Compression of raw PTF Multi-Extension FITS formatted images and Impacts on Image Co-addition

F. Masci, 3/21/2016, v2.0

1 Introduction

We used the *fpack* utility program to explore the amount of *lossy* compression that can be achieved for raw Multi-Extension FITS (MEF)-formatted image files from the Palomar Transient Factory (PTF) without significantly affecting their information content. These files consist of a primary header and 12 FITS-image extensions. These extensions contain CCD-specific headers and pixel data for each of the 12 CCDs in the PTF camera. The pixel data are represented in *signed* 16-bit format. The presence of BSCALE=1 and BZERO=32768 in the headers converts these to *unsigned* 16-bit on output where the range is -2^{15} + BZERO to 2^{15} – 1 + BZERO, or 0 to 65535 ADU.

The statistical properties of the images following decompression relative to the originals are explored in order to find an optimal set of compression parameters. This is performed for the light-sensitive CCD pixels and their accompanying over-scan regions. The latter are used to correct for floating biases during processing.

The FITS compression program (with references) is described in http://heasarc.gsfc.nasa.gov/docs/software/fitsio/fpack/

The intent here is to use this compression scheme on similar MEF-formatted images from the Zwicky Transient Facility (ZTF) in order to account for the limited bandwidth in data-transfers from Palomar to Caltech. The goal is to obtain an overall compression factor of ~3 in the file packet relative to its original size from the ZTF camera, all over the sky. We therefore explored two fields with varying source density (or complexity): one at low galactic latitude, and another at high latitude.

2 PTF Test Images

Figure 1 shows a zoom-in on two CCD images from two MEF files used in this study. The CCDs in the MEF files are surrounded by over-scan regions ~14 to 20 pixels wide. In this study, the latter are not detached from their CCDs and compressed separately. They are compressed together with the light-sensitive CCD pixels. For ZTF, these over-scan regions will be packaged into separate extensions and hence compressed separately. See below for details on why this is important.

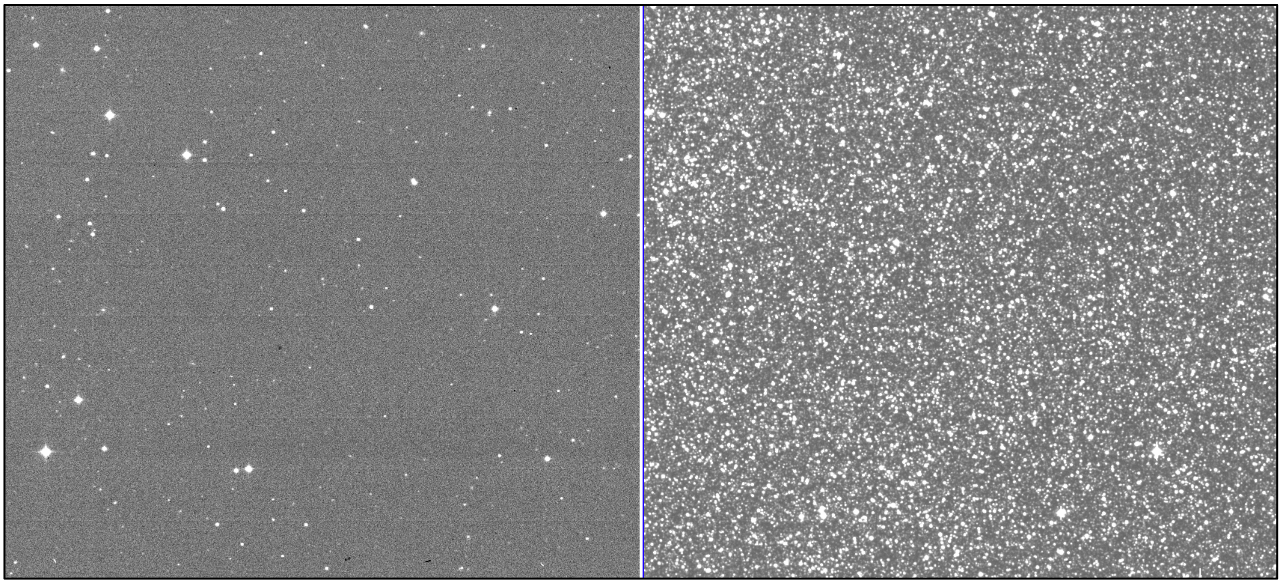

Figure 1: sample CCD sections from two raw PTF camera-image exposures used in this study. Both are CCD #4 from MEF files: `PTF201401010915_2_o_44987.fits` (left) and `PTF201408112176_2_o_40958.fits` (right). These represent fields observed at high and low galactic latitude respectively.

3 Optimal "fpack" execution

For PTF (and presumably ZTF) we find that the *fpack* call below is most optimal in terms of preserving the minimum number of noise-bits required for most science applications and for achieving desired compression factors of >~ 3. This is applicable to (i) all light-sensitive CCD-pixels regardless of sky location, and (ii) all accompanying over-scan regions, packaged as separate images (or FITS extensions).

```
fpack -i2f -r -w -q 4 input_MEF_file.fits
```

In brief, the command-line arguments correspond to:

- `-i2f` **:** internally convert images from integer (e.g., 16-bit) values to floating point (32-bit) values. This allows one of the floating-point based, *lossy* compression algorithms to be applied. Not specifying this will lead to lossless compression and lower compression ratios. Note: if the input pixel values are already in floating-point format, the `-i2f` switch can be omitted.
- `-r` **:** use the Rice compression method.
- `-w` **:** compress entire image (in each separate FITS extension) as one single tile. Applying the Rice algorithm on image sections (sub-tiles) will lead to lower compression factors
- `-q` <*n*>**:** *fixed* quantization level *n* when internally converting floating-point values to scaled integers to control the number of noise-bits to retain prior to compression. The average number of *incompressible* noise-bits preserved per pixel assuming Gaussian noise is predicted to be:

$N_{bits} \sim \log_2[q] + \log_2[\sqrt{12}]$. (1)

For $q = 4$, $N_{bits} \sim 3.8$. This implies $\sim 2^{round(3.8)} = 16$ possible discrete values for representing the pixel noise variation prior to compression. Increasing q will reduce the compression factor and make the compressed image pixel values more consistent with the original values following uncompression. For input (raw) 16-bit data, the compression ratio (input/output) using Rice is theoretically

$R \sim 16 / [N_{bits} + 1.2]$, (2)

where N_{bits} was defined in Eq. (1). For $q = 4$, this predicts $R \sim 3.2$.

Furthermore, a consequence of quantizing the noise is an increase in the overall noise due to "quantization error". The increase [in %] in pixel noise is predicted to be approximately:

$\textit{Noise increase} \sim 100*[\sqrt{(1 - 1/\{12q^2\})} - 1]\%$. (3)

For $q = 4$, we therefore expect an overall increase in the pixel noise following uncompression (relative to the original) of ~ 0.26 %.

The predictions from the above scaling relations are compared with those obtained for real PTF images below. Further details and command-line options can be found in the *fpack* User's Guide: http://heasarc.gsfc.nasa.gov/FTP/software/fitsio/c/docs/fpackguide.pdf

The output file from the above *fpack* call will be *input_MEF_file.fits.fz*, i.e., with *.fz* appended. It's important to note that only the pixel data are compressed across all extensions, not the headers. This file can then be subsequently uncompressed using the *funpack* utility.

4 Individual Image Statistics: original versus uncompressed

The following procedure was used to test the impact of three different noise-quantization levels (q = 4, 5, 6) for each field shown in Figure 1:

(i) entire raw PTF camera image MEF file was compressed using the *fpack* call in Sec. 3;
(ii) the compressed MEF file from (i) was uncompressed using the *funpack* utility;
(iii) when uncompressed, the CCD-specific pixel data in the MEF extensions will be in 32-bit floating point format with approximate dynamic range: -32768.0 to 32767.0. We split this uncompressed MEF file into its 12 constituent CCD FITS images and also add BZERO (= 32768) to the pixel values to place them on their original (pre-compressed) dynamic range: ~ 0 to 65535 ADU. This is to facilitate one-to-one statistical comparisons;
(iv) the original raw MEF file is also split into its 12 constituent CCD FITS images;
(v) the uncompressed pixel data from (iii) are compared to their original (pre-compressed) values from (iv) on a per CCD-basis using statistical metrics. These metrics are computed separately for the light-sensitive and over-scan CCD pixels.

Tables 1 and 2 summarize some of the relevant pixel metrics before and after *lossy* compression applied (with subsequent *uncompression*) for three different noise-quantization levels *q* (Section 3). Each table refers to the two source-density cases shown in Figure 1: low or nominal source density (Table 1) and high density (Table 2).

The columns in Tables 1 and 2 are defined as follows:

qlevel: quantization parameter *q*
cpr: overall MEF-file compression ratio (original size/compressed size)
nbits: average number of noise-bits preserved per pixel predicted by Eq. 1
ccd: zero-based PTF chip number
davgccd: average of pixel difference: "original – compressed" for light-sensitive CCD region
davgbias: average of pixel difference: "original – compressed" for optimal over-scan region in CCD
sigccd: robust spatial pixel sigma using percentile difference: 0.5*(84.13% – 15.85%) for light-sensitive CCD region *before compression*
sigbias: robust spatial sigma using percentile difference: 0.5*(84.13% – 15.85%) for optimal over-scan region in CCD *before compression*
psigccd: percentage change in **sigccd** metric: (compressed – original) / original; values > 0 imply an increase in the pixel noise
psigbias: percentage change in **sigbias** metric: (compressed – original) / original; values > 0 imply an increase in the pixel noise

Table 1: Statistics for nominal/low source-density field

qlevel	cpr	nbits	ccd	davgccd	davgbias	sigccd	sigbias	psigccd	psigbias
				DN	DN	DN	DN	%	%
4	3.22	3.79	00	3.43e-04	2.62e-03	35.00	7.5	7.90e-01	3.80e+00
			01	-6.21e-04	3.41e-03	32.50	7.0	9.92e-02	2.02e+00
			02	2.42e-03	1.28e-03	25.00	6.0	-1.07e+00	3.52e+00
			03	2.81e-03	1.70e-02	97.00	94.0	1.65e-01	1.86e-02
			04	-8.28e-04	1.57e-02	38.00	10.0	3.73e-01	5.99e+00
			05	-3.34e-05	-1.09e-02	38.00	5.0	1.36e-01	9.62e+00
			06	1.05e-03	-6.60e-03	33.50	10.0	9.47e-01	1.34e+00
			07	3.76e-04	2.38e-02	29.50	9.0	1.32e-02	-8.36e-01
			08	2.33e-04	2.16e-03	30.50	6.0	2.91e-01	1.14e+01
			09	4.07e-05	1.76e-02	28.50	7.0	4.83e-01	9.80e-01
			10	4.93e-04	6.69e-03	28.00	10.0	1.19e+00	2.38e-01
			11	-2.94e-04	-1.13e-02	35.50	6.0	2.94e-01	8.49e+00
5	3.04	4.11	00	7.39e-04	-1.98e-03	35.00	7.5	6.89e-01	2.88e+00
			01	-1.46e-03	2.37e-03	32.50	7.0	-9.31e-02	1.12e-01
			02	-4.10e-04	7.99e-03	25.00	6.0	-9.96e-01	1.92e+00
			03	-7.81e-05	-2.07e-02	97.00	94.0	8.36e-02	-2.97e-02
			04	2.45e-04	-2.70e-02	38.00	10.0	3.50e-01	4.19e+00
			05	-1.12e-03	8.33e-03	38.00	5.0	4.63e-02	5.64e+00
			06	-2.22e-04	9.02e-04	33.50	10.0	8.42e-01	1.14e+00
			07	-2.29e-04	-5.38e-03	29.50	9.0	-1.13e-01	-1.95e+00
			08	-6.33e-04	-6.23e-03	30.50	6.0	3.33e-01	9.83e+00
			09	-1.95e-04	-4.83e-03	28.50	7.0	4.87e-01	-8.63e-01
			10	-3.39e-04	-5.80e-03	28.00	10.0	1.06e+00	-5.54e-02
			11	-1.09e-03	-8.22e-03	35.50	6.0	2.59e-01	7.48e+00
6	2.89	4.37	00	-3.49e-04	2.92e-03	35.00	7.5	6.31e-01	7.90e-01
			01	-4.31e-04	4.12e-03	32.50	7.0	-8.71e-02	-9.50e-01
			02	6.96e-04	6.97e-04	25.00	6.0	-1.17e+00	3.74e-01
			03	7.61e-04	1.80e-03	97.00	94.0	7.25e-02	-6.31e-02
			04	1.01e-03	6.88e-03	38.00	10.0	2.90e-01	3.43e+00
			05	1.46e-03	7.82e-03	38.00	5.0	4.63e-02	3.85e+00
			06	5.33e-04	5.36e-03	33.50	10.0	9.04e-01	5.86e-01
			07	-1.16e-03	2.51e-03	29.50	9.0	-1.42e-01	-2.04e+00
			08	-2.01e-04	4.78e-03	30.50	6.0	2.95e-01	8.15e+00
			09	-3.72e-04	-3.19e-03	28.50	7.0	4.25e-01	-1.59e+00
			10	-2.39e-05	9.17e-03	28.00	10.0	1.05e+00	-5.05e-01
			11	6.25e-04	-5.31e-03	35.50	6.0	1.49e-01	5.43e+00

Table 2: Statistics for high source-density field

qlevel	cpr	nbits	ccd	davgccd	davgbias	sigccd	sigbias	psigccd	psigbias
				DN	DN	DN	DN	%	%
4	3.05	3.79	00	1.08e-03	-1.89e-02	590.50	7.0	1.74e-01	2.05e+02
			01	-5.45e-03	-8.34e-02	533.50	10.0	1.11e-01	1.16e+02
			02	9.21e-03	2.58e-02	511.00	5.0	8.22e-02	3.35e+02
			03	-5.46e-03	-4.35e-02	97.00	92.0	-1.70e-01	7.88e-01
			04	2.28e-02	5.88e-02	586.50	9.0	1.65e-01	3.07e+02
			05	3.31e-03	-1.13e-01	543.50	4.5	1.79e-01	5.53e+02
			06	-7.91e-04	7.38e-02	625.00	9.0	1.02e-01	1.89e+02
			07	1.14e-02	1.91e-01	605.50	8.0	3.65e-02	1.97e+02
			08	2.70e-03	9.51e-02	661.50	5.0	8.00e-02	5.54e+02
			09	6.96e-03	1.56e-05	540.50	5.0	1.17e-01	4.66e+02
			10	5.45e-03	1.66e-01	434.00	7.5	3.94e-02	1.68e+02
			11	9.40e-03	-5.79e-02	613.00	5.5	7.39e-02	3.14e+02
5	2.88	4.11	00	5.39e-04	-3.37e-02	590.50	7.0	1.25e-01	1.50e+02
			01	9.39e-04	1.42e-01	533.50	10.0	6.22e-02	8.08e+01
			02	-5.64e-03	-2.03e-02	511.00	5.0	6.29e-02	2.50e+02
			03	-1.69e-03	2.87e-02	97.00	92.0	-1.71e-01	6.77e-01
			04	-1.59e-02	-1.16e-01	586.50	9.0	1.41e-01	2.28e+02
			05	5.38e-04	-2.76e-01	543.50	4.5	8.51e-02	4.20e+02
			06	-3.37e-03	3.70e-02	625.00	9.0	5.14e-02	1.37e+02
			07	-1.02e-02	-7.96e-02	605.50	8.0	-8.66e-04	1.41e+02
			08	-2.83e-03	-1.07e-01	661.50	5.0	7.78e-02	4.21e+02
			09	-5.23e-03	1.92e-01	540.50	5.0	7.61e-02	3.54e+02
			10	2.96e-03	1.63e-02	434.00	7.5	3.49e-02	1.22e+02
			11	-2.74e-03	-2.02e-01	613.00	5.5	3.14e-02	2.32e+02

6	2.75	4.37	00	-1.94e-04	-2.10e-03	590.50	7.0	1.14e-01	1.14e+02
			01	-2.79e-03	6.56e-02	533.50	10.0	1.81e-02	6.29e+01
			02	7.35e-03	-3.68e-02	511.00	5.0	5.56e-02	1.95e+02
			03	1.09e-03	-2.72e-02	97.00	92.0	-2.62e-01	5.16e-01
			04	-4.31e-03	1.38e-01	586.50	9.0	1.06e-01	1.75e+02
			05	-1.92e-03	-2.45e-02	543.50	4.5	5.16e-02	3.34e+02
			06	4.62e-03	3.06e-02	625.00	9.0	4.08e-02	1.03e+02
			07	4.80e-03	6.06e-02	605.50	8.0	-4.84e-04	1.06e+02
			08	3.67e-03	1.33e-01	661.50	5.0	2.51e-02	3.34e+02
			09	-1.60e-03	-1.82e-01	540.50	5.0	9.61e-02	2.76e+02
			10	7.17e-04	1.18e-01	434.00	7.5	1.08e-02	9.27e+01
			11	-2.19e-03	-2.51e-02	613.00	5.5	7.65e-03	1.81e+02

5 Impact on Image Co-addition

With the sacrifice of a number of noise bits in the Rice compression of individual images, one may suspect that this leads to "irrecoverable information" (signal) when co-adding images to improve the pixel signal-to-noise ratio. We performed an experiment by co-adding a series of nine PTF CCD images of the same field, each with different numbers of inputs. A zoomed-in montage of these co-adds with increasing depth is shown in Figure 2. The co-add regions shown here were made from original images with no compression.

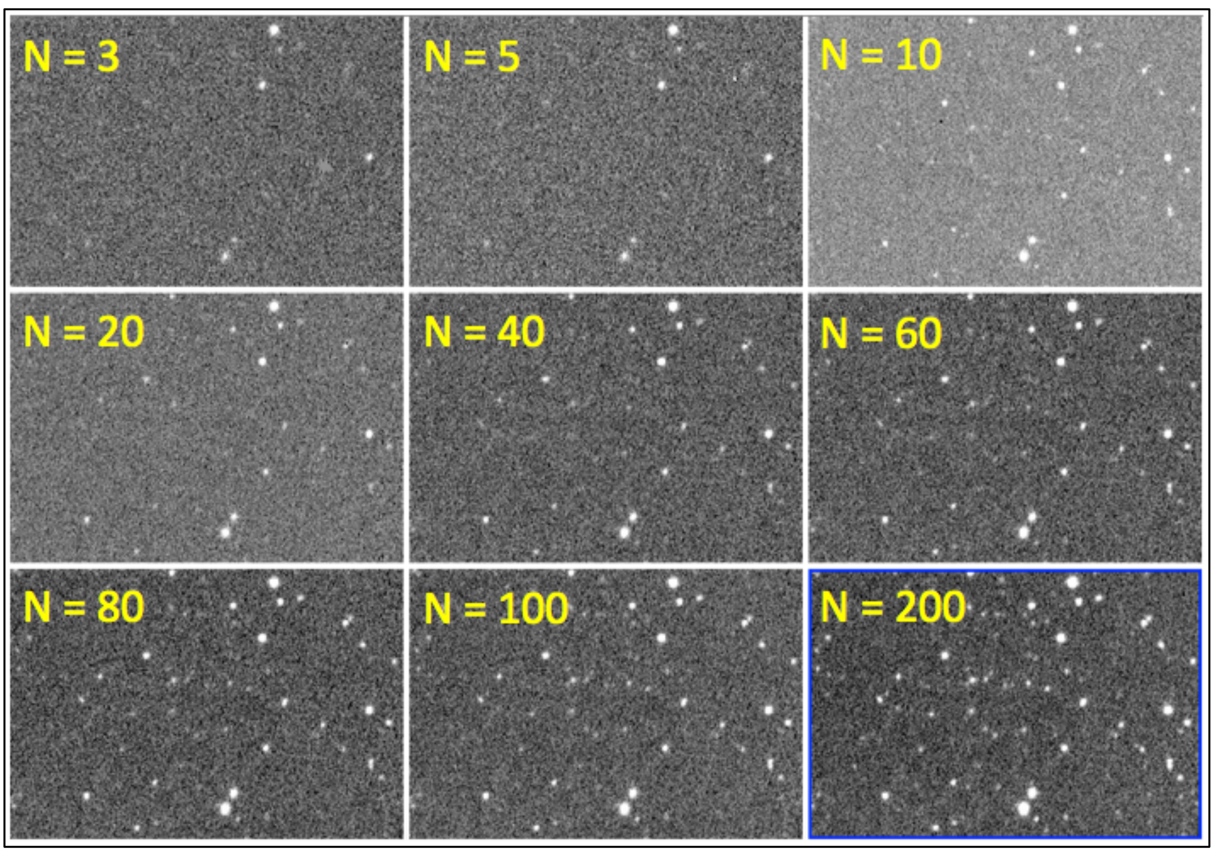

Figure 2: zoomed-snapshot of a series of co-adds with different numbers of frames: $N = 3$ to $N = 200$. All are from the same PTF CCD and field.

The same input images were then compressed using the "q=4" *fpack* execution call from Section 3 and then re-coadded (following decompression). Image statistics were then computed on the nine original co-adds and their input-compressed counterparts. These statistics comprise the following:

(i) robust *spatial* pixel-sigma using percentile difference: 0.5*(84.13% – 15.85%) on a relatively source-free region, i.e., the region shown in Figure 2.
(ii) median of trimmed pixel-*stack* standard-deviation in this same region over all pixels.

The statistic in (i) is expected to follow a $\sqrt{N}$ dependence assuming negligible residual spatial correlations, including source confusion; (ii) gives an estimate of the overall pixel uncertainty at an individual epoch, or rather the degree of temporal repeatability. The statistics for the original and uncompressed-input co-add cases are shown as a function of N in Figures 3 and 4.

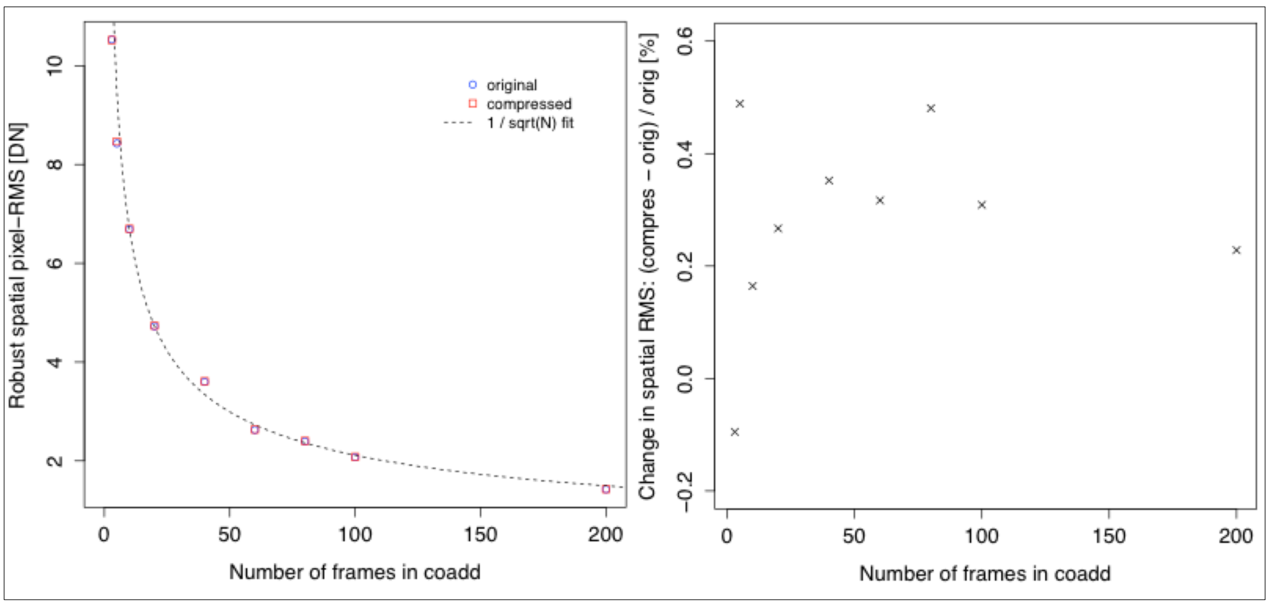

Figure 3: *left* – robust statial pixel RMS; the dashed line is a by-eye fit of the 1/√N dependence; *right* – relative change in the spatial RMS going from original to compressed image inputs prior to co-addition.

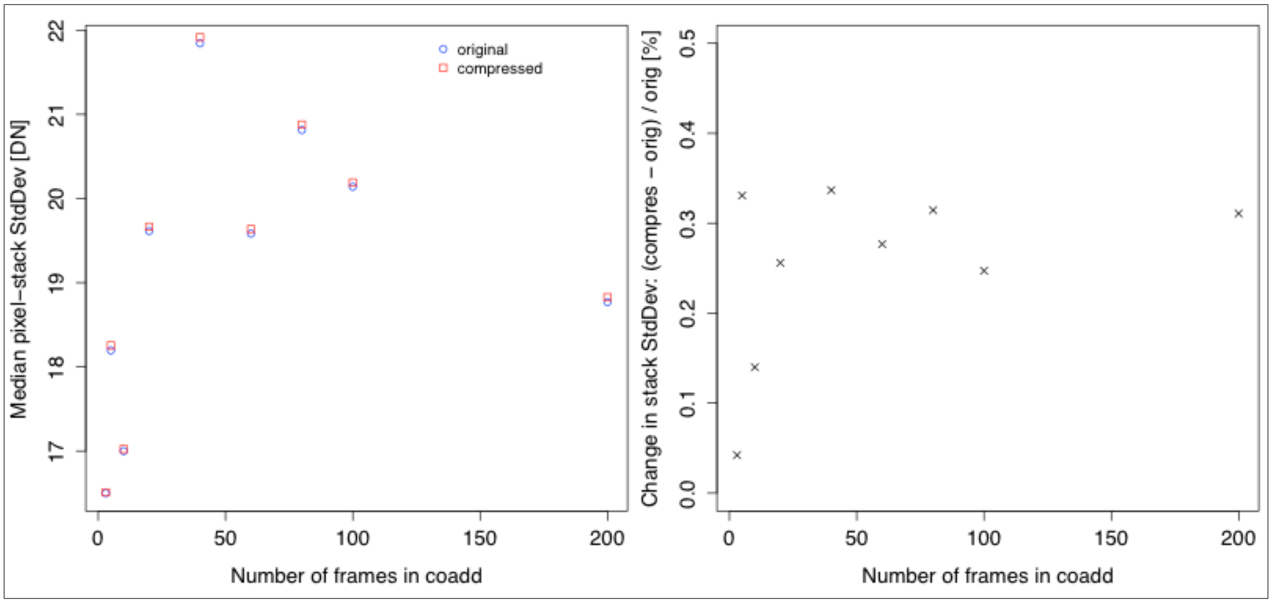

Figure 4: *left* – median of trimmed pixel-stack standard-deviation; *right* – relative change in this quantity going from original to compressed image inputs prior to co-addition.

6 Summary and Conclusions

- For low to nominal source-density fields typical of the PTF survey, quantization levels of q = 4 and 5 achieve compression ratios of ~ 3.22 and ~3.04 respectively. The pixels here are mostly dominated by background noise (including read-noise). For these cases, the increase in the pixel noise due to quantization error for the light-sensitive regions is ~ 0.2 to 1.0% across most CCDs (excluding outliers). q = 6 results in a compression ratio of ~ 2.89 and is below our goal of 3. We omit q = 6 from further consideration.
- For complex/high source-density fields, quantization levels of q = 4 and 5 achieve compression ratios of ~ 3.05 and 2.88 respectively. The pixel noise here is dominated by source-confusion and Poisson noise, and the latter is responsible for the lower compressibility in general. Nonetheless, q = 4 still satisfies our goal of >~ 3 while q = 5 does not. For q = 4, the increase in the pixel noise due to quantization error for the light-sensitive CCD regions is at most ~ 0.2%. This increase is smaller than the low source-density case because structure associated with source confusion dominates the spatial pixel noise. This confusion noise is considered information and compresses rather well.
- Changes in the overall average levels (pixel values) in the light-sensitive and over-scan regions are consistent with zero. I.e., compression and subsequent decompression results in no net bias offset relative to the original values.
- Relatively large increases in the pixel noise are observed for the over-scan regions across all CCDs following compression. For the low source-density field, this can be up to ~10% while for the high source-density field, the increase is generally > 100%. These large increases are due to the fact that the over-scan regions remained attached to the noisier light-sensitive pixels when they were compressed. The entire CCDs (both light-sensitive and over-scan pixels) were compressed as single tiles and

the Rice algorithm was dominated by information in the more numerous light-sensitive pixels. For optimal compression of over-scan regions, these must be packaged and compressed separately, i.e., as separate FITS extensions.

- In conclusion, q = 4 appears to provide the best compromise in compressing all CCD pixels to factors >~ 3 regardless of sky-location without overly inflating the pixel noise due to quantization error. This case preserves ~ 4 noise bits per pixel in general and based on PTF data, leads to an increase in the pixel noise of ~ 0.2 to 1% for background-dominated noise in single-image exposures. For co-addition, the increase in the overall pixel noise measured on co-adds using both spatial and stack (temporal) statistics is ~ 0.3% and is independent of the number of images co-added. It remains to be seen if the ZTF CCDs will have similar noise properties and limiting behaviors following co-addition.