Hybrid Approach to Mean-Variance and Photon Transfer Measurement

Blake C. Jacquot*, Brett M. Bolla, Sean Maguire
The Aerospace Corporation, 2310 E. El Segundo Blvd., El Segundo, CA 90245-4609

ABSTRACT

This paper presents a hybrid technique for measuring conversion gain that blends spatial and temporal information, allowing users to calculate an accurate conversion gain with little knowledge of sensor defects. It blends a single pixel method with multiple pixel methods. We present measured data from a visible CMOS image sensor using two multiple pixel methods and the hybrid method. Additionally, we provide arguments for validity of the hybrid method. To our knowledge, this is the first report of this technique.

Conversion gain (e-/DN) directly relates measured digital numbers (DN) to input-referred electrons (e-) for an image sensor. Conversion gain can be directly measured by considering the sensor under varying illumination states in coordination with Poisson statistics. Typically, there are two approaches: measure a single pixel over time or measure a group of pixels at one point in time after correcting for gain non-uniformity. The plotted statistics from these measurements are called either mean-variance or photon-transfer curves. The measurement of a single pixel is relatively straightforward and requires collection of many consecutive frames to get meaningful statistics not dominated by thermal noise. The data volume for an accurate single-pixel measurement can become unwieldy in terms of number of frames required. This is especially true for large format image sensors. In contrast, the measurement of a group of pixels requires fewer consecutive frames, but needs non-uniformity adjustments to correctly calculate statistics.

1. INTRODUCTION

1.1 Motivation

Conversion gain allows figures of merit such as read noise and full well to be cast from measured values of digital numbers (DN) or counts into electrons (e-). The value for conversion gain can be modeled accurately with knowledge of the unit cell and analog-to-digital converter (ADC), but such knowledge is typically unavailable to end users. Usually users must conduct an experiment.

For visible imagers, Fe_{55} sources can be used. Here, a single photon from the source results in a known amount of electrons. Comparing collected values at pixel level in DN against expected electrons allows direct calculation of conversion gain. This technique is not used for infrared sensors due to differences in detector material. Other methods may also include radiometric-based calculation of conversion gain.

The most common technique by far for experimentally finding conversion gain of both visible and infrared sensors relies on photon transfer curves which are also known as mean-variance curves. These curves rely on the discrete nature of photon arrival and the resulting Poisson statistics to empirically determine conversion gain. To extract conversion gain from a mean-variance plot, data must have fixed pattern noise (FPN) removed.

1.2 Contributions

We present a new approach to using mean-variance data for calculating conversion gain. This method differs from others primarily in that it does not require explicitly correcting for FPN. Rather, it relies on the Central Limit Theorem to implicitly correct for FPN. It does not require selection of a well-behaved region of interest (ROI). It blends single pixel methods and multiple pixel methods.

Keywords: system gain, conversion gain, mean-variance, shot noise, Poisson distribution, photon transfer
* Blake.C.Jacquot@aero.org, blakejacquot@gmail.com; phone 1 (310) 336-1510; http://www.aerospace.org/

Image Sensing Technologies: Materials, Devices, Systems, and Applications II,
edited by Nibir K. Dhar, Achyut K. Dutta, Proc. of SPIE Vol. 9481, 94810D
© 2015 SPIE · CCC code: 0277-786X/15/$18 · doi: 10.1117/12.2176115

Downloaded From: http://proceedings.spiedigitallibrary.org/ on 05/18/2015 Terms of Use: http://spiedl.org/terms

In this paper, we show results using mean of sample mean and mean of sample variance calculations because derivations and computations are relatively easy. However, using median of sample mean and median of sample variance gives slightly better performance due to ability to reject large temporal excursions (*e.g.* flicker noise). In practice we typically use the median method with excellent results.

The hybrid method can be used as a primary or supplemental method in calculating mean-variance statistics. In supplemental form, it can be used to cross-check statistics calculated by other methods.

1.3 Related Work

Photon transfer and mean variance methods were perhaps first developed to extract data from Vidicon tubes[1]. It has since been applied to solid-state image sensors, including both charge coupled devices (CCDs) and CMOS image sensors. All photon transfer methods should arrive at the same statistics but by different paths.

As a matter of convention, photon transfer curves are plotted with standard deviation against mean values on a log-log plot. Mean-variance curves plot the same data as variance against mean values on a linear-linear plot. Conversion gain can be extracted from either plot, but fitting is typically done on a mean-variance plot because it requires only a linear fit. A typical photon transfer curve is shown in Figure 1.

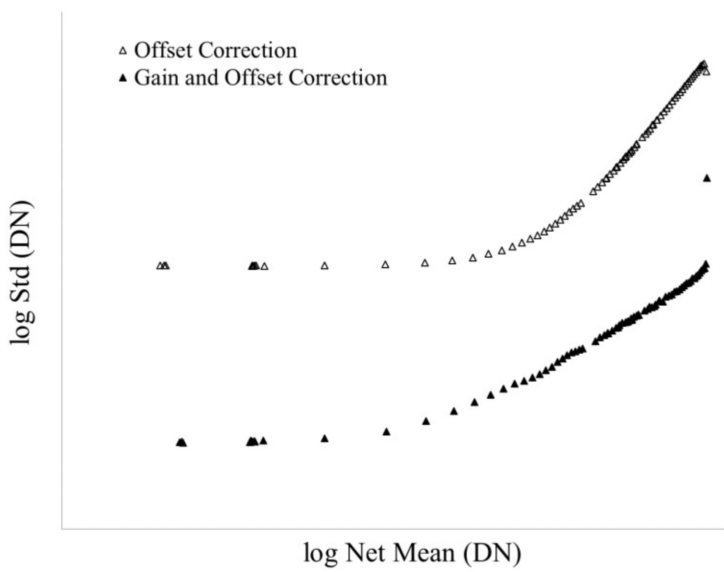

Figure 1. Photon transfer curve of a sensor with constant integration time and varying illumination. Calculations are made from spatial data.

1.4 Limitations and Benefits

We have not yet definitively proven why our technique gives good results. Rather, we give suggestive mathematical arguments and empirical data. We believe the underlying data for per-pixel response to flux is independent and similarly distributed. This similarity allows us to use central limit theorem and associated estimates of standard error. Additionally, estimates of mean(s) and variances(s) may be computed from a mixture of spatial and temporal data.

This method may fail if fixed pattern noise is large enough and "similarity" is violated. We have not quantified what level of FPN might cause the method to fail because it has succeeded on all image sensors tested that lacked significant inter-pixel capacitance, which exhibit non-Poisson behavior. Note, all methods give similarly incorrect answers in such cases.

Downloaded From: http://proceedings.spiedigitallibrary.org/ on 05/18/2015 Terms of Use: http://spiedl.org/terms

2. METHOD

2.1 Single Pixel Method

A conceptually simple way to generate a data point for mean-variance curves is to measure a single pixel's temporal values for a given integration time and illumination level. The measurement is repeated while varying either integration time or illumination. The temporal sample mean and sample variance make a single data point for the curve. In this method there is no need to account for FPN or gain non-uniformities since the pixel has the same photo-response from frame to frame and same analog signal chain. However in practice the single pixel approach is difficult because of the large number of frames needed to yield an accurate measure of standard deviation. The single pixel method negates the need to correct for FPN but comes at the expense of onerous amounts frames to collect.

2.2 Multiple Pixel Methods

In contrast to single pixel method, multiple pixel methods negate the need to collect many frames, but come at the expense of having to correct for FPN. This expense is tolerated because data volumes shrink dramatically. A region as small as 200 x 200 pixels yields 40,000 collects, which gives low standard error of both mean and variance. Correction for FPN can take several forms, but all remove FPN from one or more frames before calculating noise statistics. Removing FPN leaves read and shot noise (light and dark) and effectively makes all pixels identical from a sampling perspective. This implies pixel values over time and space are interchangeable under constant illumination.

There are three broad forms of the multiple pixel method. Typically 25-100 consecutive frames are collected for a variety of flood illumination levels spanning dark through full well.

Total spatial noise in a given frame adds as root sum square (RSS) and will equal

$$s_{Total}^2 = s_{FPN}^2 + s_{Shot}^2 + s_{\mathrm{R}ead}^2 \tag{1}$$

Here s_{Total} is total noise, s_{FPN} is fixed pattern noise, s_{Shot} is shot noise (photon and dark), and s_{Read} is read noise.

2.3 Method 1: Differencing single frames

In this method, two single frames (usually consecutive) with identical exposure conditions are differenced to yield an image lacking FPN but containing double the power of read and shot noise. Results are used to calculate variance statistics for the entire frame or a region of interest. Frame mean is calculated from any individual frame. The variance of the differenced frame will have noise modeled as

$$s_{\text{Difference Total}}^2 = 2s_{Shot}^2 + 2s_{\mathrm{R}ead}^2 \tag{2}$$

To calculate variance data for the mean-variance curve, the total variance of differenced frame must be divided by two to account for noise power doubling. This equation results from fact that FPN is constant from frame to frame but read and shot noise are not correlated frame to frame. There is no need to work with separate offset frames for this approach because the differencing removes it.

2.4 Method 2: Differencing single frame with mean frame

In this method, the user collects many consecutive frames (e.g. 25 to 100) and finds the mean or median pixel response over time to construct an image that has temporal effects (e.g. shot and read noise) removed but FPN retained (the more samples collected, the better temporal effect removal). This frame can then be differenced with any single frame to give an image that lacks FPN but retains read and shot noise statistics from the single frame. Variance is calculated from the differenced frame. Since noise power did not double in this method, there is no need to divide variance by two. All temporal noise is due to the single frame. Frame mean is calculated from the composite frame or any individual frame.

2.5 Method 3: 2-Point Correction

This method effectively performs a 2-point correction (gain and offset) on all images prior to calculating statistics. The user makes the above-mentioned Method 2 frame lacking temporal effects for a setting where the pixel values are between ~35% and 50% of full well. The resulting image is dominated by FPN. A similarly-constructed offset (dark) frame is subtracted from this FPN-dominated frame (i.e. offset correction). The differenced frame is used to calculate a single mean

Downloaded From: http://proceedings.spiedigitallibrary.org/ on 05/18/2015 Terms of Use: http://spiedl.org/terms

value for the entire 2-D region of interest. All pixels are then divided by this mean value to find relative normalized photo-response (i.e. gain). The reciprocal of pixels relative normalized photo-response can be used to adjust for photo-response non-uniformity and FPN in subsequent images (i.e. gain correction). To do this, a single frame is corrected for offset and then divided on pixel-by-pixel basis with the normalized photo-response frame. The resulting frame lacks FPN but retains read and shot noise.

In this paper we do not present results from Method 3 because we did not collect an accurate record of system offset.

2.6 Hybrid Method

The Hybrid Method relies on a large number of spatial samples and a much smaller number of temporal samples to attempt to increase sample number (and thus measurement accuracy). Figure 2 shows how the Hybrid Method approaches the statistics of the image sensor samples and shows sample defects and non-uniformities. For every pixel in an M x N grid, the temporal sample mean and sample variance are calculated. This results in M x N mean and variance values, which should have normal distribution according to the central limit theorem. The mean of the mean distribution is selected as the mean value for plotting. The mean value of variance distribution is selected as the variance value for plotting. Alternatively, the median of either distribution could be chosen.

The Hybrid Method implicitly corrects for FPN by selecting sample mean and sample variance from middling values of the normal distribution. Values from center of distribution are very nearly identical, indicating very similar pixel responses. In fact, one can argue these central values have characteristics of ideal pixels from the designer's perspective and lack non-idealities. In the absence of non-idealities (i.e. identical or very similar pixels), sampling over space and time are equivalent with respect to Poisson statistics and the sample mean and variance values accurately reflect the underlying distribution. For identical pixels, randomness in photon arrival over space has the same characteristic as over time and both result in the same values.

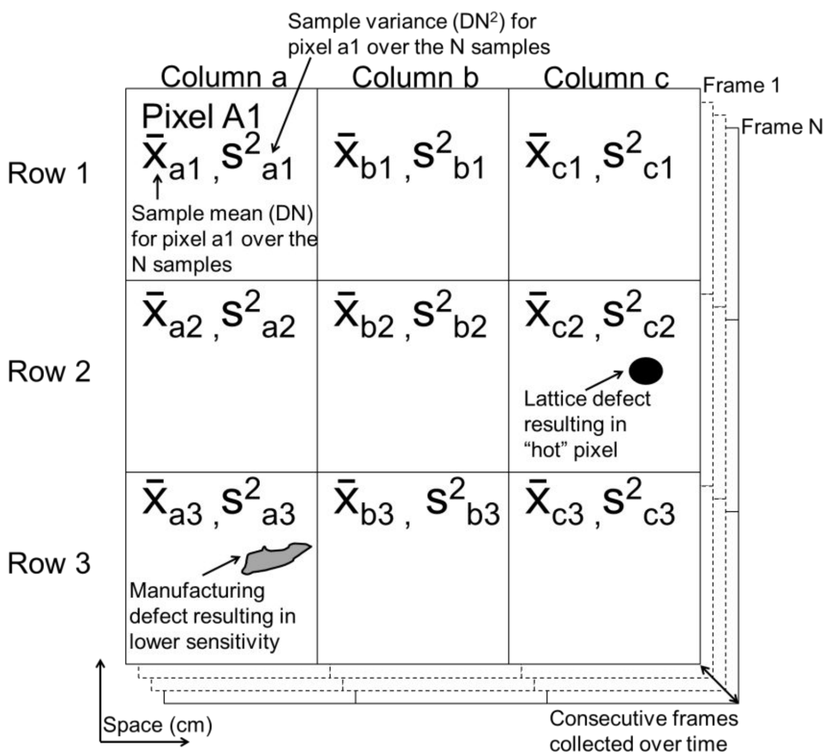

Figure 2. Sample statistics calculated from pixel array used in Hybrid Method.

2.7 Derivation of Hybrid Method Using Means

Assume we make N samples of M pixels for a total number of samples equal to NM. The mean for one pixel is

Downloaded From: http://proceedings.spiedigitallibrary.org/ on 05/18/2015 Terms of Use: http://spiedl.org/terms

$$\bar{x} = \frac{1}{N}\sum_{n=1}^{N} x_n \tag{3}$$

And the mean over all samples in all pixels is

$$\bar{y} = \frac{1}{M}\sum_{m=1}^{M}\left(\frac{1}{N}\sum_{n=1}^{N} x_n\right)_m \tag{4}$$

By properties of addition commuting and multiplication distributing we find

$$\bar{y} = \frac{1}{MN}\sum_{m=1}^{M}\left(\sum_{n=1}^{N} x_n\right)_m \tag{5}$$

Here, calculating the mean of each pixel and then the mean of these means is equivalent to simply summing up all the spatial and temporal samples and dividing by the total number of samples *MN*.

The unbiased sample variance for one pixel is

$$S^2 = \frac{1}{N-1}\sum_{n=1}^{N}(x_n - \bar{x})^2 \tag{6}$$

The mean of the sample variances over all pixels is

$$\bar{S}^2 = \frac{1}{M}\sum_{m=1}^{M}\left(S^2\right)_m \tag{7}$$

$$\bar{z} = \frac{1}{M}\sum_{m=1}^{M}\left(\frac{1}{N-1}\sum_{n=1}^{N}(x_n - \bar{x})^2\right)_m \tag{8}$$

$$\bar{z} = \frac{1}{M(N-1)}\sum_{m=1}^{M}\left(\sum_{n=1}^{N}(x_n - \bar{x})^2\right)_m \tag{9}$$

Again, the result simply sums up all the individual variances from pixels across space. The result is more samples for the statistics and less error in the estimate of the mean of the sample variances. Note that because the statistic of interest for the grid of sample variances is the mean, the procedure to compute the error via the standard error still applies.

3. IMPLEMENTATION

Data was acquired from CMOS image sensor with 6T unit-cell architecture, 14um pixel pitch, and 1280 x 1024 format[2]. Presented data were collected with 385 Hz frame rate and 2msec integration time. Sensor settings were held constant while flood illumination was varied. Flood illumination was supplied by high-power 530nm LED and 12” integrating sphere. For every illumination level 70 consecutive frames were collected.

4. RESULTS

4.1 Performance of methods applied to restricted pixel set

The three methods (1, 2, and Hybrid) were compared on 315x280 restricted pixel region. The restricted region was selected as a less stressing case for the method and has lower photo-response non-uniformity (PRNU) than the array as a whole. Here PRNU = 4.9%. Photon transfer results are shown in Figure 3.

Downloaded From: http://proceedings.spiedigitallibrary.org/ on 05/18/2015 Terms of Use: http://spiedl.org/terms

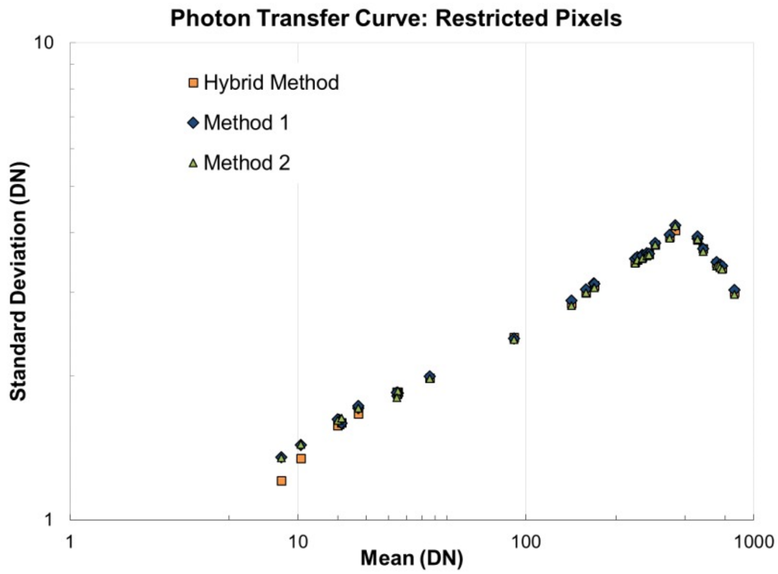

Figure 3. Photon transfer curves of restricted pixel region.

Mean-variance results of values between 10% and 90% of full well for the restricted pixel region are shown in Figure 4. The derivation of using a mean-variance curve for extracting conversion gain can be found in the Janesick text and our prior work[1-2].

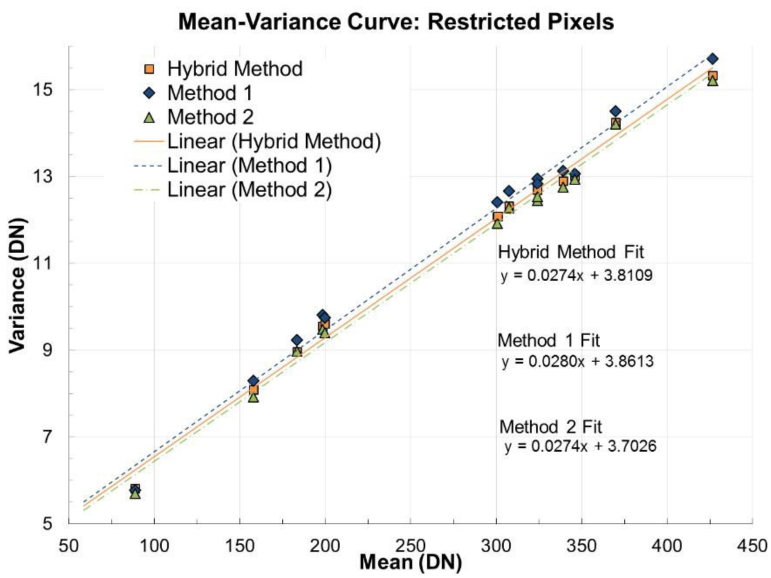

Figure 4. Mean-variance curves of restricted pixel region.

Histograms for restricted pixel region of the mean and variance values of an exposure yielding ~50% full well are shown in Figure 5. The mean value of each distribution was selected for inclusion in the photon transfer plot for Hybrid Method. Note, the variance histogram has slight positive skew.

Downloaded From: http://proceedings.spiedigitallibrary.org/ on 05/18/2015 Terms of Use: http://spiedl.org/terms

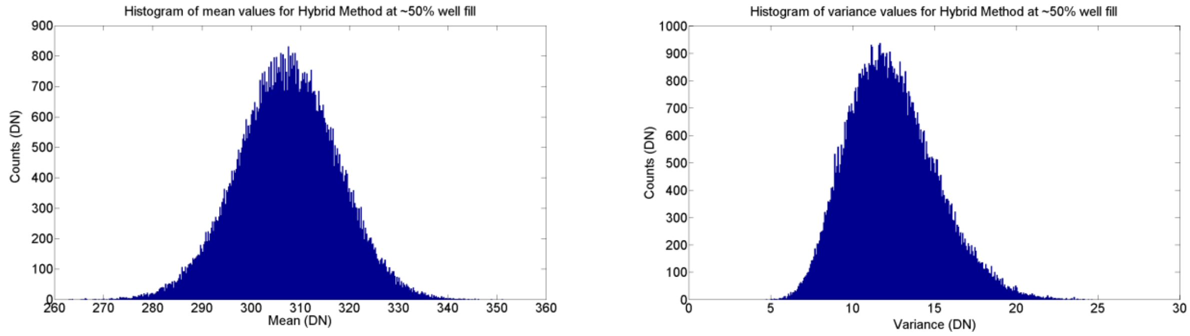

Figure 5. Histograms of sample mean and sample variance values of frame of ~50% well fill used in Hybrid Method for restricted pixels.

4.2 Performance of methods applied to entire image

The three methods (1, 2, and Hybrid) were compared on all 315x1280 pixels collected with no attempt to select an area of low FPN. Here, PRNU = 5.5%. Though PRNU is slightly higher than restricted region, results are materially the same. Photon transfer results are shown in Figure 6.

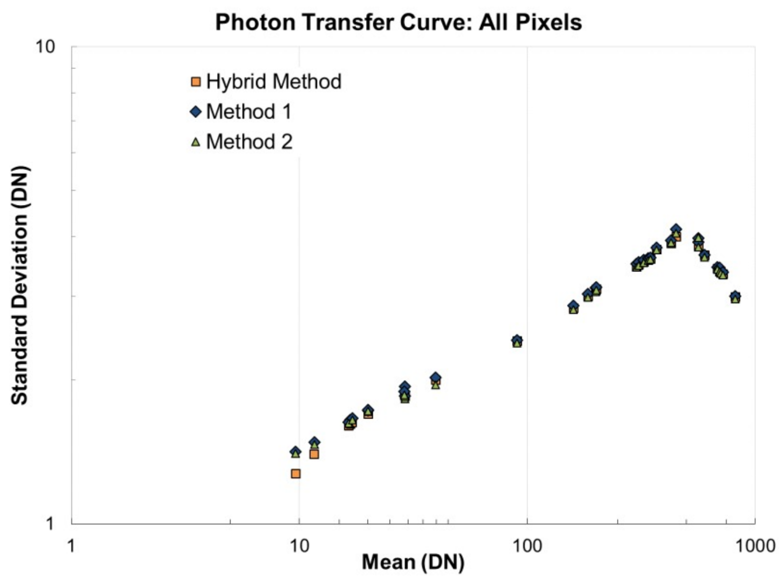

Figure 6. Photon transfer curves for all pixels.

Mean-variance results of values between 10% and 90% of full well are shown in Figure 7 and relevant histograms in Figure 8.

Downloaded From: http://proceedings.spiedigitallibrary.org/ on 05/18/2015 Terms of Use: http://spiedl.org/terms

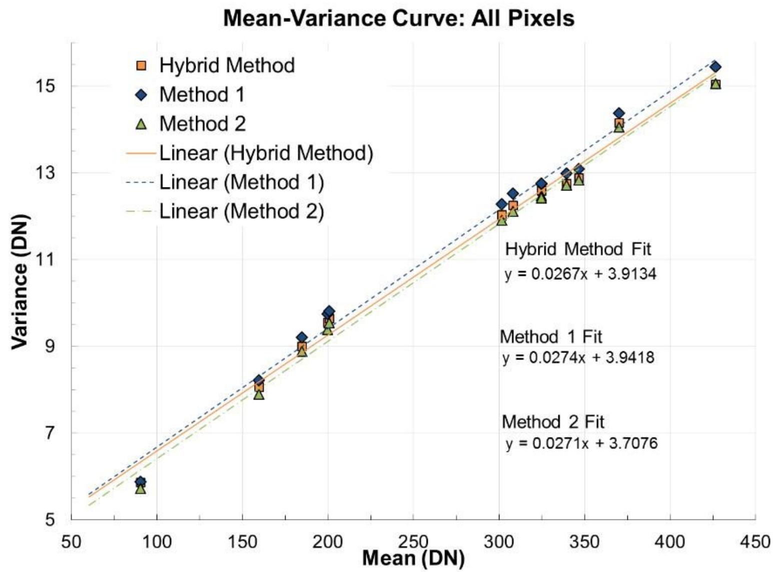

Figure 7. Mean-variance curve for all pixels.

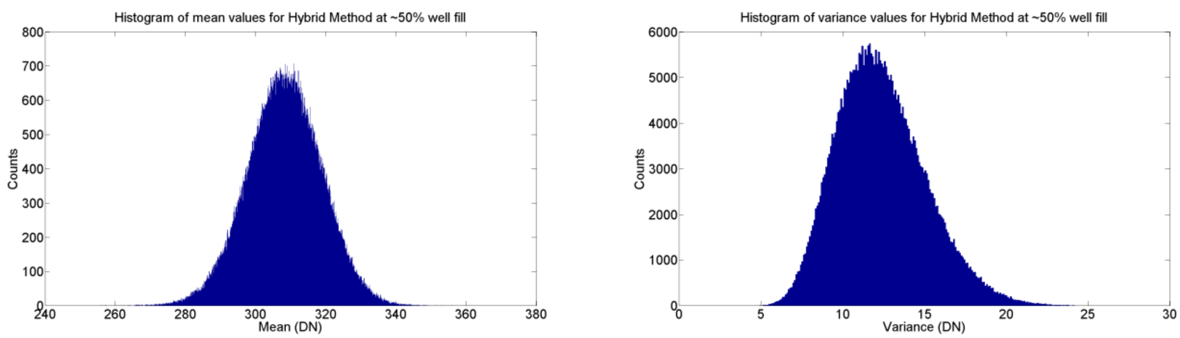

Figure 8. Histograms of sample mean and sample variance values of frame of ~50% well fill used in Hybrid Method for all pixels.

4.3 Conversion gain calculations

Table 1 shows results from the three methods, which yield slightly different measures of system conversion gain. PRNU values were calculated from composite frame of ~50% well fill using methods outlined in Appendix A.

Table 1. Conversion gain extracted from mean-variance curves.

Image Sensor Extracted Conversion Gain				
	Hybrid Method	Method 1	Method 2	Region
Slope fit (DN/e-)	0.0267	0.0274	0.0271	Restricted pixels, PRNU = 4.9%
Conversion gain (e-/DN)	37.5	36.5	36.9	Restricted pixels, PRNU = 4.9%
Slope fit (DN/e-)	0.0274	0.0280	0.0274	All pixels, PRNU = 5.5%
Conversion gain (e-/DN)	36.5	35.7	36.5	All pixels, PRNU = 5.5%

5. DISCUSSION AND ISSUES

All methods for both restricted pixels and full frame give materially similar results for conversion gain. The Hybrid Method compares well against standard methods. For restricted pixels, the Hybrid method gives relative error of 2.6% and 1.5% from Method 1 and Method 2 respectively. For all pixels, the Hybrid method gives relative error 2.2% and 0.0% from value of Method 1 and Method 2 respectively.

Downloaded From: http://proceedings.spiedigitallibrary.org/ on 05/18/2015 Terms of Use: http://spiedl.org/terms

An intuitive explanation for the success of the Hybrid Method relies on the fact that that pixels are designed to be identical. Pixel-level deviations from this design (i.e. FPN) are typically normal deviations from ideal values. The center of such a distribution accurately reflects the value of an ideal pixel (*i.e.* after FPN correction). Thus, the Hybrid Method implicitly rather than explicitly corrects for FPN by selecting sample mean and variance from pixels that most closely fit characteristics of ideal pixel. It follows that the Hybrid Method might fail if FPN is not normally distributed.

The central limit theorem is useful in showing why the method works. As expected, the histograms of sample mean and variance largely conform to normal distributions. This indicates that pixels more or less have similar responses without correcting for FPN. In the absence of FPN, a center value from the histogram accurately reflects a measure of the statistic. In the presence of moderate and normally-distributed FPN, the same is largely true.

6. CONCLUSION

With success of the hybrid method, users have another tool for calculating conversion gain that can either act as primary or supplemental method. Preliminary arguments were given for why the method worked. We presented results from mean of sample mean and mean of sample variance approach, but median of sample mean and median of sample variance yield similar data. The hybrid method gives a well-behaved mean-variance curve that compares well with other 2D methods. Future work may add rigor to these arguments and try to determine what level of PRNU might cause the method to fail.

APPENDIX A: HOW TO CALCULATE PRNU

Photo response non-uniformity (PRNU) measures the variation in response of a pixel array to flood illumination. It can be calculated as

$$PRNU = \frac{1\sigma}{FrameMean} \qquad (10)$$

Here *PRNU* is photo response non-uniformity, σ is the standard deviation of pixel response, and *FrameMean* is the mean value of the current frame. Typically PRNU is expressed as a percentage. The frame or frames used for PRNU calculation should have read and shot noise removed. This can be accomplished by averaging pixel responses over many consecutive frames for a given illumination and sensor condition (e.g. 25 to 100 frames).

ACKNOWLEDGMENT

We thank Jason Fields, Ronald Bloom, and Terry Lomheim. Funding was provided by The Aerospace Corporation through the Sustained Experimentation and Research for Program Applications (SERPA) and Internal Research and Development (IR&D) programs.

REFERENCES

[1] Janesick, J. R., Photon transfer, SPIE press San Jose (2007).

[2] Jacquot, B. C., Johnson-Williams, N., “Multiple-samples-method enabling high dynamic range imaging for high frame rate CMOS image sensor by FPGA and co-processor,” SPIE Opt. Eng. + Appl., 921609 (2014).

Downloaded From: http://proceedings.spiedigitallibrary.org/ on 05/18/2015 Terms of Use: http://spiedl.org/terms