



Scene-based non-uniformity correction: From algorithm to implementation on a smart camera



T. Toczek, F. Hamdi, B. Heyrman, J. Dubois, J. Miteran, D. Gin hac *

LE2I UMR 6306, Univ Burgundy, Dijon, France

ARTICLE INFO

Article history:

Available online 6 June 2013

Keywords:

Fixed spatial noise
Non-uniformity correction
FPGA-based smart camera
Real-time implementation

ABSTRACT

Raw output data from image sensors tends to exhibit a form of bias due to slight on-die variations between photodetectors, as well as between amplifiers. The resulting bias, called fixed pattern noise (FPN), is often corrected by subtracting its value, estimated through calibration, from the sensor's raw signal. This paper introduces an on-line scene-based technique for an improved fixed-pattern noise compensation which does not rely on calibration, and hence is more robust to the dynamic changes in the FPN which may occur slowly over time. This article first gives a quick summary of existing FPN correction methods and explains how our approach relates to them. Three different pipeline architectures for real-time implementation on a FPGA-based smart camera are then discussed. For each of them, FPGA implementations details, performance and hardware costs are provided. Experimental results on a set of seven different scenes are also depicted showing that the proposed correction chain induces little additional resource use while guarantying high quality images on a wide variety of scenes.

© 2013 Elsevier B.V. All rights reserved.

1. Introduction

CMOS image sensors are built around a grid of active pixels, each pixel containing a photodetector and an associated amplifier. The rest of the area on a CMOS sensor is comprised of readout circuitry, column amplifiers, analog to digital converters. Each element of the image sensor should ideally be identical to any other from the grid. Whenever an element differs slightly from the norm, it is likely to produce an incorrect response signal. As a result, the image sensors suffer from noise sources giving artifacts of the output data. Since the design of the first CMOS sensor, numerous works have been conducted to deeply characterize noise sources [1–4] because they limit image sensor performance, especially under low illumination and/or long exposure conditions [5]. The noise sources can be divided into two distinct families *Temporal Noise* and *Fixed Pattern Noise* [6]. Temporal noise is independent across pixels and varies from frame to frame. Sources of temporal noise include photodetector shot noise, pixel reset circuit noise, readout noise for examples. On the other hand, image sensors also suffer from Fixed Pattern Noise (FPN). FPN is the pixel-to-pixel output variation under uniform illumination due to device and interconnect mismatches across the array [7]. It includes offset FPN and gain FPN. First of all, offset FPN can be defined as a dark signal non-uniformity, which is the offset from the average dark current across the imaging array in the absence of light [8]. It is

independent of pixel signal and fixed from frame to frame for a given sensor. Offset FPN can be pixel and column noise [6,9]. Pixel FPN is due to variations in the photodetector and transistors geometry whereas column FPN is mainly caused by mismatches in the integrating amplifiers and the analog-to-digital converters common to all the photodetectors of a given column [10]. Since such defects affect all the pixels of a column, column FPN appears as stripes in the image leading to significant image quality degradation. Secondly, gain FPN is defined as the spatial variation in pixel output values under uniform illumination due to pixel-to-pixel mismatches [8]. Gain FPN is difficult to correct and for years, it was one of the major CMOS imager's disadvantages [11] because it increases with illumination and exponentially increases with temperature.

For achieving high-quality CMOS imagers, it is essential to develop good FPN compensation based on accurate models of sensors. While in the general case the FPN may be the result of any alteration of the photodetector transfer function, it can be modeled in a quite simple way. The pioneering model described in [6] is based on a detailed characterization of the silicon structure of the pixel. It focuses on offset noise and presents a detailed method for estimating the pixel and column FPN components. The same research team reports also a complementary study on gain FPN [12]. Most of the recent correction methods [10,13] consider the FPN to be characterized at each photodetector by either a gain and an offset or just an offset because it is the dominant source of FPN in most pixels [14]. Indeed, this offset (i.e. the photodetector dark current) is the most easily seen manifestation of the FPN, consider-

* Corresponding author. Tel.: +33 380393886.

E-mail address: dginhac@le2i.cnrs.fr (D. Gin hac).

ably affecting the sensor output under low luminosity and/or long exposure conditions. Gain variations, or more complex nonlinear behaviors of the photodetector, are seldom seen in practice. Lots of works have been done to integrate electronic circuitry in the sensors to significantly reduce offset FPN. The most commonly used technique is Correlated Double Sampling (CDS) that can be found in many analogue circuits. The pixel signal is sampled twice, once immediately after reset and the second time after integrating the signal. These two values are then read differentially so that offset FPN is eliminated [15]. Various other techniques have been proposed to reduce FPN by either employing pinned photodiodes [16,17] or modifying the pixel circuit itself [18,19], or optimizing the column readout circuit [20], or subtracting the response of each pixel from its response to a uniform input current generated either optically [21] or electronically [22,23]. However, despite the integration of numerous FPN-reduction circuits, CMOS image sensors still have FPN noise problems because of mismatched pixels and associating readout circuits, especially in extreme light conditions. As a result of the complex nature of sensor non-uniformities, FPN correction, also called non-uniformity correction (NUC), by means of efficient post-processing algorithms is a necessary and unavoidable task to be performed in order to achieve higher-quality images, or image sequences.

So, in this paper, we will concentrate on FPN reduction algorithms and hardware implementation of these algorithms on FPGA-based smart camera architectures. Since FPN varies slowly over time, depending mainly on the operating conditions (e.g. temperature) of the image sensor, this poses additional challenges to achieve a high quality image. Based on a state-of-the-art study of existing FPN correction methods, we present a new algorithm and different real-time hardware implementations for improved FPN compensation. This algorithm does not require any sensor calibration phase and is more robust to any dynamic changes in the FPN which may occur slowly over time. Our algorithm is naturally candidate for a real-time implementation on any FPGA-based smart camera architecture such those described in [24–27].

The rest of the paper is organized as follows. Section 2 gives a quick summary of existing FPN correction methods, and explains how our approach relates to them. Section 3 deals with our improved algorithm based on the approach developed by Harris and Chiang [28]. Three different hardware implementations on a FPGA-based platform are then discussed in Section 4. Finally, implementation details, performance, and hardware costs are described in the last section of this paper. Experimental results on a set of seven different scenes are also depicted. Comparisons in terms of image quality between our implementations and original Harris and Chiang's algorithm are also provided.

2. FPN correction algorithms

In order to significantly decrease the FPN, non-uniformity correction must be applied to images acquired by the sensor. Depending on the accuracy requirements of the candidate applications, several post-processing methods can be applied to correct non-uniformities. These techniques include basic methods such as calibration-based methods and more complex techniques such as scene-based methods [29].

2.1. Calibration-based methods

As already said, FPN is mainly due to inhomogeneity in the manufacturing of the sensor. It is also known that the resulting non-uniformity does not widely change during the service life of the sensor. So, the most common way to deal with FPN is through a calibration step at the factory in order to precisely evaluate a ref-

erence-based correction. The most basic technique is to measure the sensor output when no light reaches the photodetectors to obtain a good approximation of it [30]. Then, the resulting “dark frame” can be stored in a ROM and subtracted at each frame from the actual sensor output. While this method is easy and inexpensive, this one-time calibration has the drawback to ignore FPN variations induced by the sensor operating conditions. It is also known that FPN is not totally stationary but instead drifts slowly in time, making a one-time calibration ineffective. Of course, calibration can be done frequently to take into account the FPN variations but this unfortunately requires halting normal imaging operation during each calibration process. One way to compensate efficiently the FPN drift is to perform more than one calibration (generally two), each of them at a different temperature. Such a technique is mainly used for infrared focal-plane sensors [31] but have been also applied to more conventional image sensors [32]. The corrected image is then the one whose calibration temperature is the closest to the operating temperature, or possibly an interpolation of images whose calibration temperatures are around the operating temperature. In either case, at least twice as much ROM is required to store the calibration images, and a means to estimate the sensor operating temperature is required.

A more cost-effective solution in terms of memory is described in [33]. This better strategy for FPN estimation consists in analyzing supplementary data provided by the sensor. At the top of the sensor, a series of black and dark extra lines being shielded from light is positioned. Specifically, black lines have zero integration time whereas dark lines are shielded from incident light while having the same exposure as the image. Consequently, the data from these lines can be used to estimate a column FPN signature for the whole image. This signature is continuously averaged and updated according to the operating conditions by using the dynamic information provided by the black and dark lines and subtracted from the image data. This method only provides FPN estimation per column and is not able to compensate the variation between pixels. Additionally, despite interesting results described in [34], such a method is unfortunately not available for designers or end-users of smart cameras because it requires both the VLSI design of extra-lines of pixels in the core of the sensor and the possibility to access and read data from these extra-lines.

2.2. Dynamic FPN estimation by scene-based methods

This kind of techniques is called “scene-based” non-uniformity correction methods because they do not require specific scenes or initial conditions for the purpose of calibration. FPN is dynamically estimated and updated from the real-time analysis of the scenes captured by the sensor. Therefore, scene-based techniques are able to permanently evaluate FPN by exploiting motion-related features in image sequences. They bring the advantage to fit the drifting operating conditions without the need for any physical measurement, by adaptively using the information provided by the sensor. They can traditionally yield better results than calibration-based techniques, especially when the operating conditions of the sensor cannot be correctly approximated by a simple combination of the calibration measurements. It is obvious that such techniques increase system reliability but have higher computational complexity [35] requiring significant hardware resources to be processed at the sensor frame rate.

Scene-based non-uniformity correction (NUC) methods include two distinct categories of post-processing: (1) algebraic techniques, and (2) statistical techniques. As defined in [36], the algebraic techniques make use of global motion between the frames in the video sequence and attempt to compensate for the non-uniformity by means of algebraic methods without making statistical assumptions about the FPN [37–39]. Such approaches aim at ana-

lyzing the values of nearby pixels in order to separate the scene from the noise. Ratliff et al., in particular, have proposed an algebraical method [40] based on motion vector estimation. Such methods offer good robustness but are computationally heavy and not really suitable for real-time on-line correction.

On the other hand, statistical techniques model the FPN as a random spatial noise and estimate the statistics of the noise to remove it [41–44]. Compared with registration-based methods, statistical approaches have been more widely studied because of their relatively lower computational complexity, smaller storage demands, and better realtime performance [44]. However, they rely on hypotheses regarding the sensor output signal (and therefore the scene) in order to separate the FPN from the actual signal. It is important for those hypotheses to be reasonable enough to be satisfied by most of the scenes acquired by the sensor. Nevertheless, when these assumptions are violated, performance can become poor and ghosting artifacts are easily produced. As an example, the most well-known statistical method developed by Harris and Chiang [28] supposes the average value and standard deviation of the photodetector output is the same among all the photodetectors of the sensor.

3. Improved constant statistics method

We propose to extend the approach developed by Harris and Chiang [28] relying on the global constant statistic assumption which states that the statistics of the observed scene become constant over time. Our objective is to improve the method making it more flexible regarding the input scene, and then more efficient in any unknown visual scene. The original paper models the non-linear behavior of each photodetector as an affine function. Each photodetector k is associated with a pair of scalars a_k and b_k such that the actual luminance \tilde{I}_k received by the photodetector can be obtained from the measured one I_k through the relation:

$$I_k = a_k \tilde{I}_k + b_k \quad (1)$$

Harris and Chiang assume that the measured luminance over time has the same mean and the same variance among all the photodetectors of the sensor. Under these conditions, it can be shown that for each photodetector k , assuming m_k is its average measured luminance and σ_k the measured luminance variance, the following holds:

$$\tilde{I}_k = \frac{I_k - m_k}{\sigma_k} \quad \text{and} \quad \tilde{I}_k = A \tilde{I}_k + B \quad (2)$$

There are several easy ways to iteratively estimate m_k and σ_k – for instance, an exponential window can be used – so \tilde{I}_k can be actually and cheaply computed. The \tilde{I}_k signal mean is zero and its variance is unitary, but it is otherwise proportional to the actual luminance value plus an offset. \tilde{I}_k can be obtained by knowing the factors A and B , which are the values of the gain and the offset respectively, estimated by using the least squares method so as to minimize the error between the measured and corrected pixel values. Also, whenever the exact values of A and B do not matter, simply normalizing (\tilde{I}_k) over the dynamic range is enough (this is for example the case in IR imaging). This method has the advantage of being very simple, but it has several drawbacks in practice. The main problem is that the hypotheses regarding the mean and variance of the input signal are not satisfied in a lot of cases. Harris and Chiang recommend to move the camera or otherwise move the objects of the scene during calibration, but this implies the approach cannot be used to continuously calibrate the sensor. What's more, even when the camera is moving and a dynamic scene is being shot, the mean and variance hypotheses may not hold. For instance, as the authors of the original paper admit, most

of the time the output image is brighter at the top than at the bottom due to the fact that most light sources, both natural and artificial, illuminate the scene from above. We have observed that making accurate hypotheses regarding the FPN is much easier than regarding the scene. So, we propose to adapt this approach so as to make it work whenever the following is true for each photodetector k :

1. FPN has only an offset component (no gain component);
2. the mean value of the FPN offset is locally zero;
3. the high spatial frequency signal component at the corresponding pixel has an average value of zero over time.

Concerning the validity of the first hypothesis, we can notice that FPN is almost only seen on low luminosity and/or long exposure shots. If a significant gain component exists, its influence should be noticeable regardless of luminosity or exposure. Since it is not the case for almost any visible light sensor we studied, it means that either no FPN gain component exist at the photodetector level, or that current compensation mechanisms for this component work just fine [10,45], unlike those aiming to correct the offset component. The second hypothesis is more or less justified by assuming that manufacturing defects occur randomly. It can be partially confirmed by observing the output of the OmniVision OV9715 sensor embedded in the smart camera platform used in this study. Image capturing was performed while preventing any light from reaching the photodetectors as shown on Fig. 1(a). A zoom on a 64×64 -pixel sub-window illustrates the random distribution of noise. The absolute magnitude of its Discrete Fourier Transform is depicted on Fig. 1(b) and reveals that the column FPN contribution appears clearly as an horizontal line at $y = 0$. It can be seen that the high frequencies indeed dominate the image, but some low frequencies remain, however, and it appears that column FPN is the main source of them. The last constraint is the only one concerning the scene. Since the FPN is mostly a high frequency component, it is enough to average over time the high spatial frequencies at each pixel to determine the noise value, should this constraint be satisfied. In practice, it is often naturally satisfied, for instance whenever the camera is panning, or whenever an object traverses the field of view. While it is still possible to find pathological cases (for instance, a motionless camera shooting a completely static scene) it is much weaker and more naturally occurring than the constant mean/constant variance hypothesis used previously.

The general principle of our method is to remove the low spatial frequencies from the current frame, to apply the Harris and Chiang method on the remaining high spatial frequencies, and to add back low frequencies to obtain the final image. When applying the Harris and Chiang method, we will assume that $a_k = 1$ for each k . Benefits of our method are various: (1) range of scenes correctly corrected is much wider and, (2) the post-processing dynamic range adjustment step is no longer necessary. Indeed, for each photodetector k , let us separate the high frequency component I_k^H from low frequency one I_k^L . Using the previous notations, we will have:

$$I_k^H + I_k^L = \tilde{I}_k + b_k \quad (3)$$

Assuming $a_k = 1$ (constraint 1), using the same reasoning as in the original article but on I_k^H instead of I_k , we get:

$$\tilde{I}_k^H = I_k^H - m_k^H \quad (4)$$

where m_k^H is the average over time of the high spatial frequency component of the values of the photodetector k . The I_k^L component can be considered as noise free according to the constraint 2. Finally, let us notice that while (\tilde{I}_k^H) is zero-measured temporally just as in the original method, so is (I_k^H) , the high frequency component of the

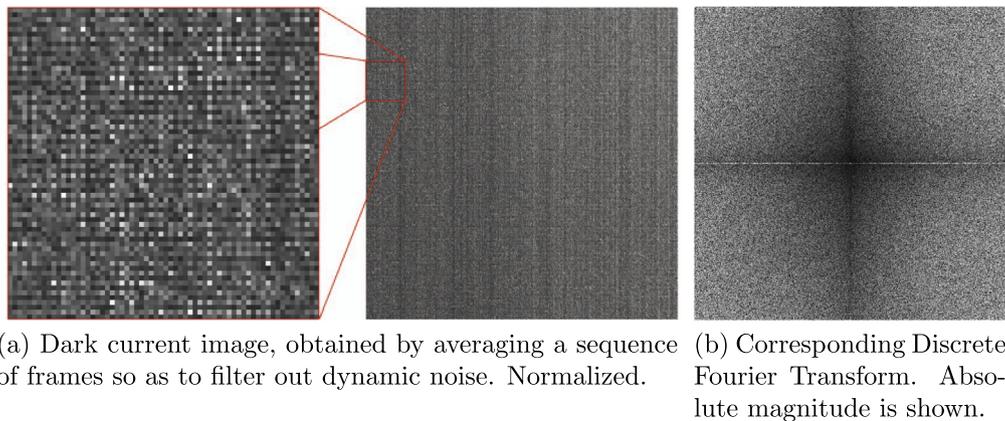


Fig. 1. Dark current FPN of an OmniVision OV9715 sensor and corresponding DFT.

actual luminance the detector would measure were it perfect. This is a consequence of the constraint 3. So, we get:

$$\tilde{I}_k = \tilde{i}_k^H + I_k^L \quad (5)$$

This can be computed as we go, confirming that no further post-processing is needed. Low-pass filtering aside, we will just need to update the (m_k^H) estimate at each frame. The cheapest method to do this is to use an exponential window [46]. We will therefore take a scalar α , such as, if $\hat{m}_k^H(n)$ is the m_k^H estimate at the n -th frame:

$$\hat{m}_k^H(n) = \alpha I_k^H + (1 - \alpha)\hat{m}_k^H(n - 1) \quad (6)$$

Alternative methods include using a sliding rectangular window, which requires to store the history of the last I_k^H values, or accounting for all the samples, which would require a division. Those methods do not yield significantly better results than the exponential window based one.

4. Generic architecture for hardware implementation in a FPGA-based smart camera

Any implementation requires to perform two main tasks: first, extracting the low frequencies of the current frame through filtering, and then correcting the individual pixel values. The first part is a well understood problem, and the second is embarrassingly parallel. What is more, both parts rely more on arithmetical performance than on bit mangling or fine grained flow control. Finally, for low values of α , using floating point arithmetics may seem to be a good idea. As a result, processor-based machines with SIMD capabilities (such as PC equipped with programmable GPUs) are quite good candidates for real-time implementations but are not easily embeddable on a smart camera. The choice of an FPGA as the implementation target should be privileged whenever the consumption is a significant problem, the FPN correction logic is to be shipped with the sensor, and/or the latency of the sensor data induced by the correction must be kept minimal.

4.1. Correction pipeline architecture

According to the principles exposed so far, the correction chain can be implemented by a simple pipeline consisting of a low-pass filter followed by a correction block working at the pixel level. The corresponding architecture is shown on Fig. 2. The high-frequency data is obtained by simply subtracting the low frequencies from the raw signal. Those low frequencies are later added back to the corrected high-frequency data. The low-pass filter and the correc-

tion block have an average throughput of one datum per cycle. The correction block has a constant latency of several cycles. Consequently, using an appropriate delay is required to synchronize the low frequency component with the high frequency one before adding them back together. Depending on its exact implementation and parameters, the low-pass filter has a latency ranging from a few cycles to a few thousands cycles (see Section 4.3 for details). In any case, its latency is much shorter than a frame duration. Since the pixel correction block needs to store the temporal average value estimate m_k^H for each pixel k , it needs an access to the external memory if the input resolution is too high. Assuming a resolution of 1280×720 and using 16 bits per estimator, 1.84 MB of memory are required, i.e. more than most FPGAs can provide internally. It is possible to prefetch the estimator values and store them back once updated in a transparent fashion, using a small on-chip scratch pad, without affecting the overall correction block latency or throughput.

While this first version gives satisfying results, it is interesting to add a block to specifically tackle column FPN correction. The column-level correction is performed similarly as the pixel-level correction. In the remainder of this paper, we will call α' the equivalent of the α coefficient for column average estimator computation. While the pixel average estimators is updated once a frame, the column estimator $m_j^{H'}$ is successively subtracted from every pixel of a given column j , and therefore successively updated by taking into account the values of each pixel of that column. Consequently, column average estimators converge more quickly since they are updated more frequently than pixel average estimators. Convergence time issues are discussed in the Harris and Chiang's original article but also in [46] in the particular case of exponential windows. Moreover, column estimators are more robust than their pixel counterparts, and column correction is less likely to produce artifacts whenever any of our three working conditions is not met. In particular, they tend to give acceptable results even for highly static scenes.

Fig. 3 shows the most natural way to integrate column-level correction the previous pipeline architecture. One should note that no SDRAM access is required this time. Indeed, for a resolution of 1280×720 and using 20 bits per estimator (see Section 5.1 for justification), all of the $(m_j^{H'})$ take about 3.2 KB. In Section 3, we have seen that the column noise is actually the noise component with the most low frequencies (cf. Fig. 1(b)), and therefore the one most likely not to satisfy our second working constraint. Given this, we propose another variant (shown on Fig. 4) of the previous architecture. In this variant, the column correction occurs before the low-pass filter, directly on the sensor data (and not on their high frequency component). Since the output of the column correc-

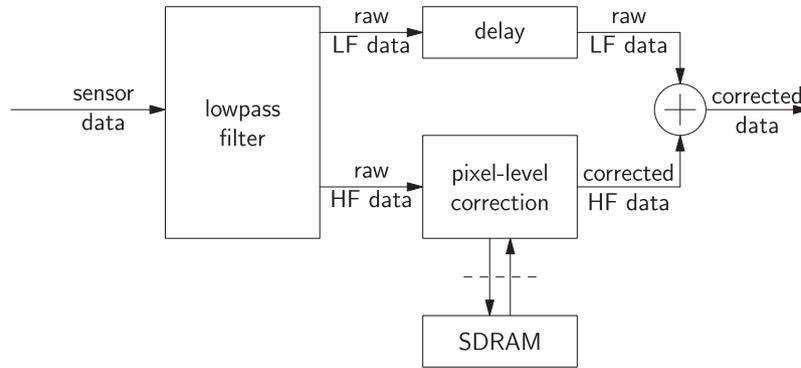


Fig. 2. Correction pipeline, variant A (pixel level correction only).

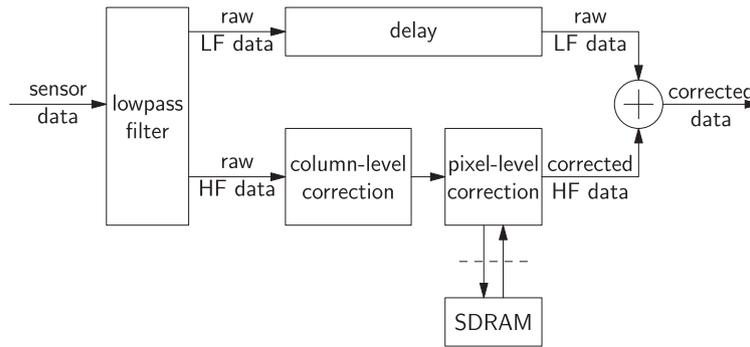


Fig. 3. Correction pipeline, variant B (column and pixel level correction).

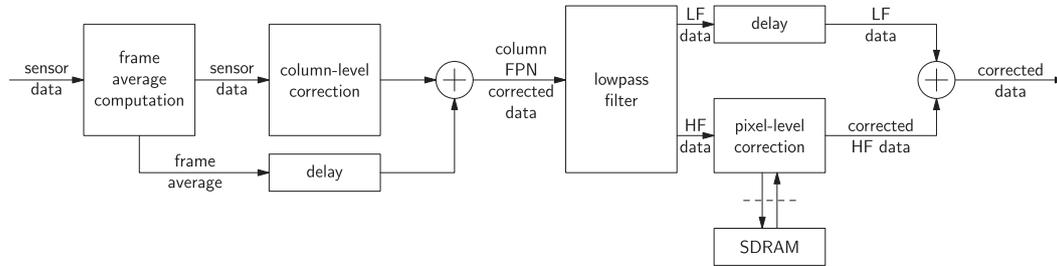


Fig. 4. Correction pipeline, variant C (column level correction before filtering, standard pixel level correction).

tion has a mean value of zero, it helps to satisfy the second working constraint, requiring a zero FPN offset. Moreover, such an approach makes also the pipeline C globally more robust regarding convergence problems because it helps the subsequent pixel-level corrector to quickly and efficiently reduce noise. This is akin to using the original approach from Harris and Chiang on whole columns instead of pixels, before using our improvements for noise correction at the pixel level. However, since the output of the column corrector is by principle zero-measured over time, the only way to obtain the correct output intensity is to compute the global spatial average pixel value of the frame before the column correction step, and to add it back behind the column correction block (for instance, directly after it). In order to compute this average, the values of all the pixels of the frame must be known. This implies a latency greater than a whole frame duration. This aspect can be heuristically mitigated by using the global average of the previous frame instead of the current one in the computations. So, there is no extra latency induced, but the corrected output will appear less precise whenever the scene luminosity varies too quickly between two successive frames.

4.2. Correction block

So far, we have not discussed the way the video data bus is organized. For the rest of this paper, we will assume it is composed of the video data itself (typically an 8-to-10 bit signal), and three one bit control signals: the active video signal indicating the video data is valid, and horizontal and vertical synchronization signals marking the beginning of a scanline or of a frame, respectively.

The overall principle of the column correction block itself is shown on Fig. 5. The main issue when implementing the correction blocks is the storage of the temporal averages for each pixel or each column. For column correction, when aiming for a throughput of one datum per cycle, it must be possible to read a coefficient and write another one during the same cycle. To avoid the use of two port on-chip RAM, we can rely on two banks: one for even-numbered columns and one for odd-numbered columns. Assuming that the correction block is fed a valid pixel value every cycle, the column whose m_j^{th} are read and written will have its index j alternating between even and odd. Since there is a FIFO before the writing port of the temporal average memory interface, it is possible to

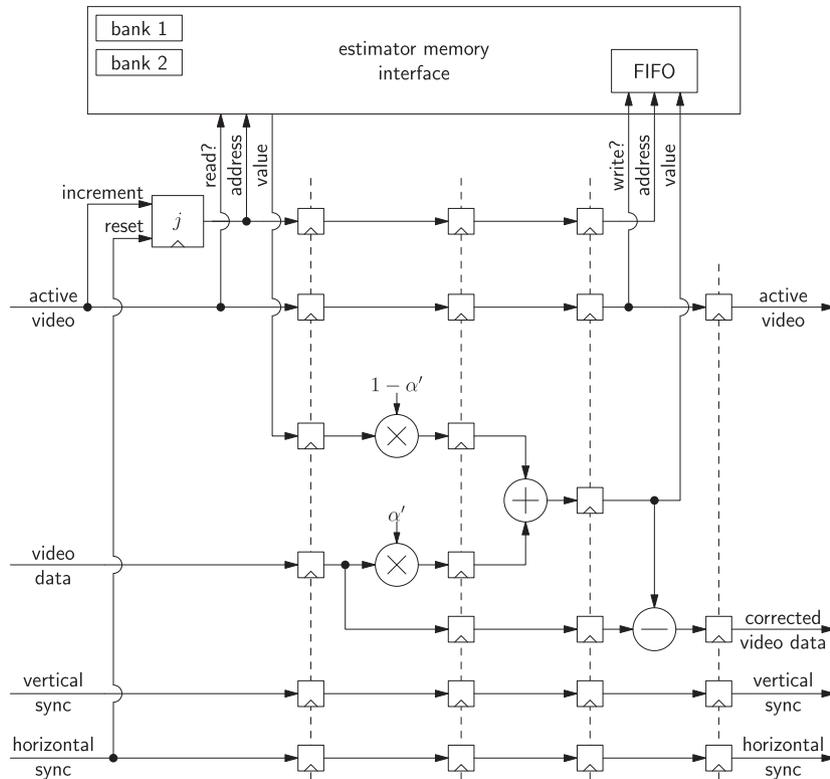


Fig. 5. Column correction block principle schematic.

perform those read and write operations simultaneously by alternating between the two banks, no matter whether the number of stages of the correction block pipeline is even or odd (Fig. 5 shows a four stage pipeline, but it can be adapted easily without loss of generality). The FIFO will contain exactly one or zero element. This works as well when there are bubbles in the pipeline. The policy is simply not to write (by leaving the value to be written in the FIFO until the next cycle) whenever a bank conflict arises. For column correction, a FIFO size of exactly one element is enough to avoid overruns.

The pixel correction block (Fig. 6) obeys essentially to the same principles, but in this case, the memory interface serves as a cache controller. Since the (m_k^H) are read and written in ascending order and that none of them is ever skipped, it is easy to prefetch them from the SDRAM, as well as store them back, in a timely manner, as long as one knows the worst case SDRAM access latency. However, to keep the throughput at one value per cycle without using double port on-chip memories, one will have to use four or more memory banks instead of two. The banks are associated by pairs of even and odd-numbered pixels (just like for columns), and at least two such bank pairs have to be used: one for the read and write operations from the correction logic, and another for storing back previously written values in the SDRAM as well as prefetching the ones to be read next. As previously, the FIFO located just before the writing port of the memory interface should have a size of at least one element for everything to work. It can nevertheless be larger, for instance when the correction block is used within a system-on-chip which contains other IPs using the SDRAM, or sharing the bus which leads to the SDRAM controller. In such systems, it may be difficult to estimate the worst-case SDRAM access latency. In such cases, it is possible to assume reasonable access times to external memory, and rely on the FIFO as a temporary store whenever the assumed deadline is not met.

4.3. Low-pass filtering

The low and high frequency components of frame data are obtained through low-pass filtering. The high frequencies are simply the difference between the original image and the low-pass filter output. There are several ways to filter an image so as to cut its high frequencies. Since we want to do noise correction on the high frequencies, we ideally have to cut all the noise frequencies while cutting as little other frequencies as possible. Since the raw signal with low frequencies is not processed and just added back to the corrected data, it can be tempting to chose a very small cut-off frequency in order to minimize the uncorrected noise in low frequencies. This can even lead to chose a filter that cuts everything but the fundamental. In such a case, the result will be similar to the simplified (no noise gain correction) implementation of the Harris and Chiang method. Despite some severe limitations (in particular in presence of high amplitude low frequency components; see for instance [40]), it is supposed to remove the FPN as expected.

So, since it is difficult to predict ahead of time the amount of FPN a sensor will generate, the prudent course of action is to use a large filtering window. The exact kernel used matters little, as long as the filter is a low-pass. It should therefore be selected for its implementation simplicity.

An obvious solution is to use a rectangle window with a width and height that are powers of two, and obtain the result by summing all the pixel values of the window before right shifting them so as to obtain the spatial average. This has the advantage to be implementable without any multiplier or divider. Let N and M be integers such as the window size is 2^N lines by 2^M columns, and let w and h be powers of two greater or equal to the frame width and height, respectively. The most complicated part is the on-chip RAM organization. We will need three memories:

- *C* needs three ports but can be implemented with four banks of single port memory; the read addresses are j and $j - M$, and the written address is j . Let us use the j parity trick for the read and write at the address j , and use another pair of banks when the M th bit, counting from the least significant to the most significant, changes. For instance, if $M = 2$ (i.e. when the window width is 4), the banks used for consecutive j would be 1, 2, 1, 2, 3, 4, 3, 4, 1, 2, 1, 2, etc. This works if the reads at j and at $j - M$ are located at the same pipeline stage.
- *R* needs two ports (if a register is used to store the last written value). By adding the extra constraint that an output is written only if j' and j have different parities, two banks of single port memory suffice.

Everything said so far does not take into account image borders. It is possible to handle them while still keeping a throughput of one value per cycle, but at the cost of extra combinatorial complexity. Alternatively, it is possible to “cut” the borders by removing the N first and last lines and the M first and last columns in order to keep the hardware simple. It is often desirable to take M such as $M > N$. This way, column noise is filtered more efficiently. An interesting alternative version of the low-pass filter can handle the extreme case where M is the image width and N is one. While this does not give the best results (see Section 5.1), it has the advantage

of using very few hardware resources, in particular no on chip memory at all, and being fit to eliminate the column component of the FPN very efficiently. Also, by taking the approximation that the average value of the previous image line is almost equal to the one of the current, it is possible to implement such a filter with a latency of just a few cycles (depending on the number of pipeline stages), instead of the latency equivalent to the duration of 2^{N-1} lines for the previous implementation (which is quite low, especially when $N < M$, but still much more than a few cycles). This new version is pretty straightforward, with just one register summing all the pixel values of the line. At every horizontal synchronization, this value is divided by the number of pixels of the line (through fixed point multiplication), and stored in a separate register. The value of this new register is used as the low frequency output, and subtracted from the original pixel value in order to obtain the high frequency output.

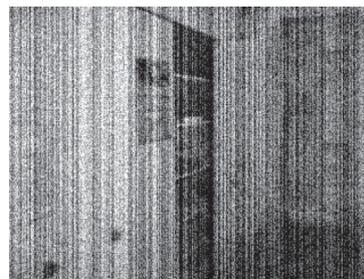
5. Results

5.1. Correction quality

We have tested the proposed method by using video sequences tainted with artificial noise. This way, an image comparison metric can be used to estimate how much the corrected version differs



(a) Original (without artificial noise)



(b) Uncorrected



(c) Corrected using pipeline A



(d) Corrected using pipeline B



(e) Corrected using pipeline C



(f) Corrected using the Harris and Chiang method (with unitary gain)

Fig. 7. 1010th frame of our test scene (same correction parameters as in Table 4).

from the original sequence. We have chosen to use the UQI [47] as such a metric. It roughly corresponds to the human perception of distance between images. The value of the UQI between two images is in the $[-1, 1]$ range, and is 1 for identical images, 0 for uncorrelated images, and -1 for completely anticorrelated images. As recommended in [47], we used a sliding window of size 8×8 pixels for UQI computation. Generally, a high UQI corresponds to a low MSE and therefore a high PSNR, but it tends to offer a higher degree of correlation with human assessment of similarity between images than those two metrics, and this is why it has been preferred to them in this article.

The artificial offset noise used is simply a uniform pixel noise within $\pm 12.5\%$ of the dynamic range, plus an uniform column noise in this range as well. The video sequence used in the tests that follow consists of a 320×240 recording with 1110 frames. The scene is a room with no mobile objects in it, but the panning of the camera across the room provides enough information for noise elimination (see Fig. 7). In order to work correctly, statistics based methods need some time for the estimator values to converge. Therefore, in the rest of this section, the evaluations of efficiency will be given by computing the average UQI on the 100 last frames only. For reference, the average UQI of the 100 last frames is 0.081 when no correction is performed (the comparison is done between the original frames and their artificially noisy versions); evaluating the mean square error on all the pixels of those 100 last frames yields a PSNR of 20.28 dB. In the rest of this section, we will refer to the average UQI between a corrected version and the original sequence on the 100 last frames as the *correction quality*.

We have assumed that there is no loss of precision in the fixed point arithmetic operations of the correction and low-pass blocks. This may or may not be true depending on the actual implementation. Of course, some precision is lost when storing the estimators in memory. Let us assume the analog-to-digital converters of the sensor produce an unsigned fixed point value between 0 and almost 1, the point being located just before the most significant bit. Similarly, each average estimator is a signed fixed point number in the range $[-1, 1]$, and the coefficients α and α' are unsigned fixed point numbers in $[0, 1]$. Taking $\alpha = 2^{-n}$ for some natural n allows to trade the two multipliers of the pixel-level correction block for a subtractor. The same can be achieved for the column-level correction block by taking $\alpha' = 2^{-n'}$. We will therefore assume α and α' to have this form.

Table 1 gives correction qualities depending on the values of α and α' . The estimators were stored as double precision floating point numbers for this measurement, so it is safe to assume that their precision was not an issue. The correction was performed using a software version of the pipeline B (Fig. 3), with a 16×16 low-pass filter. One might have thought that the value of α' should be approximately equal to $\frac{\alpha}{\text{frameheight}}$ (so as the column and pixel

estimators to converge at the same rate per frame), but in fact it is better for the column corrector to converge as early as possible. This helps the pixel mean estimators to converge into correct values (that is, values accounting for the photodetector induced noise only, excluding any noise caused by the column amplifier). We will assume $\alpha = 2^{-8}$ and $\alpha' = 2^{-12}$ for the subsequent tests. Thus, the column estimators will converge $\frac{240}{16} = 15$ times more quickly than the pixel estimators. Once acceptable values for α and α' chosen, it is possible to choose the estimator data-type. Since the required estimator precision depends on the value of α or α' (as seen on Eq. 6), it is advisable to store the column estimators with $\log_2 \frac{\alpha'}{\alpha}$ more bits than the pixel estimators. This leads to column estimators 4 bits larger than pixel estimators for the same computational precision, assuming the previously chosen values for α and α' . As previously for Table 1, a software version of pipeline B has been used to evaluate the correction quality depending on the estimator data-types and precision. Table 2 gives the correction quality by testing several candidate data-types. As we hoped, fixed point arithmetics are as suitable as floating point arithmetics, both with a value of 0.849 for correction quality. It confirms that FPGAs are good candidates for the implementation of FPN correctors. In the rest of this Section, we will assume the use of 16-bit fixed-point pixel estimators and 20-bit fixed-point column estimators.

Table 3 shows the impact of the low-pass filtering method and window size on the correction quality. Data are also obtained from the software implementation of pipeline B. One can see that the one-line filter (noted as $\infty \times 1$) gives quite honorable results compared to full-fledged rectangular window filtering, while being much simpler and not requiring any on-chip RAM at all. As expected, the horizontal filtering window size affects the correction quality more than the vertical window size, due to the column noise. The best value found is 64×32 on our test scene.

Finally, the Table 4 compares the three pipelines architectures proposed earlier, parametrized with the optimum parameters found until now (16-bit/20-bit fixed-point estimators, $\alpha = 2^{-8}$, $\alpha' = 2^{-12}$, 64×32 low-pass filter). Correction quality values in Table 4 have been evaluated from hardware simulations of the different architectures using Modelsim.

It also includes for comparison purposes a simplified (no gain correction) version of the Harris and Chiang corrector. On our test

Table 1
Correction quality depending on the pixel and column correctors exponential window coefficients (Pipeline B, 16×16 low-pass filter window, floating point estimators).

α'	α						
	2^{-4}	2^{-5}	2^{-6}	2^{-7}	2^{-8}	2^{-9}	2^{-10}
2^{-4}	0.739	0.779	0.792	0.799	0.801	0.702	0.395
2^{-5}	0.746	0.787	0.801	0.807	0.810	0.710	0.400
2^{-6}	0.751	0.791	0.806	0.813	0.815	0.715	0.402
2^{-7}	0.757	0.798	0.813	0.819	0.822	0.721	0.404
2^{-8}	0.766	0.808	0.822	0.829	0.832	0.730	0.410
2^{-9}	0.774	0.816	0.831	0.838	0.841	0.738	0.414
2^{-10}	0.779	0.821	0.836	0.843	0.846	0.742	0.418
2^{-11}	0.781	0.823	0.838	0.845	0.848	0.745	0.419
2^{-12}	0.782	0.824	0.839	0.846	0.849	0.745	0.420
2^{-13}	0.783	0.825	0.839	0.846	0.849	0.745	0.420
2^{-14}	0.739	0.787	0.806	0.819	0.832	0.738	0.418

Table 2
Correction quality depending on the estimator data-type and precision (Pipeline B, 16×16 low-pass filter window, $\alpha = 2^{-8}$, $\alpha' = 2^{-12}$).

Estimator datatype and size (pixel/column)	Correction quality
Fixed point, 12-bit/16-bit	0.527
Fixed point, 14-bit/18-bit	0.832
Fixed point, 16-bit/20-bit	0.849
Fixed point, 18-bit/22-bit	0.849
Floating point, 64-bit/64-bit	0.849

Table 3
Correction quality depending on the low-pass filter parameters (Pipeline B, 16-bit/20-bit fixed-point estimators, $\alpha = 2^{-8}$, $\alpha' = 2^{-12}$).

2^M	2^N						
	2	4	8	16	32	64	∞
1	0.143	0.251	0.415	0.600	0.737	0.805	0.905
2	0.184	0.323	0.518	0.713	0.832	0.879	
4	0.219	0.383	0.600	0.791	0.891	0.925	
8	0.242	0.423	0.648	0.833	0.922	0.950	
16	0.253	0.441	0.670	0.849	0.932	0.957	
32	0.256	0.448	0.676	0.853	0.934	0.958	
64	0.257	0.449	0.677	0.853	0.933	0.958	

Table 4

Correction quality depending on pipeline architecture (16-bit/20-bit fixed-point estimators, $\alpha = 2^{-8}$, $\alpha' = 2^{-12}$, 64×32 low-pass filter; the simplified version of Harris' and Chiang's corrector uses double precision floating point arithmetics).

Architecture	Correction quality (UQI/PSNR)
Pipeline A	0.961/36.36 dB
Pipeline B	0.958/36.18 dB
Pipeline C	0.890/28.78 dB
Simplified Harris and Chiang method	0.845/ 24.78 dB

sequence just like on the series of 6 other real world sequences, it results in incorrect low frequency distribution due to unsatisfied working hypotheses, and yields worse results than any of the three other pipelines. The Fig. 7 shows a frame near the end of the sequence for each of the correction methods, once the estimators have converged. The pipeline C may produce “ghost columns” just like the original Harris and Chiang method may produce “ghost images”. It is the case with our test scene (Fig. 7). It may give better results than the pipelines A or B when there is a significant low frequency component in the column FPN. It is not the case with our noise model, and is not generally the case with usual visible light domain sensors.

The pipeline B, which is a little more complex than the A, provides roughly the same correction quality. It converges a bit faster because of the column estimators as seen on Fig. 8. It also implies it may diverge more quickly whenever the working hypotheses of our algorithm are not met. It might be possible to work around this limitation, though: an arbiter could disable the column estimator updates when it detects insufficient variation for a given column between two consecutive frames. While the same principle could be applied to the pixel-level corrector, it seems more practical to implement it on whole columns (dynamic noise is a major problem at the pixel level). This could allow to avoid the fluctuations in the MSE when using the pipeline B and which can be seen on Fig. 8.

We have tested the three pipelines over a set of six other shorter video sequences in order to test the convergence of the proposed method (Table 5). As in Table 4, hardware simulations using Modelsim have been performed to evaluate the correction quality.

Each sequence only contains 125 frames, which is about eight times shorter than the previous test sequence. Sequences 1 and 2 use a still camera while some subjects are moving over part of its field of view. Sequence 1 has a uniform background, while sequence 2 has a very detailed still background (hence violating our third hypothesis). Sequences 3 to 6 involve a moving camera, and vary by the level of luminosity and of contrast. Sequences 3 and 4 are rather bright, while sequences 5 and 6 are almost pitch dark. High contrast elements are found on the sequences 4 and 6. The reference sequence is the one used in the previous tables, but restricted to its 125 first frames, for comparison purposes. Once

again, the pipeline A gets the best results. As expected, the sequence 2 is the least effectively corrected but the proposed method still outperforms the Harris' and Chiang's on that sequence. According to PSNR measurements, the sequences 3 and 6 are also poorly corrected, while they are satisfying using a visual criterion, such as the UQI. On the contrary, the corrected version of the sequence 4 shows visual shortcomings while having descent PSNRs for the three tested pipelines. Globally, the proposed correction method does not seem to specifically favor bright sequences over dark ones, and works for different levels of contrast.

If we observe more deeply the detailed results of the 4 different architectures, we can note the following aspects. First of all, as seen in Table 5, pipeline C always outperforms the original Harris and Chiang's method, both in terms of PSNR and UQI, except for the sequence 1 (the moving hand on an uniform background). We can also note that this specific sequence also seems to be well-adapted to the Harris method because the difference of performance between pipeline A, pipeline B and Harris and Chiang method are not significant in terms of image quality (respectively 0.895, 0.885, and 0.881). Secondly, as previously mentioned, the column estimators converge 15 times more quickly than the pixel estimators if the working conditions are met. This gives to pipeline B and pipeline C significant advantages compared to pipeline A and Harris method. Finally, the main difference between pipeline B and pipeline C is that the column-level correction of pipeline C runs before the lowpass filter. Compared to pipeline B, this approach gives similar or even lightly superior performance for sequences 4–6 in which we find poor illuminated areas. In other sequences, the performance of pipeline C is lower, and particularly in terms of PSNR, revealing that the column estimators not only filter noise but also signal.

5.2. Implementation on a smart camera

We have worked on implementing the proposed corrector using the Industrial Video Processing Kit commercialized by Avnet (part reference: AES-S6IVK-LX150T-G). It is composed of the general-purpose Avnet Spartan-6 LX150T development board (based on the Xilinx XC6SLX150T-3FGG676 FPGA) and of two FMC (FPGA Mezzanine Card) daughter cards with IO ports dedicated to video processing. It is shipped with an omnivision OV9715 image sensor [48], which can be connected to one of the daughter FMCs, and is capable of video capture at resolutions up to 1280×800 at 30 frames per second. The Fig. 9 is a photograph of the whole setup.

We have implemented the different parts of the proposed correction pipelines and assembled them into a pcore for use with the Xilinx EDK (Embedded Development Kit). Proceeding like this eases its reuse can allows its parametrization through memory-mapped registers accessible from software via the PLB (Processor

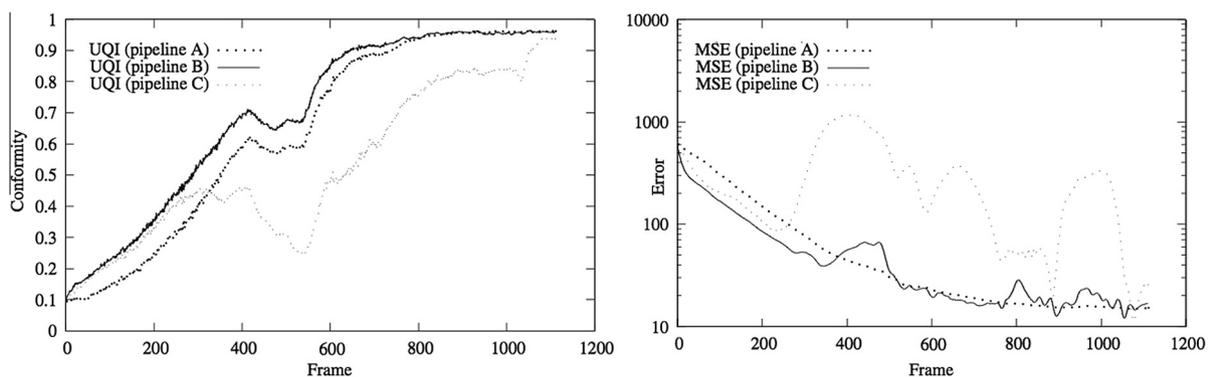


Fig. 8. UQI and mean square error depending on the frame number; same sequence and parameters as in Table 4.

Table 5

Correction quality for six shorter scenes (125 frames only, correction quality evaluated on the last 12 frames). Same parameters as in table 4, save for $\alpha = 2^{-5}$ and $\alpha' = 2^{-9}$. For comparison purposes, the previously used test sequence shortened to 125 frames appears in the column "ref. seq".

Architecture	Correction quality (UQI/PSNR)						Ref. seq.
	Seq. 1	Seq. 2	Seq. 3	Seq. 4	Seq. 5	Seq. 6	
Pipeline A	0.895 / 35.50 dB	0.254 / 20.11 dB	0.941 / 22.87 dB	0.762 / 33.01 dB	0.914 / 36.78 dB	0.833 / 22.40 dB	0.912 / 34.42 dB
Pipeline B	0.885 / 34.25 dB	0.250 / 20.10 dB	0.940 / 22.23 dB	0.700 / 31.04 dB	0.907 / 36.72 dB	0.827 / 22.22 dB	0.888 / 33.54 dB
Pipeline C	0.840 / 23.82 dB	0.220 / 18.23 dB	0.875 / 19.63 dB	0.683 / 31.92 dB	0.917 / 37.73 dB	0.823 / 21.84 dB	0.876 / 30.38 dB
Simplified Harris and Chiang method	0.881 / 27.18 dB	0.203 / 14.59 dB	0.510 / 12.88 dB	0.417 / 23.76 dB	0.901 / 34.29 dB	0.795 / 16.47 dB	0.850 / 25.89 dB



Fig. 9. The Avnet Industrial Video Processing Kit.

Local Bus) bus. The resulting pcore can be instantiated inside virtually any existing XPS (Xilinx Platform Studio) design just between the camera input and the first actual video processing IP.

The implementation itself has been done in VHDL and in Haskell. VHDL was used for the control-dominated components, for

those dealing with more than one clock domain (typically, at the interface with the memory controller), and whenever the critical path length was not a limiting factor. We have used an in-house Haskell based hardware description embedded language [49] to generate pipelined netlists from combinatorial datapaths. Most of the low pass and corrector blocks, which are arithmetics-dominated, could be implemented quite straightforwardly this way. One of the benefits of our automatized pipelining process is that arithmetical operators (such as adders, subtractors or multipliers) are not necessarily bound to a single pipeline stage, as it is the case when using a VHDL library such as IEEE `numeric_std`. On the contrary, the pipeline stages were generated in such a way that the critical paths of the different stages have approximately the same length. As a result, in our implementation the barriers between pipeline stages do not match those shown on the schematics of Fig. 5 and Fig. 6. However, as discussed later in this section, the illustrated total pipeline stage count is sensible.

A minimalist design was used for our tests as illustrated on Fig. 10. It is a system with two DMAs: one used to write at 30 frames per second the output of the sensor to the external DDR3 memory, and the other to read it and output it through a DVI port at 60 frames per second. In order to cope with these two different framerates, each processed frame is read and displayed twice on the screen. The sensor output is treated with our corrector before

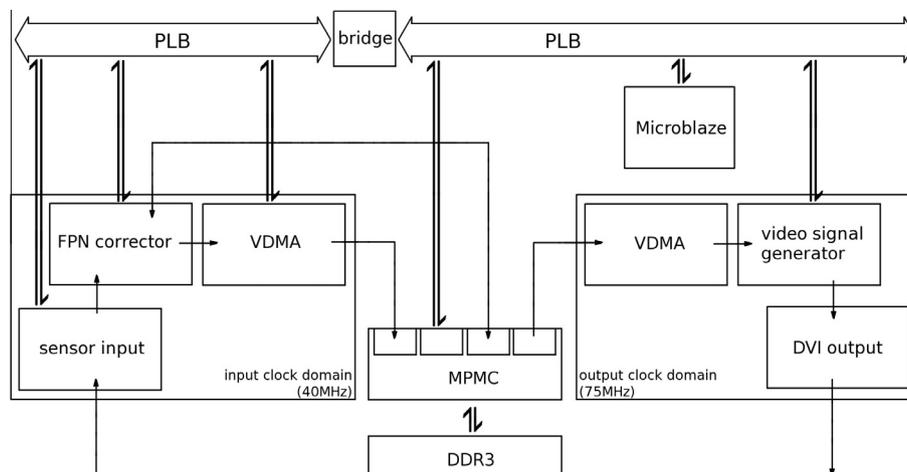


Fig. 10. Minimalist corrector pcore test design.

Table 6
FPGA implementation performance and resource usage, depending on pipeline depths.

# Pipeline stages		Frequency (MHz)	Pipeline A		Pipeline B		Pipeline C	
Lowpass	Corrector		# Registers	# LUTs	# Registers	# LUTs	# Registers	# LUTs
1	1	80.652	496	176	580	289	742	311
2	2	99.502	608	277	748	280	910	302
4	3	140.076	773	497	994	762	1156	784
5	4	149.410	844	490	1099	745	1261	767
6	4	169.319	895	475	1172	596	1334	618
7	5	129.232	946	513	1258	786	1420	808
8	6	158.278	1007	503	1348	854	1510	876
11	8	148.170	1128	517	1546	671	1708	693
22	16	134.391	1347	719	1878	1124	2040	1146

being transmitted to the DMA controller. Finally, a Xilinx Microblaze softcore is used to handle the initialization of all the peripherals, and can be used during operation to alter parameters such as sensor exposure and gain, resolution, and so on. During our tests, those parameters have been fixed and any automatic sensor behavior disabled. Nevertheless, our correction pipeline should be able to operate with automatic exposure and gain adjustments, as long as their variations are slow. It should also be possible to feed in real time the sensor gain value to the corrector for better results, especially with quickly varying gain. Those aspects are currently being investigated and will not be discussed further in this paper.

We have evaluated the hardware complexity of our pipelines in the case where they use a NPI (native port interface) point-to-point link to the Xilinx MPMC external memory controller and the one-line version of the lowpass filter. The NPI interface is the simplest of the interfaces supported by the MPMC, but requires synchronicity with the MPMC clock domain (which differs from the video input clock domain in our system). Other possible implementation choices were the use of the VFBC (video frame buffer controller) point-to-point link, or simply to access the memory controller through the PLB. However, using the NPI interface as well as the one-line filter, much simpler than a conventional rectangle window low pass filter, allows us to get an upper bound of the frequency achievable and a lower bound of the slice usage.

The Table 6 sums up the performance of the correction pipeline and its resource usage measured on the above-described test system, depending on the parameters passed to our pipelining tool. As expected, the higher the pipeline stage count, the higher the register count and the higher the slice usage, as more of them are used either as registers without accompanying logic, or simply as route-throughs. The highest reachable frequency is around 169 MHz, when using 6 pipeline stages for the low pass filter and 4 stages per corrector (the pixel and column correctors share the same architecture, so their pipelines have the same length). The maximum frequency goes down when using deeper pipelines, most likely due to routing issues as the number of slices increases. It is interesting to note that in our test system, we could have afforded to use single stage correctors, as the sensor input clockrate is 40 MHz. This is not the case for systems using better cameras: with a sensor capable of 1080p capture at 60 Hz, a working frequency of 125 MHz or higher is necessary, requiring moderately deep correction pipeline use. When comparing the resources used those available in the FPGA, the slice register usage ranges from 0.27% to 1.11%, and the slice LUT usage from 0.20% to 1.25%.

6. Conclusion

We have proposed an improved algorithm for FPN correction in visible light sensors and three associated architectural pipelines that can be easily implemented in any FPGA-based smart camera. Its hardware complexity is marginally higher than that of a simple

reference image while it significantly achieves improved quality of resulting images. Moreover, it is more robust to variations in the sensor environment or operating mode. Contrarily to the original approach we improve on, it is capable of handling a wide variety of scenes correctly by separating the signal from the noise in a more reliable fashion. This is achieved by taking into account information from the spatial neighborhood of the pixels being corrected.

Finally, our correction pipeline adds little latency in most cases and hence is suited for incorporation in applications where this is an issue. A first notable future work perspective is the automatic detection of static scenes. Indeed, a mechanism to temporarily prevent the update of estimators when two consecutive frames are too similar would allow to preserve their correctly computed values when the convergence hypotheses are not met, further increasing correction robustness. A second perspective is the implementation of this method in other FPGA smart cameras developed in our research group. We aim at implementing such a correction method to the HDR-ARTIST platform [50] dedicated to high dynamic range imaging in order to achieve the best quality possible.

Acknowledgment

The authors thank the DGCIS (French Ministry for Industry) for financial support within the framework of the project HiDRaLoN.

References

- [1] H. Tian, Noise analysis in cmos image sensors (Ph.D. thesis), 2000.
- [2] H. Tian, B. Fowler, A. Gamal, Analysis of temporal noise in cmos photodiode active pixel sensor, *IEEE Journal of Solid-State Circuits* 36 (1) (2001) 92–101.
- [3] J. Nakamura, *Image Sensors and Signal Processing for Digital Still Cameras*, CRC Press, Inc., Boca Raton, FL, USA, 2005.
- [4] S. Morfu, P. Marquie, B. Nofeie, D. Ginhac, Nonlinear systems for image processing, in: P.W. Hawkes (Ed.), *Advances in imaging and electron physics, Advances in Imaging and Electron Physics*, vol. 152, Elsevier, 2008, pp. 79–151.
- [5] HPComponentsGroup, Noise sources in CMOS image sensors, Tech. Rep., Hewlett-Packard Company, 1998.
- [6] A.E. Gamal, B. Fowler, H. Min, X. Liu, Modeling and estimation of fpn components in cmos image sensors, in: *Proc. SPIE, Solid State Sensor Arrays: Development and Applications II*, vol. 3301, 1998, pp. 168–177.
- [7] A. ElGamal, H. Eltoukhy, Cmos image sensors, *IEEE Circuits and Devices Magazine* 21 (3) (2005) 6–20.
- [8] M.V. Konnik, J.S. Welsh, On numerical simulation of high-speed ccd/cmos-based wavefront sensors in adaptive optics, *SPIE* 8149 (2011).
- [9] S.C. Kelly, R.M. Guidash, B.H. Pillman, Fixed pattern noise removal in cmos imagers across various operational conditions, 2008.
- [10] M. Schberl, C. Senel, S. Fel, H. Bloss, A. Kaup, Non-linear dark current fixed pattern noise compensation for variable frame rate moving picture cameras, in: *Proceedings of European Signal Processing Conference (EUSIPCO)*, Glasgow, Scotland, 2009.
- [11] M. Bigas, E. Cabruja, J. Forest, J. Salvi, Review of cmos image sensors, *Microelectronics Journal* 37 (5) (2006) 433–451.
- [12] B. Fowler, A.E. Gamal, D. Yang, H. Tian, A method for estimating quantum efficiency for cmos image sensors, in: *Proc. SPIE, Solid State Sensor Arrays: Development and Applications II*, vol. 3301, 1998, pp. 178–185.
- [13] M. Schoandberl, S. Foandssel, A. Kaup, Fixed pattern noise column drift compensation (cdc) for digital moving picture cameras, in: *Image Processing (ICIP)*, 2010 17th IEEE International Conference on 2010, pp. 573–576.

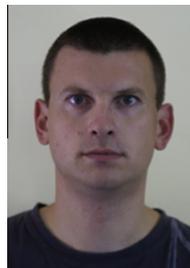
- [14] S. Otim, B. Choubey, D. Joseph, S. Collins, Characterization and simple fixed pattern noise correction in wide dynamic range logarithmic imagers, *IEEE Transactions on Instrumentation and Measurement* 56 (5) (2007) 1910–1916.
- [15] R. Nixon, S. Kemeny, B. Pain, C. Staller, E. Fossum, 256×256 cmos active pixel sensor camera-on-a-chip, *IEEE Journal of Solid-State Circuits* 31 (12) (1996) 2046–2050.
- [16] R. Guidash, T.H. Lee, P. Lee, D. Sackett, C. Drowley, M. Swenson, et al., A 0.6μm cmos pinned photodiode color imager technology, in: *IEDM'97, Electron Devices Meeting, 1997, Technical Digest 1997*, pp. 927–929.
- [17] K. Yonemoto, H. Sumi, A cmos image sensor with a simple fixed-pattern-noise-reduction technology and a hole accumulation diode, *IEEE Journal of Solid-State Circuits* 35 (12) (2000) 2038–2043.
- [18] S. Kavadias, B. Dierickx, D. Scheffer, A. Alaerts, D. Uwaerts, J. Bogaerts, A logarithmic response cmos image sensor with on-chip calibration, *IEEE Journal of Solid-State Circuits* 35 (8) (2000) 1146–1152.
- [19] M. Loose, K. Meier, J. Schemmel, A self-calibrating single-chip cmos camera with logarithmic response, *IEEE Journal of Solid-State Circuits* 36 (4) (2001) 586–596.
- [20] Y. Degerli, F. Lavernhe, P. Magan, P. Farre, Column readout circuit with global charge amplifier for cmos aps imagers, *Electronics Letters* 36 (17) (2000) 1457–1459.
- [21] B. Dierickx, B. Dierickx, D. Scheffer, D. Scheffer, G. Meynants, G. Meynants, et al., Random addressable active pixel image sensors, in: *Proceedings of the SPIE*, vol. 2950, 1996, pp. 2–7.
- [22] S. Kavadias, B. Dierickx, D. Scheffer, On-chip offset calibrated logarithmic response image sensor, in: *Proc. IEEE Workshop Charge-Coupled Devices and Advanced Image Sensors, 1999*, pp. 68–71.
- [23] L.W. Lai, C.H. Lai, Y.C. King, A novel logarithmic response cmos image sensor with high output voltage swing and in-pixel fixed-pattern noise reduction, *IEEE Sensors Journal* 4 (1) (2004) 122–126.
- [24] R. Mosqueron, J. Dubois, M. Paindavoine, High-speed smart camera with high resolution, *EURASIP Journal on Embedded Systems* 1 (2007) 024163.
- [25] R. Mosqueron, J. Dubois, M. Mattavelli, D. Mauvilet, Smart camera based on embedded hw/sw coprocessor, *EURASIP Journal on Embedded Systems* 1 (2008) 597872.
- [26] F.D. Real, F. Berry, Smart cameras: technologies and applications, in: A.N. Belbachir (Ed.), *Smart Cameras*, Springer, US, 2010, ISBN 978-1-4419-0953-4, pp. 35–50.
- [27] P.J. Lapray, B. Heyrman, M. Rosse, D. Gin hac, Smart camera design for realtime high dynamic range imaging, in: *Fifth ACM/IEEE International Conference on Distributed Smart Cameras (ICDSC 2011)*, 2011, pp. 1–7.
- [28] J. Harris, Y.M. Chiang, Nonuniformity correction using the constant-statistics constraint: analog and digital implementations, in: *Proceedings of SPIE*, vol. 3061, 1997, pp. 895–905.
- [29] B. Ratliff, M. Hayat, J. Tyo, Generalized algebraic scene-based nonuniformity correction algorithm, *Journal of the Optical Society of America A* 22 (2) (2005) 239–249.
- [30] M. Richard, G.C. Malueg, Detector array fixed-pattern noise compensation, 1976.
- [31] A. Friedenber g, I. Goldblatt, Nonuniformity two-point linear correction errors in infrared focal plane arrays, *Optical Engineering* 37 (4) (1998) 1251–1253.
- [32] J.W. Newman, Method and apparatus for image signal compensation of dark current, focal plane temperature, and electronics temperature, 2007.
- [33] A. Bosco, K. Findlater, S. Battiatto, A. Castorina, A temporal noise reduction filter based on image sensor full-frame data, in: *IEEE International Conference on Consumer Electronics, ICCE 2003, 2003*, pp. 402–403.
- [34] A. Bosco, K. Findlater, S. Battiatto, A. Castorina, A noise reduction filter for full-frame data imaging devices, *IEEE Transactions on Consumer Electronics* 49 (3) (2003).
- [35] T. Huang, L. Xiao-Lin, S. Yang, Fixed pattern noise suppression algorithm based on background modeling, *Procedia Engineering* 29 (2012) 884–888.
- [36] U. Sakoglu, R. Hardie, M. Hayat, B. Ratliff, J. Tyo, Nan algebraic restoration method for estimating fixed pattern noise in infrared imagery from a video sequence, in: *9th Annual Meeting of the SPIE: Applications of Digital Image Processing XXVII*, Denver, CO., SPIE Proc, vol. 5558, 2004, pp. 69–79.
- [37] R.C. Hardie, M.M. Hayat, E. Armstrong, B. Yasuda, Scene-based nonuniformity correction with video sequences and registration, *Applied Optics* 39 (8) (2000) 1241–1250.
- [38] P. Torle, I.A. Andersson, L. Haglund, Scene-based correction of image sensor deficiencies, *SPIE* 5074 (2003) 249–260.
- [39] C. Zuo, Q. Chen, G. Gu, X. Sui, Scene-based nonuniformity correction algorithm based on interframe registration, *Journal of the Optical Society of America A* 28 (6) (2011).
- [40] B.M. Ratliff, M.M. Hayat, R.C. Hardie, An algebraic algorithm for nonuniformity correction in focal-plane arrays, *Journal of the Optical Society of America A* 19 (9) (2002) 1737–1747.
- [41] J. Harris, Continuous-time calibration of vlsi sensors for gain and offset variations, in: *Proceedings of the SPIE: Smart Focal Plane Arrays and Focal Plane Array Testing*, vol. 2474, 1995, p. 2333.
- [42] E. Vera, R. Reeves, S. Torres, Adaptive bias compensation for non-uniformity correction on infrared focal plane array detectors, in: *HIS, 2002*, pp. 725–734.
- [43] C. Zhang, W. Zhao, Scene-based nonuniformity correction using local constant statistics, *Journal of the Optical Society of America A* 25 (6) (2008) 1444–1453.
- [44] C. Zuo, Q. Chen, G. Gu, X. Sui, W. Qian, Scene-based nonuniformity correction method using multiscale constant statistics, *Optical Engineering* 50 (8) (2011).
- [45] R. Figueras, J. Sabadell, J. Margarit, E. Martin, L. Teres, F. Serra-Graells, A 0.18 μm cmos low-power charge-integration dps for X-ray imaging, in: *Biomedical Circuits and Systems Conference, 2009, BioCAS 2009, IEEE 2009*, pp. 209–212.
- [46] A. Kumar, S. Sarkar, R. Agarwal, Fixed pattern noise correction and implementation for infrared focal plane array based staring system using scene statistics, *International Journal of Industrial and Systems Engineering* 1 (2007).
- [47] Z. Wang, A.C. Bovik, A universal image quality index, *IEEE Signal Processing Letters* 9 (3) (2002) 81–84.
- [48] Ov9715 1-megapixel product brief, OmniVision Technologies, Inc., 2011.
- [49] T. Toczek, D. Houzet, S. Mancini, Another take on functional system-level design and modeling, in: D. Borrione (Ed.), *Advances in Design Methods from Modeling Languages for Embedded Systems and SoC's*, Lecture Notes in Electrical Engineering, vol. 63, Springer, Netherlands, 2010, pp. 93–106.
- [50] P.J. Lapray, B. Heyrman, M. Rosse, D. Gin hac, Hdr-artist: High dynamic range advanced real-time imaging system, in: *IEEE International Symposium on Circuits and Systems (ISCAS 2012)*, 2012, pp. 1–6.



Tomasz Toczek is currently a postdoctoral researcher at LE2I (Laboratory of Electronics, Computing and Imaging Sciences; affiliated to the CNRS). He has obtained his PhD in Micro and Nano-electronics in 2011 from the University of Grenoble. His research interests include FPGA prototyping, functional programming for hardware description, and GPU programming.



Faouzi Hamdi is a PhD student in electronics and image processing at the laboratory Le2i (UMR CNRS 6063) - Burgundy University, France. His current research activities include embedded image processing on smart camera.



Barthelemy Heyrman received his PhD in electronics and image processing from Burgundy University, France, in 2005. He is currently associate professor at the University of Burgundy, France and member of LE2I UMR CNRS 5158 (Laboratory of Electronic, Computing and Imaging Sciences). His main research topics are system on chip smart camera and embedded image processing chips.



Julien Dubois is associated professor at the University of Burgundy since 2003. He is a member of the Laboratory Le2i (UMR CNRS 6063). His research interests include real-time implementation, smart camera, hardware design based on data-flow modeling, motion estimation and image compression. In 2001, he received PhD in Electronics from the University Jean Monnet of Saint Etienne (France) and joined EPFL based in Lausanne (Switzerland) as a project leader to develop a co-processor, based on FPGA, for a new CMOS camera.



Johel Miteran received the Ph.D. degree in image processing from the University of Burgundy, Dijon, France in 1994. Since 1996, he has been an assistant professor and since 2006 he has been professor at Le2i, University of Burgundy. He is now engaged in research on classification algorithms, face recognition, access control problem and real time implementation of these algorithms on software and hardware architecture.



Dominique Ginhac received his PhD in electronics and image processing from Clermont-Ferrand University, France, in 1999. He is currently associate professor at the University of Burgundy, France and member of LE2I UMR CNRS 5158 (Laboratory of Electronics, Computing and Imaging Sciences). His main research topics are hardware design of smart vision systems and implementation of real-time image processing applications.