Bad pixel detection for on-board data quality improvement of remote sensing instruments in CubeSats

Josua Florczak*[a], Konstantin Ntokas[b], Tom Neubert[a], Egon Zimmermann[a], Heinz Rongen[a], Uwe Clemens[a], Martin Kaufmann[b], Martin Riese[b] and Stefan van Waasen[a,c]
[a] Central Institute of Engineering, Electronics and Analytics - Electronic Systems (ZEA-2), Forschungszentrum Jülich, Germany; [b] Institute of Energy and Climate Research (IEK-7), Forschungszentrum Jülich, Germany, [c] Faculty of Engineering, Communication Systems (NTS), University of Duisburg-Essen, Germany

ABSTRACT

Large commercial of the shelf pixel arrays in current remote sensing instruments used in CubeSats make on-board processing increasingly important and enables data improvement. Therefore, we first consider the individual steps of the adapted bad pixel detection algorithm - ISMFD. In particular, we consider pixel-to-pixel variations and temporal flickering of pixels in commercial of the shelf sCMOS imaging sensors. We were able to detect an increase of bad pixels from (2.05±0.01)% to (4.1±0.1)% using real measurement images of the flying remote sensing instrument AtmoSHINE. A preliminary implemented adaptive on-board binning method was able to achieve a constant signal-to-noise-ratio on an image with a dynamic light intensity. The additional consideration of bad pixels in the binning method could demonstrate the achievement of data quality of the future remote sensing instrument AtmoLITE.
Keywords: bad pixel detection, CMOS imaging sensors, CubeSats, on-board processing, pixel binning

1. INTRODUCTION

The constantly growing demands on remote sensing instruments have led to a broader use of large pixel arrays. This is used for example in spectroscopy to measure atmospheric spectra with high spectral and spatial resolution, which can be used to retrieve temperature accurately. Remote sensing instruments in CubeSats are typically based on commercial off-the-shelf (COTS) components and take the benefit of state-of-the-art processing and high-resolution pixel arrays. Part of this image sensor family are scientific Complementary metal oxide semiconductor Image Sensors (CIS) with a high signal-to-noise ratio and highly integrated readout electronics, which are increasingly used in imaging applications[1,2]. The challenge with increasing data volumes is the transmission of all raw data due to the limited downlink capacity of CubeSats. On-board data processing is required if continuous observation is targeted. At the same time, COTS components are not specifically radiation hardened and are sensitive to displacement damage and total ionizing dose (TID)[3]. In instruments with CIS, such long-term disturbances degrade the pixels response over weeks to years, leading to an increase in bad pixels.

Bad pixels are characterized by either being unresponsive and providing a fixed value or deviating strongly from the expected average light response or dark current. The deviations in the dark current have a particularly strong effect low-light scenes[4,5]. A low signal-to-noise ratio (SNR) of pixels leads to greater uncertainties in the data products. To improve SNR methods such as increasing exposure time, average of multiple pictures or by binning of single pixels to super pixels are common solutions. This leads to reduced spatial, spectral or temporal resolution. Automated binning of pixels is supported by some state-of-the-art image sensors. Such mechanisms do not consider bad pixels lot to lot variations and the expected degradation due to radiation effects.

In this paper, we provide insight into the operation and thresholds of a bad pixel detection algorithm and introduce an adaptive pixel binning method that responds to sensor pixel degradation to improve SNR in low-light conditions. The analysis is performed for the Atmospheric Spatial Heterodyne Interferometer Next Exploration (AtmoSHINE)[6] and the upcoming Atmospheric Limb Interferometer for Temperature Exploration (AtmoLITE) missions. These instruments use spatial heterodyne interferometers to observe the spectral distribution of the O2 atmospheric A-Band signal in the mesosphere and lower thermosphere. These data are used to retrieve atmospheric temperature[1].

*j.florczak@fz-juelich.de; phone +49 2461 61-85526

2. BAD PIXEL DEFINITION AND IN-ORBIT DATA

2.1 Bad pixel

The definition of bad pixels is often mission specific and depends on the respective objectives. In this case, they can be defined by the deviation from the expected signal of the pixel to its real signal. The signal of a pixel is composed by photon driven response and the thermal dark current of the photodiode, as well as the gain of the signal[3]. Especially strong pixel deviations from the mean value in the photo response (photo response non-uniformity - PRNU), and the dark current (dark current non-uniformity - DCNU), are considered as bad pixels.
Independent of these spatial differences, pixels are also considered as bad if the dark current varies in time. This is described as dark current random telegraph signal (DC-RTS)[3] and is visible as flickering pixels on several frames. The deviations in dark current do not only come from manufacturing-related slight differences in the structure of each pixel. Also TID induced defects in the dielectric and interface states as well as defects in the silicon grid due to displacement damages, will also result cause stronger deviations of individual pixel dark currents. As a result, an increasing number of bad pixels can be observed if the detector is exposed to high energy particles. This increase in bad pixels makes it necessary to periodically monitor bad pixels in order to maintain consistent data quality.

2.2 Data

For the following discussions we use measured interferogram images from the remote sensing instruments AtmoSHINE which determine temperature profiles in the atmosphere[1] and simulated interferogram images for AtmoLITE. From AtmoSHINE we have 10 months of data collected in a sun synchronous orbit at 1000km altitude using an sCMOS imager of the type HWK1910A. With the large number of images and the long observation period, this data set is useful for studying degradation effects and flickering of bad pixels. From AtmoSHINE, we see a low-light scene in Figure 1(a), with a maximum possible sensor output of 2^{16} DN. AtmoLITE uses the GSENSE400BSI type sCMOS imager. Figure 1(b) shows a simulated interferogram image for AtmoLITE. At a maximum GSENSE output of 2^{12} DN, the simulated scene has a medium illumination. The image is a superposition of the analytical interferogram pattern of the instrument according to Harlander et al.[7]. Typical dark current for 0° C detector temperature and shot noise are superimposed. Since we determined the dark current from measurements of a GSENSE detector, Figure 1(b) also includes the dark current non-uniformity of this detector , which leads to the visible bad pixels. Figure 1 illustrates the data: the x-direction shows Fizeau Fringes (spectral information) and the y-direction shows the vertical distribution of atmospheric signal.

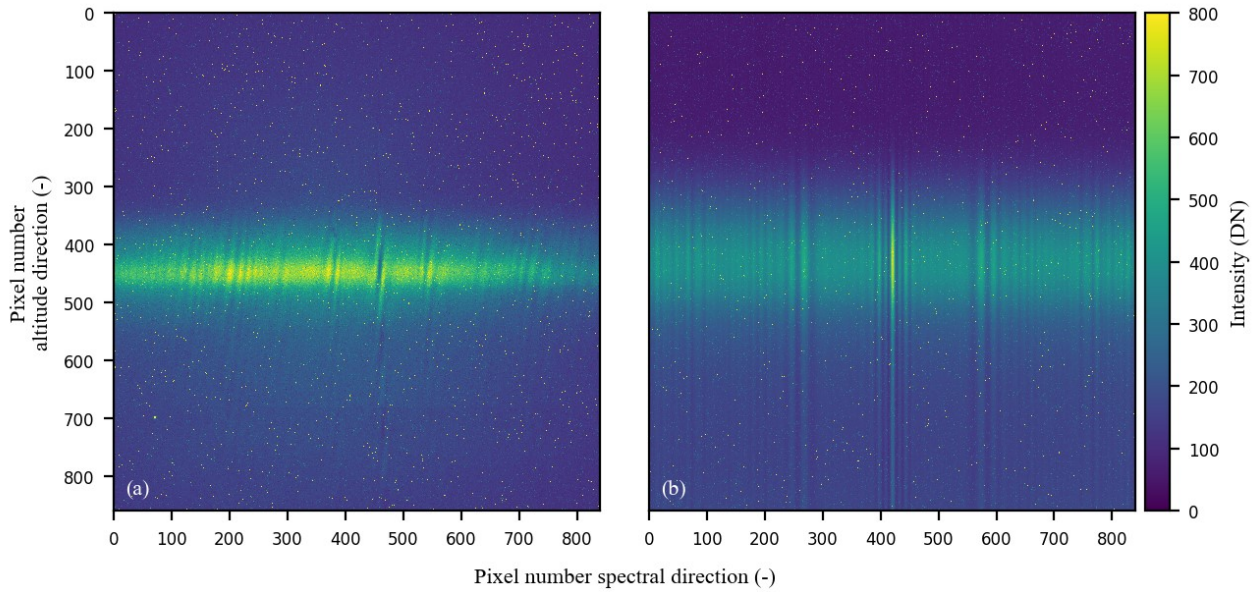

Figure 1 (a) Detector signal of the HWK sensor from the AtmoSHINE mission on Aug. 7, 2019, and (b) simulation of an image of the GSENSE sensor. Both images exhibit clearly visible bad pixels.

3. BAD PIXEL DETECTION

Because of the multiple influences of bad pixels, almost all scientific missions consider bad pixels before the main processing. The detection of flickering and slowly degrading pixels is not straight forward. Typically, bad pixels are determined under laboratory conditions on the ground. As we expect an increasing number of bad pixels during on-orbit operation, we decided to use the "Interframe Statistics Median Filtering Detection" (ISMFD) method[8]. The ISMFD method analyzes multiple frames and is able to detect flickering pixels within a measurement scene. Compared to other bad pixel detection methods[4], it performs simple arithmetic operations, is easy to adjust, and does not require special measurement routines, which prevents the need to interrupt continuous acquisition of measurement images. For the detection and correction of bad pixels during on-orbit operations, there are also special solutions, such as custom-designed CIS[9], which, however, severely limits the choice of image sensors and must be considered early in the mission planning process.

We adjusted the ISMFD for an exposure level independent detection[10] and introduced a relative gain matrix

$$r_f(i,j) = g_f(i,j)/\widetilde{m_f}(i,j) \ , \tag{1}$$

where $g_f(i,j)$ is an initial image f, as shown exemplarily in Figure 1, and $\widetilde{m_f}(i,j)$ is the associated image after a median filter. The i and j indicate the row and column position. The relative gain is independent of the exposure level, has a mean value of 1 in the ideal case and follows only the normal distribution of the shot noise. In this approach, we assume that the pixels considered in the median filter have a relatively homogeneous exposure. Because of an overall non-homogeneous exposure, we have chosen a relatively small median filter window of only 9x1 pixels along the altitude direction in order to compute the median image. However, other median filter windows can be selected, for example along the spectral direction or in multiple directions. Eq. 1 directly visualizes the relative deviation of a pixel to its neighbors. We see an example of a relative gain matrix from Eq. 1 in Figure 2(a). It can be seen that the brighter middle region deviates less from 1, which is due to the DCNU having a less strong effect in relative terms. Bad pixels, on the other hand, are visible as yellow or blue dots. Such relative deviations, together with fixed threshold values (r_{th}) are used in missions to define bad pixels. For this reason, we now define a single frame bad pixel matrix

$$N_f(i,j) = \begin{cases} 1 \text{ if } r_f(i,j) > r_{th} \\ 0 \text{ else} \end{cases} \text{ with } \ r_{th} = 1 \pm \alpha \times \sigma \ , \tag{2}$$

for each individual frame f as a function of the threshold value r_{th} and the gain matrix. We use a flexible threshold r_{th} by analyzing the histogram of $r_f(i,j)$. Three of these histograms can be seen in Figure 2(b). Here, the histogram from the HWK detector was created from 100 images taken on 7 Aug. 2019, which are similar to the one in Figure 1(a). The histogram for the GSENSE detector with noise was created from the relative gain matrix of Figure 2(a). For comparison, a histogram of a simulated interferogram signal without shot noise and dark current is shown as well. All histograms were normalized from the total number of pixels to an image of 860x860 pixels.

For the histogram of the simulated GSENSE detector image without noise, we see a very sharp and small peak around the value 1, with all values within the plotted threshold range. In comparison, the histogram from the simulated GSENSE detector image with noise has a much different characteristic. Here we see many bad pixels above the threshold range, which are significantly brighter than their surrounding pixels in Figure 1. The 1000-2000 pixels, circled in red in Figure 2(b), with a relative gain value of more than 5 already have a pixel response more than 5 times brighter than the surrounding pixels. Below the threshold range, there are a much smaller number of unresponsive pixels. The histogram for the AtmoSHINE instrument looks very similar to that of the GSENSE detector, despite the different imaging sensor and the influence of the radiation environment, which illustrates the detector independence of this method.

The two histograms with considered noise show a normal distribution around the value 1 and a standard deviation that depends on shot noise and the manufacturing variations. The threshold value is then a multiple (α) of the standard deviation (σ) and depends primarily on the sensor itself and its temperature range. The choice of α depends on the requirements for a non-bad pixel. We iterated for α a common value of 5.

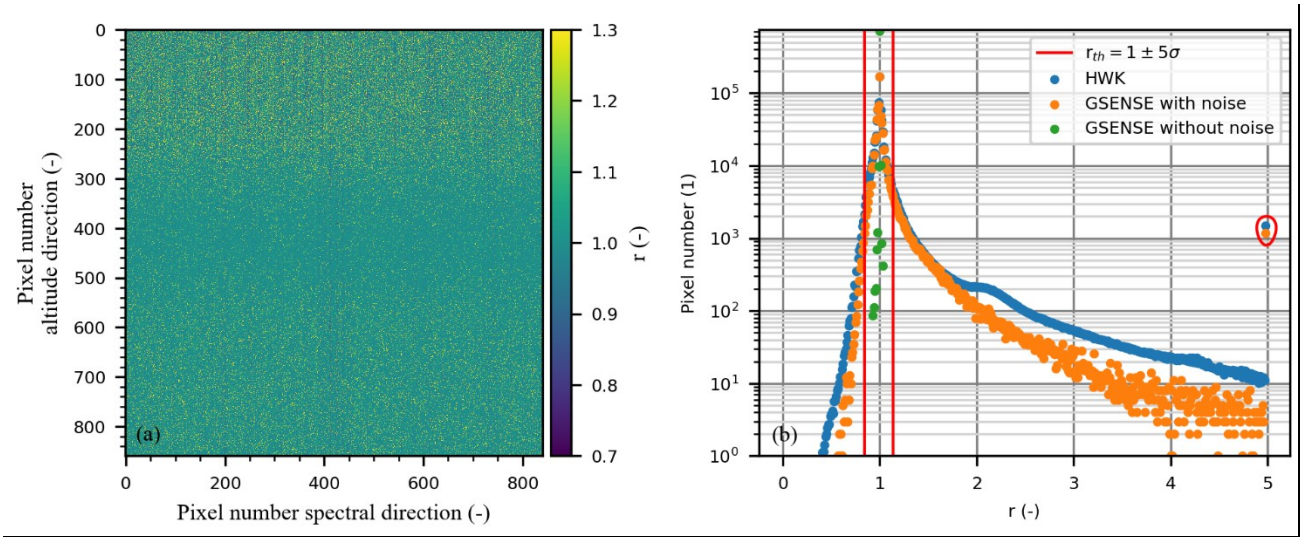

Figure 2 (a) relative gain matrix of Figure 1(b). (b) Histogram of all relative gains for real images of the HWK sensor from the AtmoSHINE mission and from a simulated image of the GSENSE sensor with and without noise.

For the detection of time-variant bad pixel mask, we analyze a time sequence of Q frames[8].

$$R(i,j) = \begin{cases} 1 \text{ if } \sum_{f=1}^{Q} N_f(i,j) \geq k_{r_{th}} \\ 0 \text{ else} \end{cases} \tag{3}$$

The matrix sum $\sum_{f=1}^{Q} N_f(i,j)$ indicates how often a pixel was detected at position i,j in a single frame bad pixel matrix. The histogram of $\sum_{f=1}^{Q} N_f(i,j)$ shows the distribution of how many pixels are flickered as bad pixels (Figure 3). The analysis involved 100 random images from the AtmoSHINE instrument. For a classical non-flicker approach of bad pixels, we would expect a distribution with a peak at 0% flicker and 100% flicker. However, we see a continuum of flickering pixels. Figure 3 indicates that only 1% of all pixels on all images were detected as bad pixels and 90% of pixels were never detected. The values between can be explained on the one hand by DC-RTS and on the other hand by pixels falsely detected as bad (false positive) in individual frames. This can occur, for example, through ionizing particles, which are visible in the image as bright individual dots, similar to bad pixels, and are therefore detected by the algorithm as bad pixels. For this reason, when determining the probability threshold $k_{r_{th}}$, above which we recognize flickering pixels as bad pixels, we must take into account how high the false positive rate of the algorithm is. For this we use a dynamic threshold in our approach which uses the knowledge from the previously analyzed histogram of relative gains and how many bad pixels we expect after Q frames[10]. We set the value $k_{r_{th}}$ to a value where the number of permanent bad pixels $\sum_{i,j} R(i,j)$ equals the average number of bad pixels per image expected from the histogram $\sum_{f,\mathrm{i,j}} N_f(i,j)\,/Q$. This probability threshold $k_{r_{th}}$ is also plotted in Figure 3.

The resulting final matrix of bad pixels $R(i,j)$ can then be used for consideration in further data processing. Since the number of bad pixels increases with TID, the determination of bad pixels should be repeated after a certain radiation dose or time period.

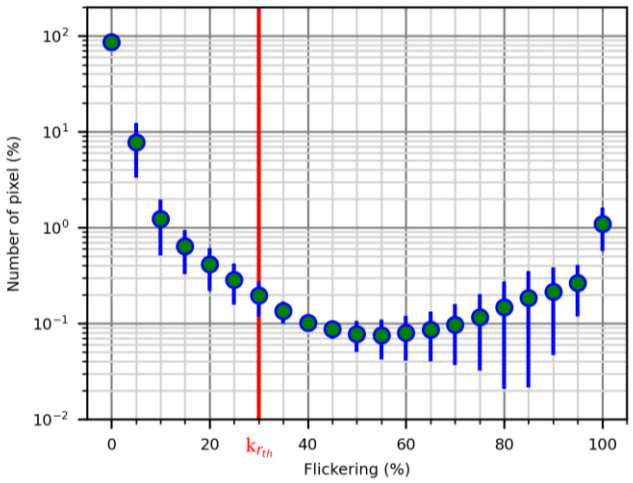

Figure 3 Histogram of flicker behavior of bad pixels. Number of pixels plotted against their frequency of being detected as bad pixels and the threshold $k_{r_{th}}$.

4. ADAPTIVE BINNING METHOD

The adaptive binning method aims to improve SNR as well as the final data product, a temperature profile for our instruments, while compressing data. The objective for AtmoLITE is a SNR > 40dB and a temperature uncertainty of less than 1K. The main sources of uncertainty, which have to be reduced, are the shot noise of the photo response, the shot noise of the dark current and the bad pixels. For the GSENSE detector, the maximum SNR specified by the vendor is between 33 and 49.5 dB depending on the gain level. There are many approaches to reduce these uncertainties. In this article we will focus on spatial binning of several pixels, as it additionally leads to the reduction of the amount of data, which enables us to measure more frequently. The values of a fixed number $m * n$ of connected pixels are averaged to form a so-called super pixel. This improves the SNR by the factor $\sqrt{m * n}$ and reduces the data by a factor of $m * n$. An extended form is the adaptive binning of pixels, illustrated in Figure 4, where different areas k of the image can be binned differently. This can respond to a dynamic lighting scene to produce a constant SNR. The compression rate depends on the light intensity of the scene and the desired improvement of the SNR.

Furthermore, with adaptive binning of pixels, bad pixels can also be masked, so that these strongly deviating values no longer effect the average value. As bad pixels are not actually noise, this does not improve the SNR, but it does improve the quality of the final data product. For adaptive binning of pixels we have implemented an algorithm in the on-board computer (programmable logic with a processing unit and a memory). A first implementation design was programmed with a high level synthesis language for the Field Programmable Gate Array (FPGA). In this implementation, we can define arbitrary areas in one direction where different levels of binning can be applied, as shown schematically with three areas in Figure 4. If bad pixels are masked within a superpixel, they are excluded from the averaging.

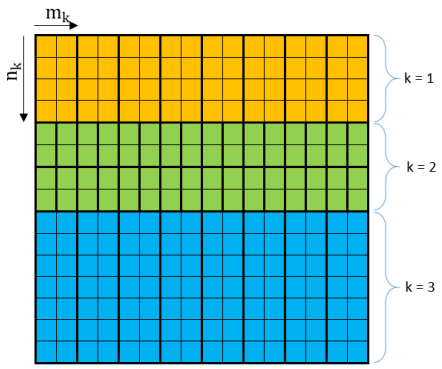

Figure 4: Schematic representation of adaptive binning. Each color marks an area that bins different numbers of pixels (thin lines) to superpixels (thick lines).

5. RESULTS

5.1 Bad pixel detection with dynamic threshold

One approach to improve data quality is to use a dynamic threshold for pixel detection that is adjusted over the entire operating time of the detector. Figure 5(a) shows the number of bad pixels detected for four days of AtmoSHINE. For each day, 10 runs of the algorithm with 100 random images were used, with recalculated thresholds r_{th} and $k_{r_{th}}$ for each run. The mean values and standard deviations for the detected bad pixels on the four days are shown in Figure 5(b). An individual threshold value r_{th} was determined for each run as described in chapter 3. In this study, we have assumed that the statistical values change only negligibly in one day. Over the period of the mission, the number of bad pixels seems to increase from (2.05±0.01)% to (4.1±0.1)%. Since the AtmoSHINE mission was a secondary payload on a satellite, the images were acquired on only a few days. Therefore, a more detailed discussion of Figure 6(b) is currently not possible.

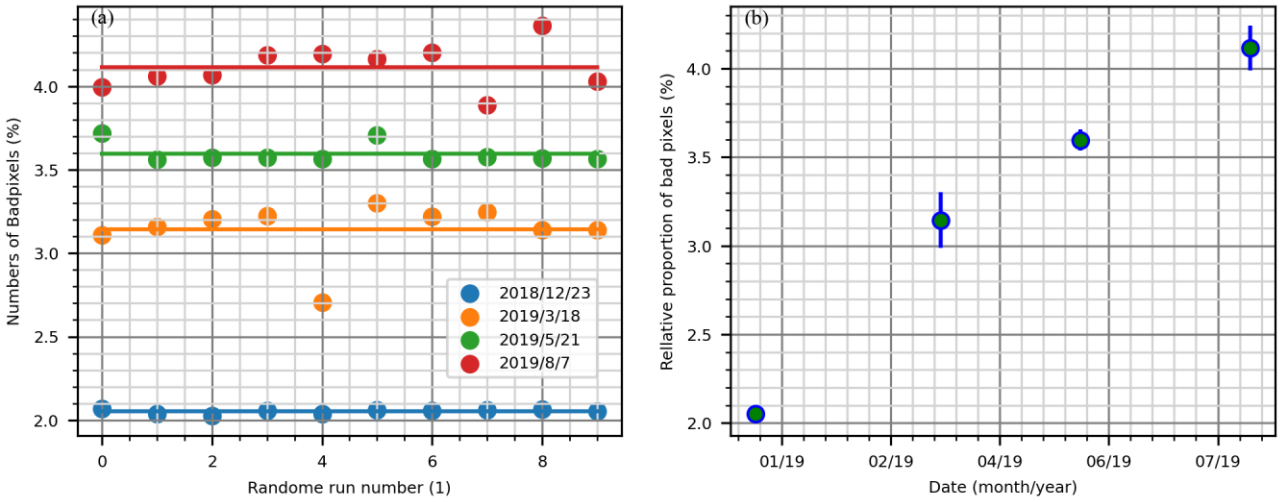

Figure 5 (a) Uncertainty in bad pixel detection for 10 different image sets from AtmoSHINE, for four days and (b) the mean and standard deviation of the detected bad pixels for each day.

5.2 Adaptive binning method

For a clear differentiation of the individual effects, we only considered simulated images of the AtmoLITE instrument for the improvement of the data quality, as can exemplarily be seen in Figure 1(b). For these simulated GSENSE detector images, no flickering pixels could be analyzed because the based dark current consisted of averaged measurements and included only the dark current non-uniformity of the measured detector, but not the DC-RTS. Therefore, the bad pixels were only analyzed for a single frame, or expressed as a threshold, $k_{r_{th}}$ was equal to 100%. This resulted in 7.8% bad pixels, which were considered in this chapter.

For error estimation, the fluctuations over 100 simulated images were determined and both shot noise and dark current were considered. In the analysis of the SNR, a constant binning of 1x20 pixels was performed once, as well as an adaptive binning, where the SNR was summarized for each altitude line of the image in Figure 6(a). For constant binning (blue), the SNR will increase by a factor of 4.5 due to a binning of 20 pixels. It can be seen that there is an increased SNR for the brighter middle rows of the image. The target SNR of 40dB is clearly exceeded for a wide altitude range. With adaptive binning (orange), a constant SNR of over 40dB was achieved over the entire altitude range. The points on the two lines in the image show the height value of an image line after binning. For constant binning, these are equidistant in contrast to adaptive binning, where significantly more data points could be resolved in the light-intensive area of the image. Both methods achieve the specifications for the SNR. In comparison, constant binning is less complex and more stable with a changed brightness distribution. Adaptive binning, on the other hand, bins a minimal number of pixels to achieve a specified SNR, which makes the data reduction dependent on the exposure scene and the targeted SNR. For our mission, constant binning would reduce the data rate by a factor of 20 and adaptive binning would reduce it by a factor of 11. This could also be increased to 20 by using a higher constant SNR.

The SNR is an important reference to transform the interferogram images[1] into a temperature profile with an uncertainty of less than 1K. Figure 6(b) shows the reverse calculated temperature (without regularization applied)for the respective altitude from the interferogram images after different binning methods. The red line indicates the initial temperature profile, which served as the basis for the simulated interferogram image and should come out as undisturbed as possible during the reverse calculation. We see that a standard 2x2 binning without considering bad pixels (blue dots) still gives strong temperature differences. The constant binning of 20 pixels without considering bad pixels (orange dots) gives a much better result. But here the temperature differences are often above 4K. If we extend the last binning with the masking of bad pixels (green dots), we achieve a temperature uncertainty of less than 1K over the whole altitude range. This shows that for accurate measurements and low uncertainties, consideration of bad pixels can be essential for achieving mission objectives. In particular, for pixel binning on-board a satellite, bad pixels can degrade data quality and cannot be removed in post-processing.

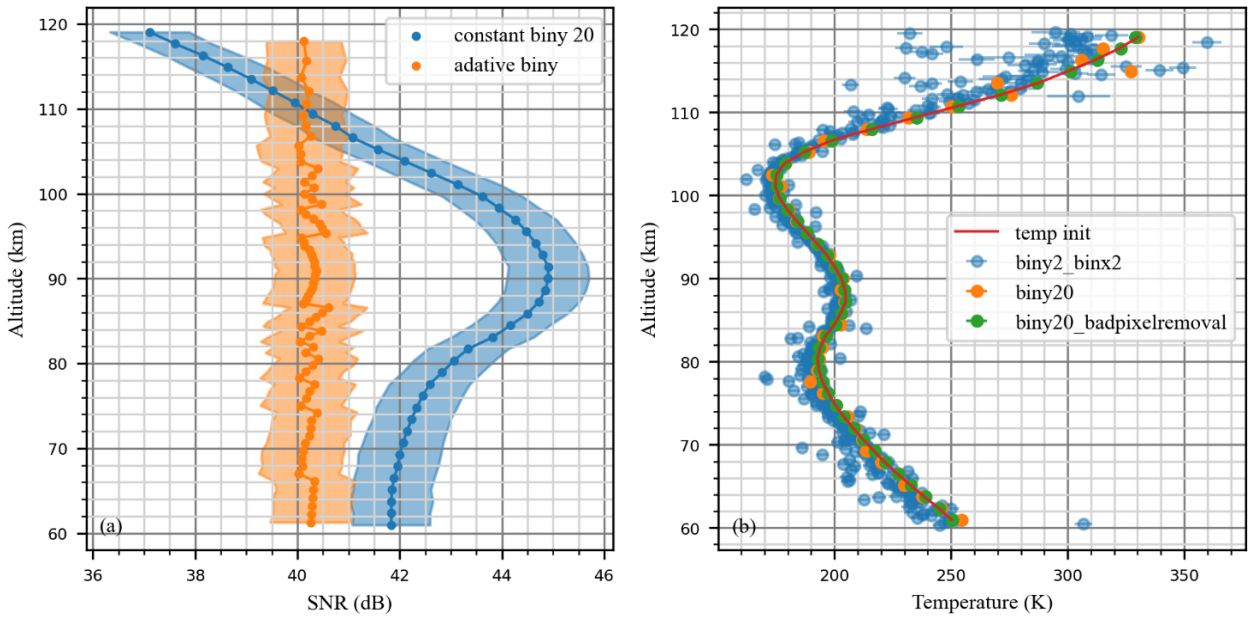

Figure 6 Analysis of a simulated GSENSE detector image from AtmoLITE including shot noise and dark current, similar to Figure 1(b). (a) Achieved SNR after uniform binning of the image (blue) and adaptive binning (orange) , including error range (semitransparent). (b) Resulting temperature profiles for different binning methods, as well as the initial temperature profile (red) for the respective altitude in the atmosphere.

5.3 Performance analysis

The introduced binning method, described in chapter 4, was implemented on a System-On-Chip (SoC) architecture, the ZYNQ7000 platform. We implemented the algorithm in software (C code) to be executed on the processing system and in VIVADO HLS (High Level Synthesis) to be placed in the fabric of the SoC. The Table 1 summarize the runtime tests of the PL (FPGA fabric) and PS (ARM CPU processing system) implementation in miliseconds. In both implementations we can see a performance increase with increasing X/Y Binning factors. Due to the image size of over 1 megapixel, we cannot hold the image inside the PL block memory. Therefore, the PL implementation was designed with an AXI Master Bus interface. We can see that the processing time of PL implementation for binning up to 8x1 is similar to the PS implementation without using the cache structure. The main bottleneck for the algorithm is the memory access from and to the main DDR based memory. Using the ARM PS cache structure speeds up the processing time by more than a factor of 10. The runtime difference between the 8x1 and 1x8 pixel binning is very low for all three implementations and the 8x1 binning always has a slightly lower value due to the line-by-line readout of the CIS. A large runtime difference between the 8x1 and 8x8 binning only results in the PL implementation, whereas the PS implementation shows almost no runtime gain.

We have also analyzed the FPGA resources used for the PL implementation without and with consideration of bad pixels and summarized them in

Table *2*. The use of 16-17% slices kept approximately constant, as well as the 6-7% DSPs. The largest increase was seen in the block RAMs, which can be explained by the additional storage of the bad pixel matrix. We finally decided to use the PL implementation because the resource consumption of the FPGA is acceptable, and offers a reliable implementation with triplicate logic.

Table 1: Runtime analysis for a PL and PS implementation, with different binning sizes.

X-Bin	Y-Bin	SoC PL runtime (FPGA Fabric) (ms)	SoC PS runtime (Arm CPU) (ms)	
			Without Cache	With Cache
1	1	22000	44700	3990
2	2	11060	12500	1070
1	8	3250	7940	620
8	1	2800	7110	580
8	8	590	6050	340

Table 2: Comparison of FPGA resources used for implementation with and without bad pixel consideration.

Implementation	Slices (13,300)	Block RAM (140)	DSPs (220)
Without bad pixel consideration	2145	11	15
With bad pixel consideration	2262	15	14

6. CONCLUSION

We obtained a detailed insight into the operation of the adapted detection algorithm for bad pixels ISMFD and discussed it on the basis of images from the sCMOS imager HWK on the AtmoSHINE mission. We measured relatively strongly deviating pixels compared to the median of their neighbors, as well as temporally flickering bad pixels and analyzed the detection uncertainty for different data from the AtmoSHINE mission.

The fact that bad pixels had to be considered was demonstrated by a planned on-board binning for the future AtmoLITE mission using simulated images from the sCMOS imager GSENSE. Thereby, adaptive binning was used to achieve a constant SNR in a non-uniform light scene, which allows stronger binning in darker areas and higher resolution in brighter areas. The binning method showed the need to consider bad pixels in order to achieve scientific mission goals, such as those of AtmoLITE. In contrast, vendor binning methods without considering bad pixels proved insufficient.

The increase of bad pixels is a critical reason why continuous monitoring of bad pixels is necessary. Thus, during the 10 months of the AtmoSHINE mission, the number of bad pixels doubled from the original (2.05±0.01)% to (4.1±0.1)%. A first implementation of the on-board consideration of bad pixels during binning shows its feasibility and is currently being optimized for use with respect to FPGA resources. In the future, on-board data compression can take place, while a continuously adapted consideration of bad pixels further improves the data quality.

REFERENCES

[1] Kaufmann, M., Olschewski, F., Mantel, K., Solheim, B., Shepherd, G., Deiml, M., Liu, J., Song, R., Chen, Q., Wroblowski, O., Wei, D., Zhu, Y., Wagner, F., Loosen, F., Froehlich, D., Neubert, T., Rongen, H., Knieling, P., Toumpas, P., et al., "A highly miniaturized satellite payload based on a spatial heterodyne spectrometer for atmospheric temperature measurements in the mesosphere and lower thermosphere," Atmos. Meas. Tech. 11(7), 3861–3870 (2018).

[2] Neubert, T., Rongen, H., Froehlich, D., Schardt, G., Dick, M., Nysten, T., Zimmermann, E., Kaufmann, M., Olschewski, F. and van Waasen, S., "System-on-module-based long-life electronics for remote sensing imaging with CubeSats in low-earth-orbits," J. Appl. Rem. Sens. 13(03), 1–17 (2019).

[3] Goiffon, V., "Radiation Effects on CMOS Active Pixel Image Sensors," 295–332 (2015).

[4] Isoz, W., Svensson, T. and Renhorn, I., “Nonuniformity correction of infrared focal plane arrays,” Infrared Technology and Applications XXXI 5783, B. F. Andresen and G. F. Fulop, Eds., 949–960, SPIE, Orlando, Florida, USA (2005).
[5] Chapman, G. H., Thomas, R., Koren, Z. and Koren, I., “Empirical formula for rates of hot pixel defects based on pixel size, sensor area, and ISO,” 19 February 2013, Burlingame, California, USA, 86590C.
[6] Chen, Q., “Retrieval of atmospheric quantities from remote sensing measurements of nightglow emissions in the MLT region,” Dr. (2020).
[7] Harlander, J., Reynolds, R. J. and Roesler, F. L., “Spatial heterodyne spectroscopy for the exploration of diffuse interstellar emission lines at far-ultraviolet wavelengths,” ApJ 396, 730 (1992).
[8] Yang Cao, Weiqi Jin, Chongliang Liu, and Xiu Liu., “Scene-based bad pixel dynamic correction and evaluation for IRFPA device,” Advances in Optoelectronics and Micro/nano-optics, 1–4, IEEE, Guangzhou, China (2010).
[9] Cao, Y., Tang, F., Bermak, A. and Le, T., “A Smart CMOS Image Sensor with On-chip Hot Pixel Correcting Readout Circuit for Biomedical Applications,” 2010 Fifth IEEE International Symposium on Electronic Design, Test & Applications, 103–107, IEEE, Ho Chi Minh City, Vietnam (2010).
[10] Florczak, J., Neubert, T., El Maghawry, K., Zimmermann, E., Rongen, H., Kaufmann, M., Olschewski, F. and Van Waasen, S., “Radiation Monitor Extension for CMOS Imaging Instruments in Nanosatellites,” IEEE Trans. Nucl. Sci., 1–1 (2022).