High-Frame Rate Low-Noise Global Shutter CMOS Image Sensor for High-Speed Machine Vision

Abhinav Agarwal[1], Jatin Hansrani[1], Kazuhisa Suzuki[2], Karthik Venkatesan[4], Wilson Law[1], Varun Shah[4], Kai Ling Ong[1], Danny Marine[1], Oleksandr Rytov[4], Tim Lu[1], Neil Kumar[1], Edward Enriquez[1], Liviu Oniciuc[1], Sam Bagwell[4], Loc Truong[1], Anders Andersson[4], and Radu Corlan[3]

[1]Forza Silicon (AMETEK Inc.), Pasadena, California, USA
[2]Forza Silicon (AMETEK Inc.), Tokyo, Japan
[3]Vision Research (AMETEK Inc.), Bucharest, Romania
[4]Not affiliated with Forza Silicon now
[1]**Corresponding Author:** abhinav.agarwal@ametek.com

***Abstract*—We present a low-noise, high-frame-rate global shutter CMOS image sensor with UHD resolution (3840×2160), targeting high-speed machine vision applications. The sensor supports video capture at up to 1100 FPS at 12-bits and 1600 FPS at 8-bits at full resolution. Fabricated in a 65 nm, 4-metal BSI process, the imager features a 5μm voltage-domain global shutter pixel with dual-gain capability for improved dynamic range and a read noise of 3e$^-$ in high-gain mode at maximum frame rate operation. For compact camera integration and low power consumption, the sensor is designed to stream video through 16 CML data ports, each operating at 7.44 Gbps, achieving a total aggregate throughput of 119 Gbps.**

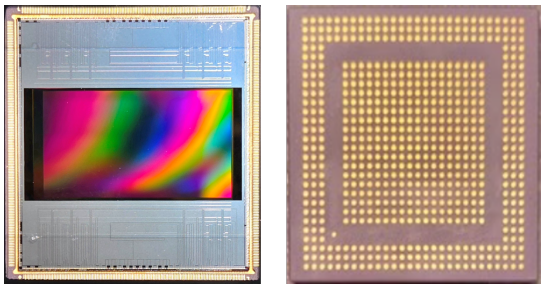

Fig. 1. Image Sensor Die Photo with 268-Pin LGA Package

I. Introduction

High-speed machine vision systems are increasingly deployed in applications such as industrial inspection, automotive safety testing, advanced microscopy, and slow-motion sports analysis, where capturing fast-moving subjects with high temporal resolution and minimal distortion is essential. A key enabler of these systems is the image sensor. In particular, a high-frame-rate, low-noise, global shutter CMOS image sensor is critical for accurately capturing transient events without motion artifacts or rolling shutter distortions.

While traditional rolling shutter CMOS image sensors offer advantages such as lower temporal noise and simpler pixel architecture, they suffer from temporal skew and motion blur—limitations that reduce their effectiveness in scenes with rapid motion. In contrast, global shutter (GS) pixel architectures enable true snapshot imaging by simultaneously exposing all pixels, thus eliminating temporal distortion. However, these architectures often introduce trade-offs, including increased dark temporal noise, reduced fill factor, and diminished dynamic range, primarily due to the added complexity and in-pixel storage elements required for global transfer operation.

Recent advancements in pixel design, low-noise analog front ends, and high-speed ADCs have significantly closed the performance gap between global and rolling shutter image sensors. These innovations now allow GS CMOS image sensors to retain their temporal advantages while achieving competitive image quality, thereby expanding their applicability in high-performance machine vision systems where both speed and image quality are critical.

In this work, we present a high-frame-rate, low-noise global shutter CMOS image sensor optimized for high-speed machine vision applications (Figure 1). The design focuses on maximizing output frame rate while maintaining image quality through low noise, minimized horizontal smearing, and support for short integration times ($< 2\,\mu$s).

To further enhance versatility, the sensor includes column and row windowing as well as port concentration modes, which enable data throughput through a reduced number of output ports. Multiple spatial subsampling modes, such as Bayer skipping and Bayer sub-sampling, are supported to increase frame rate at lower resolutions. The readout architecture is configurable for both correlated double sampling (CDS) and non-CDS modes, allowing applications to balance noise and speed, especially during windowed readout. The sensor also supports selectable output bit depths—8-bit, 10-bit, and 12-bit—enabling additional optimization of frame rate based on application-specific requirements.

II. Readout Architecture

The sensor block diagram (Figure 2) illustrates a top-bottom readout architecture in which six rows are read out concurrently. The pixel array readout is divided across two halves, with each side further partitioned into 12 superblocks, where each superblock reads 320 columns by 3 rows. The sensor employs a 12T voltage-domain global shutter pixel architecture with dual-gain capability. Within each pixel, two high-density deep trench capacitors store the reset and signal levels

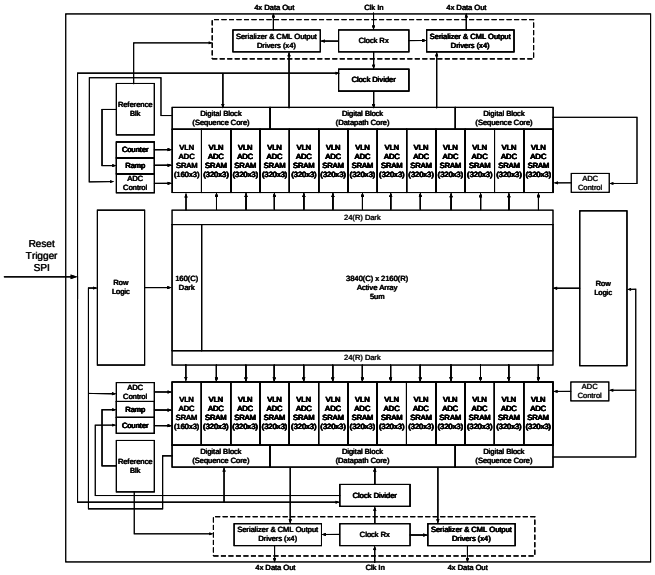

Fig. 2. Image Sensor Block Diagram

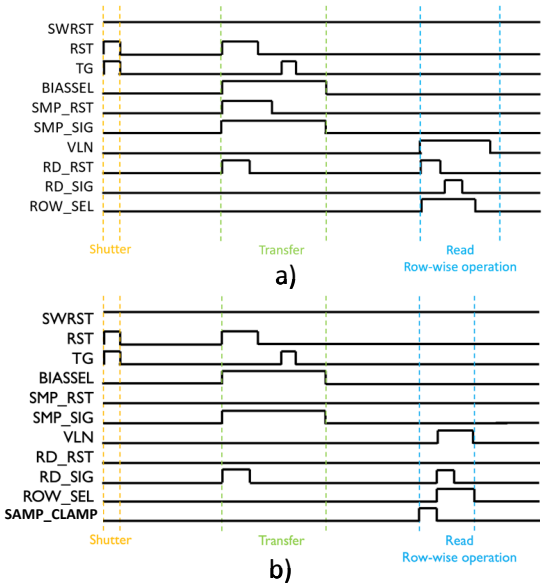

Fig. 3. Simplified Pixel Timing for a) CDS and b) non-CDS Operating Modes

required for CDS readout. The capacitor sizes are carefully optimized to minimize kTC noise—a dominant component of overall pixel noise—while maintaining sufficient fill factor to preserve QE. A common source follower configuration is used for pixel readout to suppress FPN and reduce the number of pixel output lines.

To enable high frame rates, the row time must be minimized; this is achieved using a ping-pong single-slope ADC architecture, which pipelines sampling and ADC conversion. Figure 4 illustrates the analog readout signal chain of the image sensor. In this implementation, the row time is primarily limited by the time required to sample both the reset and signal levels into the ADC during rolling readout. To trade off dark temporal noise for increased frame rate, a non-CDS readout mode is also supported, wherein a local ADC clamp can auto-zero the preamplifier and comparator—thereby reducing the reset sampling time. A high level timing diagram illustrating pixel operation in CDS and non-CDS modes is shown in Figure 3.

To support concurrent readout of six rows, each pixel column includes six pixlines. By default, even-numbered pixlines connect to the bottom readout chain, while odd-numbered pixlines connect to the top chain. To enable both dual-sided and single-sided readout modes—useful for port concentration—dynamically programmable switches are inserted along the even and odd pixlines. These switches allow each pixline to be routed to either the ADC or VLN, depending on the selected configuration.

In dual-sided readout, the VLN associated with a given pixline resides on the opposite side from its ADC. In contrast, during single-sided readout, the programmable switches enable all pixlines to access both the ADC and VLN from the same side of the readout chain. One half of the readout circuitry can be completely powered down, resulting in significant power savings.

III. High-Speed Interface Design

Operating at UHD resolution at high frame rates, the image sensor generates a substantial volume of data, necessitating careful design, simulation, and optimization of the entire high-speed signal chain. To ensure seamless integration into the final camera system, several product-level constraints—such as a compact form factor, a reduced number of FPGAs in the data downlink, and low-power operation—shaped the high-speed clock distribution and SerDes architecture. The number of high-speed data ports was limited to 16 to align with system-level FPGA constraints, making it critical to maximize data throughput per port. Additionally, implementation on a 65nm process node with only four metal layers imposed further limitations on routing resources, driving the need for precise impedance control and signal integrity optimization across the high-speed paths.

The high-speed signal chain comprises several key components, including the clock receiver, clock distribution network, serializer, and CML output driver (Figure 5a). A high-speed differential input clock operating at 3.72 GHz is supplied to both the top and bottom sides of the sensor readout. This clock is distributed through a high-speed CML-based clock distribution network to all eight CML output data ports—four located on the left and four on the right—ensuring synchronized data transmission across the entire high-speed interface.

To achieve UHD resolution at 12-bit depth and 1100 FPS, each of the 16 CML output ports (eight on the top and eight on the bottom) operate at 7.44 Gbps. The clock receiver is a CML buffer comprising multiple amplification stages, designed to boost the small-swing differential input clock to a level sufficient for the clock distribution network. This receiver is placed at the center of a T-shaped clock tree, with each branch spanning 3.4 mm.

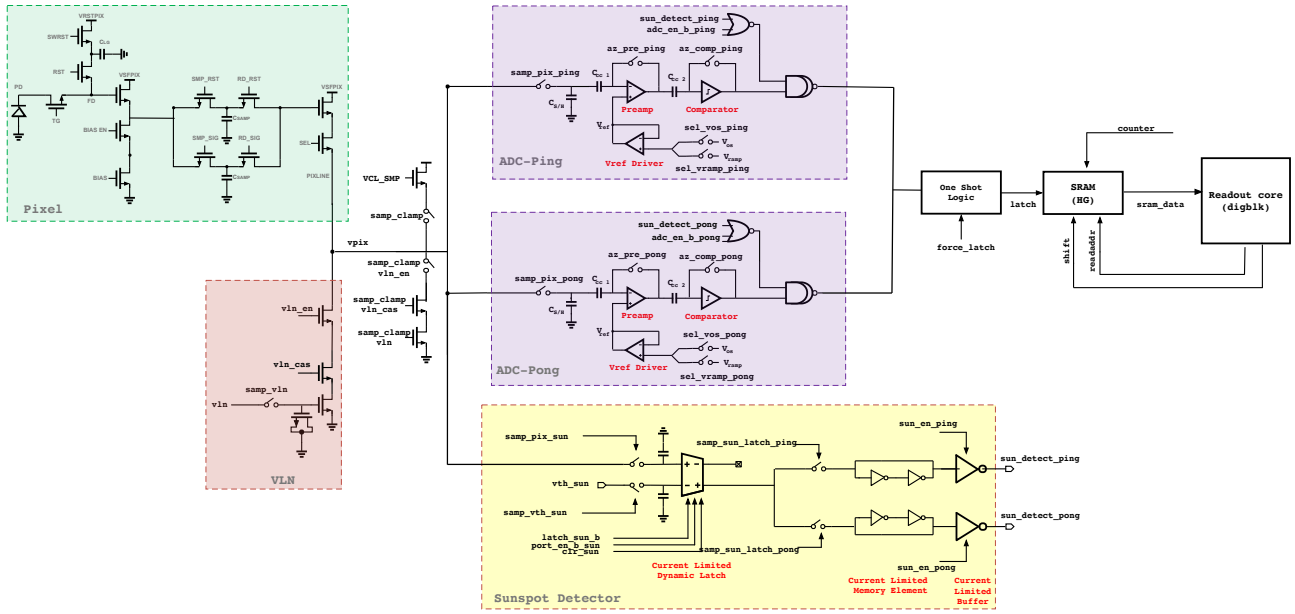

Fig. 4. Complete Analog Signal Chain (from pixel to readout core digital block)

The number and placement of CML buffers along the clock distribution network are optimized through theoretical analysis—accounting for attenuation per unit trace length—and validated by post-layout simulations to ensure clock fidelity at each stage. A key challenge is common-mode imbalance in the differential clock, caused by random device mismatch and trace impedance variations. If uncorrected, this imbalance can degrade the differential signal, potentially leading to clock failure after several stages.

To mitigate this, AC coupling is introduced periodically in the clock path. While effective in restoring common-mode balance and filtering low-frequency noise from earlier stages, AC coupling introduces additional parasitic capacitance due to the coupling capacitor (realized using MOS devices) and the input capacitance of the CML stage. This imposes a practical limit on the number of AC coupling stages. The high-pass filter formed by the AC capacitor and bias resistor is designed with a cutoff frequency at one-fourth of the operating clock frequency. However, larger capacitors add parasitic loading, and higher-value resistors increase thermal noise, necessitating careful sizing trade-offs to balance signal integrity and noise performance.

To further minimize substrate-related uncertainties, additional metal shielding is routed beneath the high-speed signal traces to provide controlled coupling and reduce undesired coupling to the substrate. While layout extraction tools often model the substrate as a zero-impedance ground, in practice, it exhibits a distributed RC behavior that depends heavily on layout geometry and process parameters. To mitigate these effects, the thick top metal layer—available in the 4-metal process—is utilized for routing critical signals, thereby reducing resistance and parasitic capacitance. This approach helps maintain stable duty cycles, which is essential for robust DDR operation of the serializer.

As shown in Figure 5a and 5b, the clock generation block

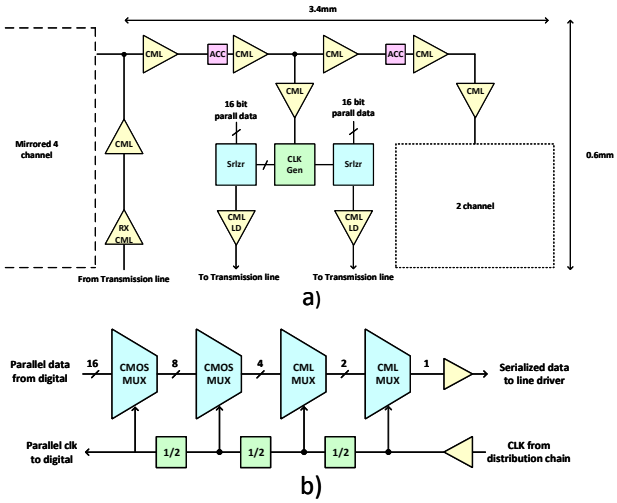

Fig. 5. a) High-Speed Signal Path (from clock receiver to serialized data output), b) High-Speed 16-to-1 Serializer Architecture

produces a divided clock for the 16:1 serializer, shared between two serializer units placed symmetrically on the left and right sides. High-speed clock domains utilize CML-type dividers, whereas CMOS TSPC-based dividers are employed for low-speed clock generation. In stages 1 and 2, the divided clock signals are converted to CMOS levels. Stage 1 clock is also used as a parallel clock routed to the digital block. This ensures synchronization between the digital parallel data and the serializer input.

The serializer is structured as a multi-stage multiplexer (MUX) tree optimized for both performance and area efficiency:

- **CML Stages:** The first stage consists of eight units of 2:1 CMOS MUXes implemented using TSPC logic for low-

TABLE I
FRAME RATE SCALING VS BIT-DEPTH ACROSS VARIOUS PORT CONFIGURATIONS (CDS MODE)

Resolution	Bit-Depth	8top+8bot	6top+6bot	4top+4bot	2top+2bot	8[one-side]	6[one-side]	4[one-side]	2[one-side]
3840×2160 (UHD)	*12-bit*	1,114	840	563	283	561	423	283	142
	10-bit	1,348	1,018	680	340	679	513	343	171
	8-bit	1,680	1,267	850	426	847	639	428	214
2560×1440 (QHD)	*12-bit*	2,172	1,853	1,250	630	1,099	938	632	318
	10-bit	3,000	2,266	1,514	760	1,518	1,147	766	385
	8-bit	3,000	2,812	1,886	948	1,518	1,423	954	480
1920×1080 (Full HD)	*12-bit*	2,873	2,873	2,192	1,108	1,460	1,460	1,114	563
	10-bit	3,968	3,968	2,653	1,339	2,016	2,016	1,348	680
	8-bit	3,968	3,968	3,306	1,673	2,016	2,016	1,680	850
1280×720 (HD)	*12-bit*	4,241	4,241	4,241	2,440	2,172	2,172	2,172	1,250
	10-bit	5,857	5,857	5,857	2,957	3,000	3,000	3,000	1,514
	8-bit	5,857	5,857	5,857	3,683	3,000	3,000	3,000	1,886
640×480 (VGA)	*12-bit*	6,214	6,214	6,214	6,214	3,219	3,219	3,219	3,219
	10-bit	8,582	8,582	8,582	8,582	4,446	4,446	4,446	4,446
	8-bit	8,582	8,582	8,582	8,582	4,446	4,446	4,446	4,446

power, compact operation at low frequencies followed by second stage with four units of 2:1 CMOS MUXes to further serialize data.

- **CMOS Stages:** Third stage has two units 2:1 CML MUXes followed by the fourth stage consisting of just one unit of 2:1 CML MUX with a one-clock delayed data path generated to support post-emphasis functionality, which improves signal integrity at high data rates.

The serializer output connects to a CML output driver with a 50 Ω resistive load for impedance matching with the transmission line. An integrated post-emphasis circuit with programmable coefficients compensates for channel loss, enhancing high-frequency signal quality. To ensure robustness, reverse-biased diodes are placed at each differential output for ESD protection. Additionally, RLC elements are inserted in the supply and ground paths during post-layout simulations to model bonding wire effects. The transmission channel is characterized using S-parameter data to accurately capture high-speed behavior and guarantee signal integrity.

IV. RESULTS AND SUMMARY

Table I summarizes the scaling of maximum achievable frame rates for various standard video resolutions for different bit depths and various port configurations. The sensor's readout architecture is designed to support higher frame rates when operating at reduced resolution formats. At UHD resolution, frame rate scales with bit depth—from 12-bit to 10-bit to 8-bit—due to limitations imposed by the total output bandwidth. For lower resolution formats, reducing bit depth from 12-bit to 10-bit results in increased frame rates; however, there is no frame rate improvement when moving from 10-bit to 8-bit, as the row time becomes limited by the signal sampling duration rather than output bandwidth. To optimize power consumption for application-specific frame rate targets, the number of active output data ports can also be reduced without sacrificing performance. Detailed specifications of the CIS sensor are presented in Table II and a full-resolution 12-bit monochrome image captured at 1100 FPS is shown in Figure 6.

TABLE II
SPECIFICATION TABLE

Parameter	Specification
Technology	1P4M, 65nm BSI
Resolution	3840×2160 (Mono and Color)
Pixel Pitch	5 μm
Shutter Type	Voltage Domain Global Shutter
Linear Full Well	8100 e^- (HG), 41 300 e^- (LG)
Conversion Gain	170 μV/e^- (HG)
Total Dark Temporal Noise	3.04 e^- (HG)
Row Noise Correction	On-Chip
Max Frame Rate	73,809 FPS (640×6)
ADC Resolution	8b / 10b / 12b
PRNU	2.2% (HG), 2.1% (LG)
DSNU	14.6 e^- (HG), 128 e^- (LG)
Dynamic Range	68 dB (HG)
Max SNR	39 dB (HG)
Quantum Efficiency	67%@430 nm (blue) 73%@530 nm (green) 61%@620 nm (red)
PLS	< -100 dB
Power Consumption	5.5 W
Data Ports	16×outputs @ 7.44 Gbps
Die Size	23.55 mm (H) × 25.6 mm (V)
Trigger Control	Internal and External
Serial Interface	SPI

Fig. 6. Monochromatic Capture at 12-bit Resolution at 1100FPS