PUBLICATIONS OF THE ASTRONOMICAL SOCIETY OF THE PACIFIC
103: 122–130, January 1991

SIGNAL-TO-NOISE CONSIDERATIONS FOR SKY-SUBTRACTED CCD DATA

MICHAEL V. NEWBERRY

Steward Observatory, University of Arizona, Tucson, Arizona 85712

Received 1990 June 25, revised 1990 October 10

ABSTRACT

The "standard" equation for calculating the uncertainty of photometry obtained from CCDs does not correctly consider the random errors, or "noise", introduced into observations by procedures used in reducing the data. We present a thorough derivation of the theoretical error equation that considers the contributions from all internal noise sources in the signal-to-noise ratio ("*S/N*") of a sky-subtracted image. A simplified version used for estimating the internal errors from empirical data is also derived. The propagation of noise through the data-reduction process is illustrated through a series of equations for the change in *S/N* that results from a variety of different operations performed on a CCD frame. Comparing these effects with the results expected for an observation made with an ideal detector suggests a number of ways to improve the precision of photometry through the practices employed in obtaining and reducing the observations.

Key words: areal detectors–data analysis–photometry

1. Introduction

In order to understand the ways in which observational technique and methods of data reduction can influence the precision of photometry obtained from CCDs, it is important to understand how various factors affect the final signal-to-noise ratio ("*S/N*") of the measurement. Unfortunately, anecdotal versions of the "CCD error equation" in common use (see, for example, Stetson 1987 §III; Schoening 1987; Schroeder 1987; Howell 1989; Harris 1990) do not correctly consider the noise added by sky subtraction, do not (with the exception of Stetson 1987) suggest how to empirically determine the precision of the measurement, and, because the equation is usually presented without derivation, obscure the nature of the contributions from various sources of noise. Other such noise sources include that added by applying bias, dark count, preflash, flat field, and possibly other "calibration" frames to the data during the reduction process. These procedures can appreciably degrade the *S/N* of the data before sky subtraction is performed, resulting in photometric measurements of lower precision than the raw observations were capable of providing. These points are seldom considered in planning an observing strategy or in choosing methods for reducing the data. The present paper will focus on the factors which affect the *S/N* and will not discuss either the operation of CCDs or the general array of methods used for reducing data obtained with them. Excellent reviews of CCD fundamentals can instead be found in Kristian 1982, Mackay 1986, and Stetson 1987, and methods of data reduction are discussed by Gilliland & Brown 1988, Holtzman 1990, and by the various manuals and "cookbooks" associated with IRAF[1] and other image-processing software.

In Section 2 the *S/N* equation for sky-subtracted data will be derived with some rigor. An empirical version of the theoretical equation will also be presented which permits the random errors of photometric measurements to be estimated correctly from the sky-subtracted data. The theoretical *S/N* equation has profound implications for observational technique and methods of data reduction, and these will be discussed in Section 3.1. In Section 3.2 we will explore the effects that data reduction can have on the *S/N* of the reduced CCD frame. A summary is given in Section 4. Although the material presented herein is directed toward two-dimensional photometry, it has general application to all forms of two-dimensional data acquisition, including spectroscopy.

2. Random Errors in the CCD Image

By the time that the magnitude of an object is estimated from its image in a CCD frame, a number of different random errors will have contributed to this measurement. In this section we will discuss all of the internal random errors that propagate into this estimate. The complete equation for the *S/N* of a sky-subtracted CCD image will be derived.

Consider that the image of an object subtends n total pixels of area on the surface of a two-dimensional detector

[1]IRAF is distributed by National Optical Astronomy Observatories, which is operated by the Association of Universities for Research in Astronomy, Inc., under cooperative agreement with the National Science Foundation.

(here taken to be a CCD) or, more appropriately, that the brightness of the object is measured by summing the count over n pixels of detector area. The total number of counts attributable to the object, C_0, is therefore given by

$$C_0 = \sum_{i=1}^{n} [C_{0+s}(i) - C_{s,e}(i)] \ , \qquad (1)$$

where $C_{0+s}(i)$ is the number of counts in pixel i that are attributable to the object plus those attributable to the background ("sky"), and $C_{s,e}$ is the *estimated* number of counts in the background of pixel i. Note that n need not be integral (it is almost never integral in the case of aperture photometry). However, for the sake of simplicity, we will hereafter treat n as if it were indeed integral. Also, the statistics of counting depends on the number of photons that are converted to electrons by the CCD and *not* the analog-to-digital units ("ADUs") present in the CCD image. Therefore, all quantities will hereafter be considered in units of electrons through the standard relationship $g = C_{\text{electrons}}/C_{\text{ADU}}$ for a gain factor g and counts C in units of electrons and ADUs, respectively.

The random fluctuations in the count, or the "noise" N, from each of j different sources are uncorrelated so that they combine in quadrature, i.e., as $\sum N_j^2$. Hence, the total number of noise electrons, N_i, associated with pixel i after sky subtraction is given by

$$N_i^2 = \sigma^2 (C_{0+s}(i)) + \sigma^2 (C_{s,e}(i)) \ , \qquad (2)$$

where the terms σ^2 denote the variance of their respective quantity. The total noise in the object count is therefore

$$N^2 = \sum_{i=1}^{n} N_i^2 = \sum_{i=1}^{n} \sigma^2 (C_{0+s}(i)) + \sum_{i=1}^{n} \sigma^2 (C_{s,e}(i)) \ , \qquad (3)$$

and, hence,

$$N^2 = \sum_{i=1}^{n} \sigma^2 (C_0) + \sum_{i=1}^{n} \sigma^2 (C_s(i)) + \sum_{i=1}^{n} \sigma^2 (C_{s,e}(i)) \ , \qquad (4)$$

where we have separated the total noise before sky subtraction, C_{0+s}, into its object and sky components. There are four internal sources of noise which comprise the terms of equations (3) and (4). These are:

1. Stochastic noise. Photons are converted into electrons randomly in time with a time-averaged efficiency that is the "quantum efficiency" of the detector. Therefore, the count in each pixel and hence the total count over n pixels is a Poisson-distributed random variable with variance equal to its zeroth moment (i.e., its total count; Sachs 1982). For example, the stochastic contribution to $\sum \sigma^2 (C_s(i))$ is just C_s.

2. Readout noise characterized by variance R^2 in each pixel.

3. Truncation noise attributable to binning, or "truncating" the number of stored electrons into a smaller number of ADUs (counts) at readout, or during the data reduction if a frame is converted from real to integer format (usually to save disk space). This truncation, or "digitization", is described by a uniformly distributed random variable having a domain equal to the gain, g (the number of electrons in one integral ADU), for the readout, and the number of precision elements in the digitization for conversion to integer format. The second moment of this rectangular distribution gives the variance to be $T^2 = (g^2 - 1)/12$, in units of electrons^2 (see Sachs 1982, p. 84). In the units of the target distribution, this is $T \approx 1/\sqrt{12} \approx 0.29$ count. The truncation noise associated with reading the CCD is normally included in the specified readout noise and, with a low gain amplifier, is usually small in comparison.

4. Processing noise contributed by all stages of the data reduction, including corrections for bias, dark count, preflashing, and interpixel variations in detector sensitivity ("field flattening"). With the exception of the bias, the calibration frames are themselves processed to some degree before they are combined with the data frames and before the background is subtracted. The processing noise, which we will denote by P, is discussed more fully in Section 3.2.

Combining items (2) through (4), the total "base-level" noise, B, in each pixel is then given by $B^2 = R^2 + T^2 + P^2$. Since R and T are fixed for a given instrumental setup and P can be made relatively constant across a frame (see Section 3.2), we will treat the base-level noise as being constant over the region of interest on the CCD frame. Over the n pixels of the object summation, the second term on the right-hand side of equation (4) then becomes

$$\sum \sigma^2 (C_s(i)) = C_s + nB^2 = n(f_s + B^2)$$

in terms of the *average* sky brightness per pixel, f_s. Note that this holds for the *internal* error even if the sky is not strictly constant under the object profile. The sky level underneath the object is estimated by sampling the local background in some "appropriate" way using some number, p, of background pixels (see Holtzman 1990 and Stetson 1990 for a discussion of the complexity of extracting a stellar magnitude from above the local background). The variance in the sky level estimated for each of the n pixels underneath the object is therefore reduced by a factor of $1/p$ relative to its variance in one pixel. This is the same relationship that exists between the standard deviation of a random sample and the standard deviation of the sample mean. Thus we have for the total internal noise in the object count

$$N^2 = C_0 + n(f_s + B^2) + n[\frac{1}{p}(f_s + B^2)] \ , \qquad (5)$$

which can be rearranged to give

$$N^2 = C_0 + n(f_s + B^2)\,(1 + \frac{1}{p}) \ . \qquad (6)$$

Since C_0 is the total "signal" for the object, its S/N is given by

$$S/N = \frac{C_0}{[C_0 + n(f_s + B^2)\,(1 + \frac{1}{p})]^{1/2}} = \frac{C_0^{1/2}}{[1 + nC_0^{-1}(f_s + B^2)\,(1 + \frac{1}{p})]^{1/2}} \ , \qquad (7)$$

where all quantities are expressed in *electrons*. One application of equation (7) is to explore the optimum subtraction or extraction area—by definition, that area over which to sum the counts from the object so as to maximize the S/N of the resulting aperture magnitude (see, for example, Pritchet & Kline 1981; Howell 1989; Stetson 1990) or extracted 1-dimensional spectrum (Horne 1986; Newberry 1987). It also finds common use in predicting exposure times needed to obtain data of some desired precision using a given instrumental setup. Note that equation (7) contains a factor $(1 + 1/p)$ which multiplies the base-level noise and the sky brightness in each pixel.

If we make the following definition,

$$\xi = nC_0^{-1}(f_s + B^2)\,(1 + 1/p) \ , \qquad (8)$$

then equation (7) reduces to the form

$$S/N = \frac{C_0^{1/2}}{[1 + \xi]^{1/2}} \ . \qquad (9)$$

The theoretical limiting S/N of an ideal detector for which the object photons are the only noise source and there is no background subtraction follows from the variance of the Poisson distribution: $(S/N)_0 = C_0/C_0^{1/2} = C_0^{1/2}$. Since this is the asymptotic limit of equation (9) for $\xi \to 0$, it is clear that in order to approach the ideal minimum noise situation—the so-called "photon-noise-limited" case—we must minimize the factors contributing to ξ in equation (8). As will be shown in Section 3 this has important ramifications for many aspects of the observational technique and reduction of the data. For the photon-limited case, the S/N of the observation increases with the square root of the exposure time (since $C_0 = R_0 t$ for an average photon rate R_0 and exposure time t).

If we assume that p is large, then equation (7) can be written in the form

$$S/N \approx \frac{C_0^{1/2}}{[1 + (C_s/C_0 + nB^2/C_0)]^{1/2}} \ . \qquad (10)$$

Two practical observational cases are illustrated by this form of the equation: the "sky-limited" case in which $\xi \approx C_s/C_0$, and the "base-noise-limited" case in which $\xi \approx nB^2/C_0$. If the total base-level noise under the object profile is small relative to the object count (the sky-limited case), then for a given sky count the S/N of the sky-subtracted image increases with the exposure time as $\sqrt{t}$. However, at any given exposure time, the S/N will be lower than for the photon-noise-limited case by a factor of approximately $(1 + C_s/C_0)^{-1/2}$. Conversely, if the object count is small relative to the total base-level noise over n pixels, then the S/N will be even lower than for the sky-limited case but will grow faster than $\sqrt{t}$ for small object count, approaching the sky-limited case and $\sqrt{t}$ growth rate when nB^2/C_0 becomes ~ 1. Hence, the undesirable base-noise-limited case can be promoted to the sky-limited case if enough photons can be collected. In either case, an extremely faint object, and especially an extended faint object, may not yield an acceptable S/N using an exposure time of any reasonable length. It is important to note that processing noise can provide a significant base-level noise component to a detector having negligible readout and truncation noise. Processing noise is discussed more fully in Section 3.2.

Although equation (7) is useful for planning an observing strategy, it is not the best way to estimate the S/N of a reduced observation because the base-level noise terms are not accurately known for a given CCD frame. The relevant noise level is that which can be measured from the frame. If the noise in each of the pixels under the image of the object could be determined, then calculating the uncertainty of the observation would be straightforward. This requires knowing the correct statistical model for the image profile over all n pixels of interest, and this is generally impossible or impractical except for unresolved images. Fortunately, the local background is usually somewhat easier to model and can be used to provide an indirect estimate of the noise in the object pixels. Since we have accounted for all internal noise sources in the above derivation, the standard deviation of the background measures the total internal noise in each pixel. In addition, there will be some nonrandom portion, ϵ^2, of the sample variance that is attributable to the images of faint objects scattered across the field of view (this is not random because it appears if the observation is repeated). Therefore, the background variance slightly overestimates the *internal* random error of the observation. Replacing the measured background ("sky") variance for the theoretical variance per sky pixel and allowing for spurious scatter ϵ^2 we have $\sigma_{bg}^2 = f_s + B^2 + \epsilon^2 = (f_s + R^2 + T^2 + P^2) + \epsilon^2$, which leads to

$$S/N \approx \frac{C_0^{1/2}}{[1 + n\sigma_{bg}^2 C_0^{-1}(1 + \frac{1}{p})]^{1/2}} \ . \qquad (11)$$

If units of ADUs are used for C_0 and σ_{bg}, then equation (10) becomes

$$S/N \approx \frac{C_0^{1/2}}{[\frac{1}{g} + n\sigma_{bg}^2 C_0^{-1}(1+\frac{1}{p})]^{1/2}} \quad (12)$$

for a gain factor of g electrons ADU^{-1}. The assumption of approximately constant processing noise is clearly now a requirement if the measured variance of the background is to correctly characterize the noise in the object pixels.

The uncertainty in magnitude is easily calculated from the S/N obtained through equations (7), (11), or (12) using the definition of a magnitude, m, corresponding to an observed count, C, viz., $m = \text{const} - 2.5 \log C$. Differentiating both sides to determine the dependence of a change in m on a change in C and substituting their respective uncertainties, $\sigma(m)$ and $\sigma(C)$, we obtain $\sigma(m) \approx 1.0857[\sigma(C)/C]$. Since the S/N is the inverse of the relative error in the net object count, we simply make the substitution $S/N = C/\sigma(C)$ and obtain

$$\sigma(m) \approx 1.0857\,(S/N)^{-1} \ , \quad (13)$$

where S/N is determined as above.

In the above derivation no *external* factors, such as scintillation noise and the uncertainty of the transformation to the standard photometric system (Beckert & Newberry 1989), were considered. Likewise, the digitization imposed by rounding to integer pixel values during data reduction can result in slightly enhanced scatter and systematic errors that increase with decreasing object brightness; a rounded pixel value can round up or down by as much as 1/2 ADU, and this can subtract or add false signal to the integrated image profile. These and other such external factors must be compounded with the internal S/N given from equations (11) or (12) before the total uncertainty of a photometric measurement can be obtained.

3. Maximizing the Signal-to-Noise Ratio

3.1 *The Photon-Noise-Limited Case*

To maximize the precision of a photometric measurement, it is clearly desirable to have $\xi \to 0$ in equation (9). In this section we will describe a number of factors which can minimize ξ through equation (8). The observer is usually constrained by the scientific program at hand to observing objects of a given brightness and angular extent and, hence, cannot improve the precision of the data by observing preferentially brighter objects or those having a small angular extent. Similarly, there are a number of factors over which an observer has no direct control because only particular instruments and a specific observing site are available. However, the observer does have the ability to define the "experimental design" so as not to degrade the precision further than the apparatus and observing conditions are capable of delivering. Considering all of these points together, the S/N of an observation can be improved or at least not degraded further than necessary by minimizing the contributions to ξ in the following ways.

1. Acquire a higher object count C_0. This is assisted by the following:

(a) Using a highly efficient instrument, including a CCD with high quantum efficiency.

(b) Good atmospheric transparency.

(c) High filter transmission and/or broad bandwidth, provided that there are not compelling scientific reasons against choosing filters of wider bandpass.

(d) Increasing the exposure time. A longer exposure may also help to raise base-noise-limited observations into the sky-limited regime.

2. Reduce the average background count per pixel, f_s, relative to the object count, C_0. This is aided by

(a) a dark observing site and a moonless sky;

(b) a CCD having a low dark count rate that does not require preflashing;

(c) minimizing the scattered light from bright objects in and near the field of view of the CCD frame; and

(d) using filters that exclude night-sky emission lines.

3. Reduce the number of pixels, n, in the object sum even when the observation of fine image detail is unimportant, for example, in observing faint standard-star fields. This is accomplished by minimizing the size of the object as viewed by the detector (this does not mean using a small aperture size in aperture photometry; see Stetson, 1990). This yields the highest possible energy density in the image of the object at the focal plane and minimizes the number of pixels contributing base-level noise to it. For an unresolved object, this also reduces the number of sky photons added to the image. However, reducing n can lead to undersampled data. If the detection of faint sources is important, the image scale should always be slightly larger than critical sampling of 2 pixels FWHM^{-1} (see Harris 1990). Some ways to minimize n are as follows:

(a) Use a shorter focal-length telescope, or a given telescope at a shorter focal ratio if reimaging optics are available.

(b) When possible, superpixel the CCD.

(c) Obtain the best possible focus and achieve the best possible guiding to take advantage of the seeing conditions.

4. Minimize, wherever possible, the base-level noise B by

(a) using a CCD and amplifying electronics that give a low readout noise per pixel;

(b) reading the CCD in "superpixels" to reduce the readout noise per unit area by reducing the number of pixels in the image (note the warning given for item 3);

(c) selecting a small gain, if possible, to minimize the truncation noise. The data reduction should also be done using floating-point arithmetic of sufficient precision, and the intermediate images and calibration images should be stored in floating-point data format; and

(d) minimizing the processing noise using high S/N dark, flat-field, preflash, and bias frames (or an average line bias). See Section 3.2 below for a discussion of processing noise.

5. Maximize the number of pixels, p, over which the local background is estimated. This is subject to the difficulties of following the sky level over some large number of pixels and of modeling the point-spread functions of other objects near the object of interest. Using a very large background region makes p large but can pull the interpolated background under the object away from its proper value. Care must be taken not to sacrifice accuracy for a small gain in precision. Note that any value $p > 0$ provides one of the advantages that a 2-dimensional detector has over a 0- or 1-dimensional one: By decreasing $1/p$, the noise introduced by sky subtraction can be diminished such that $1 + 1/p \rightarrow 1$; otherwise, $1 + 1/p = 2.0$.

3.2 *S/N Degradation from Data Reduction*

Through the course of reducing a CCD frame to the point of sky subtraction, noise can be imparted to the frame through the various "calibration" (e.g., bias, dark, preflash, and flat-field) frames, and the magnitude of this noise can grow impressively through the reduction process. As shown by equation (7), increasing the base-level noise of each pixel in the reduced image will degrade the precision of the photometric measurement and the "processing" noise introduced during data reduction should thus be minimized. As a minimal amount of processing, all frames must have the bias level subtracted from them, and the data frames must have some field flattening applied. Since all calibration frames have noise, each of these operations propagates some amount of noise into the data frames. In this section we will show how processing noise is added by the reduction process.

The processing noise imparted to an image in the kth stage of data reduction includes the processing noise introduced at all prior stages of data reduction. Specifically, if frame $k - 1$ is combined in some way with frame k, then the base-level noise for some arbitrary pixel in frame k is

$$B_k^2 = R^2 + T^2 + N_{k-1}^2 \ , \tag{14}$$

and, hence, the total internal noise for this pixel in frame k is

$$N_k^2 = S_k + R^2 + T^2 + N_{k-1}^2 \ , \tag{15}$$

where S_k is the signal in the pixel that is attributable to electrons from any photon source. The term N_{k-1} is taken to be the processing noise contributed to this pixel by the corresponding pixel in frame $k-1$ (we denoted the processing noise in general terms by P in Section 2, item 4). Here, $k = 0$ refers to the bias level that must be estimated for, and then subtracted from, all frames—dark, preflash, flat-field, and other calibration frames as well as from the data frames themselves. Hence, all bias-subtracted frames contain at least the noise contributed by the uncertainty in the bias, and also include photon noise if they are exposed to photons. Since some of *these* frames must be applied to the data then, by induction, it is clear that the S/N of a final reduced data frame is degraded by each data-reduction operation used to produce it.

The effect of processing on the S/N at each stage of data reduction can be estimated by computing the total derivative of the functional form of the operation, identifying the noise terms with the differentials, and assuming that the noise sources are uncorrelated and hence add in quadrature. The following equations are easily derived and illustrate the propagation of S/N for an arbitrary pixel of signal S through the reduction process:

raw frame − bias frame:

$$S = S_{\rm raw} - S_{\rm bias}$$
$$S/N = \left[S_{\rm raw} - S_{\rm bias} + B^2 \left(1 + \frac{1}{n_{\rm bias}}\right)\right]^{-1/2} (S_{\rm raw} - S_{\rm bias}) \tag{16}$$

combining n of frame k:

$$S = \sum_{k=1}^{n} S_k$$
$$S/N = \left[\sum_{k=1}^{n} \left(\frac{S_k}{(S/N)_k} \right)^2 \right]^{-1/2} \sum S_k \ ; \tag{17}$$

$\text{frame}_1 \pm \text{frame}_2$:

$$S = S_1 \pm S_2$$
$$S/N = \left(1 \pm \frac{S_2}{S_1}\right) \left[1 + \left(\frac{S_2}{S_1}\right)^2 \left(\frac{(S/N)_1}{(S/N)_2}\right)^2 \right]^{-1/2} (S/N)_1 \ ; \tag{18}$$

frame_1/flat-field ("ff") frame:

$$S = S_1$$
$$S/N = \left[1 + \left(\frac{(S/N)_1}{(S/N)_{\rm ff}}\right)^2 \right]^{-1/2} (S/N)_1 \ ; \tag{19}$$

frame_1 ± a constant:

$$S = S_1 \pm \text{const}$$
$$S/N = \left[1 \pm \frac{\text{const}}{S_1} \right] (S/N)_1 \ ; \tag{20}$$

frame_1 multiplied by a constant:

$$S = S_1 \cdot \text{const}$$
$$S/N = (S/N)_1 \ . \tag{21}$$

The n_{bias} term in equation (16) refers to the number of independent estimates of the bias. This would be the number of bias frames that were averaged together or the number of pixels used from the underscan or overscan region of a given frame to estimate its line bias. Equations (17) through (21) assume that the bias has been removed from the frames of interest.

To illustrate the propagation of processing noise through the data reduction, let us consider a typical CCD system that has 20 electrons of readout noise and a gain of 10 electrons ADU^{-1}. If we assume that the truncation noise is not included in this estimate of the readout noise, then the base-level noise of an unprocessed frame taken with this system is $B^2 = 20^2 + (10^2 - 1)/12 \approx 408$ electrons^2. We will examine the results of reducing the following observations: one bias frame, a flat-field frame having 50,000 electrons $pixel^{-1}$, a long-exposure dark frame having 500 electrons $pixel^{-1}$, and three data frames having 10,000, 1500, and 500 electrons $pixel^{-1}$, all referenced above the bias level (note that the bias level must actually be estimated for each image). Our third case with 500 electrons $pixel^{-1}$ will have 1/3 the exposure time of the second case having 1500 electrons $pixel^{-1}$. To obtain the effective dark count for these exposure times, we divide the bias-subtracted dark frame by 5 for the first two cases and by 15 for the third case. We also assume that the flat-field frame has been normalized to an average value of 1.0. Floating-point arithmetic and data format will be considered throughout. Table 1A shows the progression of the net signal, S, and the S/N for each of the three cases through the various stages of data reduction. The first reduction column, labeled "After bias subtraction", shows the effective doubling of the base-level noise in each frame that results from subtracting a bias *frame*, pixel by pixel, from it. Further degradation of the S/N is shown after dark frame subtraction and flat fielding. In the last column of Table 1A we show the ratio of the final S/N of the reduced frame to the ideal S/N based only on the stochastic noise ($\sqrt{S}$) attributable to the astronomical signal—the latter being the best one could expect from an ideal detector. The reduction in S/N after processing is obvious for all three signal levels in the data but clearly gets worse as the signal drops. The total processing noise added to the three signal levels between their raw and reduced states amounts to 49.5, 22.4, and 20.5 electrons for the three cases of 9900, 1400, and 467 net electrons $pixel^{-1}$, respectively—we have effectively increased the readout noise by these amounts. Immediately before flat fielding, however, only 21.5, 21.5, and 20.4 electrons $pixel^{-1}$ of processing noise—about the level of the readout noise—had been added by the data reduction. The processing noise is slightly lower in case 3 because it assumes 1/3 the dark count of cases 1 and 2. It is important to note that in case 1, for an inherently high S/N observation, the greatest degradation in S/N and, hence, the greatest noise contribution, results from flat-field division. Although obscured by the S/N form of equations (16) through (21), the processing noise added by an arithmetic operation is independent of signal level in all operations except when a signal with noise is multiplied or divided into another one. In the case of flat fielding, the noise of the flat-field frame is combined into the data frame without the addition of any net signal. Although the flat-field frame in this example has a relative uncertainty of less than 0.5% ($S/N > 200$), its 50,000 electrons add a substantial amount of stochastic noise to the high count level of case 1. However, the flat-field frame has little detrimental effect on the low signals of case 2 and, especially, case 3. In these two cases the S/N of the flat field is so high in comparison that flat fielding does little to degrade the data. It is easy to see from equation (18) that the degradation caused by the flat-field frame increases as the S/N of

TABLE 1A

S/N After Processing, Using One Bias Frame and One Dark Frame

		After bias subtraction		After dark subtraction		After flat field division		
Frame	(case)	S/N	S	S/N	S	S/N	S	(S/N)/(S/N)max
Dark	...	13.78	500	...	...	...	...	
Flat	...	221.8	50000	221.8	50000	...	...	
Data	(1)	96.15	10000	94.96	9900	87.30	9900	0.877
Data	(2)	31.17	1500	28.77	1400	28.53	1400	0.762
Data	(3)	13.78	500	12.83	467	12.81	467	0.593

the data frame increases relative to it. For example, if the S/N of the flat field equals that of a given pixel, then the flat-field division will degrade the S/N of the pixel to only 0.707 of its prior value. Therefore, the processing noise added by flat fielding will be approximately independent of the signal across the profile of an object whenever the precision of the flat field is far higher than that of any pixel under the object profile. Unfortunately, if the base-level noise varies substantially across the profile, then the standard deviation of the local background used in equation (11) becomes a poor estimator of the S/N of the photometric measurement. Hence, the S/N of the master flat field must be far higher than that of any pixel to which it is applied. This can be accomplished by combining a number of high S/N flat-field frames. One additional point must be made about flat-field frames: obtaining a flat-field frame having of order 10^4 or more electrons pixel^{-1}—a common "safe" practice—does not necessarily degrade the S/N of the reduced data by only 1% or less.

In Table 1B we show the same three cases, except that "noiseless" estimates of the bias and dark count were obtained by letting $n_{\text{bias}} \to \infty$ and using equation (20) instead of equation (18). By "noiseless" we mean that, instead of subtracting the bias and dark frames pixel by pixel from the data frames, we have reduced their noise contribution by adopting a constant value for the signal of each. This situation can be approached by using a large number of independent estimates for the value of each pixel, either by averaging a large number of individual frames or by using some average value (i.e., mean, median, or mode) obtained from a large region of the frame. For the bias level this might be the underscan or overscan region of the frame. Comparing Tables 1A and 1B shows the significant improvement in the low signal data (case 3) that results from effectively increasing the S/N of the calibration frames. Provided that the assumed processing is done, Tables 1A and 1B illustrate the extremes of S/N degradation that result from the data reduction: In Table 1A both the base-level noise and photon noise are allowed to propagate freely into the data, and in Table 1B only the photon noise in the raw frames is important. Hence, even with an inherently noiseless detector, some amount of processing noise must be added during data reduction. The processing noise added through naively conceived methods of observation and data reduction can drive the observation of a faint object into the base-noise-limited realm even using a detector having zero readout noise and unit gain.

In Tables 2A and 2B we have repeated the reduction procedure shown in Tables 1A and 1B, except that we have included the effect of using a preflash frame having an exposure level of 500 electrons pixel^{-1}. Note the devastating result in case 3 of subtracting a preflash frame having a preflash signal similar to that of the astronomical signal. This effect is mitigated somewhat in Table 2B where the processing noise has been reduced by using a "noiseless" preflash frame as for Table 1B. However, the data are still poorer than they would have been had we not used preflashing. All other things being equal, a system that requires preflashing cannot yield data as good as the data obtained with a system that does not require it, especially in the case of a faint object. It is important to note from Tables 1B and 2B that using preflash, dark, and bias frames of infinite precision still degrades the S/N of the reduced data because the extra signals added by the dark and preflash exposures contribute stochastic noise but their signals must be *subtracted* during the reduction process.

4. Summary

In Section 2 a complete version of the noise equation was derived which illustrates the various sources of inter-

TABLE 1B

S/N After Processing, Using "Noiseless" Estimates of the Bias and Dark Count

		After bias subtraction		After dark subtraction		After flat field division		
Frame	(case)	S/N	S	S/N	S	S/N	S	(S/N)/(S/N)max
Dark	...	...	500	...	...	...	...	
Flat	...	222.7	50000	222.7	50000	...	...	
Data	(1)	98.02	10000	97.04	9900	88.96	9900	0.894
Data	(2)	34.34	1500	32.05	1400	31.72	1400	0.848
Data	(3)	16.59	500	15.48	467	15.44	467	0.715

TABLE 2A

S/N After Processing, Using One Bias Frame, One Dark Frame, and One Preflash Frame

Frame	(case)	After bias subtraction S/N	S	After preflash subtraction S/N	S	After dark subtraction S/N	S	After flat field division S/N	S	(S/N)/(S/N)max
Preflash	...	13.78	500	...	...	...	...	...	...	
Dark	...	13.78	500	...	...	...	...	...	...	
Flat	...	221.8	50500	216.9	50000	216.9	50000	...	...	
Data	(1)	98.70	10500	88.97	10000	87.90	9900	81.46	9900	0.819
Data	(2)	37.69	2000	23.33	1500	21.64	1400	21.53	1400	0.575
Data	(3)	23.46	1000	8.93	500	8.33	467	8.32	467	0.385

TABLE 2B

S/N After Processing, Using "Noiseless" Estimates of the Bias, Dark Count, and Preflash

Frame	(case)	After bias subtraction S/N	S	After preflash subtraction S/N	S	After dark subtraction S/N	S	After flat field division S/N	S	(S/N)/(S/N)max
Preflash	...	...	500	...	...	...	...	...	...	
Dark	...	...	500	...	...	...	...	...	...	
Flat	...	223.8	50500	221.6	50000	216.9	50000	...	...	
Data	(1)	100.5	10500	95.74	10000	94.78	9900	87.14	9900	0.876
Data	(2)	40.75	2000	30.56	1500	28.52	1400	28.29	1400	0.756
Data	(3)	26.65	1000	13.33	500	12.44	467	12.42	467	0.575

nal uncertainty inherent to a photometric measurement made with a CCD. New to the standard theoretical equation are the effects of *processing noise,* which can be substantial, *truncation noise,* which is usually very small, and *sky subtraction,* which is necessary for any photometric measurement. These effects are not usually considered when an observing program is planned. Observers often find that the empirically determined uncertainty of a reduced image is substantially greater than they estimated before the observations were obtained. It is likely that the "extra" factors considered herein account for much of the discrepancy. Comparing the correct *S/N* equation with that for an "ideal" detector and considering the noise propagation through data reduction suggests a number of ways to maximize the precision of the reduced frame and the photometric measurements obtained from it. Many of these points given in Section 3.1 are constrained by the facilities available to the observer. However, there are a number of factors over which the observer has some control, for example, choosing the optimum image scale if a reimaging camera is available. Also, a large number of flat-field and preflash frames (if pertinent) should be obtained so that they can later be averaged into very high *S/N* masters. The master flat field should have much higher *S/N* than any pixel of interest in the data in order for the background variance to accurately characterize the uncertainty of the photometric measurement. Bias should be obtained either from a master "superbias" made of many (~ 100) individual bias frames or be determined from a reasonably large underscan or overscan region of the data frame. If possible, a reasonable average (constant) value rather than an entire frame should be used to represent the dark and preflash exposures. Floating-point arithmetic should also be used for all reduction procedures and for storing intermediate images. Finally, the background should be estimated or subtracted over as large a region as is practical. In general, *any* technique which minimizes ξ, the noise above that of a photon-noise-limited observation, will lead to higher precision photometry or spectroscopy.

This paper resulted from work started while a Postdoctoral Research Associate at the University of New Mexico and was completed in conjunction with the CCD Transit Instrument (CTI) project at Steward Observatory. The CTI is funded by NSF grant AST 88-00298 and NASA grant NAGW-804 and is supported by the Data General Corporation.

REFERENCES

Beckert, D. C., & Newberry, M. V. 1989, PASP, 101, 849
Gilliland, R. L., & Brown, T. M. 1988, PASP, 100, 754
Harris, W. E. 1990, PASP, 102, 949
Holtzman, J. A. 1990, PASP, 102, 806
Horne, K. 1986, PASP, 98, 609
Howell, S. B. 1989, PASP, 101, 616
Kristian, J. H. 1982, Scientific American, 247:4, 66
Mackay, M. 1986, ARA&A, 24, 255
Newberry, M. V. 1987, Ph.D. dissertation, University of Michigan
Pritchet, C., & Kline, M. I. 1981, AJ, 86, 1859
Sachs, L. 1982, Applied Statistics (New York, Springer)
Schoening, W. 1987, Operation of the CCD Direct Imaging Camera for the No. 1 0.9 Meter Telescope at Kitt Peak National Observatory (Tucson, National Optical Astronomy Observatories)
Schroeder, D. J. 1987, Astronomical Optics (San Diego, Academic Press)
Stetson, P. B. 1987, PASP, 99, 191
———. 1990, PASP, 102, 932