

A Multi-Resolution 100 GOPS 4 Gpixels/s Programmable CMOS Image Sensor for Machine Vision

Robert Johansson, Leif Lindgren, Johan Melander, and Björn Möller

IVP Integrated Vision Products AB

Wallenbergs gata 4, SE-583 35 Linköping, SWEDEN

robert_johansson@home.se, leifl@isy.liu.se, jme@ivp.se, bjomo537@student.liu.se

Abstract—This paper presents a multi-resolution general-purpose high-speed machine vision sensor with on-chip image processing capabilities. The sensor comprises a multi-resolution sensing area, 1536 AD-converters, and a SIMD array of 1536 bit-serial processors with corresponding memory. The SIMD processor array can deliver more than 100 GOPS sustained and the on-chip pixel-analyzing rate can be as high as 4 Gpixels/s. Experimental results showing low noise and a good digital to analogue noise isolation are shown.

I. INTRODUCTION

The sensor presented is a further development and extension of the previous sensors LAPP [1], PASIC [2], MAPP2200 [3], and MAPP2500 [4]. It is a general-purpose high-speed machine vision sensor fabricated in a standard 0.35 μm triple metal CMOS process. The sensor integrates on one chip a multi-resolution sensing area, containing one area part and one high-resolution line part (HiRes), a column parallel AD-conversion, and a column parallel single-instruction multiple-data (SIMD) processor. The processor implements image oriented instructions and delivers more than 100 GOPS sustained. The processor architecture can be used for a variety of image processing tasks like filtering, template matching, edge detection, run-length coding, and line-scan shading correction. The multi-resolution concept and the high level of integration enable low cost high performance machine vision systems.

II. ARCHITECTURE

Fig. 1 depicts the high-level chip block diagram. The core of the chip consists, with a few exceptions, of a linear array of 1536 columns butted together in the x-direction. Each column contains a part of the multi-resolution sensor area, analogue readout and AD-conversion/thresholding, and processor and registers. The instruction decoding and digital control, analogue switch control (AS), and pixel array row addressing are placed on the left side of the core. The dataport logic (DL), AD-conversion logic (AL), and the DAC used for AD-conversion are placed on the right side.

The chip has a simple interface where instructions are sent on a 24-bit bus, B , synchronised with the main chip clock, E , running at 33 MHz. The chip interprets and executes instructions within a single clock cycle, in most cases. Instructions sent to the core array constitute a typical example of a SIMD processor architecture. That is, a single instruction is distributed to all columns and each column executes the instruction using its respective column data. The instruction bus, B , can be turned around by changing

the read-write signal, RW , and status information can instead be read. A fraction of the registers in the core array are augmented with a dataport capability. This gives the possibility to stream out data on a 32-bit wide bus, DP , controlled by a separate clock, S , at 1.1 Gbit/s.

A. Analogue Part

The analogue part of one column is shown in the left part of Fig. 2. The multi-resolution sensor area is made up of two different parts. The area part has 1536 columns and 512 rows with 9.5 μm square pixels and the line part (HiRes) has 3072 columns and a set of rows of a high and narrow pixel. The pixel array has 1536 vertical buses for reading out the pixel voltages. To connect the HiRes rows to the vertical buses each HiRes row has two row addresses, making it possible to first read out the even columns and then the odd columns. Each of the vertical busses in the pixel array connects to an offset compensated and digitally programmable gain amplifier (OCDPGA) circuit. The OCDPGA has four gain settings, ranging from 1 to 4. The pixels are standard three transistor active pixels. The fill factor, being the pixel-area not covered by metal and n+-diffusion, is calculated to 60 % for the array pixels and 80 % for the HiRes pixels. The voltage levels at the gate of the reset transistor are controlled from the outside of the chip making both soft reset and hard reset possible [5]. Hard reset has been used due to the demand for high speed and no image lag.

The sensor area (including HiRes) is fully addressable via the row address decoder and the column selectability of the dataport. This enables multiple windowing and region of interest readout. Further the row address decoder allows for different integration times for different rows.

The AD-conversion is column parallel and of single slope




Fig. 1. High-level block diagram.

converter type [2]. This is implemented with a comparator and 8-bit memory in each column, and one common counter that feeds a DAC and the 8-bit in-column memories. The counter counts from 0 to 255, making the DAC produce a voltage ramp. When the ramp exceeds the voltage from the OCDPGA circuit the comparator switches and the counter value is loaded into the column memory (*ADREG* in Fig. 2). The ADC resolution is programmable from 3 to 8 bits, with a fast pseudo conversion (FPC) option. When using FPC the counter step-size is doubled after reaching value 64, thereby decreasing the greyscale resolution of the highly illuminated parts of the image. The AD-conversion is clocked at 33 MHz resulting in 7.7 μ s per 8-bit conversion and 4.8 μ s per 8-bit FPC conversion. An advantage with this ADC topology is the ability to perform fast thresholding, a crucial operation in many image processing algorithms. Besides making the DAC produce a ramp, it is possible to load a digital value into the DAC and thereby produce a voltage for the threshold operation. The fast thresholding combined with the parallel processing results in an internal pixel analysing rate of 4 Gpixels/s for simple image processing tasks.

The 8-bit DAC is of thermometer decoded current steering type [6] and has been temperature stabilized. Furthermore, the DAC features a 256-step programmable gain and a 256-step programmable offset, making it possible to do an automatic calibration of the AD-conversion. If the calibration is performed under a well-defined illumination it is possible to make different sensors have the same photo response. The programmable gain and offset can also be used together with the on-chip digital processor to implement a dithering scheme making it possible to do AD-conversions with more than 8 bits of resolution. A 10-bit conversion takes 32 μ s, and a 10-bit pseudo conversion takes 20 μ s.

B. Digital Part

The processor and registers constitute a high-performance image processing unit by the use of massive parallelism and the implementation of image oriented instructions. The processor in each column is bit-serial allowing extremely high-speed binary image processing. The bit-serial approach is also flexible while it allows variable word lengths and data formats to be used for grey image processing. For example, for each row, a 3×3 Gauss filtering takes $8b$ cycles and a 3×3 Sobel filtering takes $11b$ cycles, where b is the number of bits used. Furthermore, a multiplication takes $3b^2$ cycles and can for instance be used for shading correction in a line-scan application.

A column in the digital part of the sensor is illustrated in the right part of Fig. 2 and consists of the following parts. *PD* is a 1-bit register that holds the thresholded value of the column and *ADREG* is an 8-bit register that holds the AD-converted column value. Below this is the processor part consisting of a global logical unit (*GLU*), neighbourhood logical unit (*NLU*), and a point logical unit (*PLU*) with status registers and 16 accumulators. Further below are the global feature extractions *COUNT* and *GOR*. Finally come the general registers consisting of 96 bits per column (*RREG*) and 16 bits per column (*SREG*), where *SREG* is

also part of the high-speed dataport. Each column carries a 1-bit ALU bus that is the intra-column communication channel between the different blocks in the digital part.

The simplest type of instructions is the PLU-instructions that perform Boolean operations on the column data. 16 accumulators and special instructions for addition and subtraction support high-speed greyscale image processing. Typical for image processing is the fact that the result depends on its neighbours, such as median filtering or template matching. The NLU was designed for this purpose allowing a single instruction for a three-column median filtering or template match. Note that the NLU and PLU are tightly integrated indicating that a PLU-instruction can only be performed by passing data through the NLU. The result of an NLU/PLU-instruction is stored in one of the 16 accumulators. Since all NLU/PLU-instructions can be performed in a single clock cycle the arithmetic performance is very high exceeding 100 GOPS. The GLU was added to circumvent the problems associated with global instructions and SIMD architectures. The GLU provides a set of instructions where each column result depends on all columns input data.

The global feature extraction units, *COUNT* and *GOR*, operate on accumulator 0 in all PLUs. The *COUNT* feature outputs a binary value equal to the number of ones in accumulator 0 in all columns. The *GOR* feature is a single-bit result operation performing a global-OR on accumulator 0 in all columns. The results from *GOR* and *COUNT* are available at the *B* bus together with other status information. In addition, the *COUNT*-value can be read from the dataport and the *GOR* result is available on an external pad.




Fig. 2. The column architecture.

III. IMPLEMENTATION

A chip photograph is shown in Fig. 3. The chip measures 16.8 mm × 11.2 mm and comprises 5.8 M transistors. It is implemented in a 0.35 μm 3.3 V/5 V 3-metal 2-poly standard CMOS process. Both the analogue and the digital parts of the sensor are powered with 3.3 V. The large chip dimensions, mixed signal nature, and the use of relatively high-speed synchronous control present many design challenges. Special attention has for instance been focused on mixed-signal noise isolation, large distance signalling, synchronisation, and acceleration techniques for parallel feature extraction logic.

IV. EXPERIMENTAL RESULTS

For evaluation purposes the sensor was mounted chip-on-board on an evaluation board. An integrating sphere and a stabilised DC light provide uniform illumination. Many of the test results from the AD-conversion and the array pixels are listed in Table 1. For most measurements the DAC swing was set to 1 V, and 8-bit AD-conversions were performed.

The global DAC and the column ADC was characterised separately and the DNL/INL results are given in Table 1. For the entire signal chain, i.e. from photons to digital value, an extremely low photo response non-linearity was measured.

The FPN in dark and gain FPN was measured using an ADC swing of 1 V. The results are given in Table 1. Fig. 4 shows global and local (9×9 neighbourhood) total FPN at different grey levels.

There are two possible ways to read out the reset value from a pixel. Either while the reset transistor in the pixel is still turned on, or after turning the reset transistor off. Having the reset transistor on results in lower temporal noise since the noise from the reset transistor then becomes lowpass filtered via the pixel source follower readout, but it also results in higher FPN due to charge injection and clock feed-through from the reset transistor. Fig. 5 shows the measured ADC input referred temporal noise as function of mean ADC input signal with reset on. Calculating the theoretical RMS reset noise in the pixel [5] and adding the simulated temporal noise from the pixel source follower and the comparator result in 0.81 mV compared to the 0.90 mV measured. Another observation concluded from




Fig. 3. Chip photograph.

measurements is that heavy digital activity, including fast I/O, only affects the temporal noise marginally. An increase of 0-0.03 LSB can be measured depending on readout mode. No effect on FPN was measured. This verifies that noise isolation in the sensor is very good.

The conversion gain, from collected electrons to ADC input, was calculated by applying Poisson statistics on the measured temporal noise in Fig. 5 [7], see Table 1. From the conversion gain the pixel capacitance was calculated to 7.5 fF. The absolute spectral response was measured for the pixel array using a focused light source with variable wavelength, controllable in 5 nm steps, see Fig. 6. Peak QE×FF is 27 % and occurs at 610 nm; using a FF of 60 % a peak QE of 45 % is obtained.

To validate the increased horizontal resolution of the HiRes lines compared to the area pixels the modulation transfer function (MTF) was measured [8]. The result of the MTF measurement was 0.4 at the Nyquist frequency for the HiRes lines, and 0.45 at the Nyquist frequency for the area pixels. Since the Nyquist frequency is twice as high for the HiRes lines compared to the area pixels, it can be concluded that the horizontal resolution of the HiRes lines is almost twice as high.

The HiRes pixels behave similar to the array pixels. Dark current was however 60 % higher for the HiRes pixels, compared to the normal array pixels. This is attributed to the larger perimeter and area of the HiRes photo diodes [9]. Temporal noise was slightly lower due to the larger diode capacitance in the HiRes pixels.

The power consumed for a typical high-speed application is 0.68 W. This is distributed as 24 % in the digital part (including IO), 20 % in the global DAC, and the remaining 56 % in the analogue part. It should be noted that the analogue part has been designed to work at half of the nominal bias current (set by external resistors) which effectively halves the analogue power, thus trading lower speed for lower power.

TABLE I
CHARACTERISTICS FOR THE DAC, THE ADC, AND THE ARRAY PIXELS

Parameter	Value
Global DAC DNL, (1 V)	RMS=0.043, max=0.12 LSB
Global DAC INL, (1 V)	RMS=0.045, max=0.13 LSB
ADC DNL, (1 V)	RMS=0.031, max=0.35 LSB
ADC INL, (1 V)	RMS=0.31, max=0.92 LSB
Photo response INL	RMS = 0.09 %, max = 0.16 %
FF (array pixels)	60 %
QE	45 % at 610 nm
Pixel capacitance	7.5 fF
Dynamic Range	62 dB, (65 dB reset on)
Dark signal	0.35 LSB/ms at 60 °C
Conversion gain	14.5, 62.0 $\mu\text{V}/e^-$, (OCDPGA gain: G=1, G=4)
Temporal noise (RMS, in dark)	0.90, 3.99 mV, (G=1, G=4)
Temporal noise (RMS, in dark, pixel reset on)	0.70, 2.76 mV, (G=1, G=4)
FPN (dark, entire array)	0.15, 0.22 LSB, (G=1, G=4)
Gain FPN (RMS, Global - entire array)	1 %
Gain FPN (RMS, Local - 9x9 pixels)	0.58 %
Image lag	Below measurement limit

V. CONCLUSION

A multi-resolution general-purpose high-speed machine vision sensor with on-chip image processing capabilities has been presented. The computational power of the SIMD processor array, 100 GOPS, together with the high-speed image sensor part makes the continuous internal pixel-analyzing rate as high as 4 Gpixels/s for simple machine vision applications.

The work presented shows that it is possible to integrate high-performance programmable digital image processing circuits with a high-speed CMOS image sensor, and still achieve very low noise.

ACKNOWLEDGMENT

The authors wish to acknowledge Calle Björnlert, Jacob Wikner, Jürgen Hauser, and Magnus Engström for implementing special parts of the design. Also Ola Petersson is acknowledged for the design and implementation of the sensor evaluation board.

REFERENCES

- [1] R. Forchheimer, and A. Ödmark, "A Single Chip Linear Array Picture Processor", *Proc. of SPIE*, Vol. 397, pp. 425-430, Jan. 1983.
- [2] K. Chen, M. Afghahi, P-E. Danielsson, and C. Svensson, "PASIC: A Processor-A/D converter-Sensor Integrated Circuit", *Proc. IEEE Int. Symposium on Circuits and Systems (ISCAS'90)*, Vol. 3, pp. 1705-1708, May 1990.
- [3] R. Forchheimer, P. Ingelhart, and C. Jansson, "MAPP2200 - A second generation smart optical sensor", *Proc. of SPIE*, Vol. 1659, pp. 2-11, Feb. 1992.
- [4] M. Gökstorp, and R. Forchheimer, "Smart Vision Sensors", *Proc. IEEE Int. Conference on Image Processing (ICIP'98)*, Vol. 1, pp. 479-482, Oct. 1998.
- [5] H. Tian, B. Fowler, and A. El Gamal, "Analysis of Temporal Noise in CMOS Photodiode Active Pixel Sensor", *IEEE J. Solid-State Circuits*, Vol. 36, No. 1, pp. 92-101, Jan. 2001.
- [6] J. J. Wikner, "Studies on CMOS Digital-to-Analog Converters", Ph.D. thesis No. 667, Linköping University, Sweden, 2001.
- [7] B. Fowler, A. El Gamal, D. Yang, and H. Tian, "A Method for Estimating Quantum Efficiency for CMOS Image Sensors", *Proc. of SPIE*, Vol. 3301, pp. 178-185, April 1998.
- [8] C-S. S. Lin, B. P. Mathur, and M-C. F. Chang, "Analytic Charge

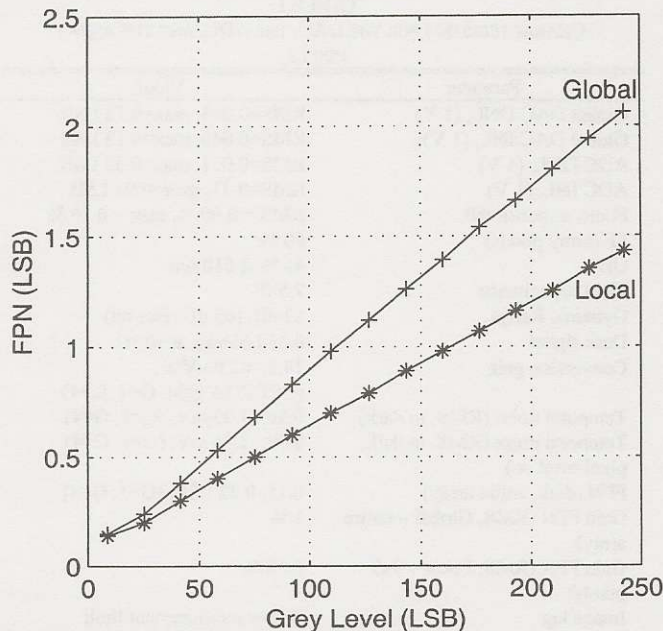


Fig. 4. Measured global and local total FPN vs. grey level.

Collection and MTF Model for Photodiode-Based CMOS Imagers",

IEEE Trans. Electron Devices, Vol. 49, No. 5, pp. 754-761, May 2002.

- [9] I. Shcherback, A. Belenky, and O. Yadid-Pecht, "Active Area Shape Influence on the Dark Current of CMOS Imagers", *Proc. of SPIE*, Vol. 4669, pp. 117-124, April 2002.




Fig. 5. Measured temporal noise vs. ADC input mean (reset on).

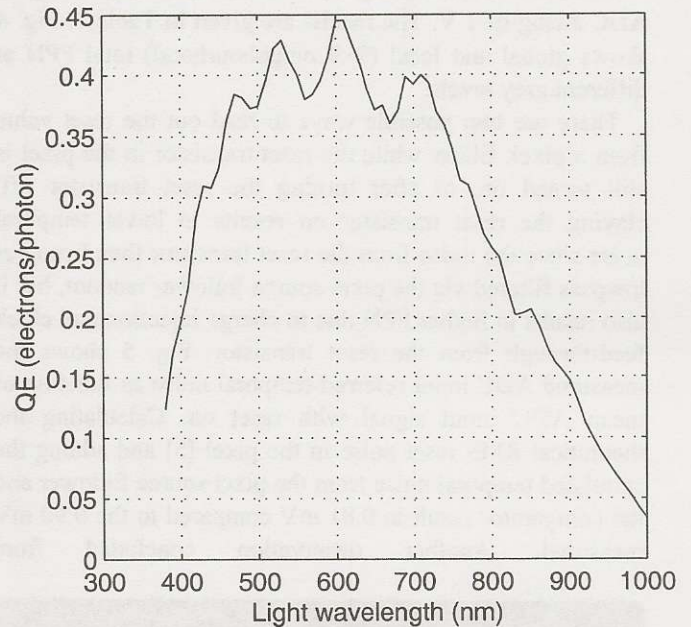


Fig. 6. Measured quantum efficiency vs. light wavelength.