Energy-Efficient Motion-Triggered IoT CMOS Image Sensor With Capacitor Array-Assisted Charge-Injection SAR ADC

Kyojin David Choo, *Member, IEEE*, Li Xu, *Student Member, IEEE*, Yejoong Kim, *Member, IEEE*, Ji-Hwan Seol, *Member, IEEE*, Xiao Wu, *Member, IEEE*, Dennis Sylvester, *Fellow, IEEE*, and David Blaauw, *Fellow, IEEE*
(Invited Paper)

***Abstract*— We propose a low-power image sensor with a motion-based triggering feature for the Internet of Things (IoT) applications. The sensor supports the near-pixel [within analog-to-digital converter (ADC)] motion-detection mode run on a heavily subsampled frame (32 × 20 pixels, while consuming only 1.7 μJ/frame) to trigger full-array capture only when a significant change has happened in the scene. Also, to maximize energy efficiency and image quality of the sensor, we propose a column-parallel capacitor array-assisted charge-injection SAR ADC that achieves 10b operation with readout noise of 226 μV_{rms}, which can be lowered to 104 μV_{rms} using multiple sampling. The ADC is highly energy efficient with best-in-class energy consumption of 63.6 pJ/frame/pixel and ADC Figure of Merit (FoM) of 14.4 μV_{rms}·nJ in full-array capture mode.**

***Index Terms*— Analog-to-digital converter (ADC), charge-injection cell (ci-cell), ciSAR, CMOS image sensors, energy efficient, Internet of Things (IoT), low-power electronics, motion detection, multiple sampling, SAR.**

I. Introduction

As THE Internet of Things (IoT) is becoming increasingly integrated into our everyday life, the demand for energy-efficient image sensors is on the rise. IoT image sensors are often size constrained and have only small batteries for energy support, yet they are expected to last weeks on a single charge. Also, their image quality determines their usefulness. Thus, it is crucial to design for both low-power operation to extend the charging cycle and high-quality signal readout to maximize their utility.

Previously reported sensors [1], [2], along with this work, adopted motion-detection (MD) triggering of full-array capture where MD is performed on a heavily subsampled frame to enable continuous low-power operation. MD effectively limits energy-hungry full-array captures to cases where the activity is detected, providing the first-order reduction in power. However, to further optimize for power consumption, the full-frame capture energy itself, which is dominated by the readout circuit, needs to be addressed.

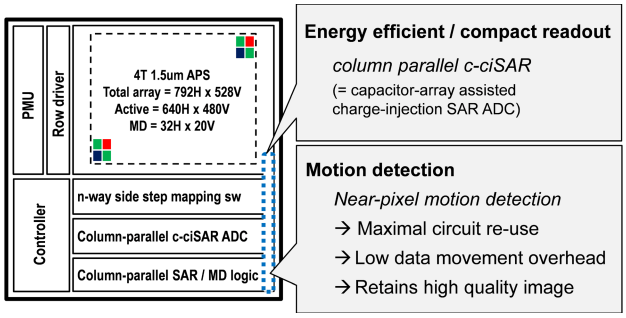

Fig. 1. Overall architecture of the proposed energy-efficient motion-triggered IoT image sensor.

Conventional image sensors use single-slope analog-to-digital converters (ADCs) with a 4T pixel structure, which provides excellent uniformity but is inherently energy-inefficient due to its linear counting operation. This is particularly a problem at a higher resolution where the energy involved in counting can quickly dominate the sensor energy. SAR ADCs are free from this problem and have superior energy efficiency but have traditionally been too large for incorporation into an image sensor column pitch. To achieve both high-quality image readout and energy-efficient low-power operation, we propose a column-parallel capacitor array-assisted charge-injection SAR ADC (c-ciSAR) structure. The ADC merges an area-efficient charge-injection cell (ci-cell) based digital-to-analog converter (DAC) [3] with a small capacitor array structure to extend its performance to 10b while improving energy efficiency. The ADC core fits into 3.1-μm column pitch even with the penalty of being fully designed using thick-oxide devices to control leakage current for low-power operation. The c-ciSAR enables significantly higher energy efficiency than a single-slope ADC, with a measured ADC Figure of Merit (FoM) [4] of 14.4 μV_{rms}·nJ (full-array mode).

In this article, we discuss the energy-saving aspects of a column-parallel c-ciSAR ADC and how it integrates into an IoT image sensor with the motion-trigger feature, shown in Fig. 1. The sensor contains video graphics array (VGA) (640 H × 480 V) active pixel array, column-parallel

Manuscript received May 20, 2019; revised August 1, 2019; accepted August 28, 2019. Date of publication September 25, 2019; date of current version October 23, 2019. This article was approved by Guest Editor Pedram Mohseni. *(Corresponding author: Kyojin David Choo.)*

The authors are with the Department of Electrical and Computer Engineering, University of Michigan, Ann Arbor, MI 48109 USA (e-mail: kjchoo@umich.edu).

Color versions of one or more of the figures in this article are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JSSC.2019.2939664

c-ciSAR, which is capable of near-pixel motion detection, peripheral power management unit (PMU), and programmable controller. The sensor behavior is software-defined via the on-chip controller to continuously monitor changes in the scene through subsampled frames and to trigger full-array single-shot capture when there is sufficient difference.

In Section II, we describe how ci-cell design expands to capacitor array-assisted charge-injection DAC (c-ciDAC) to perform energy-efficient 10b DAC operation. In Section III, we explain how c-ciDAC merges into c-ciSAR and how it benefits overall energy efficiency of image sensor readout operation. In Section IV, we cover how the motion-detection feature is implemented in c-ciSAR (near-pixel, in ADC). Then, we show measurement results and conclude this article.

II. Capacitor Array-Assisted Charge-Injection DAC

In general, SAR ADCs are considered to be highly energy efficient due to their algorithmic operation and their charge-based implementation using capacitor arrays [5]–[7]. However, they are seldom used in image sensors due to their relatively high complexity and susceptibility to mismatch compared to conventional single-slope ADCs. Also, compounded by the large layout size of their capacitor array DACs, to date, SAR ADC implementations in image sensors have rarely been truly column-parallel [8], [9].

Non-column-parallel design leads to energy inefficiency as single slice of ADC must operate multiple times faster than column-parallel design to allow for time multiplexing. This is because constant current drawing elements, namely pixel source follower and amplifier, require hyperlinear growth in current for linear speed (g_m/C) improvement. As a result, SAR ADCs in image sensors have not delivered their full energy efficiency.

For energy efficiency of the image sensor to improve, DACs in SAR ADCs must become smaller to fit into the narrow column pitch. Capacitor array structures exist that are more compact than conventional ones [10], [11]. However, even with these circuit structures, the narrow layout in image sensors renders fitting these capacitor arrays in the column pitch extremely difficult. Also, when these DACs become smaller, cap-to-cap mismatch starts to limit the accuracy. Correspondingly, capacitor array SAR ADCs must grow in size to achieve the required resolution and differential non-linearity (DNL) performance, which forces them to be drawn in multiple column pitches, reducing their energy efficiency [8], [12].

To overcome this limitation, our proposed solution takes an architectural approach to reduce the size and mismatch of DACs in SAR ADCs. We adopted ci-cell-based DACs [3].

A. Charge-Injection Cell

The ci-cell-based DAC offers reduced size and improved mismatch performance by introducing the concept of reusability to its operation, which is a great fit for area-limited image sensor application. A ci-cell functions basically the same as a current source stacked with a control switch, as shown in Fig. 2(a). When the switch is closed for a pre-determined duration to initiate charge-injection operation and then opened, a current is pulled from (or pushed out to) an output capacitor to induce a controlled voltage change. The robustness of the operation is improved by replacing the current source with a precharged capacitor [Fig. 2(b)] to introduce a tapering down current profile. This helps to reduce its sensitivity to the switch-on duration, which may be perturbed by jitter or supply noise. Both structures are essentially reusable as their internal voltages are reset when the switch is opened. Thanks to reusability, a ci-cell-based DAC can be as small as a single ci-cell because all voltage levels can be generated from the repeated operation of the ci-cell spread over time. At the same time, the DNL performance is drastically improved since those multiple charge-injection operations are virtually perfectly matched because the same ci-cell is used repeatedly.

Though reusable, ci-cells exhibit limited output impedance, which results in inconsistent charge output when the output voltage changes. As uniformity of charge injection across a wide output voltage range is crucial for obtaining the desired integral non-linearity (INL) performance, we added a cascode transistor (M2) underneath the switch transistor (M1) as shown in Fig. 2(c). Stacking the two transistors improves the output impedance of the cell greatly despite the introduction of a

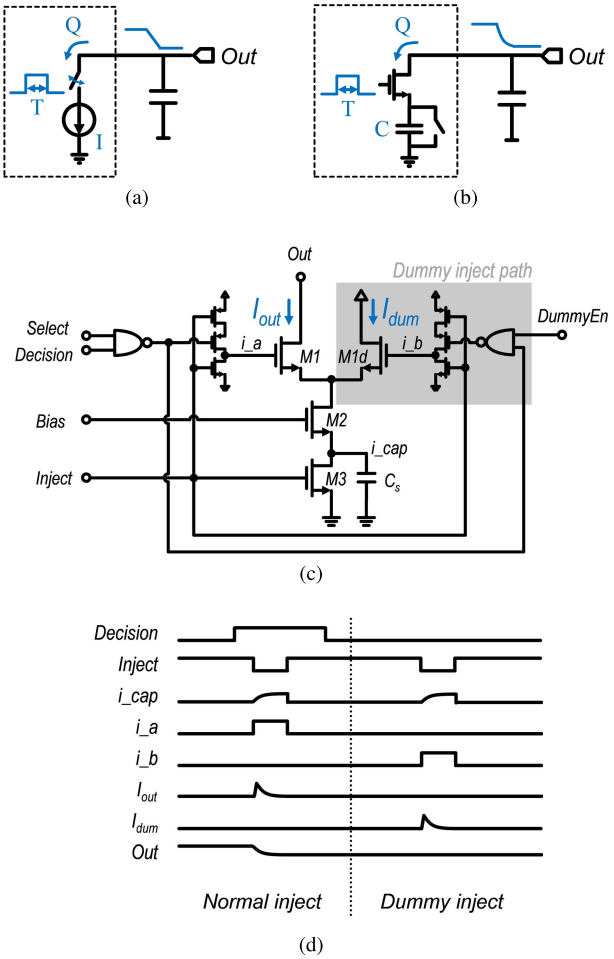

Fig. 2. Pull-down ci-cell. (a) Current source-based ci-cell. (b) Capacitor-based ci-cell. (c) Transistor-level diagram of the proposed ci-cell. (d) ci-cell operation timing diagram with Select = H and DummyEn = H. The ci-cell directs charge to Out node when Decision is H, and to VDD when Decision is L (dummy inject).

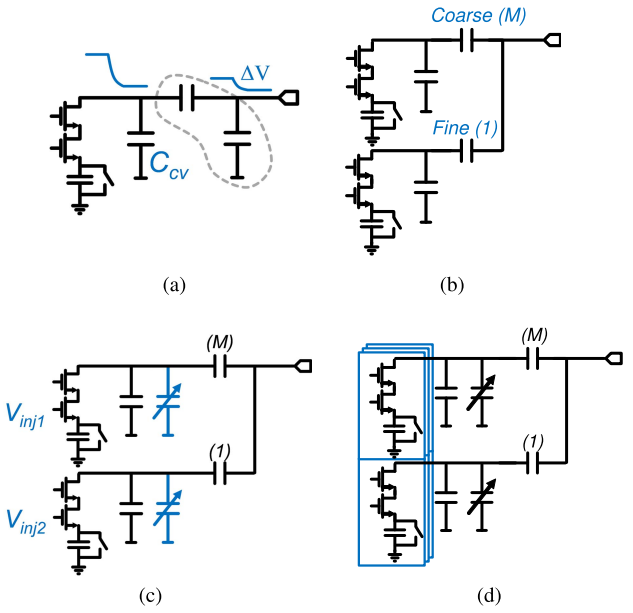

Fig. 3. ci-cell integration into c-ciDAC. (a) Added capacitor divider. (b) Added coarse/fine capacitor array. (c) Added PVT (coarse) calibration using V_{inj} and mismatch (fine) calibration using small calibration capacitor bank. (d) Added multiple ci-cells.

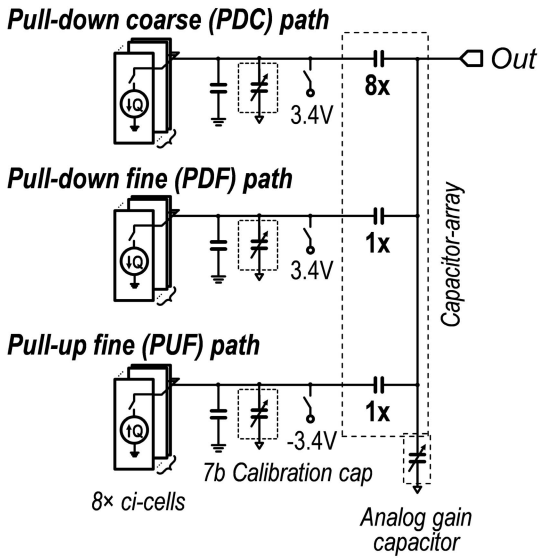

Fig. 4. Final structure of c-ciDAC.

parasitic node, which needed to be controlled to be small. The proposed ci-cell structure assumes the DAC includes multiple ci-cells, and hence extra logic gates are built in to apply gain control (Select) and the SAR decision (Decision).

Fig. 2(d) shows the detailed operation of a pull-down ci-cell. The inject signal pulses low to activate charge-injection operation. During each injection, the voltage on node i_cap rises from its reset value (Gnd) to the voltage controlled by the node Bias, whose value sets the charge-injection amount through M2. The charge created from the rising of i_cap transfers to the output capacitor and makes the voltage drop on the output node. When the ci-cell is deselected, the charge flows through dummy injection path if the DummyEn signal is high. The dummy injection path makes the current draw from the power supply to be decision-independent. This ensures a constant injection charge even with high impedance power supply regulation, which is often needed for low-power overhead in battery-operated devices.

B. Capacitor Array-Assisted Charge-Injection DAC

Though the proposed ci-cell structure, with its inherent perk of reusability, guarantees small and linear DAC implementation, there are remaining considerations to further improve the noise, energy efficiency, accuracy, and speed of the DAC. In this section, we will walk through how these performance aspects are addressed and improved in the c-ciDAC.

1) Capacitor Divider: In ci-cells, the noise charge is generated by the thermal noise on the capacitor C_s, and the SNR of the charge-injection operation can easily be described by the ratio of this noise charge to the signal output charge. The signal charge amount can be increased by controlling the ci-cell bias voltage or by increasing the size of the capacitor C_s and the SNR tends to improve with increased signal charge.

For the best SNR, the ci-cell should be set to output as much signal charge as possible. However, this would result in reduced resolution performance of the DAC since the desired LSB size increases if the output capacitor is fixed. To get around this problem, we introduced two series-connected capacitors, as shown in Fig. 3(a), to construct a capacitor voltage divider at the output. The leftmost capacitor (C_{cv}) still acts as a charge-to-voltage converter, where the swing on this internal node grows with increased signal charge. The subsequent series-connected capacitor divider reduces the output step size to the desired resolution level. Resorting to this structure is more economical than multiplying the C_{cv} size, which is already set at approximately 200 fF.

2) Coarse-Fine Capacitor Array: On the other hand, the reduced voltage swing from the capacitor divider now must meet the signal range requirement. Covering a wider voltage range conflicts with choosing a larger voltage division ratio with respect to noise and resolution performance. To resolve this conflict, we introduced an extra charge-injection path that has a division ratio close to one, as shown in the top part of Fig. 3(b), creating a coarse gain charge-injection path (ci-path). The extra coarse gain ci-path incorporates a larger (by M times) series capacitor, which also provides the necessary voltage division for the fine gain ci-path. The series capacitors are roughly sized in integer ratios of each other (M:1), and hence they compose a structure with a capacitor array merged with ci-cells, referred to as a c-ciDAC. There is no matching requirement for these capacitors as the calibration will correct for their respective gain on a per-column and per-path basis.

Introducing coarse/fine ci-paths to the c-ciDAC also reduces the DAC energy consumption, a desirable effect. If only the fine gain ci-path is used, 2^N activations are required to create 2^N DAC output levels. As each activation may consume up to several femto-Joules, which is smaller but still close to the flip-flop energy, involving an exponentially increasing number of activations per additional bit erases the gain of using

SAR ADCs over single-slope ADCs. But with a coarse/fine structure, the required number of ci-cell activations to create 2^N DAC levels scales down relative to M. This scaling enables the low-energy operation of a c-ciDAC. For our implementation, M of 8 was selected, and ci-cell activations were effectively reduced by 4.3×, including dual correlated double sampling (dual-CDS), which slightly hampers the reduction.

3) Calibration Mechanisms: Including multiple ci-paths levies calibration overhead on the c-ciDAC. Although leveraging reuse in the ci-cells reduces the impact of mismatch among themselves, they are prone to process, voltage, temperature (PVT) variations, which translate to DAC gain error. Having multiple paths can turn DAC gain error into a serious DNL error, and mismatch from the capacitor array further aggravates this issue. Therefore, we needed to incorporate a mechanism to calibrate for the ci-path gains to avoid DNL.

For calibration, we reserved two tuning knobs in the c-ciDAC design, as highlighted in Fig. 3(c). The first method is to use the biasing voltage of the ci-cells. For each ci-path, we apply different global bias voltage to correct for the collective skew of all column-parallel ci-paths, which is an effective measure against PVT variations. The second method is to add a calibration capacitor bank to the internal node of the ci-paths. As the amount of capacitance determines the conversion rate at which the injection charge will convert to voltage, controlling the capacitance has a direct impact on the ci-path gains. We designed a bank of digitally controlled 7b varactors that create small ΔC (0.25% of total capacitance per code) to fine tune this rate, which is effective at calibrating out small column-to-column mismatches. The capacitor bank is conservatively designed to cover 6σ of ci-cell local mismatch ($\sigma = 2\%$) plus array nonuniformity. Since the calibration capacitor bank only needs to provide monotonically increasing capacitance to code, it is laid out compactly with binarily sized capacitors without resorting to an area-hogging unit capacitor array. Details on how to execute these calibrations will be discussed in the next section.

4) Final Structure: For speed, we included multiple ci-cells per path to allow the DAC to operate with larger step sizes. This structure is illustrated in Fig. 3(d). In our actual implementation, as a low-power constraint on the IoT image sensor restricts the use of a fast running clock, generating control timing for ci-cells bottlenecks the overall sensor speed. Having multiple ci-cells per path helps greatly in this regard. Also, it provides the extra benefit of reducing the trimming range of the calibration capacitor bank because multiple ci-cells operating together output less mismatch charge in proportion to the signal charge than a single ci-cell.

The final c-ciDAC structure is shown in Fig. 4. Three ci-paths are implemented, pull-down course (PDC), pull-down fine (PDF), and pull-up fine (PUF), with relative path gains of 8×, 1×, and 1×, respectively. A pull-up and pull-down pair is necessary to perform binary search operation on a single-ended c-ciSAR implementation. Yet the coarse gain portion only has a pull-down ci-path because signals only tend to go down in image sensors with 4T pixels. This fact is the basis for several observed benefits when a downward linear search is performed instead of a binary search. As one of the benefits, the pull-up coarse (PUC) path can be omitted, which reduces the circuit complexity. Each path has 8 ci-cells and also 7b calibration capacitor banks. The DAC also includes an analog-gain (AG) capacitor at the output to provide the extra capacitor division ratio necessary for the AG function while preserving the DAC SNR.

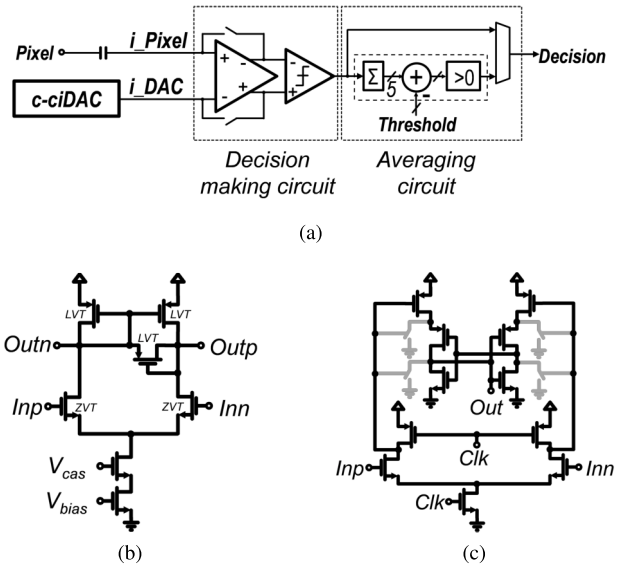

Fig. 5. c-ciSAR. (a) Circuit architecture (b) Preamplifier circuit. (c) Comparator circuit.

III. Column-Parallel c-ciDAC-Based SAR ADC for IoT Image Sensors

Fig. 5(a) shows the overall architecture of the c-ciSAR ADC. Each c-ciSAR utilizes c-ciDAC to create a local reference voltage, which is compared to the pixel voltage using the decision-making circuit for SAR operation. The decision-making circuit is comprised of a preamplifier and a clocked comparator, as shown in Fig. 5(b) and (c).

The preamplifier is an NMOS input-pair operational transconductance amplifier (OTA) structure designed to operate under 1.4-V supply and support 500-mV signal range. To achieve the required signal range, we used zero-vtNMOS devices for the input pair and low-vtPMOS as a mirror load. A clamping diode is inserted between Outn and Outp to limit the swing of the node to assist with decision speed and to eliminate significant current load change before and after the decision, where fully steered current pulls down the Outp completely to cut off the input-pair along with the tail current source.

One side of the input is from c-ciDAC, where it is a capacitive node by construction, and the other side has a capacitor inserted between itself and the pixel output. This structure decouples dc operating points of pixels and internal voltages of c-ciDAC from the dc operating point of the preamplifier, which is set by auto-zeroing. The preamplifier is designed to provide moderate gain (20 dB) to balance the settling speed and the reduction of the input referred noise of the comparator. Both Outp and Outn are connected to the differential input of the comparator to make a self-referenced digital decision. An averaging circuit is included to support an

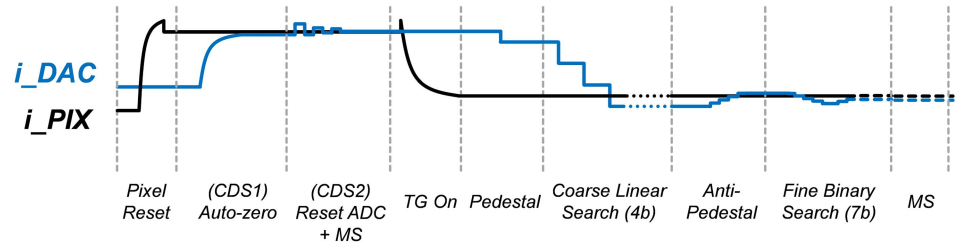

Fig. 6. c-ciSAR operation timing diagram.

TABLE I
C-CISAR SUB-BLOCKS ACTIVITY

	Pixel	c-ciDAC	Pre-amp	Comp.	Avg. circuit	SAR logic
Pixel Reset	Reset gate on	X	X	X	X	X
Auto-zero (CDS1)	O	Reset internal nodes to 3.4/-3.4V	Auto-zero	X	X	X
Reset ADC (CDS2)	O	Non-radix-of-2 progression plus few 1 LSB cycles if MS used (w/ PDF & PUF)	O	O	O if MS used	X
TG On	Transfer gate on	X	O	X	X	X
Pedestal	O	16 LSB (w/ PDF)	O	X	X	X
Coarse Lin. Search	O	64 LSB/cycle for 16 cycles (w/ PDC)	O	O	X	O
Anti-Pedestal	O	-32 LSB (w/ PUF)	O	X	X	O w/o activity
Fine Binary Search	O	Radix-of-2 progression then few redundant 1 LSB cycles (w/ PDF & PUF)	O	O	X	O
Multiple Sampling (MS)	O	Several 1 LSB cycles (w/ PDF & PUF)	O	O	O	O

O = *On or active*, X = *Off or no activity*, Text = *On and activity description*

accurate calibration decision and a multiple sampling feature, which will be discussed later in this section.

A. ADC Operation Sequence

Fig. 6 illustrates the operating sequence of the c-ciSAR converting one row of pixel values. The behavior of the c-ciSAR circuit blocks is also tabulated in Table I. The c-ciSAR performs 10b A/D conversion with analog-domain dual-CDS. The ci-cell-based DAC allows great flexibility in operating the ADC as the reference voltage generation is easily programmable by timely controlling the select signal on each of the ci-cells. By programming when and which of the 24 ci-cells in the PDC, PDF, and PUF ci-paths are activated, the c-ciSARs perform functions such as multiple sampling, pedestal/antipedestal, mixed linear/binary search, or search radix change without having to resort to hardware changes. Leveraging this flexibility, the c-ciSAR offers several benefits over conventional single-slope ADCs or SAR ADCs in terms of circuit complexity, noise, and energy efficiency. In this section, we explain how these distinctive functions are supported in the c-ciSAR and benefit the overall system performance.

1) Analog Dual-CDS: ADCs for image sensors typically support dual-CDS to remove pixel offset variation. Conventionally, the first CDS (CDS1) is done in the analog domain by auto-zeroing the decision-making circuit. The second CDS (CDS2) is performed in the digital domain after running the ADC twice, once for the reset level input and next for the signal level input. CDS2 involves extra ADC operation, which in single-slope ADCs results in a heavy energy penalty due to excess digital activities, i.e., counting, and in conventional SAR ADCs requires additional hardware for the DAC. However, with c-ciSAR, CDS2 can be performed in the analog domain without extra hardware because it can store reset level SAR result as its DAC output voltage. This allows the ADC to remember the first conversion result in analog voltage form, which the second conversion can then be based on. Hence subtraction for CDS2 is directly be performed in analog, which simplifies the ADC hardware and reduces energy overhead compared with digital computation.

Also, the absence of digital activity during dual-CDS is beneficial in low-power IoT image sensors. To lower the power consumption of the sensor, we need to lengthen the A/D conversion time, which then may increase the leakage power contribution from logic devices in SAR logic. Thanks to analog dual-CDS, the entire SAR logic can be power-gated off during a significant portion of the ADC operation to reduce its leakage power contribution.

A non-radix-of-2 SAR search is performed for reset level of CDS2. As the preamplifier consumes significant energy during the ADC operation, it is crucial to lower its tail current as much as possible while making accurate decisions. A non-radix-of-2 search helps lower the preamplifier power while enabling speedy operation since it incorporates redundancy, aiding in recovering any false decisions made along the way.

2) Coarse Linear Search and Pedestal: The signal level A/D conversion utilizes a coarse linear search to convert the first 4 bits of a 10-bit operation. A coarse linear search offers four advantages over a binary search: 1) smaller and regular changes in DAC voltage allow for the use of a slow, low-power preamplifier; 2) average ci-cell activation is reduced; 3) the PUC ci-path can be omitted; and 4) noisy PDC operation is avoided at lower light levels where its higher DAC noise (due to its high path gain) must impact less to maximize image

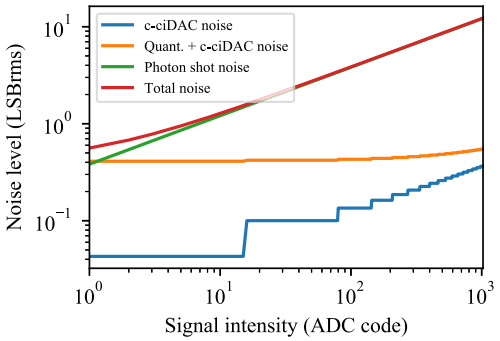

Fig. 7. Simulation result showing DAC noise accumulating over code with coarse linear search (bit depth = 10 bit, LSB = 488 μV, analog gain = 1, full-well capacity = 7000 e^-, fine steps involved at all codes = 115 steps, coarse step size = 64 LSB, $\text{Noise}_{\text{coarse-step}} = 0.09\ \text{LSB}_{\text{rms}}$, $\text{Noise}_{\text{fine-step}} = 0.004\ \text{LSB}_{\text{rms}}$, noise assumes 200 fF C_{cv}).

quality, as shown in Fig. 7. This is unlike a binary search, which would involve a constant number of noisy PDC/PUC operations regardless of the light level. Optimizing for noise at lower light conditions at the expense of increased noise at higher light conditions is a desirable tradeoff in image sensors because photon shot noise is likely to dominate the sensor SNR once the sensor is illuminated.

To further optimize for noise under dark conditions, the DAC value is intentionally shifted lower by only using the PDF ci-path before entering the coarse linear search. This behavior is referred to as pedestal, and it guarantees that there are no noisy PDC operations involved in converting signals under dark conditions.

3) Antipedestal and Fine Binary Search: After the coarse linear search, we perform a 7b fine binary search with a 1b overlap range. Prior to the fine binary search, the DAC voltage is again offset upward toward the center of the coarse search step to maximize overlap. This operation is named antipedestal. As opposed to reset level A/D conversion, the search decisions are now recorded in the SAR logic. A radix-of-2 search is preferred over non-radix-of-2 in this case because radix-2 code reconstruction is much more energy efficient than non-radix-of-2 as it involves less carry propagation. To account for any wrong decisions in the radix-of-2 search, a few more LSB decision cycles are appended at the end to provide sufficient redundancy.

The LSB cycle, when it is at the end of the whole conversion, only involves making the final decision using the decision-making circuit with no DAC voltage change. But in our application, we necessitate several more LSB cycles to be performed for redundant decisions and multiple sampling. In these cases, we must toggle the DAC voltages so that resulting voltage difference with high or low decision differs only by 1 LSB. However, our DAC is designed so that each step generated using PDF or PUF ci-path maps to this 1 LSB size that enabling both paths for LSB cycle results in 2 LSB difference. Hence we have designed the LSB cycle to include two sub-decision cycles with one involving (PDF,PUF) = (0,−1) operation and then (PDF,PUF) = (+1,0) operations.

4) Multiple Sampling: As an optional feature, the ADC can perform multiple sampling at the end of each reset and signal

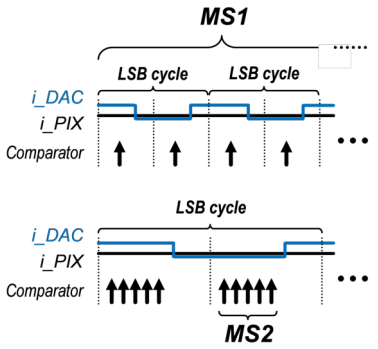

(a)

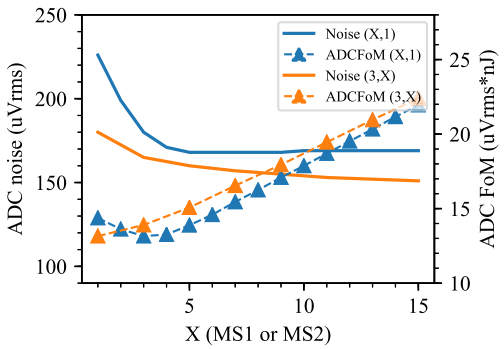

(b)

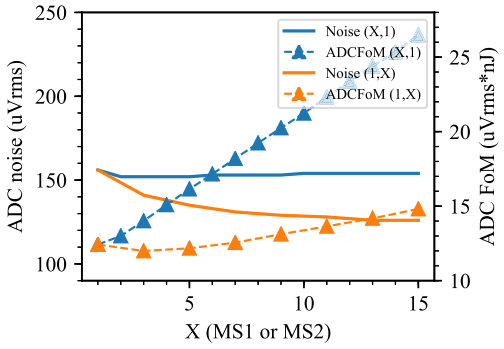

(c)

Fig. 8. (a) Multiple sampling with MS1 and MS2 methods. (b) Measurement result of noise and ADC FoM with low bias setting for preamplifier (I_{amp} = 48 nA) versus different (MS1,MS2). (c) Measurement result with high bias setting for preamplifier (I_{amp} = 108 nA).

level A/D conversion to reduce the impact of temporal noise (TN) on the system. We implemented two types of multiple sampling methods [Fig. 8(a)]: 1) repeat the LSB decision cycle for MS1 times and 2) fire the comparator decisions MS2 times within each cycles and perform majority voting. If ideal, multiple sampling can reduce the TN by $\sqrt{N}$, where N is the multiple sampling count.

At a glance, the MS1 method may not seem to yield any noise improvement if no extra statistical information is extracted [13]. However, in the presence of an input noise that is larger than an LSB bin size, repeating the LSB cycle many times will cause the DAC value to asymptotically approach the signal mean. Once it approaches the mean, it becomes

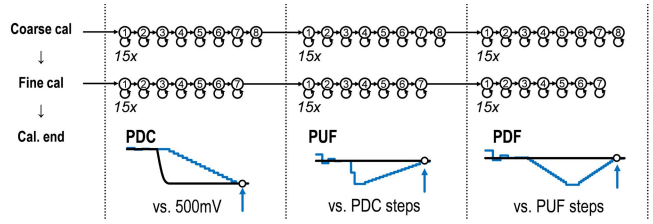

Fig. 9. Sequence for PVT (coarse) and local mismatch (fine) calibration.

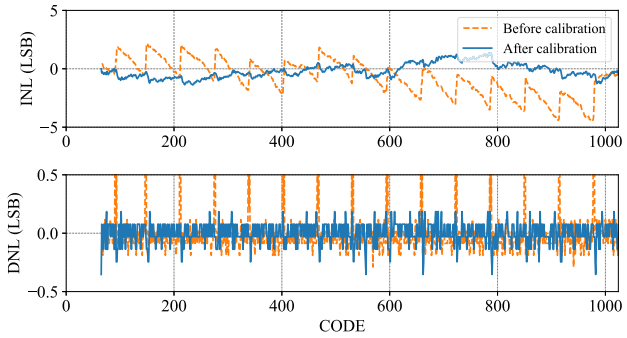

Fig. 10. INL and DNL measurement of arbitrarily selected c-ciSAR ADC column before and after calibration.

difficult to deviate from the mean since that would require an improbable streak of single-sided decisions. Effectively this mechanism forces the DAC value to be found closely around the signal mean after sufficient trials, establishing a filtering function.

For best noise reduction performance, both MS1 and MS2 must be utilized to the extreme. However, in terms of energy efficiency (ADC FoM), the optimum setting of MS1 and MS2 tends to change depending on what bandwidth of noise and how much of it is dominating the sensor, as shown in Fig. 8(b) and (c). When the readout circuit is dominated by relatively large and low-frequency noise, as in the case for Fig. 8(b), energy-optimum setting skews toward favoring MS1 than MS2. On the other hand, with smaller and higher frequency noise as in Fig. 8(c), it skews toward favoring MS2. These tendencies arise due to two reasons. First, the noise reduction of MS2 is impartial to the magnitude of noise, however, MS1 has a limit in the lowest achievable noise level that is related to LSB bin size and original noise level. Second, since MS2 packs sampling operations more densely across time, it becomes more effective with high bandwidth noise as it enjoys more uncorrelated portion of noise energy to be averaged out. For optimal energy efficiency, MS1 and MS2 must be explored in tandem.

B. Calibration

The c-ciSAR is compact and energy efficient but requires calibration to achieve high accuracy. It includes two mechanisms for calibration as mentioned in the previous section: the global ci-cell biasing voltage adjusts for PVT variations, and the calibration capacitor bank adjusts for the column-to-column mismatch. The basic idea of calibration is to match

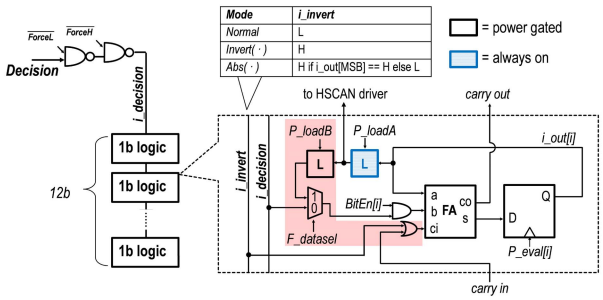

Fig. 11. SAR logic circuit with in-column computation.

the charge-injection step size to a known reference voltage in the sensor.

The calibration procedure is shown in Fig. 9. First, several PDC steps are compared against the reference voltage. Then taking PDC as a reference, the PUF steps are calibrated, and then in the same way, PDF is calibrated with PUF as a reference. This cascading calibration guarantees their relative ratio and avoids DNL. The three steps are repeated for PVT (coarse) and then for column-to-column mismatch (fine).

Each calibration step is a series of bit decisions to set the 8b biasing voltage generator (coarse) and 7b calibration capacitor bank (fine). To improve the calibration accuracy, each of the three calibration steps yields a final calibration decision after accumulating 15 iterations per bit using the averaging circuit. For coarse calibration, these final calibration decisions are streamed out and gathered at the controller to be applied to the biasing voltage generator. For fine calibration, they are collected locally in calibration capacitor banks. This entire process is done without any off-chip assistance.

With calibration, the INL and DNL performance of c-ciSAR is visibly improved, as shown in Fig. 10. As c-ciSAR is prone to gain error between each ci-path, the non-calibrated ADC exhibits noticeable spikes in the DNL plot and a ramplike pattern in the INL. After calibration, the spikes in the DNL were contained well within $+0.5/-0.5$ of the LSB boundary. Also, the maximum INL error was lower than 2 LSB with a fairly consistent profile across all c-ciSAR ADCs. This is acceptable for image sensors where pixel response and gamma correction already distort the signal to a higher degree.

C. SAR Logic

Fig. 11 shows the column SAR logic circuit. Each bit in the SAR logic is comprised of a full-adder and a D-type FF (DFF) to compute and store the reconstructed SAR result. A full-adder-based SAR logic design is chosen over a counter-based design since c-ciSAR is highly programmable in SAR search weights and equal flexibility is required in the SAR logic to reconstruct the data correctly. With this structure, the SAR logic can perform arbitrary value addition/subtraction broadcasted from outside the array.

The SAR logic is designed with leaky thin-oxide devices to reduce size. Due to leakage from the thin-oxide devices, a power gating transistor is added to each column to turn off most of the logic gates while there are no activities.

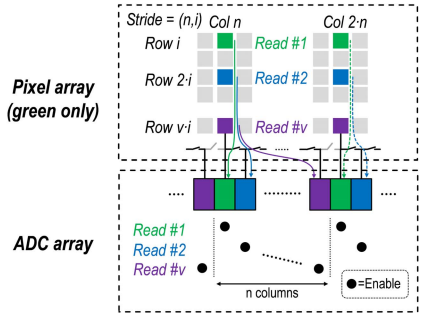

Fig. 12. n-way side-step mapping between pixel array and ADCs in motion-detection mode. Number of MD frame rows (v) is equal to column stride (n) in this illustration.

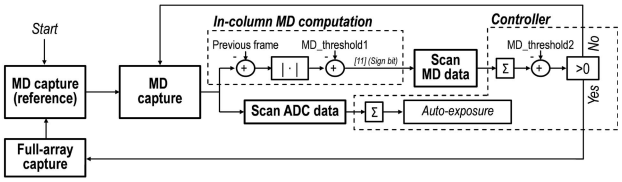

Fig. 13. Motion-detection mode signal flow.

The data retention latch, highlighted in blue in the figure, is the exception as it is required to hold the column data during speed-limited horizontal scanning or to tether on the data between two motion frames (MD frames) in MD mode.

The circuits shown in the red-shaded region are added for in-column motion-detection computation. With these additional circuits, the SAR logic can now perform bitwise inversion, absolute value function and addition/subtraction between DFF value and data retention latch value. The added circuits include an extra latch after the data retention latch to enable the data to swap positions between the data retention latch and the DFF without contention. The OR gate cuts off the carry chain and forces it to logic-H, so the full-adder always inverts the bit in a bitwise manner. The MSB of the SAR logic output is used as a sign bit and determines the polarity of the i_invert signal in absolute value mode.

IV. Motion Detection

We propose n-way side-step mapping of n MD pixel rows that re-uses the in-column hardware in the full-adder-based SAR logic. This allows motion detection without transferring the MD pixel values to external storage for MD frame differencing. Unlike prior in-pixel analog MD [1], which degraded pixel performance by -25 dB compared to typical 4T pixels, the proposed near-pixel MD approach maintains an unmodified conventional 4T pixel structure for high image quality and incurs minimal energy overhead from moving the data around.

A. N-Way Side-Step Mapping

The MD frame pixels are subsampled every nth-column and ith row from a full active pixel array. For row v, an MD pixel at coordinate $(v \times i, c \times n)$, where $c = 0, 1, \ldots$ is re-routed sideways to ADC column $(c \times n + v)$ for data retention and in-column MD computation. As readout progresses vertically across subsampled MD rows, the mapped ADCs shift horizontally, and all MD pixels are mapped to column ADC. This pixel-to-ADC mapping scheme is named n-way side-step mapping and is illustrated in Fig. 12.

This mapping scheme enables both the storage of the past MD pixel values and the differencing between two MD frames to reuse existing hardware of the c-ciSAR, thereby incurring minimal area overhead. The subsample stride (n, i) and resulting MD frame size (h, v) are configurable with the constraint that $h \times v < k$, where k is the total number of full active-array columns (640 for VGA). Skipped pixels can be enabled and binned together to minimize blind spots using three methods: floating diffusion (FD) binning on two-shared pixels (two green pixels in Bayer pattern share same FD), source-follower (SF) binning across horizontal/vertical direction, and digital binning in the column logic. For our evaluation case, $(n, i) = (20, 24)$, and FD and SF binning are used, resulting in an MD frame resolution of 32×20 pixels.

B. In-Column MD Computation

By re-purposing the SAR logic and its unit operators, we can compute whether the MD pixels have had significant changes in their code with the flow shown in Fig. 13. As the entire previous MD frame values are stored per-column in the data retention latch, and the new MD frame values are freshly computed and stored in the DFF, by bitwise inverting the new MD frame value and adding the previous frame value to it, we can take the difference between the two MD frames. Then we apply an absolute value function to erase the polarity of the difference. Next, we subtract threshold value (MD_threshold1) asserted from outside the array and then invert the result, which renders sign bit (MD flag) to turn logic-H if the MD pixel difference is higher than the threshold. By streaming out MD flags from all active columns and accumulating them at the controller side, we can determine how many MD pixels have experienced significant changes in their code. This count is again thresholded (MD_threshold2) to determine whether a meaningful amount of change has happened. If so, the sensor progresses to start integration for the full-array image and then readout, which completes a single instance of motion-triggered sensor operation. The operation then loops again for the next motion-triggering event. Utilizing the in-column computation hardware, the MD procedure is software-defined through an on-chip programmable controller and it takes less than 200 instruction cycles running at 5 MHz.

Alongside the MD computation, the controller is also programmed to intermittently perform simple auto-exposure (AE) algorithm on MD frame. AE program takes the average value of the MD frame as it is streamed out, using a dedicated accumulator, and adjusts the exposure to maintain the value to be within a preset range.

Fig. 14 shows captured images from the motion-triggered full-array capture operation.

V. Measurements

Table II provides a comparison of this work with other recently published image sensors with MD capability and with

TABLE II
PERFORMANCE SUMMARY

	This work		[2]	[1]	[12]
Technology	TPSCo 65 nm CIS		90 nm CIS + 40 nm	130 nm	180nm
Pixel structure	4T 1.5 μm		4T 1.5 μm	3T 6.4 μm	3T
Pixel array H x V	792 x 528		2560 x 1536	128 x 128	320 x 240
Frame-rate	5.6 fps[*1]		60 fps	19 fps	130 fps
ADC type	c-ciSAR ADC		Single-slope ADC	Single-slope ADC	SAR ADC
Column-parallel	O		O	O	X
Resolution	10 b		10 b	8 b	10 b
Full signal range	500 mV		435 mV	-	-
ADC noise[*2]	226 μV_{rms}		100.4 μV_{rms}	-	450 μV_{rms}[*3]
Dynamic range	64.1 dB		67 dB	38.5 dB	-
MD mode support	O		O	O	X
MD location	Near pixel		Outside array	In pixel	-
MD frame H x V	32 x 20		16 x 5	48 x 16	-
MD frame-rate	170 fps[*1]		10 fps	30 fps	-
MD mode ADC res.	10 b		8 b	3 level	-
	ADC + Pixel	**Sensor**	**Sensor**	**Sensor**	**ADC + Pixel**
E/frame [μJ]	26.6	70.1	1583	1.5	16.9
E/frame/pix [pJ]	63.6	167.5	400	610	220
ADC FoM [μV_{rms}·nJ]	14.4	37.9	40.1	-	99[*3]
ADC FoM2 [μV_{rms}·pJ/DRU]	15.7	41.4	73.5	-	-
MD E/frame [μJ]	0.26	1.7	110	0.036	-
MD E/frame/pix [pJ]	400	2650	1375000	46.3	-

[*1] Maximum frame-rate based on frame-to-frame interval including two row-readout time for integration at the beginning of frame readout

[*2] Measured at dark condition, ADC only, w/o PMU and AG=2.5

[*3] Includes pixel source follower noise

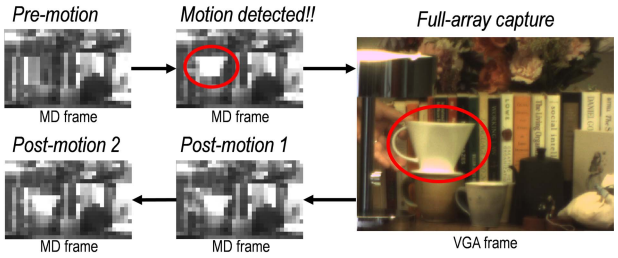

Fig. 14. Motion-detection mode triggering full-array capture based on scene difference in indoor lighting condition (approximately 300 Lux). For acquiring the image, the sensor was programed to insert 1 second blank time between MD frames.

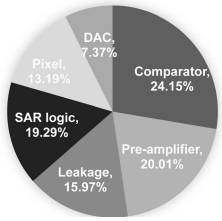

Fig. 15. Energy consumption of blocks in c-ciSAR ADC in full-array mode.

SAR ADCs. The noise-related numbers in the bottom section of the table are with an AG setting of 2.5 without multiple sampling enabled.

The energy efficiency of the complete sensor is 37.9 μV_{rms}·nJ (ADC FoM). However, this number is dominated by the fully programmable controller (for flexible testing purposes); a simple hardwired state machine could be used in a more lightweight implementation. Excluding the controller energy consumption and evaluating only the readout energy (ADC and pixel), E/frame/pix is 63.6 pJ, and the ADC FoM is 14.4 μV_{rms}·nJ. The E/frame/pix is increased to 1.7 μJ (0.26 μJ, readout only) in MD mode due to nonscaling energy (e.g., control signal driver, controller, and row driver). Yet, it is 64.7 times smaller (423 times, readout only) than that of a single-slope ADC-based sensor [2].

The proportion of energy consumed by each block in the c-ciSAR is shown in Fig. 15. The energy consumed in the pixel is well suppressed to only 13% of the overall ADC energy thanks to the column-parallel design and lowered pixel VDD (1.7 V), and the DAC energy is held to approximately 7.4% due to the c-ciDAC structure. The decision-making circuit consumes about 44% of the energy and the SAR logic consumes about 19%. Although power-gated, there is notable contribution (approx. 16%) from leakage current since SAR logic is constructed of low-vt thin-oxide devices due to limited vt choice in the process.

Table III shows measured noise performance of the sensor. With varying multiple sampling settings and the preamplifier bias current, the TN can be reduced from 226 μV_{rms} to 104 μV_{rms}, which translates to 211 μV_{rms} to 67 μV_{rms} when calculating out the quantization noise contribution. Vertical fixed pattern noise (VFPN) improves with faster preamplifier and with more redundant decisions (MS1). The sensor exhibits better ADC FoM (11.99 μV_{rms}·nJ) at the higher preamplifier bias current setting with moderate MS2, based on the extrapolated energy consumption value from Case1.

Fig. 16 shows the chip photograph. The sensor contains a 792 × 512 pixel array with VGA active array and optical black pixels. Each c-ciSAR ADC slice is designed to fit into 3.1-μm pitch, whereas it is connected to 1.5-μm pitch pixels. Smaller

TABLE III
IMAGE CHARACTERISTICS

	Case1	Case2	Case3	Case4	Case5
MS1	1	3	1	1	3
MS2	1	1	1	3	15
I_{amp} (nA)	48	48	108	108	252
TN (μV_{rms})	226	180	156	141	104
TN w/o q (μV_{rms})[*1]	211	161	134	116	67
HN (LSB_{rms})	0.043	0.030	0.026	0.024	0.019
VN (LSB_{rms})	0.067	0.042	0.035	0.032	0.024
HFPN (LSB_{rms})	0.009	0.014	0.005	0.007	0.009
VFPN (LSB_{rms})	0.380	0.235	0.154	0.151	0.152
E/frame/pix (pJ)[*2]	63.6	73.13	79.71	85.29	313.5
ADC FoM (μV_{rms}·nJ)	14.37	13.16	12.42	11.99	32.64

All measurements at dark condition, ADC only, w/o PMU and AG=2.5 (LSB=195μV)
Noise definition as found in SMIA specification
TN = Temporal noise
HN = Horizontal noise
VN = Vertical noise
HFPN = Horizontal fixed pattern noise
VFPN = Vertical fixed pattern noise
[*1] q = quantization noise
[*2] Extrapolated based on Case1

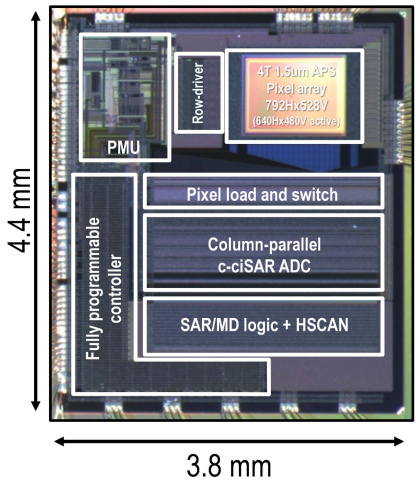

Fig. 16. Chip photograph.

pitch pixel array is selected for future integration of a small-form lens. The sensor is designed with thick-oxide devices to reduce leakage current except for column SAR logic. The ADC pitch and size can be further reduced with using high-vt thin-oxide devices if available. Also, the sensor can operate with its on-chip PMU, which generates 13 internal voltages from the external 2.5-V battery, enabling integration into an IoT sensor node. The total die area is 16.7 mm^2.

VI. CONCLUSION

In this article, an IoT image sensor is proposed that is designed with energy-efficient c-ciSAR ADCs in column-parallel form. The c-ciSAR achieves 10b performance while only consuming 63.6 pJ/frame/pix, which yields the state-of-the-art energy efficiency of 14.4-μV_{rms}·nJ ADC FoM. Also, the sensor contains motion-triggering function implemented with the near-pixel placement of previous frame data. This is done while re-using the ADC hardware already in place for c-ciSAR, which leads to a more hardware/energy-efficient solution than other sensors while maintaining high image quality (64.1 dB in dynamic range).

ACKNOWLEDGMENT

The authors would like to thank Synopsys for their tools.

REFERENCES

[1] G. Kim *et al.*, "A 467 nW CMOS visual motion sensor with temporal averaging and pixel aggregation," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2013, pp. 480–481.
[2] O. Kumagai *et al.*, "A 1/4-inch 3.9 Mpixel low-power event-driven back-illuminated stacked CMOS image sensor," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2018, pp. 86–88.
[3] K. D. Choo, J. Bell, and M. P. Flynn, "Area-efficient 1 GS/s 6b SAR ADC with charge-injection-cell-based DAC," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Jan./Feb. 2016, pp. 460–461.
[4] S. Kawahito, "Column-parallel ADCs for CMOS image sensors and their FoM-based evaluations," *IEICE Trans. Electron.*, vol. E101.C, no. 7, pp. 444–456, Jul. 2018.
[5] P. Harpe, E. Cantatore, and A. V. Roermund, "A 2.2/2.7 fJ/conversion-step 10/12b 40 kS/s SAR ADC with data-driven noise reduction," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2013, pp. 270–271.
[6] C. Liu, S. Chang, G. Huang, and Y. Lin, "A 0.92 mW 10-bit 50-MS/s SAR ADC in 0.13 μm CMOS process," in *Proc. Symp. VLSI Circuits*, Jun. 2009, pp. 236–237.
[7] L. Kull *et al.*, "A 90 GS/s 8b 667 mW 64× interleaved SAR ADC in 32 nm digital SOI CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2014, pp. 378–379.
[8] R. Funatsu *et al.*, "133 Mpixel 60 fps CMOS image sensor with 32-column shared high-speed column-parallel SAR ADCs," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2015, pp. 1–3.
[9] H. Kim *et al.*, "A delta-readout scheme for low-power CMOS image sensors with multi-column-parallel SAR ADCs," *IEEE J. Solid-State Circuits*, vol. 51, no. 10, pp. 2262–2273, Oct. 2016.
[10] J. P. Keane *et al.*, "An 8 GS/s time-interleaved SAR ADC with unresolved decision detection achieving -58 dBFS noise and 4 GHz bandwidth in 28 nm CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2017, pp. 284–285.
[11] P. Harpe, "A 0.0013 mm^2 10b 10 MS/s SAR ADC with a 0.0048 mm^2 42 dB-rejection passive fir filter," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Apr. 2018, pp. 1–4.
[12] S. Ji, J. Pu, B. C. Lim, and M. Horowitz, "A 220 pJ/pixel/frame CMOS image sensor with partial settling readout architecture," in *Proc. IEEE Symp. VLSI Circuits (VLSI-Circuits)*, Jun. 2016, pp. 1–2.
[13] L. Chen, X. Tang, A. Sanyal, Y. Yoon, J. Cong, and N. Sun, "A 0.7-V 0.6-μW 100-kS/s low-power SAR ADC with statistical estimation-based noise reduction," *IEEE J. Solid-State Circuits*, vol. 52, no. 5, pp. 1388–1398, May 2017.
[14] S.-E. Hsieh and C.-C. Hsieh, "A 0.4 V 13b 270 kS/S SAR-ISDM ADC with an opamp-less time-domain integrator," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2018, pp. 240–242.
[15] L. Kull *et al.*, "A 10b 1.5 GS/s pipelined-SAR ADC with background second-stage common-mode regulation and offset calibration in 14 nm CMOS FinFET," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2017, pp. 474–475.
[16] T. Takahashi *et al.*, "A stacked CMOS image sensor with array-parallel ADC architecture," *IEEE J. Solid-State Circuits*, vol. 53, no. 4, pp. 1061–1070, Apr. 2018.
[17] T. Haruta *et al.*, "A 1/2.3 inch 20 Mpixel 3-layer stacked CMOS image sensor with DRAM," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2017, pp. 76–77.
[18] S. Sukegawa *et al.*, "A 1/4-inch 8 Mpixel back-illuminated stacked CMOS image sensor," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2013, pp. 484–485.
[19] A. Suzuki *et al.*, "A 1/1.7-inch 20 Mpixel back-illuminated stacked CMOS image sensor for new imaging applications," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2015, pp. 1–3.

[20] T. Arai *et al.*, "A 1.1-μm 33-Mpixel 240-fps 3-D-stacked CMOS image sensor with three-stage cyclic-cyclic-SAR analog-to-digital converters," *IEEE Trans. Electron Devices*, vol. 64, no. 12, pp. 4992–5000, Dec. 2017.

[21] K. D. Choo *et al.*, "Energy-efficient low-noise CMOS image sensor with capacitor array-assisted charge-injection SAR ADC for motion-triggered low-power IoT applications," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2019, pp. 96–98.

Kyojin David Choo (M'14) received the B.S. and M.S. degrees in electrical engineering from Seoul National University, Seoul, South Korea, in 2007 and 2009, respectively, and the Ph.D. degree from the University of Michigan, Ann Arbor, MI, USA, in 2018.

From 2009 to 2013, he was with the Image Sensor Development Team, Samsung Electronics, Yongin, South Korea, where he designed signal readout chains for mobile/digital single-lens reflex (DSLR) image sensors. During his doctoral research, he interned with Apple, Cupertino, CA, USA, and consulted to Sony Electronics, San Jose, CA, USA. He is currently a Research Fellow with the University of Michigan. He holds 15 U.S. patents. His research interests include charge-domain circuits, sensor interfaces, energy converters, high-speed links/timing generators, and millimeter-scale integrated systems.

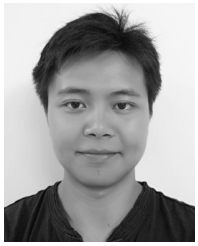

Li Xu (S'15) received the B.Eng. degree in automation from Tongji University, Shanghai, China, in 2009, and the M.S. degree in electrical and computer engineering from Northeastern University, Boston, MA, USA, in 2016. He is currently pursuing the Ph.D. degree with the University of Michigan, Ann Arbor, MI, USA.

From 2009 to 2011, he was an IC Design Engineer with Ricoh Electronic Devices Shanghai Company, Ltd., Shanghai, where he worked on low-drop out (regulators) (LDO) and dc/dc converter projects. In 2015, he was a Design Intern with Linear Technology Corporation, Colorado Springs, CO, USA. His current research interest includes energy-efficient mixed-signal circuit design.

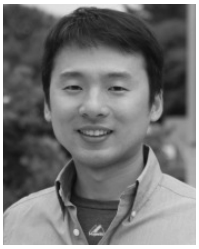

Yejoong Kim (S'08–M'15) received the bachelor's degree in electrical engineering from Yonsei University, Seoul, South Korea, in 2008, and the master's and Ph.D. degrees in electrical engineering from the University of Michigan, Ann Arbor, MI, USA, in 2012 and 2015, respectively.

He is currently a Research Fellow with the University of Michigan and the Vice President of Research and Development with CubeWorks, Inc., Ann Arbor. His research interests include subthreshold circuit designs, ultralow-power SRAM, and the design of millimeter-scale computing systems and sensor platforms.

Ji-Hwan Seol (M'15) received the B.S. degree in electrical engineering from Yonsei University, Seoul, South Korea, in 2009, and the M.S. degree in electrical engineering from KAIST, Daejeon, South Korea, in 2012. He is currently pursuing the Ph.D. degree with the University of Michigan, Ann Arbor, MI, USA, supported by the Samsung Fellowship Program.

In 2012, he joined the DRAM Design Team, Samsung Electronics, Yongin, South Korea, where he contributed to the development of mobile DRAMs including LPDDR2, LPDDR3, LDDR4, and LPDDR5. His research interests include clock generation, memory systems, and ultralow-power system design.

Xiao Wu (S'14–M'15) received the B.S. degree in electrical engineering jointly from the University of Michigan, Ann Arbor, MI, USA, and Shanghai Jiaotong University, Shanghai, China, in 2014. She is currently pursuing the Ph.D. degree with the University of Michigan.

She has experience in energy harvester, power management unit (PMU), and high-performance hardware accelerator. Her current research interests include small sensor system integration and hardware accelerator design.

Dennis Sylvester (S'95–M'00–SM'04–F'11) received the Ph.D. degree in electrical engineering from the University of California at Berkeley, Berkeley, CA, USA, in 1999, where his dissertation was recognized with the David J. Sakrison Memorial Prize as the most outstanding research in the UC-Berkeley EECS Department.

He was the Founding Director of the Michigan Integrated Circuits Laboratory (MICL), Electrical and Computer Engineering Department, University of Michigan, Ann Arbor, MI, USA, a group of 10 faculty and 70+ graduate students. He has held research staff positions at the Advanced Technology Group of Synopsys, Mountain View, CA, USA, Hewlett-Packard Laboratories, Palo Alto, CA, USA, and visiting professorships at the National University of Singapore, Singapore, and Nanyang Technological University, Singapore. He co-founded Ambiq Micro, Austin, TX, USA, a fabless semiconductor company developing ultralow-power mixed-signal solutions for compact wireless devices. He is currently a Professor of electrical engineering and computer science with the University of Michigan. He has authored or coauthored more than 500 articles along with one book and several book chapters. He holds 45 U.S. patents. His research interests include the design of millimeter-scale computing systems and energy-efficient near-threshold computing.

Dr. Sylvester currently serves on the Technical Program Committee for the IEEE International Solid-State Circuits Conference and on the Administrative Committee for the IEEE Solid-State Circuits Society. He was a recipient of the NSF CAREER Award, the Beatrice Winner Award at ISSCC, an IBM Faculty Award, an SRC Inventor Recognition Award, the University of Michigan Henry Russel Award for distinguished scholarship, and ten best paper awards and nominations. He was named one of the Top Contributing Authors at ISSCC and most prolific author at IEEE Symposium on VLSI circuits. He serves/has served as Associate Editor for the IEEE JOURNAL OF SOLID-STATE CIRCUITS, the IEEE TRANSACTIONS ON CAD, and the IEEE TRANSACTIONS ON VLSI SYSTEMS. He was an IEEE Solid-State Circuits Society Distinguished Lecturer.

David Blaauw (M'94–SM'07–F'12) received the B.S. degree in physics and computer science from Duke University, Durham, NC, USA, in 1986, and the Ph.D. degree in computer science from the University of Illinois at Urbana–Champaign, Champaign, IL, USA, in 1991.

Until August 2001, he worked for Motorola, Inc., Austin, TX, USA, where he was the Manager of the High-Performance Design Technology Group and won the Motorola Innovation Award. Since August 2001, he has been the Faculty of the University of Michigan, Michigan, Ann Arbor, USA, where he is the Kensall D. Wise Collegiate Professor of EECS. For high-end servers, his research group introduced so-called near-threshold computing, which has become a common concept in semiconductor design. Most recently, he has pursued research in cognitive computing using analog, in-memory neural-networks for edge-devices and genomics acceleration. He has authored or coauthored more than 600 articles. He holds 65 patents. He has extensive research in ultralow-power computing using subthreshold computing and analog circuits for millimeter sensor systems which was selected by the MIT Technology Review as one of the year's most significant innovations.

Dr. Blaauw currently serves on the IEEE International Solid-State Circuits Conference's technical program committee. He was a recipient of the 2016 SIA-SRC Faculty Award for lifetime research contributions to the U.S. semiconductor industry and numerous best paper awards and nominations. He was the General Chair of the IEEE International Symposium on Low Power, the Technical Program Chair of the ACM/IEEE Design Automation Conference. He is currently the Director of the Michigan Integrated Circuits Lab.