A 4-GS/s 10-ENOB 75-mW Ringamp ADC in 16-nm CMOS With Background Monitoring of Distortion

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

***Abstract*— A 4× interleaved pipelined ADC for direct-RF sampling applications is presented. It leverages the performance advantages of ring amplifiers to unlock greater architectural freedom. The first pipeline stage MDAC with a "passive-hold" mode eliminates the sub-ADC sampling path and associated problems. A high-speed ringamp topology employs digital bias control, robust common-mode feedback (CMFB), and an elegant self-resetting behavior. An asynchronous, event-driven timing control system improves several aspects of performance and enables fully dynamic power consumption and modular design re-use. A general technique is presented whereby the signal-to-distortion ratio (SDR) of any amplifier in the system can be measured in the background with an analog hardware overhead of only one comparator. In this amplifier-intensive architecture utilizing 36 ringamps, the 4-GS/s ADC fabricated in 16-nm CMOS achieves 62-dB SNDR and 75-dB SFDR at Nyquist, consumes 75 mW (including input buffer), and has a Walden figure of merit (FoM) of 18 fJ/conversion-step and a Schreier FoM of 166 dB, advancing the state of the art in direct-RF sampling ADCs by roughly an order of magnitude.**

***Index Terms*— A/D, ADC, asynchronous, direct-RF, dynamic, event driven, high speed, low power, passive-hold MDAC, pipeline, ring amplification, ring amplifier, ringamp, sample and hold, SHA, stochastic ADC.**

I. Introduction

ONE of the defining trends in the last two decades of ADC design was the shift away from architectures that rely heavily on high-performance linear amplifiers. This is primarily a consequence of the failure to gracefully scale conventional amplifier topologies into nanoscale CMOS. Several technology-related obstacles have contributed to this, such as reduced intrinsic device gain and linearity due to short-channel effects, and reduced voltage headroom due to supply voltage scaling. The evolution from planar bulk CMOS into technologies with better electrostatic control of the device channel, such as FinFET and silicon-on-insulator (SOI), has significantly alleviated the set of challenges associated with short-channel effects [1], but the trend of low supply voltages continues unabated. This persistent voltage headroom constraint requires multistage opamp topologies to be considered. Unfortunately, there are fundamental limitations to the classical approach to multistage opamps that prevent such structures from ever scaling as well as digital logic. Thus, to take full advantage of technology scaling, analog designers are presented with only two options: either eliminate linear amplifiers from circuit design altogether or invent new scalable amplifier techniques.

Below a certain ADC speed and resolution, the first option has proved to be extremely successful [2]. Architectures like SAR and techniques like time-based quantization can often avoid linear residue amplification altogether and fully capitalize on the benefits of scaling. Furthermore, architectures that can "hide" the amplification behind several MSBs of quantization also often scale well. For example, pipelined SAR and noise-shaping SAR ADCs can often utilize special-purpose low-swing residue amplifiers. There are, however, some practical limits to the speed and accuracy capabilities of these scaling-friendly architectures, and beyond certain performance boundaries, there is often no viable option. One example is in the application space of direct-RF sampling ADCs [3]–[8] located in the high-linearity, high-speed corner of the ADC performance landscape, where all reported designs to-date rely on some form of linear residue amplification [2]. The consequences of this in terms of power efficiency are severe, with designers forced to accept non-ideal solutions such as high-power telescopic single-stage class-A opamps operating on special high-voltage supplies [9], [10].

Meanwhile, in pursuit of the second option, many opamp alternatives have been proposed over the years, experimenting with full feedback [11]–[13], partial feedback [14]–[17], and no feedback (open-loop) [18]–[24]. This includes techniques, such as zero-crossing based circuits [14], [15], [25], digital amplifiers [11], inverter-based amplifiers [12], [13], [26], charge-steering amplifiers [16], [17], pulsed bucket brigade [23], capacitive charge pump [24], Gm-C "dynamic" amplifiers [18]–[20], Gm-R amplifiers [21], and open-loop with high-order non-linear calibration [22]. Many of these ideas provide valuable specialized solutions and when combined with the right architecture advance the

Manuscript received August 28, 2020; revised December 23, 2020; accepted January 16, 2021. Date of publication February 23, 2021; date of current version July 23, 2021. This article was approved by Associate Editor Shanthi Pavan. *(Corresponding author: Benjamin Hershberg.)*

The authors are with the Interuniversity Microelectronics Center (IMEC), 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.2021.3053893.

Digital Object Identifier 10.1109/JSSC.2021.3053893

state-of-the-art [18], [21], [22]. However, all of these techniques also suffer from at least one critical limitation, such as low-swing, low-linearity, low-speed, or expensive calibration. This both limits the generality of these solutions and, in many cases, also limits their ability to fully realize the original motivation of optimal technology scaling.

Ring amplification is another alternative that has been proposed [27]. In essence, a ring amplifier (ringamp) is a multistage inverter-based amplifier with a dominant output pole, typically stabilized through dynamic large-signal transient feedback mechanisms. Ringamps have been shown to be capable of wide output swing, high gain, high linearity, high-speed, dynamic switchable operation, and performance that scales with digital CMOS. Published ringamp designs have demonstrated the broad applicability of this approach, enabling state-of-the-art power efficiency in both high-accuracy [27]–[32] and high-speed ADCs [33]–[38]. Structural optimizations and techniques covering a diverse set of applications and technologies have also been demonstrated, including low-voltage [39] and high-voltage [30] environments, power management applications (LDO) [40], [41], linearity-priority applications [32], [36], [42], and speed-priority applications [36], [43]. Techniques for reliable operation across process, voltage, and temperature (PVT) based on both robust analog design [28], [29], [42], [44] and digital background tuning or calibration [33], [34], [38] have been reported.

The ADC presented in this article is an upgraded version of the design presented in [33] with the clocking system of [37]. It demonstrates how ringamps, when applied to the aforementioned application of direct-RF sampling ADCs, provide advantages both directly by eliminating the amplifier performance bottleneck and indirectly by the consequent increase in architectural freedom. Operating at 4 GS/s, the ADC consumes 75 mW including input buffer and has 62-dB SNDR and 75-dB SFDR at Nyquist, a Walden figure of merit (FoM) of 18 fJ/conversion-step, and a Schreier FoM of 166 dB. It advances the state of the art for direct-RF sampling ADCs across all three key metrics of accuracy, linearity, and power efficiency.

This article is organized as follows. Section II describes the overall system and highlights several notable architectural innovations, including an MDAC circuit that eliminates the sub-ADC sampling path and a fully asynchronous event-driven control system. Section III discusses the high-speed ringamp topology that is used to enable such an architecture. In Section IV, a method is introduced for background measurement of an amplifier's signal-to-distortion ratio (SDR) using only a single comparator and digital processing. This can be used for background optimization of ringamp biasing across PVT. Measured results are reported in Section V, and conclusions are drawn in Section VI.

II. Architecture

A. Overview

A top-level view of the ADC is shown in Fig. 1. Four 1-GS/s pipelined ADC channels are interleaved to attain the system rate of 4 GS/s. The input signal is buffered by the ac-coupled

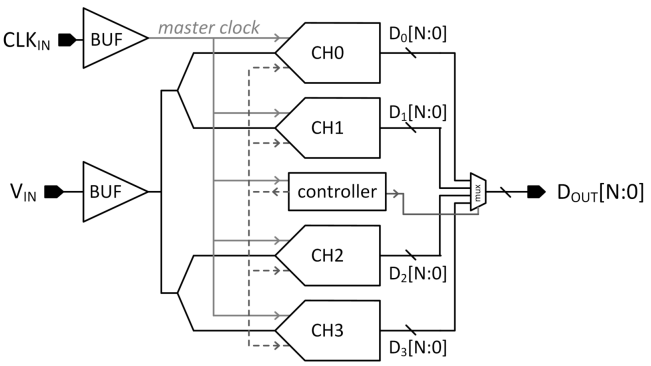

Fig. 1. Top level of the 4-GS/s ADC.

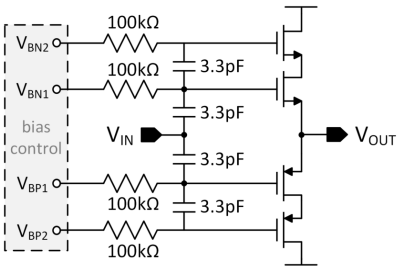

Fig. 2. Input buffer of Fig. 1.

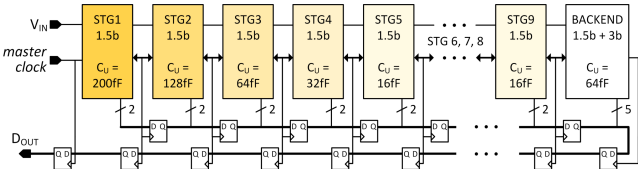

Fig. 3. Pipeline channel architecture of Fig. 1.

pseudo-differential class-AB push–pull source follower shown in Fig. 2 [5].

The channel architecture is presented in Fig. 3. It contains nine 1.5-bit MDAC stages scaled for thermal noise, followed by a 1.5-bit + 3-bit backend. Despite the lower theoretical power efficiency of 1.5-bit/stage versus multi-bit/stage [45], the availability of high-efficiency residue amplification in this design reduces the impact of this tradeoff on the overall system power budget. In return, the 1.5-bit/stage approach reduces complexity and offers more design freedom. In particular, the inherent linearity of a 1.5-bit sub-DAC allows both capacitor mismatch and finite-gain error to be corrected with a single code-independent stage gain correction factor, and the large redundancy range provides extra analog residue capacity for both systematic errors like amplifier and sub-ADC offsets, and techniques that utilize redundancy like dither-based background calibration [46] and early quantization [47], [48].

B. Quantization Techniques

1) MDAC With Passive-Hold Mode: The stage 1 MDAC of the pipeline channel is shown in Fig. 4. Traditionally, a 1.5-bit flip-around MDAC can be configured into two modes: *track* and *amplify*. Here, a third "passive-hold" mode is introduced. Unlike traditional sample-and-hold circuits, this hold function

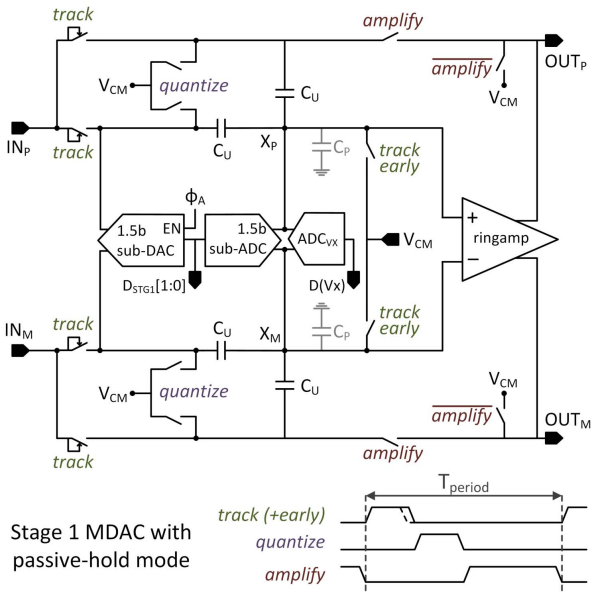

Fig. 4. Pipeline stage 1 MDAC with extra "passive-hold" mode.

does not require an amplifier. During *track*, the MDAC samples the input signal $V_{\rm IN}$ ($= V_{\rm INp} - V_{\rm INm}$) onto unit capacitors C_U. Then, during *quantize*, the top plates of the C_U capacitors are connected to $V_{\rm CM}$ and the bottom plates are left floating. The voltage V_X ($=V_{Xp} - V_{Xm}$) formed across the floating bottom plates will be

$$V_X = -V_{\rm IN} \cdot \frac{2\ C_U}{2\ C_U + C_P} \tag{1}$$

where C_P is the sum of all other capacitances present at the X nodes. In other words, V_X becomes an attenuated replica of $-V_{\rm IN}$. The sub-ADC is attached directly to the X nodes and uses this information to perform quantization. To minimize its parasitic influence, the sub-ADC flash is implemented using the source-shifted programmable threshold comparator described in [49]. This topology embeds the flash threshold offsets inside the comparators while also providing low decision delay with near-minimum sized input transistors. The large stage redundancy range easily absorbs any quantization error due to uncompensated attenuation of $V_{\rm IN}$ by C_P in (1) even when using digital "firmware default" values to program the sub-ADC comparator thresholds.

The proposed technique uses the exact same physical charge stored on $2C_U$ for both quantization and residue amplification, and this provides two key benefits. First, all skew and bandwidth-related mismatch between the MDAC and sub-ADC is eliminated, which, for high-speed applications, is critical. Second, there is no dedicated sub-ADC sampling network. This minimizes the input capacitance, relaxing the design requirements of the input buffer.

2) Early-Quantization: The first stage of the pipeline samples an unknown, unpredictable input signal. This creates a causality constraint that requires the fundamental operations *track*, *quantize*, and *amplify* to be performed sequentially. Fortunately, this fits nicely into the timing constraints of a 4× interleaved system, where track time is anyway limited to 25% duty cycle and residue amplification occupies up to 50%, leaving 25% still free for quantization. By contrast, in the latter stages of the pipeline, up to 50% track time and 50% amplify time are theoretically possible, but this leaves no time for quantization. However, it also becomes possible to parallelize the *track* and *quantize* phases in the latter stages. There are several approaches to this that have been proposed, including look-ahead quantization [50], [51] and early quantization [47], [48].

This design uses a form of early quantization. The basic principle is seen in pipeline stage 2 of Fig. 5. When stage 2 is tracking, its sub-ADC also tracks the output residue of stage 1 in parallel. Crucially, the sub-ADC samples "early" and already begins quantization, while the main signal path continues to settle. The built-in stage redundancy absorbs the extra error due to quantizing this imperfect and not fully settled residue. The sub-ADC result is ready by the time the stage 2 MDAC samples, removing *quantize* from the critical timing path and allowing stage 2 to transition directly from *track* to *amplify*. The main challenge of early quantization is sampling at the optimal "early" moment. Sampling too soon may result in too much quantization error, and sampling too late can inject error into the main signal path due to disturbance by the stage 2 sub-ADC sampler. The optimal moment is, therefore, just after slewing has finished. In this design, each ringamp circuit outputs a *slew done* signal that instructs the sub-ADC of the subsequent stage when to sample. This is discussed further in Section III-E.

3) Composite Backend: The pipeline's 1.5-bit + 3-bit backend stage in Fig. 3 combines both the passive-hold and early quantization techniques discussed earlier into a two-step procedure, with 1.5-bit early quantization followed by a 3-bit passive-hold flash. As with all other sub-ADCs in the system, the wide threshold tuning range comparator topology of [49] is used, and the threshold levels of the 3-bit flash are one-time calibrated using an on-chip reference ladder.

C. Asynchronous Event-Driven Timing Control

In "deep" pipelined ADCs that contain many stages, a conventional clock tree constitutes a highly distributed network, with parasitics and mismatch creating skew between the different branches. Sufficient margin must be included in the timing generation such that all non-overlap and causal relationships are maintained. This leads to a difficult set of design tradeoffs in terms of power, speed, jitter, reliability, and reconfigurability. High conversion speeds and interleaving of multiple channels exacerbate these tradeoffs further. To overcome these limitations, this design pursues an alternative approach. Localized state machines in each pipeline stage coordinate operation through an asynchronous event-driven communication protocol. The result is a modular, correct-by-construction timing system that minimizes routing complexity and maximizes performance. Combined with the switchable operation of ring amplifiers, it also leads to fully dynamic power consumption of the ADC core, resulting in constant power efficiency and performance independent of clock rate.

At the ADC top level in Fig. 1 a self-biased clock buffer based on [52] is used to convert an external differential sine

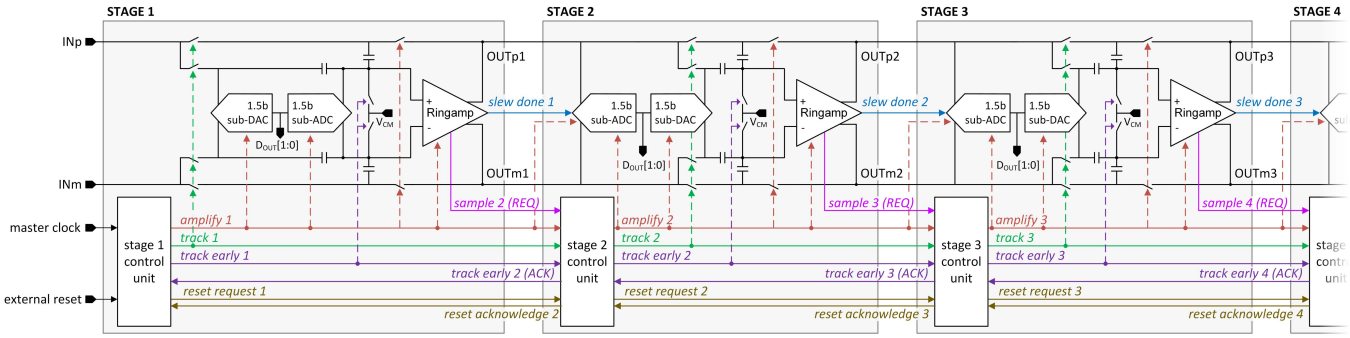

Fig. 5. Simplified view of the first three pipeline stages of Fig. 3 with asynchronous event-driven timing control.

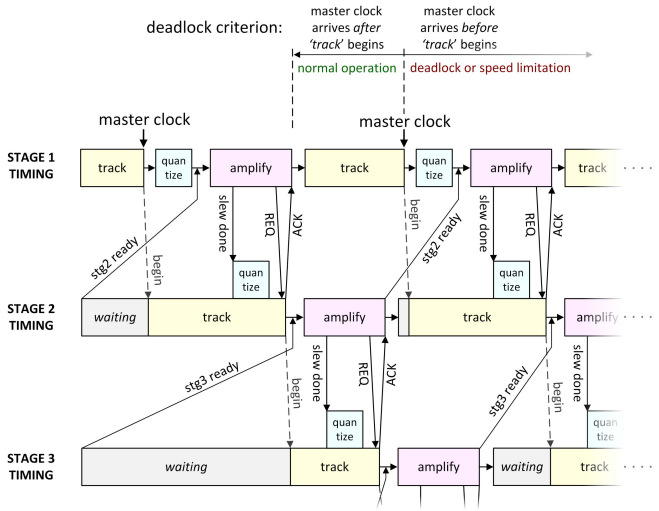

Fig. 6. Example timing control state flow for Fig. 5.

input into a single-ended square *master clock* output that is fed directly into all four channels and a controller. The controller coordinates interleaving of both the channel frontend samplers and the final digital outputs. Only the jitter of the *master clock* affects system performance, and due to the way that it is directly fed into each channel, only four gates separate the *master clock* from the critical bottom-plate sampling switch of the stage 1 MDAC.

Each stage in the pipeline communicates with its direct neighbors through an event-driven protocol, as shown in Fig. 5. The wires linking the stage control units together serve a dual purpose: they facilitate both inter-stage communication and control of local circuits. Fig. 6 shows a simplified version of the event-driven timing that occurs during normal operation. It begins in the system reset state, where all stages are either tracking or waiting. The arrival of the *master clock* initiates a chain reaction of processing events that propagate down the pipeline. First, stage 1 samples the input and begins quantizing. When quantization is complete and stage 2 indicates that it is ready, stage 1 amplifies the residue. After internal ringamp timing controls indicate that the residue is settled, *sample* 2 *(REQ)* is sent to stage 2. Stage 2 then samples and returns the acknowledgment *track early* 2 *(ACK)*, allowing stage 1 to power down its ringamp and return to the *track* state. The inherent causality of this approach makes it correct-by-construction; safe and valid phase relationships are guaranteed.

Stage 2 now begins a similar exchange with stage 3, with a few notable differences. First, all stages except stage 1 use the early quantization technique discussed earlier, where quantization occurs in parallel with tracking. Second, when stage 2 is ready to resume tracking, it will wait until stage 1 indicates that it is safe to reconnect. This prevents stage 2 from reconnecting to stage 1's output just as stage 1 samples, and avoids kickback-related errors.

The asynchronous digital outputs of each stage propagate to the end of the pipeline for alignment. As shown in Fig. 3, they are then shifted back up to stage 1 using the causally related edges of *sample (REQ)*, where they can be safely re-synchronized with the *master clock*.

Event triggers in the system are either generated with digitally controlled delay cells, derived directly from analog events, or a combination of both. Triggering on analog events requires that the analog blocks themselves function correctly, and triggering on digital events assumes that the state machines are in valid states. Under certain boundary conditions like power-ON, initial convergence of analog dc biases, or a saturated input signal, it is possible that these conditions will not be met. This can result in a deadlock where the chain reaction of processing operations halts. Fortunately, a deadlock is easily detected with the criterion shown in Fig. 6. Any deadlock or stall in the pipeline will create a backlog that will eventually propagate back to stage 1, and if a *master clock* edge arrives at stage 1 before it has begun tracking the input, we know that either the clock speed is too high for the given settings or a deadlock has occurred. When this condition is detected, a "reset" command is issued that forces each stage back into its *track* or *waiting* state. Analog blocks such as ringamp CMFB are designed so that repeatedly returning to this reset condition will eventually result in correct bias convergence.

III. Ring Amplifier

The fully differential ring amplifier topology used in this work is shown in Fig. 7. It is a refinement of the ringamp used in [33], first appearing in [34], and described in detail for the first time here.

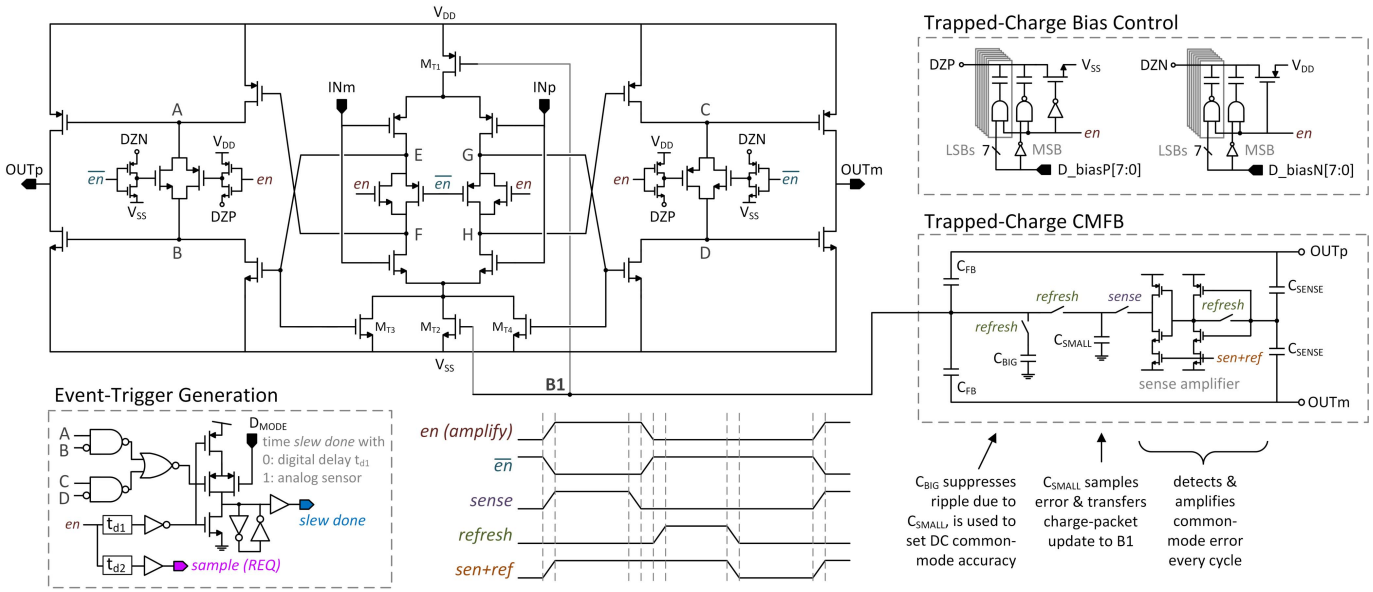

Fig. 7. Fully differential, self-resetting, bias-enhanced, digitally biased ring amplifier structure used in each pipeline stage.

A. Bias Control

A resistor-based dead-zone embedding approach determines the ringamp's dynamic stabilization behavior, based on the "self-biased" ringamp topology introduced in [39]. This is implemented with tunable CMOS resistors spanning nodes A/B and C/D [35]. In the OFF-state (en = "0"), the gates of these nMOS/pMOS "resistors" are reset to V_{SS}/V_{DD} and their conductance is zero, like an open switch. In the ON-state (en = "1" = *amplify*), the digitally controlled DACs labeled "Trapped-Charge Bias Control" in Fig. 7 synthesize voltages at nodes DZN and DZP using charge redistribution. Simultaneously, these nodes also connect to the gate nodes of the CMOS resistors. Charge sharing occurs between the bias DACs and the pre-charged transistor gates, and this results in some voltage attenuation of V_{DZN} and V_{DZP}. To counteract this, the control logic for the MSB capacitor of each DAC is configured differently than the LSBs so that it expands the voltage tuning range of the DAC to compensate. This makes it possible to apply the maximum overdrive voltage to the CMOS resistor, which in turn minimizes its required size, helping to minimize internal ringamp parasitics and thus overall ringamp speed and efficiency.

B. Common-Mode Feedback

Robust common-mode rejection is ensured with the combination of three feedback paths. Two of these paths are contained in the "Trapped-Charge CMFB" block of Fig. 7. The first is a high-speed feedback path through the two C_{FB} capacitors. This provides instantaneous suppression of per-cycle errors. About 4 dB of rejection per pipeline stage is provided through this path, which is found to be sufficient for robust operation while also allowing the size of C_{FB} to be kept moderate, minimizing output loading of the ringamp.

The second CMFB path is a low-speed discrete-time path that sets the dc bias of node B1. When the ringamp is powered down (en = "0") and its outputs OUTp and OUTm are shorted to VCM (see Fig. 4), a small sense amplifier auto-zeroes itself with respect to VCM by charging the C_{SENSE} capacitors. During the ringamp's ON-state, this sense amplifier outputs the polarity of the common-mode level at the ringamp output with respect to VCM (i.e., the common-mode error). Just before the ringamp ON-state ends, this information is sampled onto C_{SMALL}. Obtaining an "early" clock edge that can do this is simply a matter of selecting a slightly earlier transition in the event-driven timing chain of the control system. The charge information sampled on C_{SMALL} is then passively shared with all capacitances attached to node B1, including the charge reservoir C_{BIG}. C_{BIG} enforces an upper limit on the voltage ripple amplitude that is possible at node B1 due to charge packet updates from C_{SMALL}. When the ringamp transitions from its OFF-state to its ON-state, transient charge redistribution at B1 will occur due to coupling through C_{FB}, M_{T1}, and M_{T2} into nodes whose ON-state voltages are different from their OFF-state voltages. This results in a common-mode offset because the charge packet updates from C_{SMALL} are delivered during the OFF-state but used to set the bias for the ON-state. However, the gain of the sense amplifier is large enough to suppress this CM offset to negligible levels. The bias information at B1 is applied to the current sources M_{T1} and M_{T2}. This complementary current source approach provides a significantly wider common-mode rejection range than a single-sided current source biasing scheme [28] and eliminates the need for a static current reference, enabling fully dynamic operation. Ultimately, the key advantage of the trapped-charge approach is that the bias at B1 re-settles instantly, allowing the ringamp to be rapidly powered ON and OFF.

The third CMFB path is a local loop applied to the first stage of the ringamp using M_{T3} and M_{T4} [28]. This high-speed active feedback path suppresses peaking in the transfer function of the primary CMFB path (C_{FB}). In simulation, it is

found to improve the steady-state ac common-mode phase margin of the main signal path from roughly 15° to 100°.

C. Performance Optimization

In direct-RF sampling applications, where concerns of intermodulation and harmonic interference often elevate SFDR above even SNR as the most important performance spec, a high linearity residue amplifier is particularly useful. The simulated steady-state open-loop gain of the ringamp in this design is roughly 62 dB, but the higher order settled linearity observed for a wide-swing output signal is better than 90-dB SFDR, due to the inherent compression immunity of the ringamp [35], [36]. In other words, although 62-dB open-loop gain will cause a non-negligible finite-gain error, the variation of this error with respect to ringamp output voltage is very small, and 90-dB linearity can be achieved with only a first-order gain error correction. Applying known background calibration techniques that utilize the large available redundancy range of the 1.5-bit pipeline stage, it is straightforward to correct both capacitor mismatch and first-order gain error with a single shared stage-gain correction factor [46]. The remaining higher order errors are non-limiting at the system level, well below the 78-dB HD3 linearity bottleneck imposed by the ADC input buffer at $f_{\text{Nyquist}} = 2$ GHz.

However, high linearity is only guaranteed for a well-settled output. To ensure this, the ringamp is carefully optimized for speed, which ultimately revolves around two core design principles: minimum capacitance at nodes A–H (for maximum internal pole frequencies) and unrestricted dynamic current steering. The topology of Fig. 7 is specifically chosen with this in mind, using the CMOS resistors at nodes E–H to increase the speed with the bias-enhancement technique introduced in [36], [43], [53]. This technique also reduces parasitics, for reasons explained in Section III-D. The ringamp follows a roughly 4:2:1 transistor sizing scheme for internal stages 1–3, which provides a good tradeoff between power efficiency and speed. Finally, special priority in the layout is given to nodes A–H to minimize interconnect parasitics.

D. Self-Resetting

The presence of switchable CMOS resistors in both the first and second stages of the ringamp leads to an elegant self-resetting behavior [36]. During the OFF-state, when the MDAC is tracking, these CMOS resistors act like open switches, cutting OFF-current flow through each branch. The voltage presented to INp and INm is roughly mid-rail, and nodes A, C, E, and G will charge to V_{DD}, whereas nodes B, D, F, and H will discharge to V_{SS}. Consequently, the output stage is also disabled, placing the ringamp in a full signal-independent shutdown. There are three key benefits of this scheme. First, all pull-up, pull-down, and power-gating functions are implemented by the core transistors themselves. This minimizes parasitics at the critical nodes A–H. Second, the CMOS resistors that provide current cutoff (power-gating) functionality in the OFF-state must also generate fairly large (>100 mV) IR drops in the ON-state. This means that small transistors can be used to implement the CMOS resistors, with zero headroom penalty. By contrast, in other schemes with dedicated power-gating switches, these switches do not serve a functional purpose in the ON-state and must be sized very large to minimize their IR drop impact on voltage headroom [27]–[29], [31]–[33]. Large switches also place a heavy burden on the clock driver, leading to increased power consumption and supply noise. Third, contrary to other schemes that use single-sided power gating, in this topology, charge kickback at the ringamp input is mostly self-canceled. The pMOS and nMOS input pairs kick charge with opposite polarity when resetting, and this minimizes the influence of ringamp power cycling on common-mode levels in the MDAC.

E. Timing Control

The "Event-Trigger Generation" block in Fig. 7 produces the two event triggers exported by each ringamp in Fig. 5. Namely, the *slew done* signal used for early quantization that was discussed in Section II-B2 and *sample (REQ)* that initiates the end of *amplify*, as discussed in Section II-C. Two options for generating *slew done* are implemented. When $D_{\text{MODE}} =$ "0", *slew done* occurs a fixed time after the start of *amplify*, set by the digitally tunable delay t_{d1}. When $D_{\text{MODE}} =$ "1", t_{d1} is set to its minimum delay and *slew done* is instead determined by a simple analog sensor built with digital gates. Consider the binary value of the voltages at nodes A–D in Fig. 7. During slewing, the logic equation AB'CD' will never equal "1" because A = B and C = D. During settling, the dynamic formation of a dead-zone voltage across the CMOS resistor forces the four nodes apart such that A = C = "1" and B = D = "0" and the logic equation evaluates to "1." Waveforms further illustrating this analog slew-done detection technique are presented in [37].

IV. SDR Estimation

The digital bias control capability of the CMOS resistors in the ringamp of Fig. 7 enables the possibility of run-time bias adjustment. In the simplest use case, this tunability can be combined with static firmware presets to compensate for unanticipated parasitics and non-idealities of a systematic nature. However, if we wish to track PVT variation and/or optimize performance of individual ringamps, an additional element is required: adaptive bias control.

In this section, a technique for background estimation of SDR in a feedback circuit is introduced. Although we consider it from the perspective of ringamp bias control, the technique itself is general and can be applied to any switched capacitor feedback circuit or residue amplifier of interest.

A. Basic Principle

The technique stems from an observation that the summing node voltage V_x at node X in Fig. 8 contains the residual error due to non-ideal feedback. In other words, in an ideal feedback system where the amplifier is noiseless, infinite gain, infinite bandwidth, and zero-offset, V_x will be zero. In a real system, any non-zero value of V_x indicates that a proportional error will also be present in V_o at node OUT.

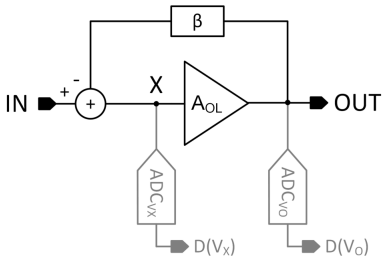

Fig. 8. Basic setup for measuring the SNDR of an amplifier in feedback.

A portion of this error is due to the finite gain of the amplifier, A_{ol}, and in [54], it was shown how using the two auxiliary ADCs in Fig. 8, it is possible to estimate and correct this error. Naively, one might try to do this by averaging multiple samples

$$\hat{A}_{ol} = \frac{1}{n}\sum_{i=0}^{n}\frac{\hat{V}_{oi}}{\hat{V}_{xi}} \tag{2}$$

where $\hat{V}_{xi}$ and $\hat{V}_{oi}$ are digitized estimates that include noise and distortion injected by ADC_{vx} and ADC_{vo} and n is the number of samples averaged. However, the presence of error and noise in $\hat{V}_{xi}$ and $\hat{V}_{oi}$ will result in the denominator of some samples having a value of zero or near zero, resulting in large digital noise and poor estimation accuracy and convergence. In [54], this problem is addressed by using an LMS estimator instead. In this work, for reasons that will become apparent later, we solve it by computing A_{ol} with a linear regression

$$\hat{A}_{ol} = \frac{\sum_{i=0}^{n}\hat{V}_{oi}^{2}}{\sum_{i=0}^{n}\hat{V}_{oi}\cdot\hat{V}_{xi}}. \tag{3}$$

With this knowledge, the effect of linear finite gain on the total closed-loop stage gain can be corrected. However, there are other well-known background calibration techniques that can accurately correct for not only first-order finite gain but also random capacitor mismatch [46], so in this work, we are not directly concerned with the V_o/A_{ol} component of V_x. Rather, we are interested in the component of V_x that contains all other noise and distortion components of the feedback system. This error V_{nd} can be extracted from V_x as

$$V_{nd} = V_x - V_{dc} - \frac{V_o}{A_{ol}}. \tag{4}$$

The value of static offset V_{dc} is simply mean(V_x) and thus easily determined. The input-referred SNDR of the feedback circuit can be estimated by comparing the error power contained in V_{nd} to the signal power

$$\widehat{\text{SNDR}}_{dB} = 10\cdot\log_{10}\left(\frac{\beta^2\cdot\text{Var}(V_o)}{\text{Var}(V_{nd})}\right) \tag{5}$$

where β is the loop feedback factor (Fig. 8). Note that because V_o is not purely signal but rather a combination of signal and output-referred V_{nd}, this estimator is not exact. However, in typical cases where $\text{Var}(V_o) \gg \text{Var}(V_{nd})$, the resulting error will be very small and can be ignored. In this work, we seek to use the estimator $\widehat{\text{SNDR}}$ as an objective function for optimizing ringamp settling behavior.

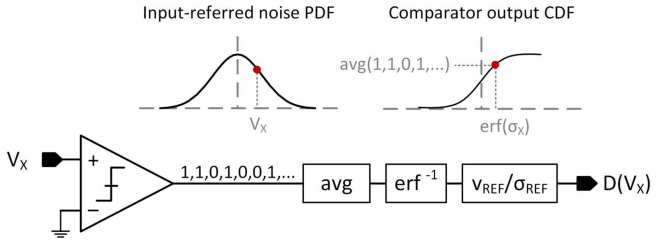

Fig. 9. Principle of operation for a 1-bit stochastic ADC.

B. Stochastic ADC to Measure V_x

To apply this estimation concept, the practical implementation of ADC_{vx} and ADC_{vo} in Fig. 8 must be considered. In the case of a pipelined ADC, a digital representation $D(V_o)$ of V_o with sufficiently high resolution is conveniently already produced by the ADC backend, providing ADC_{vo} implicitly. For ADC_{vx}, obtaining an accurate digitization $D(V_x)$ of V_x requires extra hardware and is not trivial. Any corruption of the charge information stored on node X will inject additional errors, and the small amplitude of V_x (typically on the order of 1 mV) makes high-resolution, low-noise capture difficult. A direct implementation approach comes with many drawbacks. In [54], ADC_{vx} involves a large 128-fF input capacitance that loads node X and degrades loop gain, a 64× gain pre-amplifier stage, and an area intensive 13-bit auxiliary ADC. The performance and complexity overhead that this imposes severely constrains the applicability of the technique. In this work, we wish to monitor many different feedback loops in the system, so the analog hardware must be simple, small, and noninvasive.

An elegant alternative is to use a single-comparator stochastic ADC [55]. The basic principle is shown in Fig. 9. Consider when the input V_x is stationary and its amplitude is small enough to fall within the central region of the comparator's thermal noise probability distribution function (PDF). After many sequential comparisons, the average value of the output bit stream will describe V_x in terms of the comparator's cumulative density function (CDF). This can be linearized using the inverse error function (erf^{-1}) into units of probability distribution sigma (σ). Finally, if the relationship between absolute voltage and relative sigma is known, the digital output $D(V_x)$ can be scaled in terms of absolute voltage.

The appeal of this approach is that it transforms thermal noise from foe to friend, allowing a small-amplitude signal to be quantized with high precision. It can be implemented with a small analog footprint, shifting complexity into digital processing blocks that do not impact main signal path performance and can easily be shared by multiple monitor circuits.

C. Practical Algorithm

A practical algorithm for estimating SDR in the background that accounts for the unique aspects of the 1-bit stochastic ADC_{vx} is shown in Fig. 10. For reasons explained later, the procedure low-pass filters random noise, and thus, $\widehat{\text{SNDR}}$ becomes $\widehat{\text{SDR}}$. The algorithm improves upon [33], with additional refinements to reduce computational complexity and

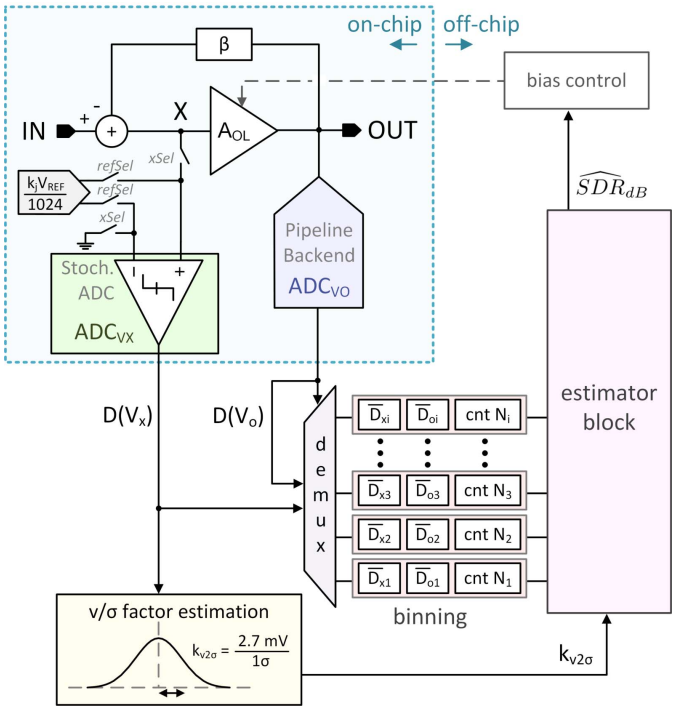

Fig. 10. Conceptual depiction of the practical algorithm used to perform SDR estimation.

increase robustness. It consists of three main steps: preparation, data collection, and estimation.

1) Preparation: Before SDR estimation can occur, two preparation steps are necessary. First, the input-referred offset of the stochastic ADC must be adjusted so that its noise PDF is approximately centered around the active signal at node X. Multiple circuits in the MDAC contribute to the cumulative offset, including ADC_{vx} itself and the residue amplifier. Offset cancellation is performed by connecting ADC_{vx} to X ($xSel$ = "1", $refSel$ = "0" in Fig. 10), while the pipeline is operating under normal conditions and sampling V_x at the end of each cycle with ADC_{vx}. The resulting output data stream is used to tune the internal offset of ADC_{vx} until the 1-b output from ADC_{vx} has an average value of 0.5. Once an initial digital offset compensation value is determined, it can be continuously monitored and adjusted in parallel with the main algorithm described later, since it uses an identical setup.

In a second preparation step, a volts-per-sigma unit conversion constant $k_{v2\sigma}$ is determined. This requires applying a set of known voltages to ADC_{vx} and observing their corresponding locations on the ADC's output CDF. An on-chip reference ladder, shared by multiple stages, is included for this purpose. To ensure rejection of systematic offsets, a minimum of two references must be tested, and adding additional test points further increases accuracy by reducing sensitivity to DNL errors of the reference ladder. In this work, we apply seven test voltages

$$V_{\text{test}j} = \frac{k_j \cdot 2(V_{\text{REFP}} - V_{\text{REFM}})}{1024} \Big| k_j = \{3, 2, \ldots, -3\}. \quad (6)$$

For each $V_{\text{test}j}$ applied, the average value of ADC_{vx} identifies its location on the noise CDF. This can be linearized in terms of PDF σ

$$\hat{\sigma}_j = \sqrt{2} \cdot \text{erf}^{-1}(2 \cdot \text{mean}\big(D(V_{xj})\big) - 1). \quad (7)$$

Performing a linear regression on the set of test inputs and outputs $\{V_{\text{test}j}, \sigma_j\}$, a volts-per-sigma conversion factor $k_{v2\sigma}$ is obtained. With this factor, it is possible to normalize the digital outputs of both ADC_{vx} and ADC_{vo} to a common known reference, V_{REF}. As explained later in Section IV-D, with a thoughtful implementation of ADC_{vx}, it is possible to periodically repeat this procedure in the background during normal operation in order to track changes in $k_{v2\sigma}$ with respect to PVT variation.

2) Data Collection (Binning): The 1-bit stochastic ADC approach provides a solution to the challenge of quantizing a small amplitude signal to high accuracy by averaging the thermal noise of the comparator. However, when this is done by making many successive measurements of a single held sample of V_x, any noise or error in that single sample will also be quantized, including kT/C thermal noise. Considering the small amplitude of the signal of interest, this requires a very large sampling capacitor to achieve a decent SNR. The stochastic ADC implementation in [55] is conveniently able to quantize a signal held on a 1-pF capacitor in the foreground. However, in this system where we desire transparent background operation, a large sampling capacitor that can directly meet such strict kT/C noise requirements is impractical.

Fortunately, there is an alternative approach to stochastic quantization that eliminates this problem. Rather than making multiple measurements of a single sample, it is possible to combine single measurements of multiple related samples, provided that each sample in the set captures the same input voltage. The information needed to identify which cycles of ADC_{vx} contain the same (or very similar) sampled input voltage is already available in the output codes of ADC_{vo}. All cycles with the same $D(V_o)$ code are assumed to have also sampled the same value of V_x.[1]

In this alternate "multi-sample averaging" approach, the input-referred noise PDF of ADC_{vx} simply becomes the sum of all noise sources, including both the comparator and the input sampler. Thus, kT/C noise becomes uncorrelated with the signal, and capacitor sizing requirements disappear entirely. Noise contained in V_x is also averaged by this procedure, and thus, V_{nd} of (4) is stripped of random noise and (5) becomes an estimator of SDR rather than SNDR.

Practical implementation of this new approach to stochastic quantization requires sorting $D(V_x)$ and $D(V_o)$ data into "bins." In a naive implementation, one bin is allocated to each unique $D(V_o)$ code, and whenever that code is encountered, the corresponding $D(V_x)$ and $D(V_o)$ data from the given conversion cycle are sorted into the bin. As shown in Fig. 10, each bin index "i" consists of three accumulators producing scalar outputs: the average values of $D_i(V_{oi})$ and $D_i(V_{xi})$, and the number of samples N_i that have been sorted into bin_i.

[1]This assumes that there is only one transient settling characteristic for each output voltage, regardless of other factors such as data-dependent sub-DAC switch connections in the pipeline stage MDAC. Empirically, we find this to be an acceptable assumption.

Mode	ADC_{vx} Behavior
disable	entirely disabled, never sample, never convert
idle	sample X nodes every cycle, but never convert
normal	sample X nodes every cycle, and convert every N^{th} cycle (programmable)
prepare	connect to reference DAC, and only sample and convert every N^{th} cycle (programmable)

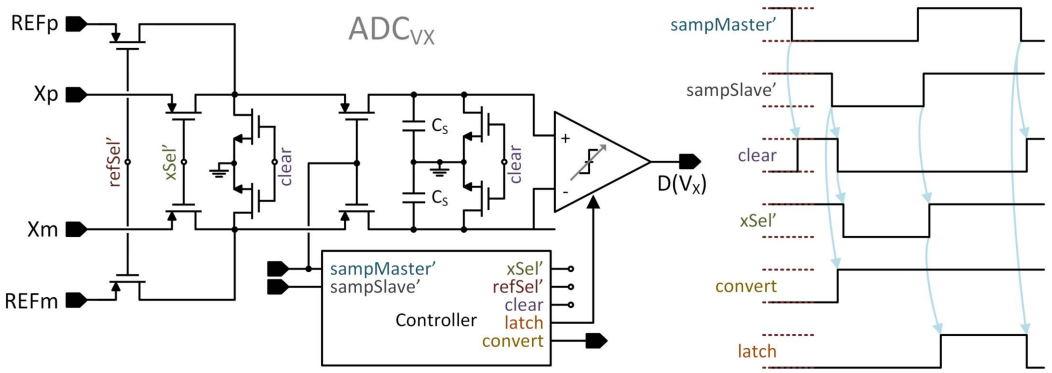

Fig. 11. Schematic of ADC_{vx}, with modes of operation (left) and key control signals during a conversion cycle when the circuit is operating in the "normal" mode (right).

To obtain acceptable quantization accuracy of V_{xi}, a sufficient number of samples must be accumulated in each bin_i [55]. If one bin is allocated for each code of $D(V_o)$, there will be many bins (thousands) and it will take a long time to fill each bin_i with a statistically significant amount of data. This problem is solved by grouping contiguous ranges of $D(V_o)$ codes into shared "macro" bins. Empirically, it is found that even for low numbers of these shared bins (16, 32, or 64 bins in total), the estimation accuracy is minimally affected. Meanwhile, digital complexity and acquisition speed are greatly improved. As shown in Fig. 10, practical implementation of macro-bin sorting is simply a matter passing the $D(V_x)$ and $D(V_o)$ data streams into a demux controlled by the MSBs of $D(V_o)$.

3) SDR Estimation: For each bin_i, the scalar outputs $\overline{D}_{xi}$, $\overline{D}_{oi}$, and N_i are fed into the estimator block of Fig. 10. First, digital values are converted into absolute units of volts. For $\overline{D}_{oi}$, this is a straightforward linear scaling

$$\hat{V}_{oi} = \overline{D}_{oi} \cdot 2(V_{\text{REFP}} - V_{\text{REFM}}). \tag{8}$$

The other, $\overline{D}_{xi}$, represents a point on the noise CDF curve of ADC_{vx}. Thus, it must be linearized into a PDF σ and then converted to volts using the conversion factor $k_{v2\sigma}$ obtained during the preparation steps

$$\hat{V}_{xi} = k_{v2\sigma}\sqrt{2} \cdot erf^{-1}(2 \cdot \overline{D}_{xi} - 1). \tag{9}$$

With estimates of $\hat{V}_{xi}$ and $\hat{V}_{oi}$ for each bin_i in hand, the next step is to estimate A_{ol}. To account for the signal-dependent, nonuniform sorting of $D(V_o)$ codes into different bins, (3) is modified by weighting the data from each bin by factor N_i, the number of samples averaged into bin_i

$$\hat{A}_{ol} = \frac{\sum_{i=a}^{b} N_i \cdot \hat{V}_{oi}^2}{\sum_{i=a}^{b} N_i \cdot \hat{V}_{oi} \cdot \hat{V}_{xi}}. \tag{10}$$

The summations iterate bin_i from indexes a to b, which usually are chosen to be the minimum and maximum bin indexes. Note that the computational cost of this calculation is a function of the number of bins used, not the number of data samples captured, keeping digital power costs low.

Likewise, variance can be calculated in a similarly compact form that also accounts for the nonuniform, signal-dependent sorting of data into the bins

$$\text{Var}(x, N) = \frac{\sum_{i=a}^{b} N_i \cdot (x_i - \mu_x)^2}{\sum_{i=a}^{b} N_i} \tag{11}$$

where μ_x is the N_i-weighted mean value of x

$$\mu_x = \frac{\sum_{i=a}^{b} N_i \cdot x_i}{\sum_{i=a}^{b} N_i}. \tag{12}$$

Finally, applying (10) and (11) to (5), we arrive at the desired estimator equation

$$\widehat{\text{SDR}}_{dB} = 10 \cdot \log_{10}\left(\frac{\beta^2 \cdot \text{Var}(\hat{V}_o, N_i)}{\text{Var}(\hat{V}_{xi} - \hat{V}_{oi}/\hat{A}_{ol}, N_i)}\right). \tag{13}$$

A by-product of this processing procedure is that the dynamic INL of the amplifier's closed-loop transfer function is also estimated. This is further discussed and visualized later in Section V.

D. Circuit Implementation

A simplified schematic of ADC_{vx} is drawn in Fig. 11. The 1-bit stochastic ADC consists of a sampling network and a comparator. The comparator is a two-stage structure with a digitally tunable threshold similar to [56, Fig. 5]. The design requirements of the comparator are quite relaxed: offset is canceled digitally, noise is not only tolerated but desired, and decision speed is not critical. Rather, the main challenge in implementation is the sampling network. ADC_{vx} connects to nodes X_P and X_M in Fig. 4, where the "golden" information of the main signal path resides as trapped charge. It is critical that ADC_{vx} samples in a way that does not corrupt this information. Two things accomplish this. First, pre-clearing ensures that when ADC_{vx} connects to X_P and X_M, the charge sharing that occurs will be signal-independent. Second, ADC_{vx} samples at exactly the same moment as the main signal path's sampler in the subsequent pipeline stage. This is done by driving the gates of the sampling switches in ADC_{vx} with the same physical wire that drives the gates of the next-stage sampler (i.e., *sampMaster'* in Fig. 11 is *track early'*, the inverse of *track early* in Figs. 4 and 5). To support all of the modes of operation that are required, in particular the

need to be able to connect to a reference ladder, a two-level master/slave sampling approach is used. When connecting to the reference DAC, *refSel'* is set to "0." When connecting to the X nodes, *xSel'* is derived from the *amplify* phase of the given stage. Periodically taking ADC_{vx} offline to re-measure $k_{v2\sigma}$ in the presence of PVT variation can be done in parallel with normal operation, but this creates a different loading condition at node X in Fig. 10 compared to normal operation. Including a second "dummy" sampling network in ADC_{vx} that connects to the X nodes when ADC_{vx} is offline for $k_{v2\sigma}$ estimation can overcome this.

A size of 5 fF is chosen for the sampling capacitor C_S. This both minimizes the loading of the X nodes and helps to generate noise for the stochastic conversion process. However, with such a small capacitance, clock-feedthrough from the sampling switches causes considerable common-mode voltage shift in the sampled signal. In the case of the pMOS sampler shown here, this can be handled by using an nMOS-input comparator.[2]

All clock inputs to ADC_{vx} as well as all internal timing controls in ADC_{vx} follow the asynchronous, event-driven approach of the overall system. The blue arrows in the timing diagram in Fig. 11 indicate the event-driven state flow for one possible timing scenario in the "normal" mode of operation. However, due to the asynchronous nature of the system, edge arrivals can vary depending on events elsewhere in the main ADC (e.g., the relative edge relationships of *sampMaster'* and *sampSlave'*), and the controller is designed with all possibilities in mind. The controller can be programmed to only latch the comparator once every Nth cycle, allowing further power reduction if algorithm convergence speed is already sufficient. The *idle* mode of operation described in Fig. 11 presents the same loading conditions to the X nodes as the *normal* mode, allowing for transparent background operation regardless of whether SDR estimation is enabled or not. The only impact on the main signal path of ADC_{vx} sampling the X nodes is a small decrease in amplifier loop gain and bandwidth and a slight shift in the common-mode level, but compared to other contributors to C_P in Fig. 4 such as the ringamp input capacitance and the MDAC feedback switch, the influence of ADC_{vx} is negligible.

V. Measured Performance

The ADC described in this work is fabricated in a 16-nm CMOS technology and occupies an active area of 0.194 mm^2 (360 μm $\times$ 540 μm), excluding decoupling. A die photograph is shown in Fig. 12.

A. System Performance

System-level performance is summarized in Fig. 13. For a -1 dBFS input tone, the ADC produces a low-frequency SNDR and SFDR of 64.1 and 78.5 dB, respectively [see Fig. 13(c)], and a 2-GHz Nyquist SNDR and SFDR of 61.9 and 75.2 dB, respectively [see Fig. 13(d)]. A single, fixed configuration of digital settings and stage gain correction coefficients (determined off-chip) are used for all measurements reported. At the maximum sampling rate of 4 GS/s, total power consumption is 75 mW, with 11.9 mW consumed by the input buffer, 2.4 mW by the clock buffer, and 15.2 mW per channel. The ADC voltage references are $V_{\text{REFM}} = 50$ mV and $V_{\text{REFP}} = 850$ mV, allowing the ringamps to utilize 89% of the available 900-mV core supply. The input buffer operates on a 1.8-V supply. Interleaving spurs due to channel mismatch are one-time foreground calibrated using precision tunable delay elements inserted into each channel's clock path. Degraded SFDR observed at certain frequencies in the input frequency sweep of Fig. 13(f) is dominated by HD2 and HD3. It is determined that this distortion is not generated in the ADC core. The exact root cause is not certain, but evidence suggests that it is related to aspects of the chip interface boundary, either the off-chip packaging and PCB, on-chip signal and clock buffers, or some combination of both.

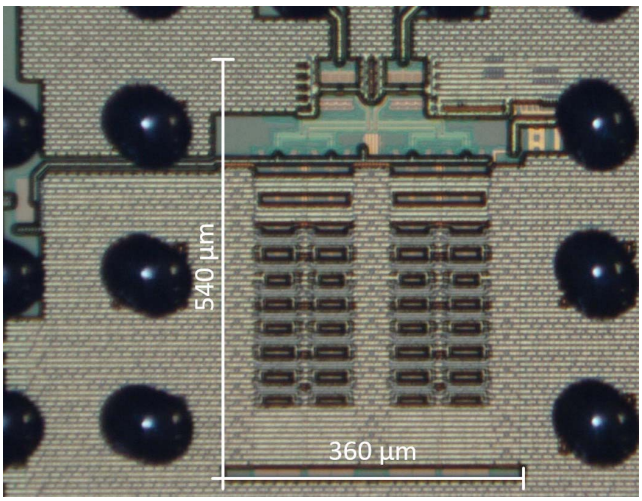

Fig. 12. Die photograph.

Fig. 14 provides insight into the ringamp tuning landscape, where the two digital bias controls D_biasP and D_biasM of Fig. 7 in the channel 0, stage 1 ringamp are independently swept and the SNDR of channel 0 is measured. Notably, the line defined by D_biasN = D_biasP is guaranteed to intersect the diagonal band of optimum or near-optimum biasing codes in Fig. 14. This means that by enforcing D_biasN = D_biasP, control complexity can be reduced to a simple 1-D tuning problem.

Table I shows a comparison to the state-of-the-art in high-resolution multi-gigasample direct-RF sampling ADCs. This design reports the best performance in all three key metrics of accuracy, linearity, and power efficiency. Taken together, this represents a roughly order-of-magnitude improvement in the state of the art. This is made possible not only by the direct performance advantages of ring amplifiers but also the indirect advantage of greater design freedom when amplification is relaxed as an architectural constraint.

[2]In the implemented circuit, for reasons of design reuse, a pMOS-input comparator was used. To shift the common-mode back to approximately mid-level, C_S is split into two halves, and the "ground" terminal of one of the halves is down-shifted from V_{DD} to V_{SS} before converting. For simplicity, we do not show this in Fig. 11 as it is not an essential aspect of implementation.

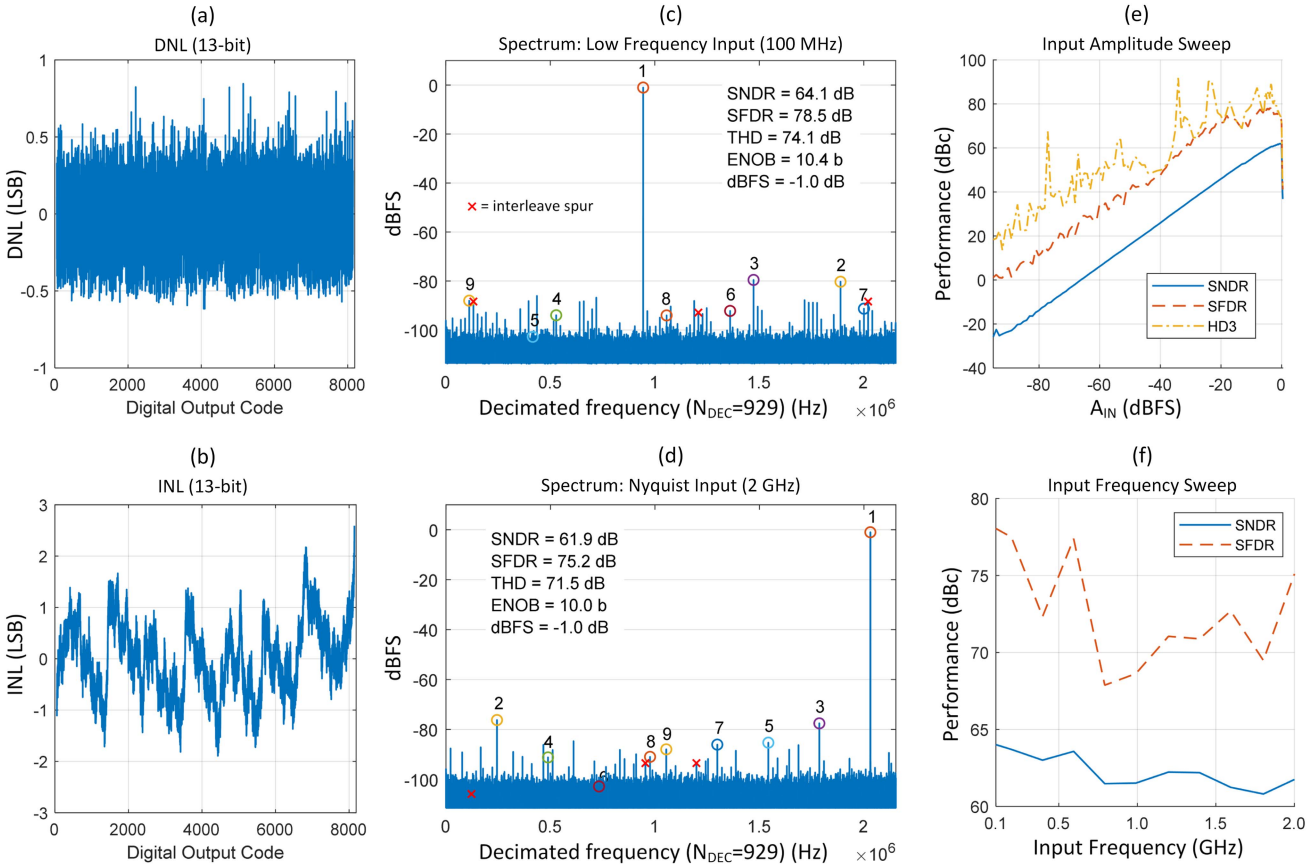

Fig. 13. ADC performance summary.

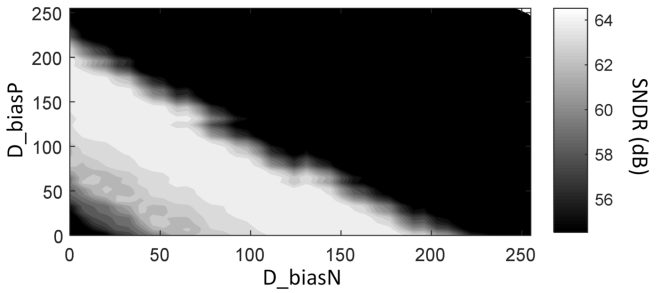

Fig. 14. 2-D Sweep of ringamp bias controls in channel 0, stage 1.

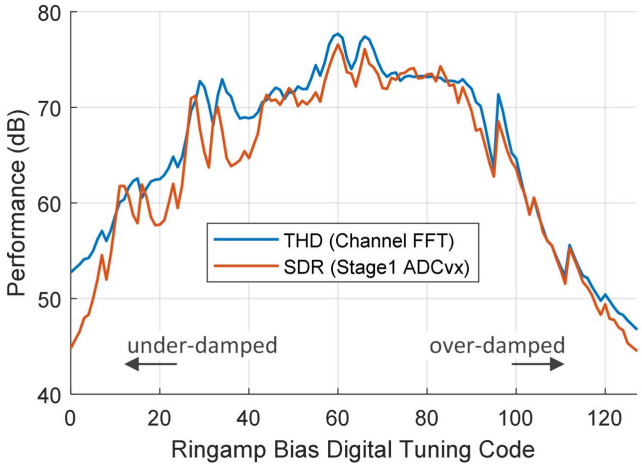

Fig. 15. Sweep of SDR estimated by ADC_{vx} in pipeline channel 0, stage 1 compared to THD of the channel 0 digital output spectrum.

B. SDR Monitoring

An ADC_{vx} SDR monitor is included in stages 1–3 of each channel but omitted in the latter stages where accuracy requirements are relaxed and adaptive bias control is less useful. Fig. 15 shows the estimated SDR of channel 0, stage 1 across a 1-D sweep of ringamp digital bias code by setting D_biasP = D_biasN in Fig. 7. For comparison, the THD of channel 0 obtained from an FFT of the channel's digital output spectrum is also plotted. Considering that THD and SDR are not identical quantities and that the THD also includes distortion from additional sources such as the input buffer and the other pipeline stages, the correspondence between these two curves is very close, suggesting accurate estimation. Also notable is the broad range of bias codes that give near-optimal SDR, indicating that coarse tuning control is sufficient for PVT-robust operation. Most of the bumps and jumps in the curves of Fig. 15 are due to DNL breaks of the ringamp bias DACs, designed to have a redundant sub-binary tuning characteristic. This can be smoothed out with a bias DAC re-design that prioritizes monotonicity.

TABLE I

COMPARISON WITH STATE-OF-THE-ART MULTIGIGASAMPLE DIRECT-RF SAMPLING ADCS

	This Work	Ramkaj ISSCC 2019 [3]	Vaz ISSCC 2017 [4]	Devarajan ISSCC 2017 [5]	Straayer ISSCC 2016 [6]	Wu ISSCC 2016 [7]	Ali VLSI 2016 [8]
Architecture	**Pipeline**	Pipe-SAR	Pipe-SAR	Pipeline	Pipeline	Pipeline	Pipeline
Sampling rate [Gsps]	**4**	5	4	10	4	4	5
Number of channels	**4**	8	8	8	16	4	2
Channel rate [Msps]	**1000**	625	500	1250	250	1000	2500
Residue Amplifier	**Ringamp**	Open-loop	Open-loop	Opamp	ZCD	Opamp	Opamp
Resolution [bit]	**13**	12	13	12	14	13	14
Technology [nm]	**16**	28	16	28	65	16	28
Supply [V]	**0.9, 1.8**	1.0	- , - [3]	1.0, - [3]	1.0, 1.8	0.8, 1.8	0.9, 1.8
ENOB LF input [bit]	**10.4**	10.1	10.3	9.4	-	9.7 *	10.2
ENOB Nyquist [bit]	**10.0**	9.4	9.2	8.8	8.9	9.0	9.3 *
SNDR LF input [dB]	**64.1**	62.4	63.8	58.5 *	-	60 *	63
SNDR Nyquist [dB]	**61.9**	58.5	57.3	55	55.5	56	58 *
SFDR LF input [dB]	**78.5**	75.2	75 *	70 *	-	72 *	80
SFDR Nyquist [dB]	**75.2**	65.4	67.0	64	64.0	68.0	70 *
Power [mW]	**75**	159	513	2900	2214	300	2300
FoM_{Walden} [fJ/c.step]	**18**	46	214	631	1130	145	709
$FoM_{Schreier}$ [dB]	**166**	161	153	147	145	154	148
Area [mm²]	**0.194** [1]	1.76 [1]	1.04 [2]	20.2	11.0 [1]	0.34 [1]	14.4 [2]

* = estimated from figure
- = information not provided
[1] = active area only
[2] = total area w/ decap
[3] = buffer supply exists but not listed

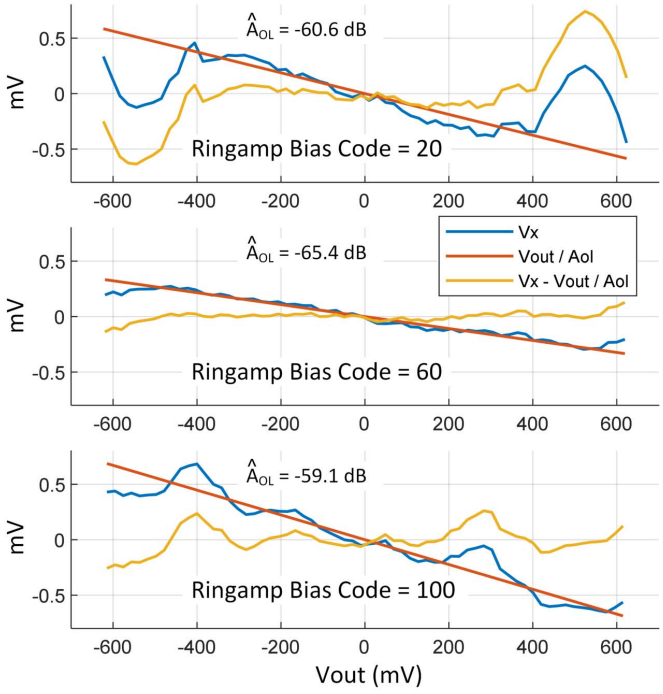

Fig. 16. Binned $\hat{V}_x$ and $\hat{V}_o$ data used to estimate SDR for three bias codes contained in the sweep shown in Fig. 15.

Fig. 16 shows the binned $\hat{V}_x$ and $\hat{V}_o$ data (number of bins = 64) used to estimate SDR for three different bias codes contained in the sweep shown in Fig. 15. The three traces in each plot visualize the primary components of (4): V_x stripped of dc offset, first-order gain error, and higher order distortion. In addition to estimating SDR, this information also provides a measurement of "dynamic INL" of the ringamp as a function of output voltage. Clear differences can be observed in the higher order distortion characteristics of these three bias conditions, representing under-damped (code 20), critically damped (code 60), and over-damped (code 100) settling behavior. This provides unique insight into the way that incomplete settling and its impact on distortion vary as a function of both output voltage and bias condition.

VI. CONCLUSION

Linear amplifiers remain a fundamental tool of analog signal processing, despite their challenges in scaling well with technology. They have clearly imposed a performance bottleneck on many circuits and systems, and efforts to circumvent this bottleneck have profoundly shaped the design landscape. Conversely, a general solution (or diverse set of specific solutions) for scalable amplification offers the possibility of an equally profound restructuring of design tradeoffs. The rehabilitation of existing amplifier-intensive architectures like pipelined ADCs is just one aspect of this reshaping. An even larger potential for innovation may lie in the increased design freedom offered to all architectures, even ones where amplifiers are not typically used.

With regard to ringamps in particular, two key aspects of future work can help to enable this restructuring. First, continued innovation and refinement of robustness and industrialization techniques will enable ringamps to serve as a simple, designer-friendly, commodity block that can be used without hesitation. Just as industrialization techniques for dynamic Gm-C and Gm-R amplifiers have been demonstrated using both analog and digital approaches [21], [57], [58], a diverse set of both analog and digital ringamp robustness techniques will provide designers with more options for tackling specific applications and design challenges. Good progress has already been made in this regard, but there is still room for further improvement of both analog [28], [29], [42], [44] and digital [33], [34], [38] approaches.

Second, opportunity for future work lies in the expansion of ringamps into the wider world of circuits. It is important to remember that ring amplification is a generic set of design principles, not a specific topology. Beyond just applying existing ringamp topologies to other applications, there are opportunities to re-imagine ringamps from the ground-up, drawing from the core concepts of dominant output pole, high-frequency internal poles, inverter-based, dynamic stabilization, and large-signal feedback mechanisms. We already see this topological diversification beginning to occur as ringamps move into other applications such as noise-shaping SAR [44] and power management [40]. There is no inherent restriction binding ringamps to CMOS technology, and certain non-CMOS and beyond-CMOS technologies may also hold opportunities for future work.

Ultimately, a number of recent publications do indeed seem to indicate a resurgence of linear amplifiers in circuit design. The ringamp-based ADC presented in this article significantly advances the state of the art in its target application of high-linearity direct-RF sampling. Likewise, ADCs utilizing

ringamps in other application spaces are finding similar success in pushing the boundaries of performance [28]–[30], [32], [34]. This encouraging trend suggests that an era of circuits and systems enhanced by scalable amplification may now be upon us.

REFERENCES

[1] V. Subramanian *et al.*, "Planar bulk MOSFETs versus FinFETs: An analog/RF perspective," *IEEE Trans. Electron Devices*, vol. 53, no. 12, pp. 3071–3079, Dec. 2006.

[2] B. Murmann. *ADC Performance Survey 1997–2020*. Accessed: Aug. 2020. [Online]. Available: http://web.stanford.edu/~murmann/adcsurvey.html

[3] A. Ramkaj *et al.*, "A 5 GS/s 158.6 mW 12b passive-sampling 8×-interleaved hybrid ADC with 9.4 ENOB and 160.5 dB FoMS in 28 nm CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2019, pp. 62–64.

[4] B. Vaz *et al.*, "A 13b 4 GS/s digitally assisted dynamic 3-stage asynchronous pipelined-SAR ADC," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2017, pp. 276–277.

[5] S. Devarajan *et al.*, "A 12b 10 GS/s interleaved pipeline ADC in 28 nm CMOS technology," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2017, pp. 288–289.

[6] M. Straayer *et al.*, "A 4 GS/s time-interleaved RF ADC in 65 nm CMOS with 4 GHz input bandwidth," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Jan. 2016, pp. 464–465.

[7] J. Wu *et al.*, "A 4 GS/s 13b pipelined ADC with capacitor and amplifier sharing in 16 nm CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Jan. 2016, pp. 466–467.

[8] A. M. A. Ali *et al.*, "A 14-bit 2.5 GS/s and 5 GS/s RF sampling ADC with background calibration and dither," in *Proc. IEEE Symp. VLSI Circuits (VLSI-Circuits)*, Jun. 2016, pp. 1–2.

[9] A. M. A. Ali *et al.*, "A 14 bit 1 GS/s RF sampling pipelined ADC with background calibration," *IEEE J. Solid-State Circuits*, vol. 49, no. 12, pp. 2857–2867, Dec. 2014.

[10] J. Wu *et al.*, "A 5.4 GS/s 12b 500 mW pipeline ADC in 28 nm CMOS," in *Proc. Symp. VLSI Circuits*, Jun. 2013, pp. C92–C93.

[11] K. Yoshioka *et al.*, "Digital amplifier: A power-efficient and process-scaling amplifier for switched capacitor circuits," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 27, no. 11, pp. 2575–2586, Nov. 2019.

[12] M. Figueiredo, R. Santos-Tavares, E. Santin, J. Ferreira, G. Evans, and J. Goes, "A two-stage fully differential inverter-based self-biased CMOS amplifier with high efficiency," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 58, no. 7, pp. 1591–1603, Jul. 2011.

[13] Y. Chae and G. Han, "Low voltage, low power, inverter-based switched-capacitor delta-sigma modulator," *IEEE J. Solid-State Circuits*, vol. 44, no. 2, pp. 458–472, Feb. 2009.

[14] J. K. Fiorenza, T. Sepke, P. Holloway, C. G. Sodini, and H.-S. Lee, "Comparator-based switched-capacitor circuits for scaled CMOS technologies," *IEEE J. Solid-State Circuits*, vol. 41, no. 12, pp. 2658–2668, Dec. 2006.

[15] L. Brooks and H.-S. Lee, "A zero-crossing-based 8-bit 200 MS/s pipelined ADC," *IEEE J. Solid-State Circuits*, vol. 42, no. 12, pp. 2677–2687, Dec. 2007.

[16] B. Razavi, "Charge steering: A low-power design paradigm," in *Proc. IEEE Custom Integr. Circuits Conf.*, Sep. 2013, pp. 1–8.

[17] S.-H.-W. Chiang, H. Sun, and B. Razavi, "A 10-bit 800-MHz 19-mW CMOS ADC," *IEEE J. Solid-State Circuits*, vol. 49, no. 4, pp. 935–949, Apr. 2014.

[18] B. Verbruggen, M. Iriguchi, and J. Craninckx, "A 1.7 mW 11b 250 MS/s 2× interleaved fully dynamic pipelined SAR ADC in 40 nm digital CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2012, pp. 466–468.

[19] M. S. Akter, K. A. A. Makinwa, and K. Bult, "A capacitively degenerated 100-dB linear 20–150 MS/s dynamic amplifier," *IEEE J. Solid-State Circuits*, vol. 53, no. 4, pp. 1115–1126, Apr. 2018.

[20] 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.

[21] W. Jiang, Y. Zhu, M. Zhang, C.-H. Chan, and R. P. Martins, "A 7.6 mW 1 GS/s 60 dB 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.

[22] A. M. A. Ali *et al.*, "A 12b 18 GS/s RF sampling ADC with an integrated wideband track-and-hold amplifier and background calibration," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2020, pp. 250–252.

[23] N. Dolev, M. Kramer, and B. Murmann, "A 12-bit, 200-MS/s, 11.5-mW pipeline ADC using a pulsed bucket brigade front-end," in *Proc. Symp. VLSI Circuits*, Jun. 2013, pp. C98–C99.

[24] I. Ahmed, J. Mulder, and D. A. Johns, "A low-power capacitive charge pump based pipelined ADC," *IEEE J. Solid-State Circuits*, vol. 45, no. 5, pp. 1016–1027, May 2010.

[25] L. Brooks and H.-S. Lee, "A 12b, 50 MS/s, fully differential zero-crossing based pipelined ADC," *IEEE J. Solid-State Circuits*, vol. 44, no. 12, pp. 3329–3343, Dec. 2009.

[26] J. K.-R. Kim and B. Murmann, "A 12-b, 30-MS/s, 2.95-mW pipelined ADC using single-stage class-AB amplifiers and deterministic background calibration," *IEEE J. Solid-State Circuits*, vol. 47, no. 9, pp. 2141–2151, Sep. 2012.

[27] 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.

[28] 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.

[29] 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.

[30] 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, Dec. 2019.

[31] B. Hershberg and U.-K. Moon, "A 75.9 dB-SNDR 2.96 mW 29 fJ/conv-step ringamp-only pipelined ADC," in *Proc. Symp. VLSI Circuits (VLSIC)*, Jun. 2013, pp. C94–C95.

[32] T.-C. Hung, J.-C. Wang, and T.-H. Kuo, "A calibration-free 71.7 dB SNDR 100 MS/s 0.7 mW 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.

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

[34] B. Hershberg, N. Markulic, J. Lagos, E. Martens, D. Dermit, and J. Craninckx, "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," in *Proc. IEEE Symp. VLSI Circuits*, Jun. 2020, pp. 1–2.

[35] J. Lagos, B. Hershberg, E. Martens, P. Wambacq, and J. Craninckx, "A single-channel, 600-MS/s, 12-b, ringamp-based pipelined ADC in 28-nm CMOS," *IEEE J. Solid-State Circuits*, vol. 54, no. 2, pp. 403–416, Feb. 2019.

[36] J. Lagos, B. P. Hershberg, E. Martens, P. Wambacq, and J. Craninckx, "A 1-GS/s, 12-b, single-channel pipelined ADC with dead-zone-degenerated ring amplifiers," *IEEE J. Solid-State Circuits*, vol. 54, no. 3, pp. 646–658, Mar. 2019.

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

[38] Y. Chen *et al.*, "A 625 MS/s, 12-bit, SAR assisted pipeline ADC with effective gain analysis for inter-stage ringamps," in *Proc. IEEE 45th Eur. Solid State Circuits Conf. (ESSCIRC)*, Sep. 2019, pp. 197–200.

[39] Y. Lim and M. P. Flynn, "A 100 MS/s, 10.5 bit, 2.46 mW comparator-less pipeline ADC using self-biased ring amplifiers," *IEEE J. Solid-State Circuits*, vol. 50, no. 10, pp. 2331–2341, Oct. 2015.

[40] J.-E. Park, J. Hwang, J. Oh, and D.-K. Jeong, "A 0.4-to-1.2 V 0.0057 mm^2 55fs-transient-FoM ring-amplifier-based low-dropout regulator with replica-based PSR enhancement," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2020, pp. 492–494.

[41] B. Xiao *et al.*, "An 80 mA capacitor-less LDO with 6.5 μA quiescent current and no frequency compensation using adaptive-deadzone ring amplifier," in *Proc. IEEE Asian Solid-State Circuits Conf. (A-SSCC)*, Nov. 2019, pp. 39–42.

[42] T.-C. Hung and T.-H. 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, May 2019.

[43] 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.
[44] X. Tang *et al.*, "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.
[45] D. W. Cline and P. R. Gray, "A power optimized 13-b 5 Msamples/s pipelined analog-to-digital converter in 1.2 μm CMOS," *IEEE J. Solid-State Circuits*, vol. 31, no. 3, pp. 294–303, Mar. 1996.
[46] J. Li and U.-K. Moon, "Background calibration techniques for multistage pipelined ADCs with digital redundancy," *IEEE Trans. Circuits Syst. II, Analog Digit. Signal Process.*, vol. 50, no. 9, pp. 531–538, Sep. 2003.
[47] B.-M. Min, P. Kim, F. W. Bowman, D. M. Boisvert, and A. J. Aude, "A 69-mW 10-bit 80-MSample/s pipelined CMOS ADC," *IEEE J. Solid-State Circuits*, vol. 38, no. 12, pp. 2031–2039, Dec. 2003.
[48] T. Sundstrom, C. Svensson, and A. Alvandpour, "A 2.4 GS/s, single-channel, 31.3 dB SNDR at Nyquist, pipeline ADC in 65 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 46, no. 7, pp. 1575–1584, Jul. 2011.
[49] E. Martens, B. Hershberg, and J. Craninckx, "Wide-tuning range programmable threshold comparator using capacitive source-voltage shifting," *Electron. Lett.*, vol. 54, no. 25, pp. 1417–1418, Dec. 2018.
[50] T. Matsuura *et al.*, "A 240-mbps, 1-W CMOS EPRML read-channel LSI chip using an interleaved subranging pipeline A/D converter," *IEEE J. Solid-State Circuits*, vol. 33, no. 11, pp. 1840–1850, Nov. 1998.
[51] O. Stroeble, V. Dias, and C. Schwoerer, "An 80 MHz 10b pipeline ADC with dynamic range doubling and dynamic reference selection," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, vol. 1, Feb. 2004, pp. 462–539.
[52] X. Yang, Y. Zhu, C.-H. Chan, S.-P. U, and R. P. Martins, "Analysis of common-mode interference and jitter of clock receiver circuits with improved topology," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 65, no. 6, pp. 1819–1829, Jun. 2018.
[53] Y. Dong, W. Yang, R. Schreier, A. Sheikholeslami, and S. Korrapati, "A continuous-time 0–3 MASH ADC achieving 88 dB DR with 53 MHz BW in 28 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 49, no. 12, pp. 2868–2877, Dec. 2014.
[54] A. M. A. Ali *et al.*, "A 16-bit 250-MS/s IF sampling pipelined ADC with background calibration," *IEEE J. Solid-State Circuits*, vol. 45, no. 12, pp. 2602–2612, Dec. 2010.
[55] 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.
[56] 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.
[57] S. Liu, H. Han, Y. Shen, and Z. Zhu, "A 12-bit 100-MS/s pipelined-SAR ADC with PVT-insensitive and gain-folding dynamic amplifier," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 67, no. 8, pp. 2602–2611, Aug. 2020.
[58] B. Verbruggen *et al.*, "A 2.1 mW 11b 410 MS/s dynamic pipelined SAR ADC with background calibration in 28 nm digital CMOS," in *Proc. Symp. VLSI Circuits*, Jun. 2013, pp. C268–C269.

Benjamin Hershberg (Member, IEEE) received the H.B.S. degrees in electrical engineering and computer engineering from Oregon State University, Corvallis, OR, USA, in 2006, and the Ph.D. degree in electrical engineering from Oregon State University, in 2012, for work focusing on ADCs and scalable, low-power switched-capacitor amplification solutions, including the invention of ring amplification.

In 2013, he joined the Interuniversity Microelectronics Center (IMEC), Leuven, Belgium. He is 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 International Solid-State Circuits Conference (ISSCC) 2019 Lewis Winner Award for outstanding paper.

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, Nice, France, working on integrated high-speed digital phase-locked loop (PLL) frequency synthesizers in 65-nm CMOS for cable-TV. From 2010 to 2015, he was an Analog and RF ICs Designer with STMicroelectronics, Grenoble, 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 multi-standard 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 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 the 2019 International Solid-State Circuits Conference (ISSCC).

Barend van Liempd (Member, IEEE) received the B.Sc. and M.Sc. degrees in electrical engineering from the Eindhoven University of Technology, Eindhoven, The Netherlands, in 2009 and 2011, respectively, and the Ph.D. degree from the Vrije Universiteit Brussel, Brussels, Belgium, in 2017, in collaboration with IMEC, Heverlee, Belgium. His Ph.D. dissertation concerned tunable, highly integrated RF front-end circuits and modules in silicon-on-insulator (SOI) CMOS.

In 2011, he joined IMEC, working as an R&D Engineer on multi-standard transceivers, until 2014, when he became a Ph.D. Researcher and later a Senior Researcher in 2017. He was appointed as a Program Manager of Radar ICs in 2018 and now leads IMEC's radar IC R&D activities. His research interests are analog, RF, and mm-wave circuits for wireless and sensing applications. He has authored or coauthored more than 30 articles, patents, and patent applications.

Dr. van Liempd was a co-recipient of the 2015 NXP Prize at the European Microwave IC (EuMIC) Conference and the 2019 Lewis Winner Award for outstanding paper at the International Solid-State Circuits Conference (ISSCC).

Ewout Martens (Member, IEEE) received the M.Sc. degree in electronic engineering and the Ph.D. degree *(summa cum laude)* from the Katholieke Universiteit Leuven (KU 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 the Chief Scientist focused on analog design automation. In 2010, he joined IMEC, Leuven, as a Researcher for WLAN transceivers and is currently working as a Principal Member of Technical Staff within IMEC's IoT Department. His interests and work are related to the development of innovative architectures for RF receiver front-ends, transceiver building blocks like phase-locked loop (PLL), filters and LNA, and ADCs for various applications, including image sensors and wireless receivers.

Dr. Martens served as a Technical Program Committee Member for the Symposium on Very Large Scale Integration (VLSI) Circuits from 2017 to 2020.

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 Interuniversity Microelectronics Center (IMEC), Leuven, Belgium, on digital subsampling phase-locked loops (PLLs) and polar transmitters.

He is currently a Research Scientist at 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 currently serves for the Technical Program Committee of the IEEE Very Large Scale Integration (VLSI) Circuits Symposium. He was a co-recipient of the International Solid-State Circuits Conference (ISSCC) 2019 Lewis Winner Award for outstanding paper.

Jorge Lagos (Member, IEEE) received the B.Sc. degree *(summa cum laude)* in electronics engineering from the 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 very-large-scale integration (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 was a co-recipient of the 2019 International Solid-State Circuits Conference (ISSCC) Lewis Winner Award for outstanding paper.

Jan Craninckx (Fellow, IEEE) received the M.S. and Ph.D. *(summa cum laude)* degrees in microelectronics from the 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 voltage-controlled oscillators (VCOs) and phase-locked loops (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 an IMEC fellow. He has authored and coauthored more than 200 articles, book chapters, and patents. 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.

Dr. Craninckx is/was a regular member of the Technical Program Committee for several Solid State Circuits Society (SSCS) conferences, was the chair of the SSCS Benelux chapter (2006–2011), SSCS Distinguished Lecturer (2012–2013), and elected SSCS AdCom member (2017–2019). He was an Associate Editor (2009–2016) and Editor-in-Chief (2016–2019) of the IEEE JOURNAL OF SOLID-STATE CIRCUITS. He currently serves as the RF Subcommittee Chair for the IEEE International Solid-State Circuits Conference.