A 1-MS/s to 1-GS/s Ringamp-Based Pipelined ADC With Fully Dynamic Reference Regulation and Stochastic Scope-on-Chip Background Monitoring in 16 nm

Benjamin Hershberg, *Member, IEEE*, Nereo Markulić, *Member, IEEE*, Jorge Lagos, *Member, IEEE*, Ewout Martens, *Member, IEEE*, Davide Dermit, *Member, IEEE*, and Jan Craninckx, *Fellow, IEEE*

***Abstract*— This article presents a fully dynamic ringamp-based pipelined ADC with integrated reference buffer that operates from 1-MS/s to 1-GS/s and maintains a Walden Figure-of-Merit (FoM) of 14 fJ/conversion-step across this range. A "split-reference" regulation technique is introduced, which provides multiple buffered replicas with varying accuracies and output impedances to the core ADC circuitry, relaxing overall buffer design requirements and improving efficiency. The regulator blocks are implemented with fully dynamic discrete-time loops. Furthermore, a technique for background reconstruction of residue amplifier settling behavior is also described. The "scope-on-chip" captures high-resolution transient waveforms using a 1-bit stochastic ADC. It is shown how these waveform data can be used for optimization of ringamp biasing and PVT tracking. The ADC is fabricated in a 16-nm CMOS technology and at 1 GS/s with a Nyquist input achieves 59.5-dB SNDR, 75.9-dB SFDR, and 10.9-mW total power consumption with only 8% consumed by the reference regulation.**

***Index Terms*— A/D, ADC, asynchronous, dynamic, event driven, fully dynamic, high speed, low power, pipeline, pipelined ADC, reference buffer, reference regulation, ring amplifier, ring amplification, ringamp, scope on chip, stochastic ADC, voltage reference, voltage regulation, waveform capture.**

I. Introduction

A STABLE analog reference is a key component of almost all ADCs. It provides a known voltage or current reference (or set of references) against which to compare the analog input. The precision and noise of this reference is critical; errors injected into the reference can be just as problematic as errors injected into the main signal. In many popular ADC architectures, the core circuitry pulls charge from the reference, requiring a sufficiently low-impedance source to maintain voltage accuracy and minimize reference ripple. For ADCs in higher speed and resolution categories, the amount of on-chip decoupling capacitor that would be required to satisfy these requirements consumes too much chip area to be cost effective (nanofarads). Instead, an additional buffering stage is often added to convert a high-impedance voltage reference provided by a bandgap or external source into a low-impedance regulated replica that can be used by the ADC core. The conventional solution to this buffering stage is to use a continuous-time voltage regulation loop stabilized by a more practically sized decoupling capacitor located at the regulator output.

Many alternative techniques and enhancements for reference regulation have also been proposed, particularly in the domain of medium- and high-resolution, high-speed SAR ADCs. In such systems, the power efficiency of the ADC core is often quite good, but charge pull from the reference is also relatively high and contains strongly data-dependent high-frequency switching transients. This places a burden on the regulator, and it often consumes a large portion of the overall power budget (e.g., four times the ADC core power in [1]). The proposed methods for reducing regulator requirements include techniques based on using redundancy to correct MSB-related reference errors [2], minimizing or eliminating signal-dependent charge pull from the reference with an internal charge reservoir and/or special switching techniques [3]–[6], active reference error cancellation [7], and digital correction [8], [9]. Notably, many of these techniques exploit specific aspects of the SAR algorithm itself. In other common ADC architectures such as pipeline and delta–sigma, charge is pulled from the references by switched-capacitor residue generation circuits such as MDACs or integrators, and this type of charge pull is not always compatible with SAR-based techniques.

There thus exists an unmet need for a more general class of reference regulation techniques. This is emphasized by the recent appearance of high-performance reconfigurable pipelined ADCs [10], [11]. The ADC cores of these designs are capable of high-speed operation with state-of-the-art power efficiency in the medium- and high-resolution design space. Furthermore, they support fully dynamic operation with constant energy per conversion independent of clock rate. To maintain this reconfigurability advantage at the system level, the reference regulation subsystem must be fully dynamic too. Unfortunately, conventional regulator designs

Manuscript received August 25, 2020; revised October 28, 2020; accepted December 3, 2020. Date of publication January 29, 2021; date of current version March 26, 2021. This article was approved by Associate Editor Brian Ginsburg. *(Corresponding author: Benjamin Hershberg.)*

The authors are with the Interuniversity Microelectronics Center, 3001 Leuven, Belgium (e-mail: benjamin@hershberg.com).

Color versions of one or more figures in this article are available at https://doi.org/10.1109/JSSC.2020.3044831.

Digital Object Identifier 10.1109/JSSC.2020.3044831

involve significant static power consumption. A fully dynamic alternative was proposed in [12], but is only practical for lower resolution ADCs. In this article, we present a fully dynamic discrete-time regulator that can be used in high-accuracy applications and is generalizable to several classes of ADC.

In recent years, ring amplifiers (ringamps) have provided designers with a general-purpose scalable switched-capacitor residue amplification tool [13], enabling state-of-the-art ADCs in various architecture types and performance categories, including the fully dynamic pipelined ADCs mentioned earlier [10], [11], [14]–[18]. Ringamps operate in closed-loop, with feedback helping to suppress variation due to process, voltage, and temperature (PVT). Thus, the challenge of PVT robustness is essentially a matter of ensuring stability and clean settling. There are several approaches to this: robust analog design with sufficient built-in margin [14], [15], [17], [19], [20], adaptive control using either analog [21] or digital [10] compensation, or calibration [22]. Of these options, adaptive bias control is well suited for pushing the limits of performance and efficiency, at the cost of added complexity. In this article, we demonstrate a new approach that monitors the transient settling behavior of a ringamp in the background with a "scope-on-chip" and uses this information to generate an objective function that can be used for adaptive bias control. Waveform capture is done using a very compact 1-bit stochastic ADC circuit followed by simple digital processing.

These two key techniques of fully dynamic reference regulation and stochastic waveform capture are demonstrated in an 11-bit 1.5-bit/stage pipelined ADC fabricated in a 16-nm CMOS technology. At a maximum sampling rate of 1 GS/s with a Nyquist input tone, the ADC achieves 59.5-dB SNDR and 75.9-dB SFDR with regulation enabled. Total power consumption is 10.9 mW with only 8% of this consumed by the reference regulator. The ADC's fully dynamic event-driven method of operation maintains near-constant performance and a Walden Figure-of-Merit (FoM) of 14 fJ/conversion-step from 1-MS/s to 1-GS/s.

This article is organized as follows. Section II introduces the proposed regulation technique and provides details of implementation. Section III explains the stochastic waveform capture system. Measurements are reported in Section IV. Conclusions and a discussion of future work are drawn in Section V.

II. Split-Reference Regulation

A. Basic Principle

The proposed regulation technique stems from a simple but fundamental insight; at the local circuit level, the accuracy and charge-sourcing capabilities required from a voltage reference are often mutually exclusive with respect to time. That is, when a circuit pulls a large amount of charge from a reference, it rarely requires high accuracy, and when it requires high accuracy, it usually pulls very little charge.

Take, for example, the conventional 1.5-bit flip-around MDAC structure of Fig. 1(a) and the simulated residue amplifier (RA) settling behavior and total integrated reference charge waveforms in Fig. 2. More than 98% of the total

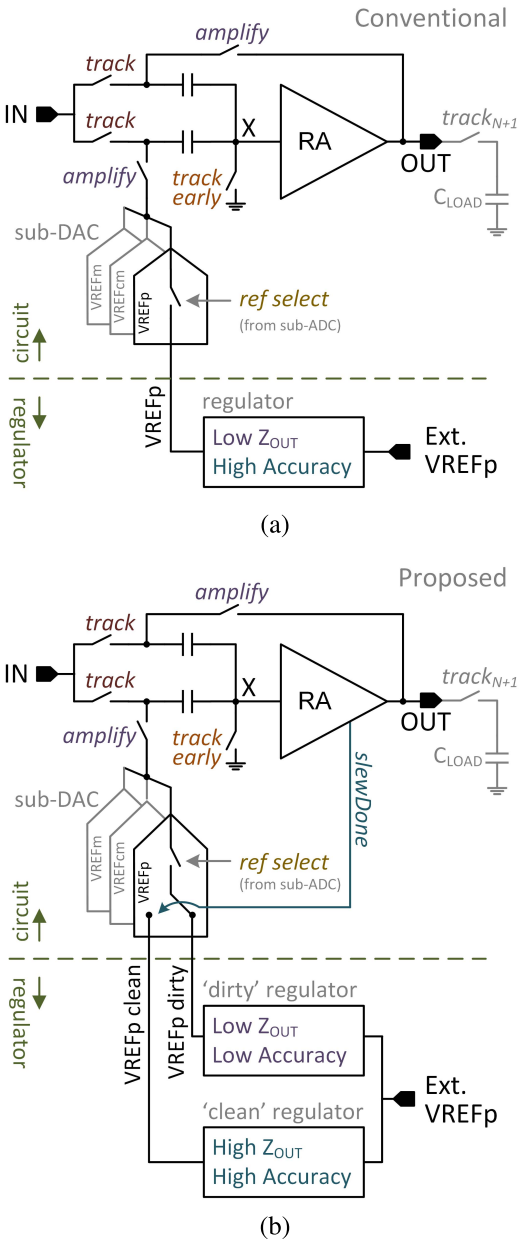

Fig. 1. Simplified 1.5-bit flip-around MDAC with (a) conventional reference regulation and (b) split-reference regulation.

integrated charge is pulled from the reference during the initial connection to C_{LOAD} and subsequent slewing. Because the amplifier is only coarse charging during this time, the accuracy requirement of the reference is temporarily reduced. As long as the reference error is less than the unsettled error of the main feedback loop, it will not significantly affect the fine settling that follows. During fine settling, less than 2% of the total integrated charge is pulled from the reference. However, the reference accuracy requirement increases to its maximum. Any error in the reference at the end of the *amplify* phase will also appear in the final sampled residue.

Fig. 1(b) shows the technique of "split-reference" regulation that exploits the above observations to relax the overall design requirements of the regulator subsystem. The external reference is split into two regulated replicas, a "dirty" replica with low output impedance but also low (i.e., relaxed) accuracy

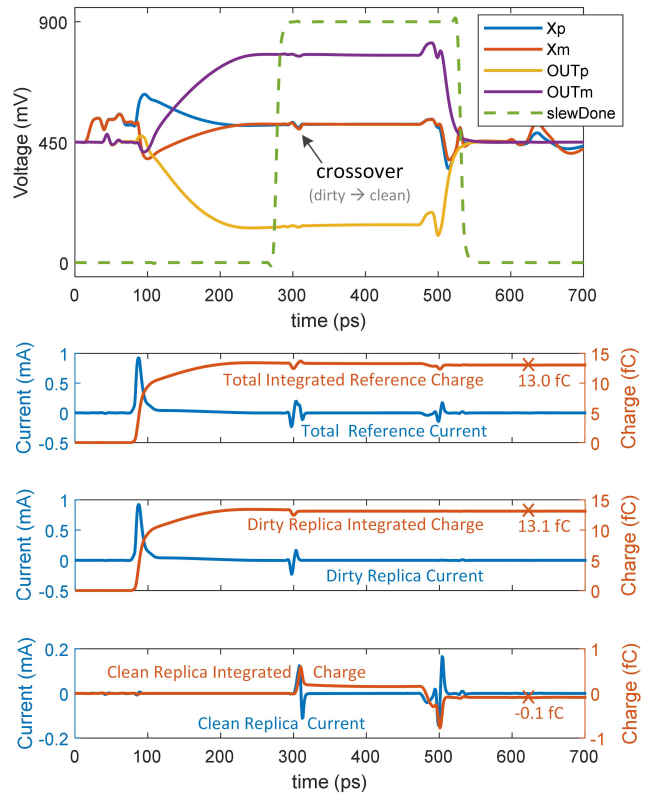

Fig. 2. Example residue amplification waveforms for a fully differential equivalent of the MDAC circuit of Fig. 1(b).

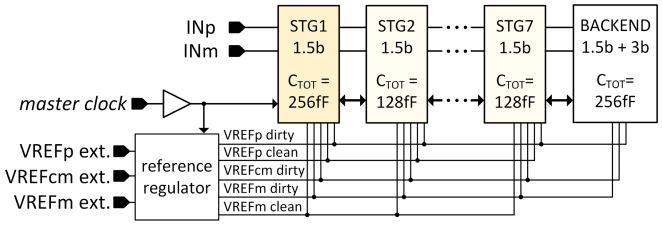

Fig. 3. System overview. A single set of regulated voltages are shared by all stages.

requirements and a "clean" replica with high accuracy requirements but also high (i.e., relaxed) output impedance. During slewing, the dirty replica is connected to the feedback circuit, and during fine settling, the clean replica is connected.

The simulated waveforms in Fig. 2 depict an example residue amplification for the split-reference MDAC of Fig. 1(b). It can be seen from the integrated charge waveform of the dirty replica that it provides an initial signal-dependent spike of 9 fC at roughly 100 ps when the MDAC connects to the output load and another 4 fC during slewing. Most of the charge pulled from the clean replica during fine settling is only temporarily borrowed by the sub-DAC switches, and after these switches are opened, a final integrated charge of -0.1 fC is observed.

B. System-Level Considerations

The split-reference technique involves three key components: the regulator subsystem that generates the replicas (i.e., transmitter), the local circuits in the ADC that pull charge from them (i.e., receiver), and how they interface at the system level (i.e., interconnect). We will begin this discussion at the system level and then dive into the implementations of the regulator and local circuitry in the later subsections.

A top-level overview of the ADC is given in Fig. 3. The 11-bit pipelined ADC uses the asynchronous event-driven timing control, fully dynamic operation, early quantization, and slew-done detection techniques described in [11], the self-resetting ringamp topology of [23], and the passive-hold quantization technique described in [10]. The 1.5-bit stage architecture requires three reference voltages for operation: VREFp, VREFcm, and VREFm. The split-reference regulation subsystem provides five regulated replicas in total, as shown in Fig. 3. Due to the relaxed accuracy requirements of the common-mode reference in this fully differential system, VREFcm does not require a clean replica.

These five replicas are shared by all pipeline stages. It is important to note that the localized switching between dirty/clean replicas at the circuit level solves a critical system-level problem related to shared replicas; when "stage X" is fine settling, "stage Y" might be simultaneously coarse charging and pulling large amounts of charge from the replica. This is particularly true for event-driven pipelines like this one where intra-stage timing relationships are asynchronous [11]. In the conventional single-replica case, this scenario leads to the problem of reference-ripple crosstalk errors, placing a heavy burden on the regulator. The split-reference technique solves this problem by ensuring that all large charge pull is confined to the dirty replicas only, and the clean replicas remain quiet and accurate with minimal amounts of crosstalk.

The key regulator design requirements that must be determined based on system-level needs are the minimum acceptable charge-sourcing capability and the maximum allowable reference error. With regard to charge sourcing, each replica must be able to handle the maximum charge-pull scenario. Charge equations describing local circuit operation can be used to solve for this worst case per-cycle charge pull. However, details of circuit implementation and parasitic modeling can alter this result significantly, so simulation-based extraction of the value is ultimately necessary. Repeating the analysis for each circuit that is connected to the shared global replicas and summing all values together provides a spec for the minimum amount of charge per ADC conversion that each regulated replica must be capable of delivering.

Errors in the clean replica have a direct conversion path to the ADC output, and determining a spec for the maximum allowable reference error is fairly straightforward. The left plot of Fig. 4 shows the result of an experiment where the ADC is provided with ideal references except for an added Gaussian noise source superimposed onto the clean replica. Then, the un-quantized analog residue of the pipeline backend is combined with the digital outputs of the ADC to give an estimate of total SNDR with SQNR and thermal noise excluded. From this, we see that to achieve a regulator-related SNDR target of 70 dB, the total rms error power of less than 100 μV is required.

Determining an accuracy spec for the dirty replica is more complicated due to multiple filtering mechanisms that are

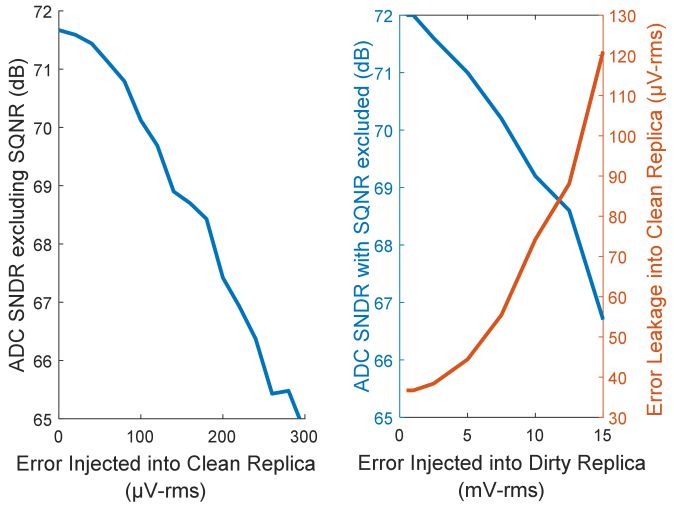

Fig. 4. Influence of reference error in clean (left) and dirty (right) replicas on ADC performance. An SNDR ceiling near 72 dB exists due to other unrelated error sources in the ADC.

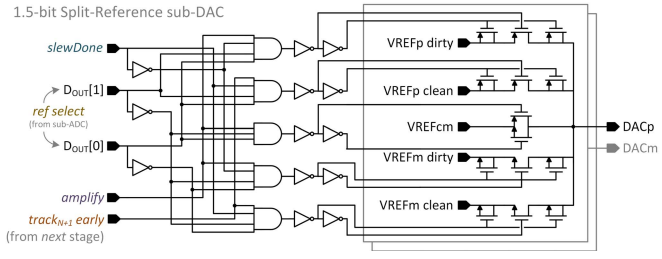

Fig. 5. Detailed implementation of the 1.5-bit sub-DAC in Fig. 1(b).

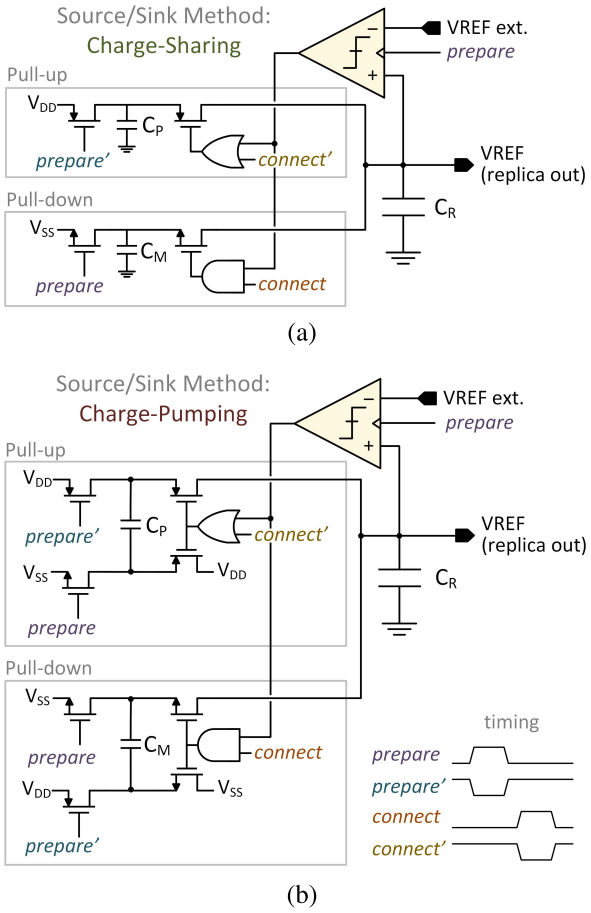

Fig. 6. Discrete-time fully dynamic regulators using (a) charge-sharing and (b) charge-pumping methods of charge delivery.

involved. The two primary transfer paths of dirty replica error to the ADC output are: (a) charge-transfer into the clean replica through the MDAC capacitors and (b) error in the dirty replica causing the ringamp to settle toward the wrong target value during the coarse-charge phase, consequently creating a larger settling burden for the residue amplifier in the fine-settling phase and thus more incomplete settling error. The right plot of Fig. 4 shows a second experiment where the dirty replica is an ideal voltage source superimposed with a Gaussian noise source, and the clean replica is a realistic regulator implementation that meets its own 100 μV accuracy spec. Approximately 8 mV-rms of error in the dirty replica produces 70-dB SNDR at the ADC output. Notably, this occurs with less than 60 μV of error leakage observed on the clean replica due to error transfer path "(a)". This is well below the 100-μV level of clean replica noise determined to correspond to 70-dB SNDR and implies that error transfer path "(b)" is more dominant. To improve SNDR for this particular scenario, it would therefore be necessary to increase the duration of fine settling.

C. ADC Integration

Integration of the split-reference technique into the ADC core consists of modifying all reference injection points to support dual injection. The 1.5-bit sub-DAC in Fig. 1(b) is the injection point of this architecture. It is shown in more detail in Fig. 5. Switching from the dirty to clean replica is governed by the *slewDone* signal. This signal can be generated with either a fixed time delay or a slew-done detector. Both options are conveniently already available in this system [11]. To avoid shorting dirty/clean replicas together, the crossover must be nonoverlapping. The sub-DAC selection logic implements this with a nonoverlap window of only one gate delay. Clock feedthrough glitches during crossover are canceled by complementary dummy switches at the sub-DAC outputs. This minimizes the size of the crossover glitch visible in Fig. 2 and its impact on settling. The main penalty associated with split-reference integration is about 30 ps of lost settling time when the MDAC feedback is cut. To isolate the clean replicas from large voltage swings in the MDAC at the end of *amplify* due to ringamp power-down, the sub-DAC is first disconnected using the slightly earlier $track_{N+1}$ *early* control line of the subsequent pipeline stage [11].

D. Fully-Dynamic Regulator

The fully dynamic event-driven operation of this ADC offers several advantages for both design- and run-time reconfigurability [11]. To maintain this advantage at the system level, the regulation system must be fully dynamic too. This can be achieved using a discrete-time regulator similar to [12]. Consider the implementation drawn in Fig. 6(a). The comparator measures the polarity of error between the regulated replica's output and an external reference voltage

at discrete time intervals. Based on the comparator decision, corrective feedback in the form of charge-packet updates are sourced/sunk by capacitors C_P/C_M.

1) Deterministic Ripple: During the elapsed time between each comparator decision, there is no feedback. This guarantees loop stability [12] but also increases reference ripple. All charge pulled by the ADC during each interval is taken from the reservoir C_R and is only replenished once per cycle. As discussed earlier in Section II-B, the size of the charge packet that must be delivered each interval is set by the charge-pull requirements of the ADC. Therefore, the only way to reduce ripple is to either increase C_R or decrease the interval between updates. In a low-resolution ADC, this ripple constraint is manageable, due to the combination of low charge-pull and low accuracy requirements. In [12], a 6.2 ENOB ADC is shown to be possible using a single-replica discrete-time regulation system. However, as system accuracy targets are increased, this approach quickly becomes impractical. To meet the 70-dB reference accuracy targets of the design presented here, decoupling on the order of nanofarads is ultimately required for a single-replica solution. The split-reference regulation technique overcomes this practical limitation, enabling high-accuracy discrete-time reference regulation.

Based on the system-level error requirements discussed previously in Section II-B and Fig. 4, the dirty replica can tolerate a relatively large amount of deterministic ripple without impacting system accuracy. This relaxed design spec can be used to limit the size of C_R to a practical value despite the large charge-sourcing requirement. For example, in the ADC described here, a value of $C_R = 35$ pF in the dirty replicas results in a very manageable ripple of less than 1 mV-rms. In the case of the clean replica, the relaxed charge-sourcing requirements make it much easier to meet the higher accuracy requirements, and in this ADC, a value of $C_R = 15$ pF in the clean replicas gives less than 100-μV-rms ripple. By comparison, it would require more than 370 pF to achieve the same 100-μV-rms ripple using a conventional single-replica system.

2) Comparator Noise: While deterministic ripple is the dominant error source in the dirty replica, comparator noise typically dominates in the clean replica. Fig. 7 shows the relationship between the input-referred noise of the comparator versus total rms error of random noise plus deterministic ripple at the regulated replica output. For the clean replica (i.e., left plot), a comparator with 200-μV noise and $C_R = 45$-pF results in 76-μV-rms total error due to noise plus deterministic ripple, leaving a healthy design margin for additional dirty-to-clean error leakage and PVT variability in the total 100 μV budget. For the dirty replica (i.e., right plot), comparator noise is overshadowed by the relatively large deterministic ripple error. Using a comparator with 200 μV input-referred noise and $C_R = 45$-pF results in 850-μV-rms total noise. Increasing the comparator noise by five times to 1 mV only raises the total combined noise to 1 mV-rms. The regulator comparator of the dirty replica can therefore be designed for low power consumption. When these results are interpreted from the perspective of a conventional single-replica regulator, it is found that to achieve the same final 76 μV-rms total error, a comparator with 200 μV noise would require $C_R > 1.1$ nF.

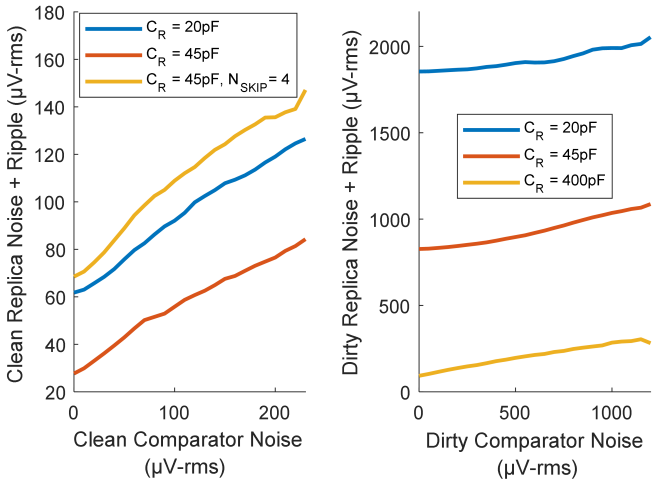

Fig. 7. Influence of comparator noise (input-referred) on the total error due to noise and deterministic ripple at the regulated replica output.

3) Charge-Delivery Method: Fig. 6 shows two different approaches to sourcing and/or sinking charge. Fig. 6(a) uses a passive charge-sharing approach. In the *prepare* phase capacitors C_P/C_M are precharged to V_{DD}/V_{SS}. Then, in the *connect* phase, charge is transferred to C_R. To deliver the required amount of charge per cycle, Q_{PACKET}, the required source/sink capacitor sizes are given by the following simplified equations that assume $C_R \gg C_P, C_M$:

$$C_P \geq \frac{Q_{PACKET}}{V_{DD} - VREF}, \quad C_M \geq \frac{-Q_{PACKET}}{VREF}. \tag{1}$$

Note that the required size of C_P asymptotically reaches infinity at VREF $= V_{DD}$, as does C_M at VREF $= V_{SS} = 0$V. In systems requiring reference voltages close to the supplies, this means that for constant capacitor values, the charge delivered to the load (and ripple it injects) varies significantly as a function of VREF. This is a key drawback of passive charge-sharing. It can be overcome with the charge-pumping implementation of Fig. 6(b). Again, using the simplifying assumption of $C_R \gg C_P, C_M$, the capacitor design equations are

$$C_P \geq \frac{Q_{PACKET}}{2V_{DD} - VREF}, \quad C_M \geq \frac{-Q_{PACKET}}{V_{DD} + VREF}. \tag{2}$$

The additional $C \cdot V_{DD}$ of charge transferred by this scheme doubles the supported voltage range. For VREF values close to the supplies, the required sizes of C_P and C_M are significantly reduced and their optimal values are much less sensitive to the variation of VREF. Overall, this provides a more robust and reconfigurable implementation than charge sharing.[1]

[1]In this design a charge-sharing implementation is used, but due to these important advantages, the charge-pumping scheme is recommended for future work. In many cases, it is even worth considering a hybrid scheme where one capacitor charge-shares and the other charge-pumps; closer inspection of (1) reveals that the more one capacitor approaches the asymptotic limit, the further the other one moves away from it.

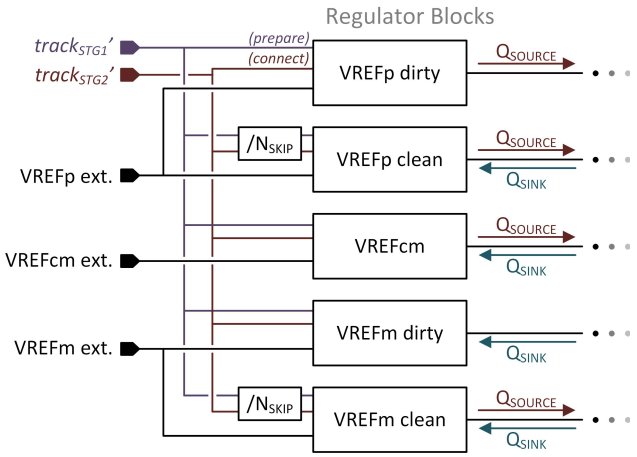

Fig. 8. Overview of the complete regulation subsystem.

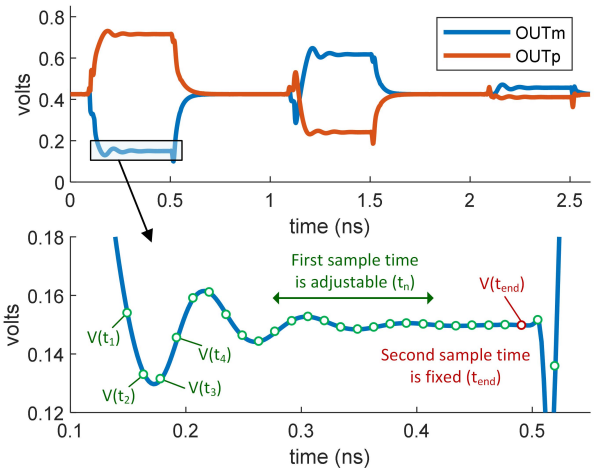

Fig. 9. Self-referenced deterministic sampling concept. By comparing the waveform to its own final value, only the relative variation is sampled.

4) Comparator: The regulator comparator is a two-stage structure with a digitally tunable threshold similar to [24, Fig. 5]. Depending on the regulated level (VREFp, VREFcm, and VREFm), either a PMOS- or NMOS-input implementation is used. The tunable threshold plus a programmable option to short the comparator inputs (not shown in Fig. 6) allows the comparator offset to be nulled, maintaining absolute accuracy of the reference and also eliminating static offset between dirty and clean replicas.

5) Complete Regulator Subsystem: The complete regulator subsystem is shown in Fig. 8. The input impedance seen by the external references is very high because they only connect to the CMOS gates of the regulator comparators. This allows for any type of external reference to be provided without buffering, including a bandgap or resistor divider. The nature of operation in the MDACs of the pipeline is such that charge is always pulled from VREFp dirty each cycle and always dumped into VREFm dirty. Considering this, the dirty regulators are configured to operate with only uni-directional feedback by disabling one of the charge-delivery paths in Fig. 6. This leads to less total deterministic ripple error on those replicas compared to bi-directional feedback. The other three replicas (VREFp clean, VREFm clean, and VREFcm) operate with bi-directional feedback.

The maximum regulator refresh rate is the same as the ADC *master clock* rate. The *prepare* phase of Fig. 6 begins at the end of the pipeline stage 1 *track* phase, and the *connect* phase begins at the end of the pipeline stage 2 *track* phase, as shown in Fig. 8. This timing scheme is chosen to provide stage 1 of the pipeline with the cleanest reference possible because it is the most critical stage for accuracy. Additional details pertaining to the timing control of this design are described in [11].

A clock decimation feature shown in Fig. 8 is included for the clean regulators such that the refresh rate can be reduced by a factor N_{SKIP}. This tunable parameter provides an additional degree of freedom for optimizing the tradeoff between accuracy, power consumption, and the size of C_R.

III. Transient Waveform Capture

The time-domain settling behavior of a feedback-based residue amplifier is a function of its stability and phase margin. In the design and validation of ring amplifiers, transient methods of analysis are frequently used in simulation due to the inherently transient and large-signal aspects of ringamp operation. In physical systems, however, such methods of analysis are rarely used, even though they would be useful in a variety of applications, including test, debug, adaptive control, and calibration. This is mainly due to the impractical size, complexity, and power required to provide a so-called "scope-on-chip" that can probe on-chip nodes. In this section, we describe a scope-on-chip solution specialized for high-resolution background capture of low-amplitude signals that manage to overcome these hurdles. It has a low-complexity implementation with compact layout. Although the specific implementation presented here utilizes ADC backend hardware out of convenience, we show how a fully self-contained implementation is also possible, allowing it to be used in any type of discrete-time switched-capacitor system. Later, in Section IV, we will show how captured waveform data can be used for PVT-robust adaptive bias control for ringamps.

A. Key Challenges (and Solutions)

There are three main obstacles preventing the practical implementation of a scope-on-chip for high-speed residue amplifier settling waveform capture. In this section, we will describe each problem and a corresponding solution, building up to the implementation in Fig. 10. Each solution is made possible by customizing the capture system to the specific requirements of low-amplitude deterministic waveforms.

1) Input Amplitude and Resolution Requirements: Consider the simulated output residue waveforms in the top half of Fig. 9 that show three consecutive periods of ringamp operation. Each *amplify* phase is about 420 ps wide, sufficient for a 1 GS/s timing budget. Ultimately, it is the portion of the waveform that occurs after initial slewing that provides the key

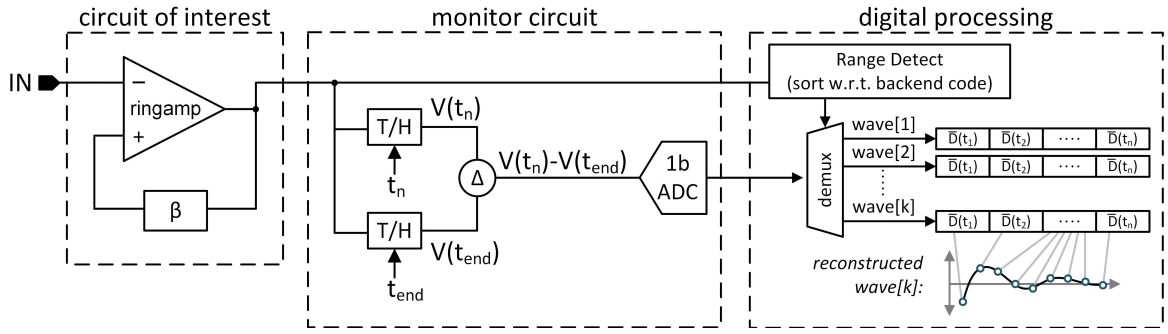

Fig. 10. Concept of the waveform capture and reconstruction procedure. A small monitor circuit quantizes many samples that are sorted into bins and averaged in digital domain with respect to backend code range.

information regarding settling behavior and stability. We wish to know this settling behavior with sub-millivolt precision. Due to the large voltage swing of these residue waveforms, to acquire this information directly would require a capture-ADC with very high resolution.[2]

The solution, shown in the lower half of Fig. 9, is to decrease the capture-ADC input range requirement. This is achieved by using a self-referenced sampling scheme to narrow in on the small voltage range of interest. For each time-series capture point "n", the capture-ADC is fed the difference between the voltage sampled at time point t_n and a fixed reference sample: $V_{\rm IN} = V(t_n) - V(t_{\rm end})$. In theory, any sample can be used for the reference value $V(t_{\rm end})$, although, for practical reasons, using a sample near the end of the settling time is best. Effectively, self-referenced sampling removes the much larger "main" residue signal, and only the residual settling variation remains. This significantly reduces the input range requirement of the capture-ADC and, consequently, the resolution that is needed for sub-millivolt precision.

2) Sampling Rate Requirements: In the 420-ps *amplify* phase during which the residue is generated, roughly 320 ps is used for settling. To sample 64 time-series data points in this window, the capture-ADC would need to operate at an impractical rate of 200 GS/s. The solution to this problem stems from the insight that although random noise does influence the ringamp settling waveform to a limited extent, it is still mostly a deterministic process. In other words, all cycles with the same final value will also have roughly the same settling waveform. Thus, rather than measure a single waveform at 64 discrete-time points, it is approximately the same to measure one time point for 64 separate but identical waveforms. This allows the capture-ADC sampling rate to be reduced to arbitrarily low speeds because cycles can also be skipped. Grouping similar waveforms together is simply a matter of sorting with respect to the final value. However, this sorting information is the same information that was removed through self-referenced sampling, so an auxiliary circuit is needed to extract it. In a pipelined ADC, the digital output of the backend is conveniently available to serve this function. In circuits where no backend is available, a simple range detector circuit built using two comparators with thresholds above and below the target level can be used. Notably, the sorting range can be made quite large (e.g., tens or even hundreds of millivolts) and still lead to useful waveform capture. This is because the close-in settling behavior of residues with similar final value (i.e., backend code) tends to also be similar. The self-referenced sampling removes the offsets from all of these similar residues, allowing their data to be combined without major loss of information.

3) Noise, Input Capacitance, Complexity, and Area: The solutions listed above overcome resolution and speed obstacles to make a practical implementation of the capture-ADC possible, but they are still not enough to enable a lightweight design that can "invisibly" probe many nodes in the system. The small input range provided by self-referenced sampling still does not reduce thermal noise requirements such as sampled kT/C noise, and to achieve acceptable SNR, a large sampling capacitor is still required. Input-referred noise of the capture-ADC itself also remains a problem, and this leads to a challenging, power-hungry design despite the low resolution. Ideally, a capture-ADC with negligibly small input capacitance, design complexity, area, power, and effect on system performance is desired.

This can be achieved by implementing the capture-ADC with a 1-bit stochastic ADC [25]. The stochastic approach offers an elegant solution for multi-bit quantization of low-amplitude signals by utilizing random noise as a tool for quantization. The basic idea is to use a single comparator to quantize a small sampled input voltage multiple times. Due to noise in the sampler and comparator, the 1-bit ADC output will be a stream of "1"s and "0"s whose average value describes the location of the input voltage on the ADC's noise cumulative density function (CDF). Through additional processing steps, this value can be linearized and converted into units of volts [10]. However, for this application, we skip these extra processing steps to keep digital complexity low; we are only interested in the general characteristic of the settling waveform and can tolerate some distortion. Thus, a multi-bit quantization of the input is obtained simply by calculating the average of the comparator output.

Performing this stochastic quantization process on a single input sample as in [25] solves the multi-bit quantization challenge, but it still does not solve the kT/C problem. This second challenge can be solved with the multi-sample

[2]In certain architectures where a high-resolution backend is available, this can be accomplished by re-using the backend hardware [18]. However, this only works as a foreground technique and the capture resolution will likely still be quite coarse.

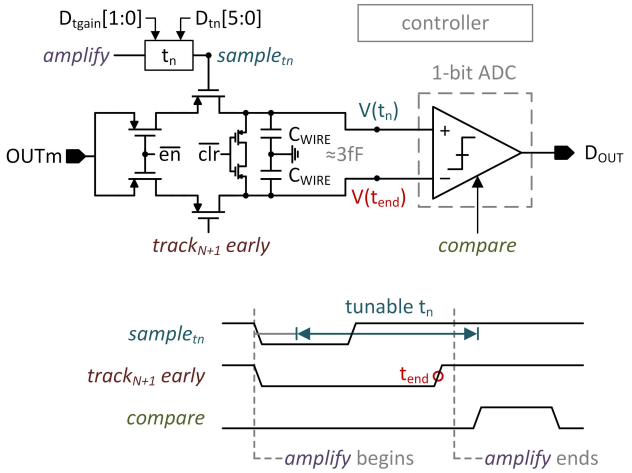

Fig. 11. Simplified schematic of the monitor circuit.

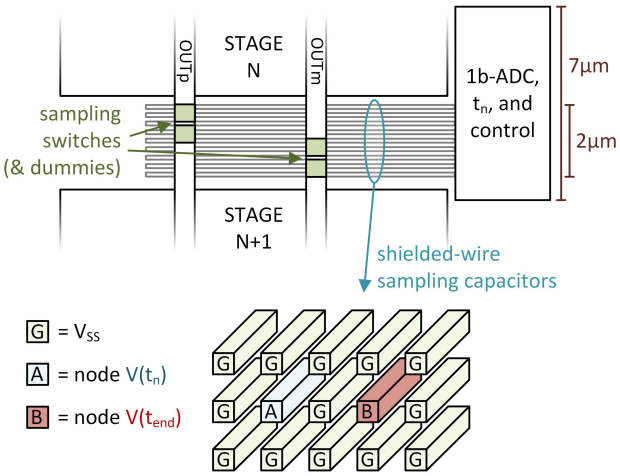

Fig. 12. Layout details of Fig. 11 monitor circuit.

averaging approach mentioned in [10]. Rather than make multiple comparisons of a single sample, it is equivalent to make a single comparison for multiple samples, provided that each sample is of the same input voltage. In this way, kT/C noise gets rolled into the ADC's total input-referred noise CDF that is used for quantization, and the input capacitance can be made arbitrarily small. This multi-sample averaging approach is inherently compatible with the multi-sample sorting procedure described earlier. Details about how these two techniques are combined is described next.

B. Practical Implementation

A block diagram of the complete capture system is shown in Fig. 10. The monitor circuit (i.e., capture-ADC) consists of a self-referenced sampler and a comparator. The 1-bit ADC's output is fed to a digital processing block and sorted. The "Range Detect" block sorts data into rows with respect to ringamp output voltage. This allows multiple waveforms to be reconstructed in parallel, corresponding to different ringamp output levels. However, if a very low digital complexity is desired, just a single row is sufficient to store the capture data for a single subrange of interest. In this implementation, sorting is done by demuxing the data according to the MSBs quantized by the pipeline backend. The columns in each row (bins) correspond to the discrete steps (64 in this design) of the programmable time delay t_n. The stream of 1-bit values that are fed into each bin is averaged to obtain the final multi-bit scalar quantization result.

A simplified schematic of the monitor circuit block in Fig. 10 is drawn in Fig. 11. Although a fully differential implementation is possible, for this design, we choose a single-ended implementation for simplicity, attaching the monitor to the negative ringamp output. To provide symmetric loading of the ringamp, a dummy sampling network is also connected to the ringamp positive output. Note that the sampling network uses only PMOS switches, so it can only accurately sample waves corresponding to the lower half of the ringamp's output residue range. For our uses, this is sufficient. The sampler can be disabled by asserting $\overline{en}$. The comparator with PMOS input pair and digitally tunable threshold is similar to the structure in [24, Fig. 5]. The comparator threshold can be calibrated by de-asserting $\overline{clr}$. The beginning of the pipeline stage's *amplify* phase is used as the reference edge for triggering the variable-delay sampling clock used to acquire the $V(t_n)$ samples. The tunable delay cell has a 6-bit main linear tuning control plus three coarse gain settings that adjust the LSB size to support the capture of different settling-time durations. To ensure that the $V(t_{\text{end}})$ sample does not disturb residue settling or the main MDAC sampler, it is sampled using the same physical wire that samples the residue in the following pipeline stage (*track*$_{N+1}$ *early*).

Due to the multi-sample averaging technique, the kT/C requirement of the C_{WIRE} sampling capacitors in Fig. 11 can be reduced to any desired amount. They should, therefore, be made as small as possible to minimize capacitive loading of the circuit node that the monitor probes. As shown in the layout floorplan of Fig. 12, C_{WIRE} is formed from the parasitic coupling of shielded routing wires that connect the sampling switch matrix placed below the residue amplifier outputs to the comparator placed on the right edge of the stage. The total distributed capacitance of C_{WIRE}, including the parasitic contributions from switches and comparator, is roughly 3 fF. The simplicity of this circuit leads to a very compact layout. The switches and C_{WIRE} routings are only 2 μm high and the main layout block consisting of the comparator, tunable delay cell, and digital controller is only 7 μm high by 22 μm wide. Power consumption of the monitor circuit is negligible due to the lightweight design, and even more so, because it is possible to operate at a reduced rate by skipping cycles. The 1-bit output from the monitor circuit can be multiplexed with other monitor circuits in the system into a single shared digital processing block. In this design, the digital processing block in Fig. 10 is implemented off-chip.

IV. Measured Performance

A. General Performance

The ADC is fabricated in a 16-nm CMOS FinFET technology, occupying an active area of 0.095 mm^2

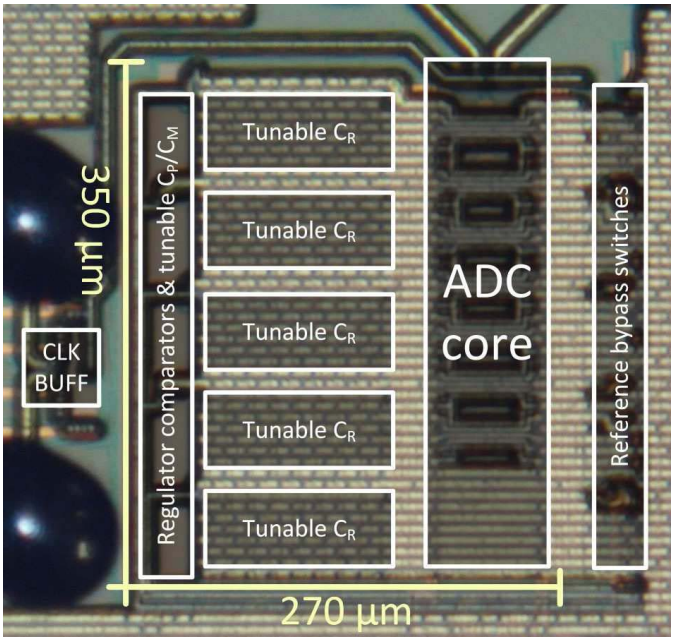

Fig. 13. Die photograph.

(270 μm × 350 μm). A die photograph is shown in Fig. 13. General system performance is summarized in Fig. 14. The 11-bit ADC operates from a single 900-mV supply with VREFm = 50 mV and VREFp = 850 mV. Except where otherwise noted, all measurements are performed with reference regulation enabled, using the values tabulated in Fig. 15(a). A fixed set of digital firmware settings and linear stage-gain correction coefficients are used for all measurements, except where otherwise noted.

As seen in the decimated digital output spectrum ($N_{\mathrm{DEC}} = 457$) of Fig. 14(d), at the maximum sampling rate of 1 GS/s with a −0.3 dBFS, 500-MHz input tone, the ADC achieves 59.5-dB SNDR, 75.9-dB SFDR, and 69.9-dB THD. The full-scale input range is 1.6-V pk-pk differential. Total power consumption is 10.9 mW with 10.0 mW (92%) consumed by the core ADC and 0.9 mW (8%) by the regulator subsystem. The sweep of clock frequency in Fig. 14(e) demonstrates that the ADC, including all reference regulation, exhibits fully dynamic power consumption. This results in a near-constant energy efficiency and performance profile from 1-MS/s to 1-GS/s and Walden and Schreier figures of merit of 14.1 fJ/conversion-step and 166.1 dB, respectively.

Table I compares this work with the state-of-the-art for medium-resolution ADCs operating in the vicinity of 1 GS/s [11], [26]–[31]. This design is the only one to include integrated reference regulation and achieves highly competitive performance across all metrics. The comparison also illustrates how recent advances in scalable amplification have helped to make high-speed single-channel pipeline architectures competitive with interleaved SAR ADCs, without the associated drawbacks of interleaving.

TABLE I
COMPARISON WITH RELATED STATE-OF-THE-ART ADCS

	This Work	Jiang ISSCC 2019 [26]	Hershberg ISSCC 2019 [11]	Lagos CICC 2018 [27]	Venca ISSCC 2016 [28]	Nam VLSI 2016 [29]	Lien VLSI 2016 [30]	Sung ISSCC 2015 [31]
Architecture	**Pipeline**	Pipe-SAR	Pipeline	Pipeline	TI SAR ΔΣ	TI SAR ΔΣ	TI Pipe-SAR	TI FATI-SAR
Sampling Rate [MS/s]	**1000**	1000	600	1000	600	1600	800	1600
Number of Channels	**1**	1	1	1	4	8	4	12
Per-Channel Rate [MS/s]	**1000**	1000	600	1000	150	200	200	133
Technology [nm]	**16**	28	16	28	28	65	28	65
ENOB Nyquist [bit]	**9.6**	9.7	9.7	9.1	9.3	10.6	9.8	9.0
SNDR Nyquist [dB]	**59.5**	60.0	60.2	56.6	58.0	65.3	60.8	56.1
SFDR Nyquist [dB]	**75.9**	74.6	78.3	73.1	66.0	65.3	75.0	61.2
Power [mW]	**10.9**	7.6	6.0	24.8	26.5	37.7	14.6	17.3
$\mathrm{FoM_{Walden}}$ [fJ/c.step]	**14.1**	9.3	12.0	45.0	68.0	17.8	20.0	21.0
$\mathrm{FoM_{Schreier}}$ [dB]	**166.1**	168.2	167.2	159.6	158.5	168.6	165.2	162.8
Core Area [mm²]	**0.095**	0.009	0.037	0.540	0.042	0.900	0.093	0.360
Calibration	**gain**	gain	gain	gain	offset, gain, mismatch	offset, gain, timing skew	offset, gain	offset
Reference Regulation	**yes**	no	no	no	no	no	no	no

B. Regulation Subsystem

The total impact of the regulator on overall ADC performance is a 0.3 dB reduction in SNDR, 0.5 dB increase in THD, and no measurable change in SFDR.[3] The *slewDone* signal used for initiating sub-DAC crossover is generated using a statically configured time-delay control. Measurements reveal that the other option of using a dynamically generated crossover with a slew-done detector produces a slightly less optimal timing and degrades SNDR by 2.5 dB at 1 GS/s. However, when the system clock rate is reduced to 800 MS/s and the ringamps are given more time to settle, both crossover generation approaches exhibit the same system-level performance.

All C_P, C_M, and C_R capacitors in the regulator subsystem are made tunable. This significantly increases the regulator layout area but provides valuable test capabilities. Fig. 15 reports several experiments where one control variable is swept at a time, with all others held at the values of Fig. 15(a). The charge pulled from the references by the ADC varies as a function of input signal frequency and amplitude. To best capture this signal-dependent effect, all sweeps are conducted at one of the worst case input conditions of −0.3 dBFS Nyquist input tone where a large charge-intake beat-frequency is present, placing the heaviest burden on the active ripple suppression of the regulator.

Fig. 15(b) shows the impact of the clean replica decimation factor on performance, showing a clear power-saving benefit; for the nominal setting of $N_{\mathrm{SKIP}} = 12$, only 11% of total regulator power is consumed by the clean replicas. For decimation values above this, the SNDR begins to degrade because the clean replicas can no longer deliver enough charge for full ripple suppression. This can be overcome by increasing the sizes of C_P and C_M in the clean replicas, but the sweep of the clean replica charge delivery capacitor sizes in Fig. 15(c) shows the tradeoff that comes along with this. The optimal sizes for the clean replica capacitors are values

[3]This performance impact is determined by comparing the default mode to a full bypass mode where the regulator is disabled, the MDACs operate from single reference lines with no switching from dirty to clean, and all replicas are connected to off-chip references complemented by a large array of on-chip decoupling.

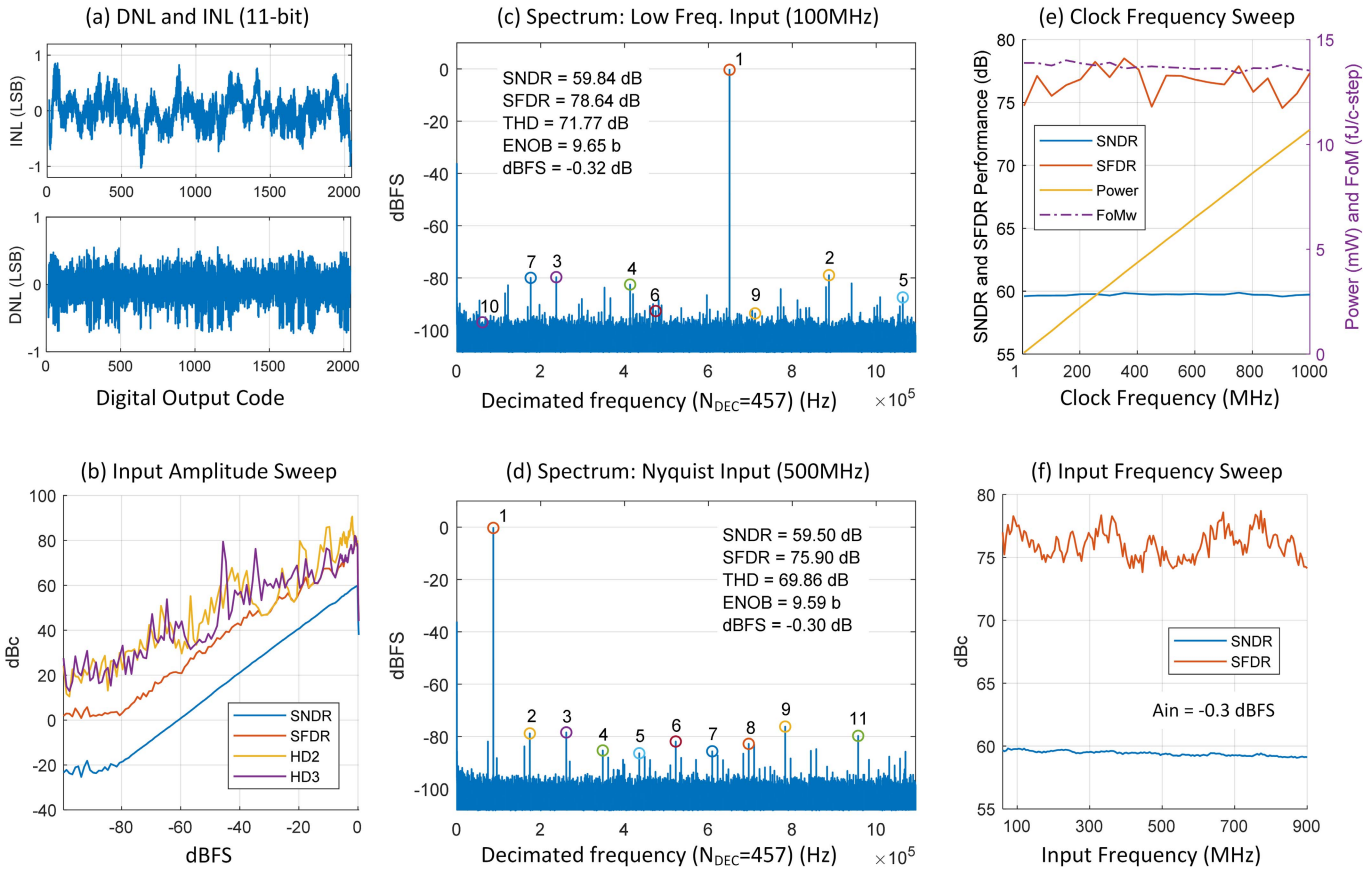

Fig. 14. Summary of measured ADC performance.

that are sufficient to deliver the required reference charge, but nothing more. Above this optimum value, performance can be seen to slowly decline because of the larger charge-packet updates that are generated, creating larger deterministic ripple. This trend can also be observed in the dirty replicas [visible in the zoomed-in view in Fig. 15(d)], but the effect is less pronounced due to the inherent filtering of dirty replica errors. Fig. 15(d) also shows the only two parameters in the system where a steep roll-off in performance can be observed; if the dirty replicas cannot source/sink sufficient charge, both dirty and clean replicas will begin to collapse toward mid-rail.[4]

Fig. 15(f) shows how a larger C_R improves error filtering. The same sweeps are repeated in Fig. 15(e) with decimation disabled by setting $N_{\text{SKIP}} = 1$. Comparison of Fig. 15(e) and (f) (note the different y-scales) plus Fig. 15(b) illustrates the three-way tradeoff between accuracy, efficiency, and the size of C_R. It is possible to reduce the size of C_R by either spending more power or sacrificing some accuracy.

[4]A key reason why all other parameters produce gentle performance roll-off is because the change transfer path between clean and dirty replicas inside the pipeline stage MDAC creates a switched-RC low-pass filter that powers the clean replica with charge from the dirty, whether the clean replica regulation is enabled or not. This filtering effect can even be deliberately engineered. However, the filter does not reject content that falls inside the filter passband (i.e., signal-dependent beat frequencies) and ultimately either active suppression (as used here) or a more complex feedback control in the dirty replica to filter out low-frequency content is necessary for robust operation.

In Section II-D1, it was determined that $C_R > 35$ pF would provide a very large design margin for the dirty replicas. Moreover, in Section II-D2, it was determined that $C_R = 45$ pF for the clean replicas provides a comfortable design margin when $N_{\text{SKIP}} = 1$. These two expectations align well with the measurements of Fig. 15. Sweeping C_R of the dirty replicas in Fig. 15(e) and (f) shows no measurable impact on performance, indicating that 45 pF is much larger than necessary. Fig. 15(e) shows that $C_R = 45$ pF is sufficient for the clean replicas when $N_{\text{SKIP}} = 1$ but leads to some performance degradation in Fig. 15(f) when $N_{\text{SKIP}} = 12$. Ultimately, it is clear from these results that the sizes of C_R and overall layout area of the regulator subsystem can be dramatically reduced in future implementations.

C. Transient Waveform Capture

Monitor circuits probe all seven residue amplifier outputs in the pipeline. Fig. 16 shows four waveforms captured by the monitor circuit located at the output of the first pipeline stage. Each waveform corresponds to a different ringamp bias condition: unstable, under-damped, critically damped, and over-damped. The backend-code sorting range is set to 15% of full scale (240 mV) for these measurements. This wide sorting range indicates that the digital Range Detect block in Fig. 10 can also be implemented as a simple analog detector in systems where a backend is not available, increasing the

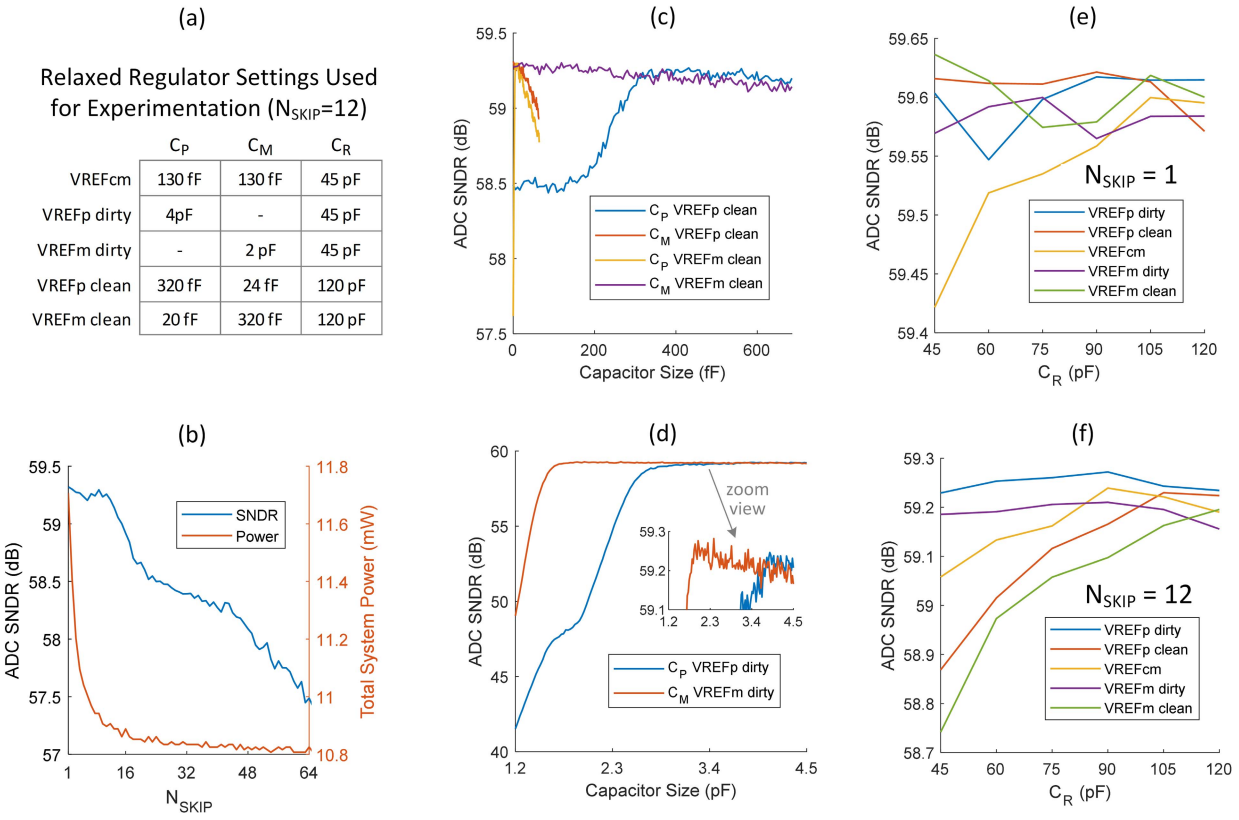

	C_P	C_M	C_R
VREFcm	130 fF	130 fF	45 pF
VREFp dirty	4pF	-	45 pF
VREFm dirty	-	2 pF	45 pF
VREFp clean	320 fF	24 fF	120 pF
VREFm clean	20 fF	320 fF	120 pF

Fig. 15. Characterization of the reference regulation subsystem.

generality of the technique. The trace denoted "sample flag" is a flag-bit generated by the controller in Fig. 11 that compares the $sample_{tn}$ and $track_{N+1}$ *early* clock edges, indicating which index in the time-series capture data corresponds to the next pipeline stage's sampling moment. There is a slight time delay in this detection circuit, so the fixed-length "chop" range indicated in Fig. 16 can be used to reliably locate the sample corresponding to $V(t_{\text{end}})$. Beyond t_{end}, kickback of the next-stage sampler appears in the capture data, so knowledge of where the "clean" part of the waveform ends is useful for additional postprocessing.

From visual inspection of Fig. 16, it is fairly obvious to the trained human eye which of the four waveforms is the closest to critically damped settling. Similarly, it is possible to define an objective function that can extract this feature and use it for adaptive bias control. Consider the following objective function that compares two sub-vectors of the time-series waveform data[1 : 64] shown in Fig. 16:

$$\text{dampingOptimality}(\text{data}, a, b) = \frac{\text{Stdev}(\text{data}[1:b])}{\text{Stdev}(\text{data}[a:b])}. \quad (3)$$

For this particular implementation, a is chosen to be 20% of b. The output of this function versus ringamp bias control code, plotted in Fig. 17, correctly identifies the optimum ringamp biasing point, as evidenced by the overlay of SNDR at the ADC output. Essentially, (3) calculates what percentage of the waveform's total signal energy occurs at the beginning of the settling time. The highest initial concentration occurs in the critically damped case. Intuitively, this also makes sense; faster settling means a faster rate of amplitude decay.

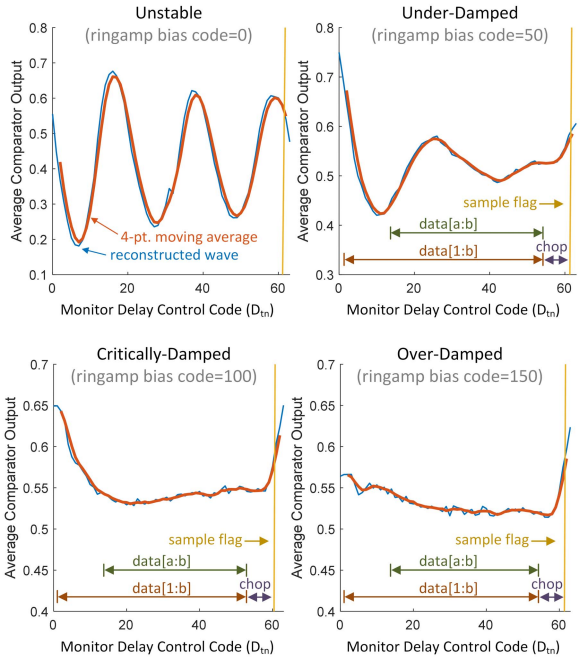

Fig. 16. Reconstructed waveforms captured by the monitor circuit placed at the output of the first pipeline stage. The y-axis is the averaged monitor output, bounded from 0 to 1 (note different y-scales and waveform amplitudes of each subplot).

By making a relative comparison between two parts of the same data set (i.e., expressing as a ratio), (3) ensures immunity to fluctuations in absolute waveform amplitude that can result from either changes in the ringamp when the bias is adjusted or

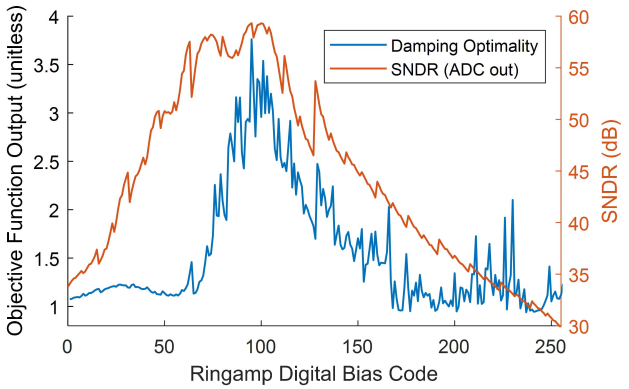

Fig. 17. Output of (3) damping optimality estimator, with SNDR of the ADC digital output provided for comparison.

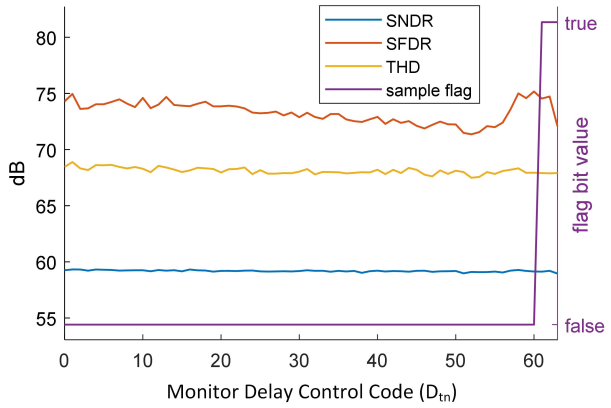

Fig. 18. Influence of monitor circuit background operation on the main ADC performance.

PVT variation affecting the voltage-to-sigma conversion gain of the stochastic ADC.

Fig. 18 characterizes the impact of background operation on overall ADC performance for the monitor circuit at the output of the first pipeline stage. With respect to delay control code D_{tn} of Fig. 11, SNDR and THD are negligibly affected. It can be seen that SFDR degrades by about 3 dB in the worst case. Acquisition of $V(t_{\text{end}})$ has no impact on performance because it samples at the same instant as the main signal path sampler. The acquisition of the $V(t_n)$ sample, however, does disturb the residue settling slightly due to charge injection and clock feedthrough of the sampler kicking into the ringamp output. The closer this occurs to the end of the settling time, the less time the ringamp has to re-settle this kick. Improving the sampler design with well-known charge injection and clock feedthrough cancellation techniques is therefore recommended for future work.

V. Conclusion

This article has explored two main topics. First: the technique of splitting each reference voltage into multiple regulated replicas with different accuracy and current-sourcing capabilities, and switching between them at the circuit level so as to relax the overall regulator design requirements. Implementing this concept using discrete-time regulator blocks extends the benefits of a fully dynamic ADC core to the system level. Consuming only 8% of total power, this approach to regulation provides a power-efficient solution that is generalizable to several classes of ADC. Possibilities for future work on the topic include methods for reducing total required reservoir capacitance (and area) and methods for switching between references with zero time penalty.

The second key contribution of this article is a technique for high-resolution waveform capture of small-amplitude transient signals. It offers a lightweight, noninvasive tool for gaining unique insight into time-domain effects in physical systems. It is shown how this information can be processed to form an objective function that is useful for adaptive bias control and PVT tracking. Other possible applications include characterization and debug features for SoCs, such as a power-supply glitch monitor that could be used for debugging deterministic supply noise. Possibilities for future work include a fully differential monitor circuit implementation and a kickback-free sampler for cleaner background operation.

From a system perspective, the design reported here demonstrates how high-efficiency residue amplification using ring amplifiers enables pipeline architectures that offer a competitive alternative to time-interleaved SAR ADCs. The higher per-channel speed of the pipeline offers a compelling advantage when pushing to higher sampling rates due to a lower interleave factor and relaxation of key frontend sampling design challenges.

References

[1] C.-C. Liu, M.-C. Huang, and Y.-H. Tu, "A 12 bit 100 MS/s SAR-assisted digital-slope ADC," *IEEE J. Solid-State Circuits*, vol. 51, no. 12, pp. 2941–2950, Aug. 2016.

[2] F. van der Goes *et al.*, "A 1.5 mW 68 dB SNDR 80 Ms/s 2 × interleaved pipelined SAR ADC in 28 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 49, no. 12, pp. 2835–2845, Dec. 2014.

[3] R. Kapusta, J. Shen, S. Decker, H. Li, E. Ibaragi, and H. Zhu, "A 14b 80 MS/s SAR ADC With 73.6 dB SNDR in 65 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 48, no. 12, pp. 3059–3066, Dec. 2013.

[4] B. Chen, M. Maddox, M. C. W. Coln, Y. Lu, and L. D. Fernando, "Precision passive-charge-sharing SAR ADC: Analysis, design, and measurement results," *IEEE J. Solid-State Circuits*, vol. 53, no. 5, pp. 1481–1492, May 2018.

[5] E. Martens, B. Hershberg, and J. Craninckx, "A 69-dB SNDR 300-MS/s two-time interleaved pipelined SAR ADC in 16-nm CMOS FinFET with capacitive reference stabilization," *IEEE J. Solid-State Circuits*, vol. 53, no. 4, pp. 1161–1171, Apr. 2018.

[6] J. Kuppambatti and P. R. Kinget, "Current reference pre-charging techniques for low-power zero-crossing pipeline-SAR ADCs," *IEEE J. Solid-State Circuits*, vol. 49, no. 3, pp. 683–694, Mar. 2014.

[7] Y. Shen *et al.*, "A 10-bit 120-MS/s SAR ADC with reference ripple cancellation technique," *IEEE J. Solid-State Circuits*, vol. 55, no. 3, pp. 680–692, Oct. 2019.

[8] C. Chan *et al.*, "A 0.011mm^2 60dB SNDR 100MS/s reference error calibrated SAR ADC with 3pF decoupling capacitance for reference voltages," in *Proc. IEEE Asian Solid-State Circuits Conf. (A-SSCC)*, Nov. 2016, pp. 145–148.

[9] W.-H. Tseng, W.-L. Lee, C.-Y. Huang, and P.-C. Chiu, "A 12-bit 104 MS/s SAR ADC in 28 nm CMOS for digitally-assisted wireless transmitters," *IEEE J. Solid-State Circuits*, vol. 51, no. 10, pp. 2222–2231, Oct. 2016.

[10] B. Hershberg *et al.*, "A 3.2GS/s 10 ENOB 61mW Ringamp ADC in 16nm with background monitoring of distortion," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2019, pp. 58–60.

[11] B. Hershberg *et al.*, "A 6-to-600MS/s fully dynamic ringamp pipelined ADC with asynchronous event-driven clocking in 16nm," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2019, pp. 68–70.

[12] L. Kull *et al.*, "A 3.1 mW 8b 1.2 GS/s single-channel asynchronous SAR ADC with alternate comparators for enhanced speed in 32 nm digital SOI CMOS," *IEEE J. Solid-State Circuits*, vol. 48, no. 12, pp. 3049–3058, Dec. 2013.

[13] B. Hershberg, S. Weaver, K. Sobue, S. Takeuchi, K. Hamashita, and U.-K. Moon, "Ring amplifiers for switched capacitor circuits," *IEEE J. Solid-State Circuits*, vol. 47, no. 12, pp. 2928–2942, Dec. 2012.
[14] Y. Lim and M. P. Flynn, "A 1 mW 71.5 dB SNDR 50 MS/s 13 bit fully differential ring amplifier based SAR-assisted pipeline ADC," *IEEE J. Solid-State Circuits*, vol. 50, no. 12, pp. 2901–2911, Dec. 2015.
[15] Y. Lim and M. P. Flynn, "A calibration-free 2.3 mW 73.2 dB SNDR 15b 100 MS/s four-stage fully differential ring amplifier based SAR-assisted pipeline ADC," in *Proc. Symp. VLSI Circuits*, Jun. 2017, pp. C98–C99.
[16] X. Tang *et al.*, "9.5 A 13.5b-ENOB second-order noise-shaping SAR with PVT-robust closed-loop dynamic amplifier," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2020, pp. 162–164.
[17] T. Hung, J. Wang, and T. Kuo, "16.4 a calibration-free 71.7db SNDR 100ms/s 0.7mw weighted-averaging correlated level shifting pipelined SAR ADC with speed-enhancement scheme," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2020, pp. 256–258.
[18] A. ElShater *et al.*, "A 10-mW 16-b 15-MS/s two-step SAR ADC with 95-dB DR using dual-deadzone ring amplifier," *IEEE J. Solid-State Circuits*, vol. 54, no. 12, pp. 3410–3420, Oct. 2019.
[19] T. Hung and T. Kuo, "A 75.3-dB SNDR 24-MS/s ring amplifier-based pipelined ADC using averaging correlated level shifting and reference swapping for reducing errors from finite opamp gain and capacitor mismatch," *IEEE J. Solid-State Circuits*, vol. 54, no. 5, pp. 1425–1435, Feb. 2019.
[20] M. Kinyua and E. Soenen, "A 72.6 dB SNDR 14b 100 MSPS ring amplifier based pipelined SAR ADC with dynamic deadzone control in 16 nm CMOS," in *Proc. IEEE Custom Integr. Circuits Conf. (CICC)*, Mar. 2020, pp. 1–4.
[21] P. K. Venkatachala *et al.*, "Process invariant biasing of ring amplifiers using deadzone regulation circuit," in *Proc. IEEE Int. Symp. Circuits Syst. (ISCAS)*, May 2018, pp. 1–5.
[22] Y. Chen, J. Wang, H. Hu, F. Ye, and J. Ren, "A time-interleaved SAR assisted pipeline ADC with a bias-enhanced ring amplifier," *IEEE Trans. Circuits Syst. II, Exp. Briefs*, vol. 65, no. 11, pp. 1584–1588, Nov. 2018.
[23] B. Hershberg, N. Markulic, J. Lagos, E. Martens, D. Dermit, and J. Craninckx, "A 1MS/s to 1GS/s ringamp-based pipelined ADC with fully dynamic reference regulation and stochastic scope-on-chip background monitoring in 16nm," in *Proc. Symp. VLSI Circuits*, Jun. 2020, pp. 1–2.
[24] B. Verbruggen, J. Craninckx, M. Kuijk, P. Wambacq, and G. Van der Plas, "A 2.6 mW 6 bit 2.2 GS/s fully dynamic pipeline ADC in 40 nm digital CMOS," *IEEE J. Solid-State Circuits*, vol. 45, no. 10, pp. 2080–2090, Oct. 2010.
[25] B. Verbruggen, J. Tsouhlarakis, T. Yamamoto, M. Iriguchi, E. Martens, and J. Craninckx, "A 60 dB SNDR 35 MS/s SAR ADC with comparator-noise-based stochastic residue estimation," *IEEE J. Solid-State Circuits*, vol. 50, no. 9, pp. 2002–2011, Sep. 2015.
[26] W. Jiang, Y. Zhu, M. Zhang, C. H. Chan, and R. P. Martins, "3.2 A 7.6mW 1GS/s 60dB SNDR single-channel SAR-assisted pipelined ADC with temperature-compensated dynamic Gm-R-based amplifier," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2019, pp. 60–62.
[27] J. Lagos, B. Hershberg, E. Martens, P. Wambacq, and J. Craninckx, "A 1Gsps, 12-bit, single-channel pipelined ADC with dead-zone-degenerated ring amplifiers," in *Proc. IEEE Custom Integr. Circuits Conf. (CICC)*, Apr. 2018, pp. 1–4.
[28] A. Venca, N. Ghittori, A. Bosi, and C. Nani, "27.8 A 0.076mm2 12b 26.5mW 600MS/s 4×—Interleaved subranging SAR-ΔΣ ADC with on-chip buffer in 28nm CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2016, pp. 470–472.
[29] J.-W. Nam, M. Hassanpourghadi, A. Zhang, and S.-W. Chen, "A 12-bit 1.6 gs/s interleaved sar adc with dual reference shifting and interpolation achieving 17.8 fj/conv-step in 65nm CMOS," in *Proc. IEEE Symp. VLSI Circuits (VLSI-Circuits)*, Jun. 2016, pp. 1–2.
[30] Y.-C. Lien, "A 14.6mW 12b 800MS/s 4×time-interleaved pipelined SAR ADC achieving 60.8dB SNDR with Nyquist input and sampling timing skew of 60fsrms without calibration," in *Proc. IEEE Symp. VLSI Circuits (VLSI-Circuits)*, Jun. 2016, pp. 1–2.
[31] B. Sung *et al.*, "26.4 A 21fJ/conv-step 9 ENOB 1.6GS/S 2× time-interleaved FATI SAR ADC with background offset and timing-skew calibration in 45nm CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2015, pp. 1–3.

Benjamin Hershberg (Member, IEEE) received the H.B.S. degree in electrical engineering and computer engineering and the Ph.D. degree in electrical engineering, with a focus on ADCs and scalable, low-power switched-capacitor amplification solutions, including the invention of Ring Amplification, from Oregon State University, Corvallis, OR, USA, in 2006 and 2012, respectively.

In 2013, he joined Interuniversity Microelectronics Center (imec), Leuven, Belgium. He is currently a Principal Member of Technical Staff and conducts research on a variety of topics spanning the receiver chain, from the RF frontend down to the ADC. He has authored or coauthored more than 50 articles, book chapters, and patents.

Dr. Hershberg was a recipient of the 2010 Analog Devices Student Design Award, a winner of the 2012 Broadcom University Research Competition, and a co-recipient of the ISSCC 2019 Lewis Winner Award for Outstanding Paper.

Nereo Markulić (Member, IEEE) received the M.Sc. degree in electrical engineering from the University of Zagreb, Zagreb, Croatia, in 2012, and the Ph.D. degree *(summa cum laude)* from Vrije Universiteit Brussel, Brussels, Belgium, in 2018. His Ph.D. work was in collaboration with the Interuniversity Microelectronics Center (imec), Leuven, Belgium, on digital subsampling phase-locked loops (PLLs) and polar transmitters.

He is currently a Research Scientist with imec, working on RF and mixed-signal circuits for radar applications and next-generation connectivity. He has authored and coauthored publications and patents on PLLs and analog-to-digital converters, a book on frequency synthesis.

Dr. Markulić is a co-recipient of the ISSCC 2019 Lewis Winner Award for Outstanding Paper. He currently serves on the Technical Program Committee of the IEEE VLSI Circuits Symposium.

Jorge Lagos (Member, IEEE) received the B.Sc. degree *(summa cum laude)* in electronics engineering from Pontificia Universidad Católica del Perú (PUCP), Lima, Peru, in 2003, the M.Sc. degree *(cum laude)* in nanotechnologies for integrated systems from the Politecnico di Torino, Turin, Italy, Institut Polytechnique de Grenoble, Grenoble, France, and École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland, in 2013, and the Ph.D. degree *(summa cum laude)* from Vrije Universiteit Brussel, Brussels, Belgium, in 2019, for his work on high-performance, power-efficient analog-to-digital converters, in collaboration with imec, Leuven, Belgium.

From 2004 to 2014, he occupied several Research Assistant positions with a focus on various topics. From 2004 to 2006, he was with PUCP, researching analog VLSI networks for channel decoding. From 2006 to 2011, he was with the Politecnico di Torino, working on chip-level electromagnetic compatibility and system-on-chip functional test. From 2012 to 2014, he was with the Fraunhofer Institute for Integrated Circuits, Erlangen, Germany, focusing on front-end design for MEMS sensors. In 2018, he joined imec, where he is currently a Senior Researcher in analog and mixed-signal IC design.

Dr. Lagos is a co-recipient of the 2019 ISSCC Lewis Winner Award for Outstanding Paper.

Ewout Martens (Member, IEEE) received the M.Sc. degree in electronic engineering and the Ph.D. degree *(summa cum laude)* from Katholieke Universiteit Leuven, Leuven, Belgium, in 2001 and 2007, respectively, for research on modeling techniques for Delta-Sigma modulators and RF front-ends.

From 2007 to 2010, he worked in a spin-off company of KU Leuven, as a Chief Scientist focused on analog design automation. He joined imec, Leuven, in 2010, as a Researcher for WLAN transceivers and is currently working as a Principal Member of Technical Staff with the IoT Department, imec. His interests and work are related to the development of innovative architectures for RF receiver front-ends, transceiver building blocks like PLL, filters and LNA, and ADCs for various applications including image sensors and wireless receivers.

Dr. Martens served as a Technical Program Committee Member of the Symposium on VLSI Circuits from 2017 to 2020.

Davide Dermit (Member, IEEE) received the B.S., M.S., and Ph.D. degrees in electrical engineering from the University of Modena and Reggio Emilia, Modena, Italy, in 2005, 2007, and 2011, respectively.

In 2008, he was an Internship Student with Entropic Communications, Sophia-Antipolis, France, working on integrated high-speed digital PLL frequency synthesizers in 65-nm CMOS for cable-TV. From 2010 to 2015, he was an Analog and RF ICs Designer with STMicroelectronics, France, working on integrated I^2C shifters in 600-nm BCD for the HDMI standard and integrated 10–28 GHz receivers in 28–14-nm CMOS for multistandard PHY and SerDes. In 2015, he joined imec, Leuven, Belgium, where is he currently working as an Analog and RF ICs Designer. In imec, he worked on integrated Cartesian Digital RF Modulators in 28-nm CMOS for Software Defined Radios while his current main research interest is the design of integrated high-speed ADCs in 16-nm FinFET.

Dr. Dermit was a co-recipient of the 2019 Lewis Award for Outstanding Paper at 2019 ISSCC.

Jan Craninckx (Fellow, IEEE) received the M.S. and Ph.D. degrees *(summa cum laude)* in microelectronics from ESAT-MICAS Laboratories, KU Leuven, Leuven, Belgium, in 1992 and 1997, respectively. His Ph.D. work was on the design of low-phase noise CMOS integrated VCOs and PLLs for frequency synthesis.

From 1997 to 2002, he worked with Alcatel Microelectronics (later part of STMicroelectronics) as a Senior RF Engineer on the integration of RF transceivers for GSM, DECT, Bluetooth, and WLAN. In 2002, he joined imec, Leuven, Belgium, as a Principal Scientist working on RF, analog, and mixed-signal circuit design. He is currently a Fellow with imec. His research focuses on the design of RF transceiver front-ends in nanoscale CMOS, covering all aspects of RF, mmwave, analog, and data converter design. He has authored and coauthored more than 200 articles, book chapters, and patents.

Dr. Craninckx was a Regular Member of the Technical Program Committee for several SSCS conferences. He was the Chair of the SSCS Benelux Chapter from 2006 to 2011, an SSCS Distinguished Lecturer from 2012 to 2013, and an elected SSCS AdCom Member from 2017 to 2019. He was an Associate Editor from 2009 to 2016 and an Editor-in-Chief from 2016 to 2019 of the IEEE JOURNAL OF SOLID-STATE CIRCUITS and currently serves as an RF Subcommittee Chair for the IEEE International Solid-State Circuits Conference.