Bit-plane Processing Techniques for Low-Light, High Speed Imaging with a SPAD-based QIS

Istvan Gyongy[1], Neale A.W. Dutton[2], Luca Parmesan[1], Amy Davies[3], Rebecca Saleeb[3], Rory Duncan[3], Colin Rickman[3], Paul Dalgarno[3], Robert K. Henderson[1]
[1]The University of Edinburgh, Institute for Integrated Micro and Nano Systems, Edinburgh, U.K.
[2]STMicroelectronics Imaging Division, 33 Pinkhill, Edinburgh, U.K., EH12 7BF
[3]Heriot-Watt University, Institute of Biological Chemistry, Biophysics and Bioengineering, Edinburgh, U.K.
Istvan.Gyongy@ed.ac.uk Tel: +44 131 650 5568

***Abstract*–Advances in SPAD sensor technology have recently yielded a >10kfps, 320x240 resolution imager [1] that can be operated as a Quanta Image Sensor (QIS) [2]. This device (labelled SPCImager) thus combines the high speed and sensitivity afforded by SPAD devices with the negligible readout noise and logarithmic compression of QIS. These attributes open the door to a number of potential applications [3]. This work explores two applications: high speed imaging of fast moving objects and low light microscopy. Appropriate signal processing schemes are considered in both cases.**

I. INTRODUCTION

The SPCImager is implemented in 0.13µm imaging CMOS; it features an 8µm pixel pitch at 26.8% fill factor, and has a peak photon detection efficiency of 35% at 450nm. When operated as a QIS, the raw output of the sensor are bit-planes composed of binary pixel values, indicating photon detections, which are then aggregated in time (and/or space) to compose grayscale images. Both rolling and global shutter modes are available, and the range of lines read out may be reduced for an increased frame rate. The imager is twinned with an FPGA board (Opal Kelly XEM6310) that controls the acquisition of image data, relaying a continuous stream of bit-planes to a PC over a USB 3.0 link. With the benefit of a complete record of bit-planes, different back processing (or aggregation) schemes can then be explored in order to identify or enhance particular features in the image data.

II. HIGH SPEED IMAGING

Figure 1 shows an example sequence of raw bit-planes from the imager, capturing a rotating fan at 10kfps. A standard approach for increasing the bit depth of a QIS sensor's output is to sum distinct groups of bit-planes (similar to the integration of light in a conventional CMOS or CCD sensor). However, the frame rate is then lowered in proportion to the number of bit-planes summed. An alternative scheme is proposed here whereby a (real-time) rolling sum is carried out across bit fields (as illustrated in the figure). The native frame rate of the sensor is therefore preserved, enabling closer tracking of moving objects. Figure 2 suggests a possible hardware implementation of the scheme (on FPGA) involving a shift register and accumulator for each pixel of the sensor. The method lends itself to extensions; for instance, the length of the summation window can be adapted independently for individual pixels (or groups of pixels) according to the variability in the measured light intensity [2]. More specifically, shorter windows are used for "dynamic" areas of a scene to reduce motion blur, and longer windows in "static" regions ensuring high bit depth.

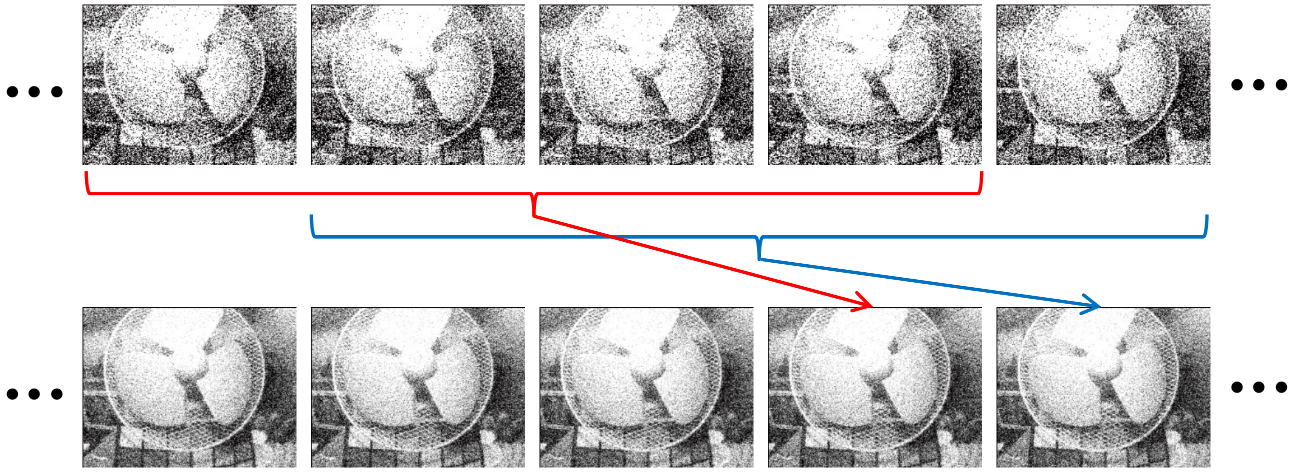

Figure 1. Rolling summation of bit-planes capturing a rotating fan. The upper sequence of images consists of consecutive raw bit-planes, each exposed for 2µs. Carrying out a moving sum of four bit-planes leads to the second set of images. The bit depth has been increased, whilst preserving the frame rate.

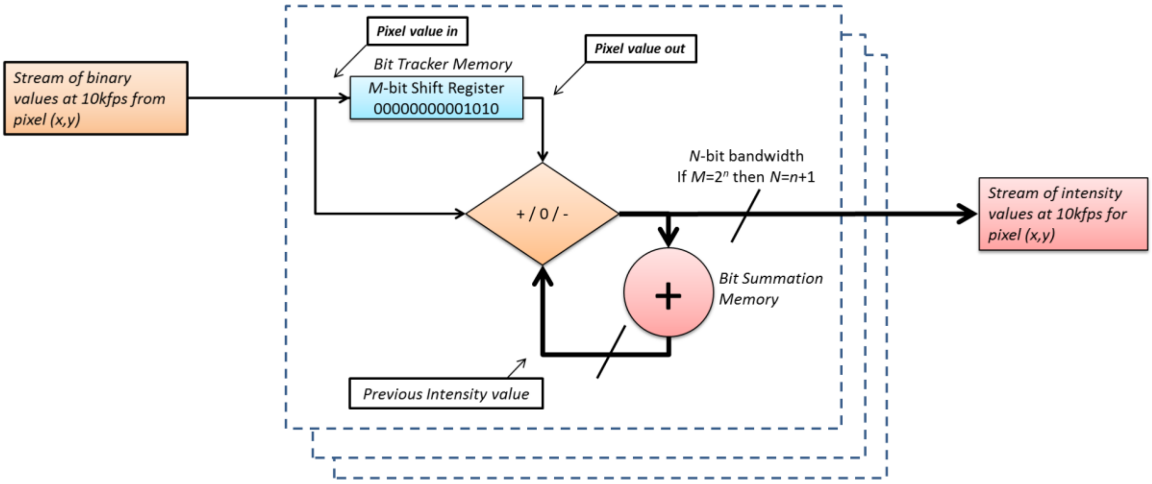

Figure 2. Hardware implementation of rolling summation. Width M of shift register corresponds to the number of bit-planes to be summed.

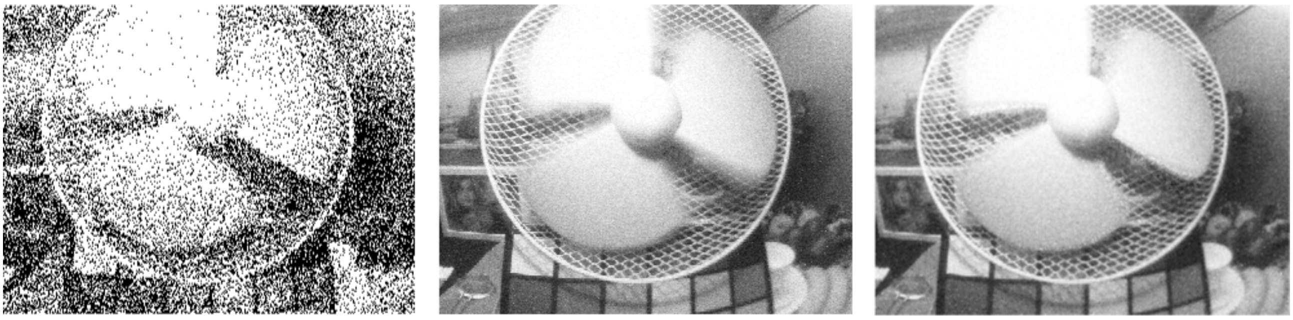

a **b** **c**

Figure 3. Rotating fan imaged using: a) Single raw bit-plane b) Rolling sum of 128 bit-planes c) Adaptive rolling sum of bit-planes

A simple version of such an algorithm calculates rolling sums over, for example, both the last 32 and last 128 binary values of each pixel. In addition, for each pixel a rolling measure of the variability of light intensity is taken. The “processed” pixel intensity then either becomes the longer or (a scaled version of) the shorter bit-sum depending on the current value of variance. Figure 3 demonstrates the effect of the suggested algorithm. On the left is a frame from the sequence of raw bit-planes; the middle image is an aggregation of 128 bit-planes. The right image shows the result of adaptive aggregation (i.e. selectively opting for 32 or 128 bit-plane summation over the pixel array); the motion blur on the fan blades is seen to be reduced, whilst the background is largely unaffected. Figure 4 indicates the pixels that, for the purpose of composing the adaptively aggregated frame of Fig. 3c, were considered to have “high variance”. Most of these pixels are shown to be concentrated around the edges of the fan blades.

Just as standard rolling aggregation, “adaptive” aggregation can be computed independently for each pixel and is therefore highly parallelisable for an FPGA implementation. An outline of the algorithm applied here is given in Figure 5.

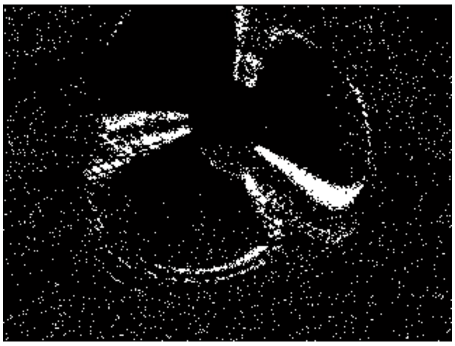

Figure 4. Identifying dynamic areas of a frame for adaptive bit-plane aggregation: shown in white are pixels where the detected light intensity is deemed to be highly varying (according to the last 128 bit-planes)

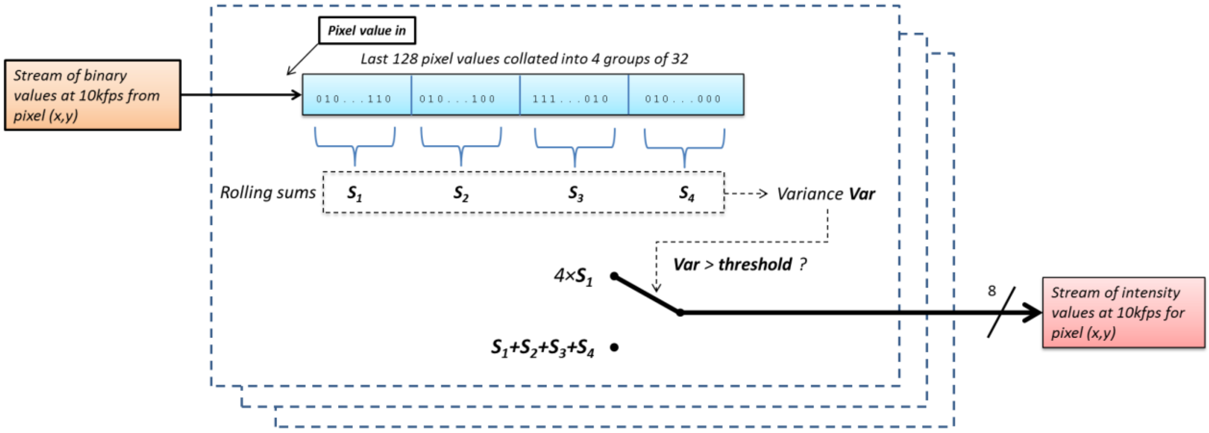

Figure 5. Example algorithm for adaptive rolling summation of pixel values

III. LOW LIGHT IMAGING

SPAD-based image sensors are, by their very nature, capable of detecting single photons, making them ideal candidates for low light microscopy. The performance of SPCImager was therefore assessed in this application. Head-to-head tests were carried out with a commercial sCMOS camera (Hamamatsu ORCA-Flash4.0), with both devices coupled to a standard 100× widefield fluorescence microscope (Olympus IX71) via a 50:50 beam splitter (see Fig. 6). In this configuration, both cameras simultaneously recorded the identical field of view. In one test, a sample of spatially dispersed, immobilized, quantum dots was studied. Both bright clusters and isolated dots were observed. Figure 7 shows the output of the cameras; in the case of SPCImager 100 bit-planes, each exposed for 100µs, have been added together, whilst the image of sCMOS is for a single 10ms exposure (making the total exposure time the same for both cameras). The noise in the standard image of SPCImager (Fig. 7a), a result of the dark count rate (DCR) of the pixels, becomes more pronounced in low light, masking some of the finer details in the sample (the median DCR at room temperature is 50Hz; no cooling is used). By taking a sequence of background bit-planes under dark conditions prior to imaging one can compensate for the DCR, and after median filtering a much “cleaner” image is obtained (Fig. 7b), comparable to the sCMOS output (Fig. 7c).

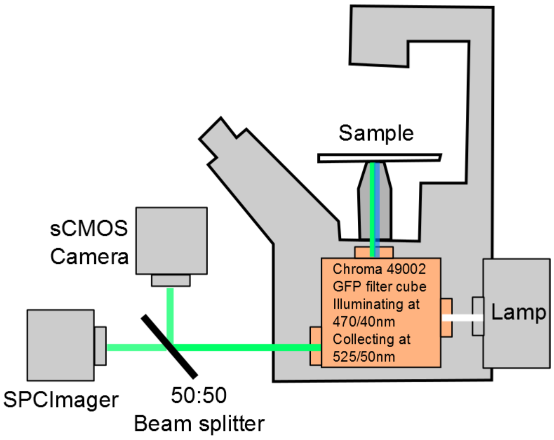

Figure 6. Diagram of setup for quantum dot microscopy experiment

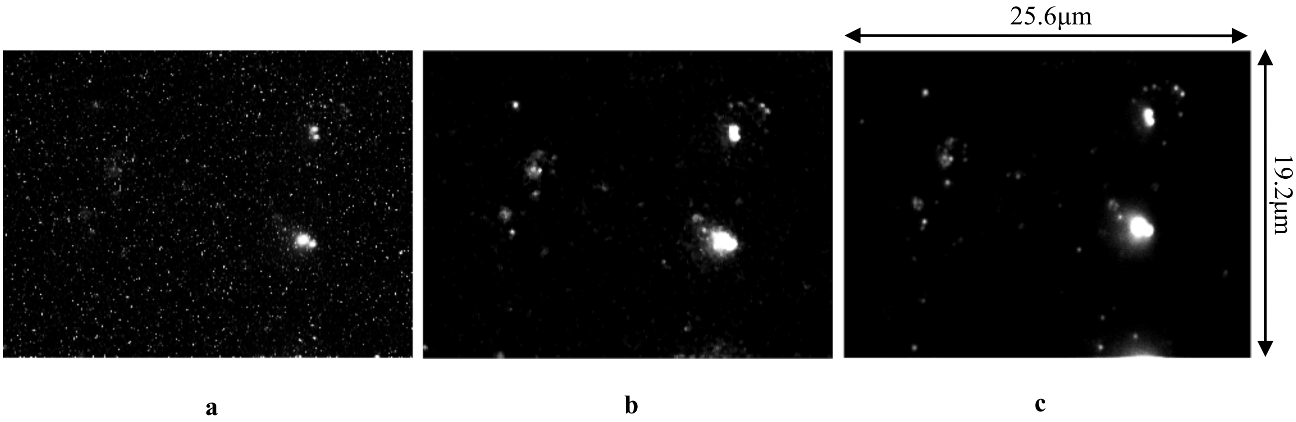

a b c

Figure 7. Quantum dots (QDot® 525 ITK™ Carboxyl) as captured by: a) Aggregated output of SPCImager (100×100µs exposures) b) Aggregated output of SPCImager after DCR compensation and median filtering c) Hamamatsu ORCA-Flash4.0 sCMOS camera (single 10ms exposure). Region of interest of sCMOS was chosen to match imaging area of SPCImager. Autocontrast has been applied to all images.

It is important to note that the sCMOS camera does offer higher sensitivity in most circumstances, owing to the higher quantum efficiency and fill factor of its pixels. Whilst SPAD image sensors have been predicted (using computer simulations) to match or even exceed the performance of sCMOS and EMCCD when used for single-molecule localisation [4], this was under the assumption of a much higher overall photon detection efficiency than currently available in SPAD imagers. The type of experiment that SPCImager may be best suited for is when dynamic behaviour occurring over short time-scales (inaccessible to sCMOS) is to be observed. To test out the device on a dynamic microscopy sample, fluorescent beads undergoing Brownian motion were imaged. Figure 8a shows a sequence of consecutive images of a single bead, each being composed of just 10 bit-planes. Despite the low total exposure time per image (0.5ms), the bead is clearly visible throughout, enabling its position to be tracked [5]. Increasing the number of bit-planes aggregated from 10 to 100 leads to an image with a smoother intensity profile, but one that suffers from motion blur (Fig. 8b). This highlights the power of being able to adapt the summing of bit-planes for particular imaging conditions, trading bit depth for increased time resolution when necessary.

In addition to capturing quantum dots and fluorescent beads, SPCImager was used to image cells. An example image is given in Figure 9 of a HeLa cell, revealing some of the cellular structure.

IV. CONCLUSIONS

The flexibility in the way the output of a QIS can be processed means that image composition can be tailored to particular applications. This paper has illustrated, in the context of a SPAD QIS, that by adapting the summation of bit-planes, whether at a per pixel level or uniformly over the whole pixel array, distinct advantages may be gained.

ACKNOWLEDGEMENTS

This research was founded by the ERC TotalPhoton grant. The authors appreciate the support of STMicroelectronics who fabricated the device. The use of the ESRIC facilities at Heriot-Watt University is also gratefully acknowledged.

REFERENCES

[1] Dutton, N.A.W. et al. “320x240 Oversampled Digital Single Photon Counting Image Sensor” IEEE VLSI Sym. 2014

[2] Fossum, E. “A Digital Film Sensor” in “Image Sensors and Signal Processing for DSCs”, Ed. Nakamura, J. (2005)

[3] Burri, S. et al. “Architecture and applications of a high resolution gated SPAD image sensor” Opt. Express 22, 17573-17589 (2014)

[4] Krishnaswami, V. et al. “Towards digital photon counting cameras for single-molecule optical nanoscopy” Optical Nanoscopy 3(1), 1 (2014)

[5] Thompson, R. et al. “Precise nanometer localization analysis for individual fluorescent probes” Biophys J. 82(5), 2775–2783 (2012)

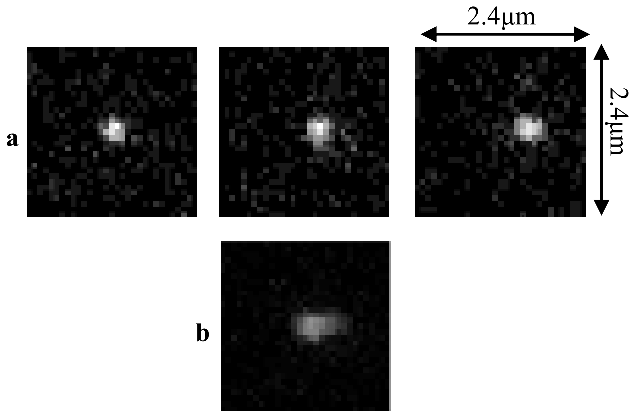

Figure 8. Images of 0.1μ fluorescent bead (580/605 Fluorosphere suspended in water): a) Sequence of images in which each image is a sum of 10×50μs exposures (total acquisition time is 1.5ms per image) b) Aggregate image of 100×50μs exposures (acquisition time 15ms). 30×30 pixel area is shown in all images.

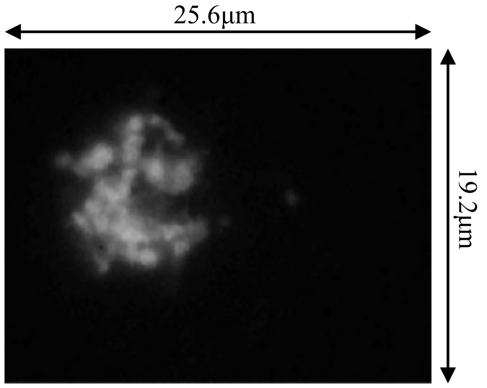

Figure 9. Image of HeLa cell (with green fluorescent protein-marked Syntaxin17 proteins) taken using SPCImager (Sum of 1000×50μs exposures, DCR compensation and median filtering has been applied)