_i of each code range increases at a constant rate.

If the raw data N are integral, then these Δ_i and Δ'_i must also be integral. If, for example, we desire $b = 1$ bit of noise and have gain $g = 0.5e \text{ ADU}^{-1}$, then $\Delta'_i = 1$ produces the desired compression algorithm. When $1/2b^2g$ is nonintegral, however, we will more generally obtain $\langle \Delta'_i \rangle = 1/2b^2g$. In fact, we can define a class of compression algorithms as follows:

1. Choose a sequence of integers $\{s_1, s_2, \dots, s_n\}$ that will repeat cyclically to give the incremental range steps $\{\Delta'_1, \Delta'_2, \dots\}$.
2. Define the parameter b via

$$b^2 = \frac{1}{2g\langle s_i \rangle}. \quad (12)$$

3. Choose a range Δ_0 for the first compression code that extends from $N = 0$ to $N = (RN)/gb$. (The codec algorithm can be extended to small negative values of N if desired using the fixed range width Δ_0 for small negative i values.)

Then the codec defined by this range-step sequence will have nearly constant number b of bits per pixel noise level under the Poisson + read noise data model. The more nearly uniform the values s_i are, the better will be the codec performance, as shown in §§ 4.2 and 4.3. For example, we can produce a codec with $b = 1$ with $g = 1$ by using code steps $\Delta'_i = \{0, 1, 0, 1, 0, 1, \dots\}$. We will refer to this as a “ $g = 1, \{0, 1\}$ codec.”

2.2. Data Model

When compressing integer data, we denote the probability of the raw pixel value being N as P_N . For notational convenience we will assume that there exists a continuous and differentiable function $P(x)$ over $x \geq 0$ which matches P_N at $x = N$. For example, the Poisson formula for P_N is easily extended to non-integral values with a gamma function.

In the CCD data the Poisson signal has Gaussian read noise added to it. We will assume the read noise has zero mean. Hence we can consider the predigitization pixel photoelectron count e to be a real number, with probability distribution $P(e)$. In the limit of high gain g , we can consider the detector output to be a real number e or N ; at finite gain, P_N is again defined only at integral N , but we can assume there exists some analytic extension to the full real domain $P(N)$.

Most of our results are general to any well-behaved $P(N)$ or $P(e)$, and we will note when the assumption of Poisson or Poisson + Gaussian noise models are important.

3. CODEC NOISE

A raw pixel value N will incur an error $\delta N = N_i - N$ upon being encoded and decoded through code i . If the input values are uniformly distributed across the range of code i , then the rms error induced by the codec is $\sigma_{\text{codec}} = \Delta_i/\sqrt{12}$. If we make the further simplifying assumptions that the coding error is

uncorrelated with the raw N and that Δ is nearly constant for all i accessed by $P(N)$, then the variance of the output is

$$\sigma_{\text{out}}^2 = \sigma_N^2 + \sigma_{\text{codec}}^2 = \sigma_N^2 \left(1 + \frac{1}{12b^2} \right). \quad (13)$$

The integration time required for a shot-noise-limited astronomical exposure to reach fixed S/N will be inflated by

$$\frac{\sigma_{\text{out}}^2}{\sigma_N^2} = 1 + \frac{1}{12b^2}. \quad (14)$$

If we want the codec to cause $<10\%$ penalty in exposure time, we hence need $b \geq 0.91$. In this article we will focus our attention on this question: *does a square-root algorithm with $b \geq 1$ cause any bias or degradation in astronomical measurements beyond the expected $<4\%$ increase in noise level?*

A floor on the bandwidth required to transmit real images is given by the bpp found for pure-noise images. The minimum bpp required for lossless compression of a sequence of codes i with probabilities p_i is the Shannon entropy $H = -\sum p_i \log_2 p_i$. For nearly-Gaussian probability distributions digitized to b bits per σ , the Shannon entropy is $H \approx \log_2 \sqrt{2\pi e} + \log_2 b$ (e.g., Romeo et al. 1999). For $b = 1$, a numerical evaluation of the Shannon entropy shows that ≥ 2.1 bpp will be required. We construct a 16-bit integer FITS image consisting of random noise with $\bar{e} = 40$ and $\text{RN} = 4$, apply square-root compression with $b = 1$, then apply lossless bzip2 or CCSDS 121B (CCSDS 1997) compression. The resulting image requires 2.4–2.5 bits pixel $^{-1}$, within 0.3–0.4 bpp of the Shannon limit. We can therefore consider 2.4 bits pixel $^{-1}$ to be a floor on the telemetry required when using a pixel-by-pixel lossy compression step, subject to $\leq 10\%$ integration-time penalty from codec errors. We note that the “raw” data digitized at $g = 1$ would have $b = 7.5$ and require 2.9 additional bits per pixel. We verify that with square-root compression enabled, the bpp for pure-noise images is largely independent of $\bar{e}$, RN , and g , as expected.

4. RECONSTRUCTION BIAS

The expectation value of the raw data $\bar{N} = \langle N \rangle$ should be unchanged after compression and decompression. We will denote by $\delta\bar{N}$ the bias in the expectation value of the data after the codec operates. In practice this bias is not a problem if it is independent of $\bar{N}$, since it is easily corrected. If $\delta\bar{N}$ varies with $\bar{N}$ and/or with the details of the $P(N)$ distribution, then the response of the detector effectively becomes nonlinear and/or ambiguous to correct, which can bias photometry and shape measurements.

Any codec bias could of course be corrected if the raw distribution P_N were known at each pixel. But the mean level $\bar{N}$ will vary from pixel to pixel due to background gradients and sources, so there is no single P_N that applies to the entire image.

We would prefer to debias the data by decompressing code i to $N_i - \delta_i$, where δ_i is some small correction that is independent of the raw probability distribution. There is no guarantee that such a solution exists but we will show that a good approximation does exist for cases when $P(N)$ spans multiple codes.

4.1. Real-valued Data

We begin with the case in which the raw value N can be considered a continuous real variable. This case is instructive if not necessarily realistic. The center and width N_i and Δ_i of the range for each code i can be set to any positive real numbers. The codec bias can be expressed as

$$\delta\bar{N} = -\sum_i \int_{-\Delta_i/2}^{\Delta_i/2} du u P(N_i + u), \quad (15)$$

where $u = N - N_i$ is the distance from the center of code step i . If we can find a set of δ_i such that

$$\sum_i \int_{-\Delta_i/2}^{\Delta_i/2} du \delta_i P(N_i + u) = \delta\bar{N} \quad (16)$$

holds for any distribution $P(N)$, then we will have a complete debiasing prescription for the codec.

When the distribution $P(N)$ is wider than a code step width Δ , it can be approximated by the Taylor expansion of $P(N)$ about N_i :

$$P(N_i + u) = P_i + P'_i u + P''_i u^2/2 + P'''_i u^3/6 + O(u/\sigma_N)^4, \quad (17)$$

with $P'_i = P'(N_i)$, etc. We will truncate the series at cubic order henceforth. The contribution of the neglected terms will scale as $(\Delta/\sigma_N)^4 = b^{-4}$. We will later perform a numerical check of the significance of the neglected terms. Substituting the Taylor expansion into equation (15) gives

$$\delta\bar{N} \approx -\sum_i \Delta_i [P'_i \Delta_i^2/12 + P'''_i \Delta_i^4/480]. \quad (18)$$

If this sum over i were an integral, we would evaluate this quantity through integration by parts. We can perform an analogous discrete process by defining $P_{i\pm 1/2} \equiv P(N_i \pm \Delta_i/2)$ and noting that to our chosen order of Taylor expansion, we have

$$\Delta_i P'_i = P_{i+1/2} - P_{i-1/2} - \frac{\Delta_i^3}{24} P'''_i. \quad (19)$$

Substituting this into equation (18) yields

$$\delta\bar{N} = -\sum_i \left[\frac{\Delta_i^2}{12} (P_{i+1/2} - P_{i-1/2}) - \frac{\Delta_i^5}{720} P'''_i \right] \quad (20)$$

$$= \sum_i \left[\frac{\Delta_i}{12} (\Delta'_i P_{i-1/2} + \Delta'_{i+1} P_{i+1/2}) + \frac{\Delta_i^5}{720} P_i'''' \right]. \quad (21)$$

Recalling that $\Delta'_i = \Delta_i - \Delta_{i-1}$, it is clear first that *the codec bias is very small if the step width Δ is constant* ($\Delta' = 0$). This would be the case if the codec were a simple decimation over the range of the data. This is the case for normal digitization of the data.

The perfect square-root compressor of real-valued raw data equation (8) has the unique property that $\Delta'_i = \Delta'$ is constant. In this special case we can simplify further and use the Taylor expansion one more time to get $P_{i+1/2} + P_{i-1/2} = 2P_i + \Delta_i^2 P_i''/4$, and now we have

$$\delta \bar{N} = \sum_i \Delta_i \left[\frac{\Delta'}{6} \left(P_i + \frac{\Delta_i^2}{8} P_i'' \right) + \frac{\Delta_i^4}{720} P_i'''' \right]. \quad (22)$$

Second, it is clear that *the square-root codec bias will be nearly independent of $P(N)$* because $\sum_i \Delta_i P_i \rightarrow 1$ in the limit of small Δ_i , so we expect there to be bias

$$\delta \bar{N} = \frac{\Delta'}{6} = \frac{1}{12b^2g}. \quad (23)$$

A constant bias is relatively benign; for example, it does not influence the appearance of a sky-subtracted image. It is also easily remedied by decompressing code i to value $N_i - \delta^{(0)}$, where $\delta^{(0)} = \Delta'/6$. After making this adjustment, the residual bias, to third order in the Taylor expansion of each code range, becomes

$$\delta \bar{N} = \sum_i \Delta_i \left[\frac{\Delta' \Delta_i^2}{72} P_i'' + \frac{\Delta_i^4}{720} P_i'''' \right]. \quad (24)$$

Further application of the “integration by parts” technique shows that the leading terms in this residual bias scale as $\int dN \Delta' P'(N)$, which is zero if Δ' is constant across N . Hence we expect the reconstruction bias to be very close to the constant $1/12b^2g$ for the reconstruction of real-valued raw data with $b \geq 1$, i.e., whenever the intrinsic noise spans multiple compression code steps.

To summarize the key results: a simple decimation codec ($\Delta_i = \text{constant}$) reproduces the mean of the input data with no bias, and the square-root codec ($\Delta'_i = \text{constant}$) biases the mean by the constant value in equation (23) independent of the distribution $P(N)$ of the raw data. These statements are very good approximations when $b \geq 1$ and the raw data are real-valued.

4.2. Integer Data

Next we consider the case when the raw data N must be integral. In this case Δ_i and Δ'_i must be integral, and N_i is integral or half-integral depending upon whether Δ_i is odd or even,

respectively. The integrations over N in equation (15) must be replaced by sums over the allowed integral values of N . We obtain

$$\delta \bar{N} = \sum_i \sum_j j P_{N_i+j}, \quad (25)$$

where the sum runs over $j \in \{-(\Delta_i - 1)/2, -(\Delta_i - 3)/2, \dots, (\Delta_i - 3)/2, (\Delta_i - 1)/2\}$.

We smoothly extend the discrete probabilities P_N to a continuous function $P(N)$ that maintains the normalization condition $\int dN P(N) = 1$, and we again perform a third-order Taylor expansion of $P(N)$ about the center N_i of each code range. The bias can now be written

$$\delta \bar{N} = - \sum_i \left[P'_i \sum_j j^2 + \frac{P_i'''}{6} \sum_j j^4 \right]. \quad (26)$$

The sums of powers of integers (or half-integers) can be found from well-known formulae.⁸ The result is

$$\delta \bar{N} = - \sum_i \frac{\Delta_i}{12} (\Delta_i^2 - 1) \left(P'_i + \frac{3\Delta_i^2 - 7}{120} P_i''' \right). \quad (27)$$

The discrete version of integration by parts transforms this into an expression identical to equation (21) up to small corrections in the P''' term:

$$\delta \bar{N} = \sum_i \frac{\Delta_i}{12} \left[\Delta'_i P_{i-1/2} + \Delta'_{i+1} P_{i+1/2} + \frac{(\Delta_i^2 - 1)(2\Delta_i^2 + 7)}{120} P_i''' \right]. \quad (28)$$

For an integer-valued codec, we cannot necessarily maintain constant Δ'_i for an arbitrary b ; for example, with $g = 1$ and $b = 1$, we must have $\langle \Delta' \rangle = 1/2$. Hence we must choose adjustments δ_i to each decompressed value that will best cancel the bias in equation (28). After further manipulation it is possible to show that the leading-order terms in the bias are canceled, independent of the form of P_N , if

$$\delta_i = \frac{\Delta'_{i+1} + \Delta'_i}{12} - \frac{\Delta'_{i+2} - \Delta'_{i+1} - \Delta'_i + \Delta'_{i-1}}{48} \quad (29)$$

$$= \frac{\Delta_{i+1} - \Delta_{i-1}}{12} - \frac{\Delta_{i+2} - 2\Delta_{i+1} + 2\Delta_{i-1} - \Delta_{i-2}}{48}. \quad (30)$$

To summarize: the compression algorithm is defined by a repeating integer sequence $\{s_i\}$ for the Δ' values, along with a gain g and read noise RN, as described in § 2.1. The compressed

⁸ For example, <http://mathworld.wolfram.com/FaulhabersFormula.html>.

values will have $b = (2g\langle s_i \rangle)^{-1/2}$ bits per σ of input noise. On decompression, each compressed code i should be replaced with the value $N_i - \delta_i$, with δ_i given by Equation (30).

One should prefer the most uniform possible sequences $\{s_i\}$. A constant value (nominally $\Delta' = 1$) will yield a “perfect” square-root compression with minimal reconstruction bias. An alternating sequence, e.g., $\{0, 1\}$, has a constant bias correction δ_i , but slightly larger biases. Longer $\{s_i\}$ sequences have more complicated and less precise bias corrections.

4.3. Numerical Validation

Figure 1 shows the noise and bias introduced by several versions of the square-root codec. In each case the raw data are drawn from a Poisson distribution with the mean number of photoelectrons plotted on the x axis, ranging from 40–100 e . The data are taken to have Gaussian read noise with zero mean and 4 e rms. Each plotted point is the result of 10^7 trials from this distribution. Each trial produces an integral raw N value and a real-valued decompressed value $N_i - \delta_i$.

We examine 4 different codecs in Figure 1. (See online version of figure for the color details described here.) The red, blue,

FIG. 1.—Performance of four square-root codecs on data with Poisson noise plus 4 e rms Gaussian read noise. *Bottom panel*, the center values N_i and the range widths Δ_i of each step for the four codes. The *solid line*, *dotted line*, and *long-dashed line* codecs each produce $b = 1$ bit per noise σ_N of the input value, with decreasing levels of smoothness in the derivative Δ' of the range width. The *short-dashed line* codec retains more noise, $b = 1.22$. *Top panel*, the increase in data rms due to the codec, which is independent of input level. *Central panel*, the bias in the reconstructed data stream, with the noise level of this measure indicated by the black error bar. Bias levels are $O(0.001e)$, in some cases consistent with zero within measurement noise. See text for further detail. See the electronic edition of the *PASP* for a color version of this figure.

and green curves are designed to produce $b = 1$ bit of noise. The red curve is a “perfect” square-root codec, with $\Delta' = 1$ and $g = 0.5e \text{ ADU}^{-1}$. The blue and green curves have $g = 1$ and $\{s_i\}$ sequences of $(0,1)$ and $(0,0,1,1)$, respectively. The magenta curve is a codec with $g = 1$ and step sequence $(0,0,1)$, which retains $b = 1.22$ bits of noise.

The bottom panel plots the mean gN_i and range $g\Delta_i$ of each code step i for each of the four codecs, expressed in units of photoelectrons. The top panel shows the rms fluctuation of the decoded data, relative to the RMS fluctuation of the input data. All of the codecs inflate the rms by almost exactly the factor $(1 + 1/12b^2)^{1/2}$ predicted in equation (13). In particular, the $b = 1$ codecs inflate the noise by 4%.

The central panel plots the mean difference between decompressed and raw data. The expected noise in this measurement from the 10^7 trials is plotted in black at left. The “perfect” codec, in red, has zero bias within errors, $|\delta\bar{N}| < 0.001$. This bias is, remarkably, $<10^{-5}$ of the input values, $<10^{-4}$ of the typical noise σ_N and code range $\Delta = 7$ –10, and also $<1\%$ of the bias $1/12b^2 = 0.083$ that the codec would have generated without any correction. The $\{0, 1\}$ -sequence, in green, shows some significant residual bias at the ± 0.002 level, and the $\{0, 0, 1, 1\}$ -sequence codec (blue) has bias of $\approx \pm 0.004$ electrons. While these biases are quite small, they do highlight that equation (30) ignores some higher-order terms that make the bias dependent on P_N , and that it is preferable to use the smoothest possible step sequences. The $\{0, 0, 1\}$ codec also has some bias at the 0.001 level, despite having slightly less compression and lower noise than the “perfect” $b = 1$ codec in red.

These numerical tests insure that square-root compression at $b = 1$ will introduce no significant biases into photometric measurements, and increases in errors will be limited to 4%.

5. SHAPE MEASUREMENT BIASES

Using simple simulated galaxies and a very basic shape-measurement routine, we demonstrate here that the codec does not induce any biases in the measurements of galaxy fluxes, sizes, or ellipticities at the part-in- 10^4 level. In this work we analyze simple galaxies with specified size and signal-to-noise ratios which span a range that we might expect in real images. A test for codec with much more realistic images is presented by Vanderfeld et al. (2010, in preparation). In that work the galaxies have sizes, appearances, fluxes, and shapes that accurately mimic those we might expect from a visible survey telescope in orbit.

5.1. Procedure

The simulated images are postage stamps constructed as follows:

1. Start with an exponential-profile galaxy, with intensity $I \propto e^{-r/r_0}$. The exponential disk is convolved with a Gaussian of width $\sigma = r_0/10$, to suppress high frequencies associated with the $r = 0$ cusp of the exponential disk.

2. This circular galaxy is then sheared to an ellipticity $e = (a^2 - b^2)/(a^2 + b^2) = 0.1, 0.3, \text{ or } 0.7$ (a and b are the major and minor axes of the galaxy). The major axis of the elliptical galaxy can have position angle $\phi = 0, 22.5^\circ, \text{ or } 45^\circ$ from the x -axis.

3. The galaxy is (de)magnified until the resulting galaxy has half-light radius of $r_{1/2} = 2, 4, \text{ or } 6$ pixels.

4. The galaxy profile is shifted by a random fraction of a pixel in x and y , then convolved with the square pixel response function and sampled on a 1 pixel grid to yield a noiseless postage-stamp image.

5. The flux of the galaxy is adjusted until its total $S/N = 15, 50, \text{ or } 300$ when noise is added in the following step. The S/N is defined as

$$(S/N)^2 = \sum_{i \in \text{pixels}} \frac{I_i^2}{\sigma_i^2} \quad (31)$$

where I_i and σ_i are the signal and noise in each image pixel.

6. A sky level of 50 electrons is added to each pixel's expected signal I_i . Then an integral "observed" photoelectron count for each pixel is determined by drawing from a Poisson distribution with mean I_i . Gaussian read noise with zero mean and $\sigma = 4$ electrons is added to give a floating-point measurement value for each pixel.

At this point two digitized versions of the postage-stamp image are produced. For the "codec" image, the photoelectron count I_i is scaled by gain $g = 0.5e \text{ ADU}^{-1}$, digitized, compressed with a sequence 1 perfect square-root codec, and then decompressed to a floating-point image following the debiasing procedures in § 4.2. Note that the codec retains $b = 1$ bits per noise σ on each pixel. For the "raw" image, we add Gaussian noise with rms of $\sigma_i/\sqrt{12b}$ to each pixel, so that the noise level of the raw and codec images will be expected to be equal. Then we scale by the gain g and digitize *without* square-root preprocessing.

Both the raw and codec versions of the postage-stamp galaxy are measured using a basic "adaptive moments" shape-measurement algorithm (Bernstein & Jarvis 2002) as follows: the pixel data are fit to an elliptical Gaussian model, with the flux, centroid, size, and ellipticity of the Gaussian allowed to vary. The sky background is fixed at $50/g = 100 \text{ ADU}$. There are hence 6 degrees of freedom in the galaxy model (and the Gaussian galaxy model is *not* an exact fit to the underlying exponential-profile galaxy).

5.2. Results

For each choice of ellipticity e , orientation ϕ , and size $r_{1/2}$ of galaxy, we create 1000 simulated galaxies at $S/N = 300, 30,000$ realizations with $S/N = 50$, and 300,000 realizations with $S/N = 15$. To test for biases induced by the codec, we measure the mean difference $e_{\text{codec}} - e_{\text{raw}}$ between the ellipticities

measured on the codec and the raw images. Similarly we measure the mean difference in the (log of) flux and size measured for the galaxies. To see if the codec induces extra shape noise, we compare the rms variation of the fitted e in the codec data set to the RMS variation in the raw fits. We also check for additional variance in flux and size measurements.

Figure 2 presents the biases and noise levels on the ellipticity, size, and flux measurements of galaxies with $r_{1/2} = 2$ pixels and $\phi = 0$. The results for larger galaxies and for different ellipticity orientations are similar and not plotted here. We find no evidence for any codec-induced bias in size, flux, or ellipticity measurements at any S/N or ellipticity. The number of trials is sufficient to determine these biases to $\approx \pm 10^{-4}$ for each configuration, approximately 10 times below the maximum permissible ellipticity measurement errors in a large WL survey (Amara & Réfrégier 2008; Huterer et al. 2006). For these simplified galaxies and shape-measurement techniques, we confirm that the square-root codec at $b = 1$ does not induce any significant bias in weak-lensing shape measurements.

FIG. 2.—Biases (left) and relative noise level (right) of measurements on codec images, relative to raw images, are plotted for several quantities measured on simulated galaxies. See text for details of the simulation and measurement process. From top to bottom, we plot the bias in the (log of) galaxy flux, (log of) galaxy size, and the two components of galaxy ellipticity: symbols with white center are for galaxies with $S/N = 15$, filled symbols have $S/N = 50$, and filled circles with black outline have $S/N = 300$, with each point plotted at its input ellipticity. All data are fully consistent with the codec images yielding no bias or change in noise level from the raw images—at accuracies of $\pm 10^{-4}$ for biases and $< 1\%$ in noise levels. (Noise levels for $S/N = 300$ cases are not plotted because there are too few trials for an accurate measure of rms noise.) See the electronic edition of the *PASP* for a color version of this figure.

For the $S/N = 50$ and $S/N = 15$ cases, we have enough trials to determine the measurement noise in e , size, and flux to $\ll 1\%$. We find the raw and codec images to have identical measurement noise levels to this accuracy, confirming that the noise added by the codec is accurately described by equation (13).

6. COMPRESSION FACTORS

6.1. Simulated Scene

We create simulated images from a space-based visible survey telescope to estimate the bits of telemetry required per pixel (bpp) needed to transmit the images after they have been compressed with the square-root algorithm followed by a lossless compression algorithm. The simulated images have three components: the expected cosmic scene; a uniform sky background; and bright linear features produced by cosmic rays (CR) in the detectors. Our test images are aimed at reproducing the appearance of images taken with an instrument resembling the CCD camera proposed for the *JDEM/IDECs* mission. The relevant characteristics assumed are: a 1.5 m diameter telescope operating in an L2 Lissajous orbit, where the sky background is dominated by zodiacal light; a 200 s integration through a bandpass in the red part of the visible spectrum, with bandwidth $\ln(\lambda_{\max}/\lambda_{\min}) = 0.3$; a fully-depleted, 200 μm -thick p-channel CCD (Roe et al. 2007); a scale of $0.14'' \text{ pixel}^{-1}$; and a point-spread function (PSF) with EE50 radius $\approx 0.14''$, after including all optical, mechanical, and detector-induced sources of image blur.

The cosmic scene is taken from deep moderate-latitude exposures taken with the ACS/WFC F606W filter aboard the *Hubble Space Telescope* (*HST*) in the course of a search for very faint outer solar-system objects (Bernstein et al. 2004). These exposures combine many *HST* exposures to produce an image whose noise is far below that expected from a single *JDEM/IDECs* exposure as taken here. The scene is scaled so that a source of magnitude $I_{AB} = 31.09$ produces 1 photoelectron. The *HST* image is blurred and repixelized to match the *JDEM/IDECs* specifications described here.

Diffuse sky emission is added assuming high-ecliptic-latitude zodiacal background at L2. We estimate ≈ 45 electrons pixel^{-1} in the detector.

6.2. Simulated Cosmic Rays

Cosmic rays in space are isotropic and uniformly impinge on a sphere. From ESA's Space Environment Information System (SPENVIS)⁹ program, the CR rate is found to be approximately 5 particles $\text{s}^{-1} \text{ cm}^2$ over 4π steradian. This is the average rate for Galactic cosmic rays, which are dominated by protons. Low-energy particles (mainly solar protons) have little integrated effect on the total rate even without shielding. We ignore periods

of enhanced low-momentum solar protons during solar storms. These episodes are known in advance and data taking is avoided.

We use SPENVIS to generate a momentum spectrum for a particular orbit (L2 in this case) at a particular solar epoch behind appropriate shielding (3 mm Al assumed here). The Monte Carlo simulation samples this spectrum and for each momentum, an average ionization rate ($e \mu\text{m}^{-1}$) is calculated for the target material (Si in this case) according to a modified Bethe-Bloch formalism (Amsler et al. 2008, § 27). For the data generated and analyzed here, a single particle momentum was used corresponding to minimizing protons—approximately 3 GeV c^{-1} , corresponding to $70e \mu\text{m}^{-1}$ in silicon.

For a given exposure time, particle flux, and detector area, the total number of cosmic rays is calculated. We generate this number of cosmic rays uniformly over 4π sr. Each ray is randomly placed over the surface of detector and tracked through the detector material. Tracing is done in uniform step lengths. For each step, a Landau distribution is sampled to determine the number of ions produced. Hard scatters and δ -rays are not simulated. For each ion generated in a step, a Gaussian lateral charge diffusion distribution is randomly sampled. The diffusion constant is based on a measured 4 μm rms for a 200 μm -thick CCD operated at 100 V bias voltage. This is linearly scaled by ion conversion height above the pixel plane. The diffused charge is projected to the pixel plane where its position is discretized and added to any previous pixel charge.

In the simulated 200 s exposure, 7% of the pixels receive $\geq 5e$ of CR-generated charge. We also simulate an image with triple this dose of cosmic rays, which is simply the sum of 3 200 s CR simulations.

6.3. Compression Results

The scene, CR, and sky components are summed, and then Poisson noise and $4e$ rms read noise are added to each pixel. The noisy images are then scaled by gain $g = 0.5e \text{ ADU}^{-1}$ and digitized, then compressed using a perfect square-root algorithm with sequence 1. The compressed images are saved as 16 bit integer FITS-format images. A segment of a simulated image, before compression, is shown in Figure 3.

The Shannon entropy of the square-root-compressed image, calculated by direct measure of the p_i of each code in the image, is 2.8 bits pixel^{-1} . Tripling the CR dose on the exposure increases the entropy to 3.6 bits pixel^{-1} , suggesting that the normal dose of cosmic rays is responsible for 0.4 bits pixel^{-1} of entropy in the image. Thus most of the 0.7 bit increase in entropy over the pure-noise image (as measured in § 3) is due to cosmic rays.

There are many possible lossless compression algorithms to apply to the square-root-compressed FITS images. This choice will, of course, not affect the biases in the data, but will affect the compression ratio and the robustness of the telemetry

⁹ At <http://www.spennis.oma.be/spennis/>.

FIG. 3.—Section of an image used to test the square-root + lossless compression scheme. This simulates a 200 s CCD exposure with a 1.5 m space telescope observing near the ecliptic pole from an L2 orbit. Most pixels are dominated by sky and read noise, and celestial objects are visible, but the CR tracks dominate the information content of the image. The scale is in ADU, and the image gain is 0.5 photoelectrons ADU^{-1} . The image requires 3.1 bits pixel^{-1} to transmit after application of square-root compression and lossless bzip2 compression.

stream. Applying the commonly available bzip2 algorithm to the FITS files yields the following results:

1. The images simulating a 200-second exposure and 200 seconds' worth of cosmic-ray tracks occupy 3.1 bits pixel^{-1} after square-root and bzip2 compression.
2. The information content of the images is dominated by the cosmic-ray component. If we triple the number of cosmic rays (mimicking a 600-second exposure, even though the cosmic scene and sky background are still 200 seconds' worth), we get a pessimistic value of 4.0 bits pixel^{-1} in the compressed images.

The bzip2 performance is again 0.3–0.4 bpp above the entropy of the image.

Full-frame bzip2 compression is not very robust: a single-bit transmission error can lead to loss of a full image. A more robust algorithm is CCSDS 121B (CCSDS 1997), which has already been implemented on >25 space missions. We find the file size after square-root preprocessing and CCSDS 121B lossless compression to be within 0.1 bits pixel^{-1} of the bzip2 results.

A different lossless compression scheme, optimized toward efficient compression of the linear cosmic-ray tracks that dominate the 2D image, might yield still better compression. If further compression were desired, the most effective approach would be to identify CR-contaminated pixels on that satellite and replace them with a constant tag value before compression, so that their entropy is not transmitted to the ground.

6.4. Comparison to Linear Preprocessing

Implementation of square-root preprocessing is very simple, requiring just a fixed lookup table to convert ADC output values to compressed codes. An even simpler algorithm is linear pre-scaling, i.e., decimation, in which the code range Δ is fixed over the full dynamic range. As has been noted, a decimation algorithm should be free of bias, just as the corrected square-root algorithm is measured to be free of bias. What are the relative merits of the two lossy codecs?

We compress our test images with a linear algorithm with code step Δ chosen to yield $b = 1$ steps per noise σ at the sky level of the images. As noted, e.g., in Pence et al. (2009), most pixels are at sky level, so a fixed b at sky level should yield similar compression performance. Indeed, for images with *no* cosmic rays, we find that linear and square-root compression yield indistinguishable bpp values after the lossless step. As expected, the square-root algorithm performs better if the number of CR-impacted, high-signal pixels is substantial. In our nominal 200 s exposure, the square-root algorithm uses 0.12 fewer bpp than decimation. For a 600 s CR dose, the square-root algorithm requires 0.38 fewer bpp than decimation. This suggests that the square-root algorithm removes about 2 bits of noise from the average CR-impacted pixel. If an algorithm to remove cosmic rays is implemented before compression, then there would be little difference in compression ratio between the square-root and linear codecs. However, such an algorithm would require more processing power than the lookup table of the square-root codec, and one would have to insure that objects of interest are not discarded with the cosmic rays.

We should expect the square-root algorithm to gain advantage also in the case of fields with many objects brighter than the sky background, e.g., fields at low Galactic latitude, or with large nearby galaxies.

The square-root algorithm has practical advantages as well, in that it will retain the desired $b = 1$ value as the sky background varies. The linear codec would need to change its step size Δ to accommodate different sky levels, e.g., from changing filters or zodiacal background. An onboard algorithm could determine the sky noise level and choose Δ appropriately for the linear codec. The square-root codec requires adjustment only when the bias level or gain of the detector ADC output change.

7. CONCLUSION

Future dark-energy surveys such as *JDEM* or *Euclid*¹⁰ propose to survey a substantial fraction of the full sky with space-based visible/NIR imaging near the diffraction limit. A data compression scheme is needed to satisfy telemetry constraints, but the weak gravitational lensing measurements place very strict requirements on the degree to which the compression algorithm can bias the measured shapes or fluxes of galaxies. For this reason, we choose to investigate lossy compression processes that operate on a pixel-by-pixel basis, fearing that algorithms examining pixel groups might induce correlated codec errors between pixels that bias the galaxy ellipticities.

Square-root preprocessing is known to be a good means for compressing shot-noise-dominated data streams while keeping a fixed number b of bits per σ of noise across the full detector dynamic range. The reduction of entropy by reducing the number of noise bits is essential to efficient subsequent lossless compression algorithms. We derive an expression for bias in the naive reconstruction of square-root-compressed data and present a correction scheme, which we numerically verify is successful to high accuracy.

Since the noise induced by roundoff in the compression process increases the required observing time by $1/12b^2$, we confine our examination to $b \geq 1$. We verify that flux and shape measurements of simple galaxy models are biased by $\lesssim 10^{-4}$ for a range of galaxy sizes, ellipticities, orientations, and S/N. This is well within the weak-lensing requirements. We verify this with simple postage-stamp galaxy images; Vanderveld et al. (2010, in preparation) bound the shape biases determined from realistic simulations of the true extragalactic sky.

At $b = 1$, the Shannon entropy of the square-root-compressed images varies from 2.1 bits pixel⁻¹ for pure noise, to 3.6 bits pixel⁻¹ for images with the expected astronomical scene and triple the expected number of cosmic rays. Lossless compression algorithms such as *bzip2* achieve compression levels within 0.3–0.4 bits of the Shannon limit. For simulated images most

closely resembling the expected *JDEM* CCD data, we find compression to 3.1 bits pixel⁻¹ using *bzip2*. The size of the transmitted data stream depends primarily on the degree of cosmic-ray contamination once $b = 1$ is set: schemes to more efficiently compress or mask cosmic rays prior to compression might lead to lower bit rates. Higher CR rates or more crowded astronomical fields, e.g., near the Galactic plane, will require more bits. Since pure-noise images require 2.4 bits pixel⁻¹ with *bzip2* at $b = 1$, we can state that compression to < 2.4 bits pixel⁻¹ will require $b < 1$, and hence a higher observing-time penalty.

We thus validate square-root lossy compression as useful for *JDEM* data and capable of meeting requirements for low bias. Other approaches would be feasible as well. For instance, simple decimation to a fixed code step Δ that equals the sky noise level should produce similar noise and bias properties. The presence of cosmic rays and/or bright objects in $\gtrsim 10\%$ of the pixels gives the square-root codec a useful advantage in mean number of required telemetry bits per pixel. The square-root algorithm also has a robust and exceptionally simple implementation—a single lookup table—meaning that there is no need for onboard processing to recognize cosmic rays or determine the sky noise level.

It is thus expected that visible/NIR dark-energy extragalactic imaging survey data can be transmitted to the ground with 2.5–4 bits pixel⁻¹, depending upon the CR flux level and the sophistication of onboard CR identification or removal. This compression can be achieved with no significant bias of observed quantities, and with 8% penalty in observing time to achieve fixed S/N.

The work of J. R. and R. A. V. was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with NASA, and was funded by JPL's Research and Technology Development Funds. G. B. is supported by grant AST-0607667 from the NSF and the US Department of Energy (DOE) grant DE-FG02-95ER40893. C. B. is supported by the Director, Office of Science, Office of High Energy Physics, of the DOE under grant DE-AC02-05CH11231. Fermilab is operated by Fermi Research Alliance, LLC under DOE grant DE-AC02-07CH11359.

REFERENCES

- Amara, A., & Réfrégier, A. 2008, *MNRAS*, 391, 228
 Amsler, C., et al. (Particle Data Group) 2008, *Phys. Lett. B*, 667, 1, <http://pdg.lbl.gov/>
 Bernstein, G., & Jarvis, M. 2002, *AJ*, 123, 583
 Bernstein, G. M., Trilling, D. E., Allen, R. L., Brown, M. E., Holman, M., & Malhotra, R. 2004, *AJ*, 128, 1364
 Candès, E., Romberg, J., & Tao, T. 2006, *IEEE Trans. Inf. Theory*, 52, 489
 Consultative Committee for Space Data Systems 1997, *Lossless Data Compression*, CCSDS 121.0-B-1 Blue Book, Washington, DC 20546, <http://www.ccsds.org>
 Frieman, J. A., Turner, M. S., & Huterer, D. 2008, *ARA&A*, 46, 385
 Gowen, R. A., & Smith, A. 2003, *Rev. Sci. Instrum.*, 74, 3853
 Huterer, D., Takada, M., Bernstein, G., & Jain, B. 2006, *MNRAS*, 366, 101
 Lin, H., & Marriner, J. 2003, SNAP collaboration memorandum
 Nieto-Santisteban, M. A., Fixsen, D. J., Offenberg, J. D., Hanisch, R. J., & Stockman, H. S. 1999, *Astronomical Data Analysis Software and Systems VIII*, 172, 137
 Pence, W. D., Seaman, R., & White, R. L. 2009, *PASP*, 121, 414
 Press, W. H. 1992, *Astronomical Data Analysis Software and Systems I*, 25, 3

¹⁰ At <http://sci.esa.int/euclid>.

- Roe, N. A., et al. 2007, *Nuclear Instruments and Methods in Physics Research A*, 572, 526
- Romeo, A., Gaztanāga, E., Barriga, J., & Elizalde, E. 1999, *Int. J. Mod. Phys. C*, 10, 687
- Seaman, R. L., White, R. L., & Pence, W. D. 2009, preprint (arxiv:0910.3733)
- Thompson, A. R., Moran, J. M., & Swenson, G. W. 2001, *Interferometry and Synthesis in Radio Astronomy*, (2nd ed.)
- Weinreb, S. 1963, MIT Research Laboratory of Electronics, Technical report 412
- White, R. L., Postman, M., & Lattanzi, M. G. 1992, in *Digitised Optical Sky Surveys* (Dordrecht, Boston: Kluwer), 174, 167