A 5GS/s 7.2 ENOB Time-Interleaved VCO-based ADC Achieving 30.5fJ/cs

Maarten Baert, *Student Member, IEEE*, and Wim Dehaene, *Senior Member, IEEE*

Abstract—This paper presents an 8-channel time-interleaved VCO-based ADC achieving 7.2 ENOB at 5 GS/s in 28nm CMOS. A high-speed ring oscillator with feedforward cross-coupling and a shared tail transistor is combined with an asynchronous counter in order to improve the resolution while minimizing the power consumption. Asynchronous double sampling is used to enable reliable sampling of the asynchronous counter state. On-chip digital calibration is used to compensate for channel mismatch and nonlinear distortion, and sampling time mismatch is corrected using tunable clock delays. With a total power consumption of just 22.7 mW it achieves a Walden figure of merit of 30.5fJ/cs.

Index Terms—ADC, VCO-based, asynchronous logic, time interleaving, digital calibration.

I. INTRODUCTION

TECHNOLOGY scaling has resulted in significant performance improvements for digital CMOS circuits in terms of both speed and power consumption. Unfortunately, analog circuits do not benefit from technology scaling to the same extent. The ever-decreasing supply voltage makes some traditional analog techniques unusable, and exacerbates noise and distortion issues. In the domain of analog-to-digital converters (ADCs), technology scaling has motivated the transition from pipeline ADCs to successive approximation register (SAR) ADCs with a capacitive DAC. The elimination of the analog amplifier allows these ADCs to scale significantly better than traditional pipeline-based designs, however the comparator and capacitive DAC are still inherently analog components susceptible to noise and mismatch, which still limits the scaling potential.

VCO-based ADCs provide an interesting alternative with great scaling potential as they replace almost the entire ADC with digital circuits [1]. The general idea is to connect the input voltage of the ADC to the control input of a voltage-controlled oscillator (VCO), and then measure the oscillation frequency of the VCO in order to derive the unknown input voltage. A conceptual block diagram of such a VCO-based ADC is shown in figure 1. The frequency measurement is accomplished by connecting the VCO to a cycle counter, and sampling this counter at regular intervals. The difference between two consecutive counter values is proportional to the frequency of the VCO, which in turn is proportional to the input voltage. Since the VCO can be implemented as a simple ring oscillator, the resulting architecture is highly digital, the only true analog component being the sampling switch. The

Fig. 1. Conceptual block diagram of a VCO-based ADC.

resolution of such a VCO-based ADC depends on both the tuning range of the VCO and the sample rate. The maximum effective number of bits (ENOB), ignoring thermal noise and distortion, is given by the following equation:

$$\text{ENOB}_{\max} \approx \log_2 \left(\frac{F_{\max} - F_{\min}}{F_{\text{sample}}} \right)$$

This equation shows the inherent trade-off between sample rate and resolution of VCO-based ADCs: by halving the sample rate, we can gain one bit in resolution. This intuitively makes sense since the VCO converts voltage-domain information to time-domain information. As a result, the resolution of VCO-based ADCs does not depend directly on the available supply voltage, but on the time available for conversion and the speed of the circuit (which may depend on the supply voltage and many other factors). As digital technology gets faster, the resolution of VCO-based ADCs will increase. [2]

In the following sections, a brief overview of existing VCO-based designs is given and we present several new techniques to improve the resolution and linearity while maintaining low power consumption. With our improvements, VCO-based ADCs become a viable choice for high-speed, medium resolution applications such as next-generation Ethernet and high-speed wireless communication.

II. RELATED WORK

The concept of using a VCO for analog-to-digital conversion was originally introduced in 1999 [3], but did not gain much popularity until a decade later, after technology scaling had made time-based techniques significantly more attractive. Initially VCOs were mainly used as quantizers in continuous-time Delta-Sigma ADCs, an application for which they are particularly well-suited due to their inherent first-order noise shaping and dynamic element matching (DEM). Although this is still their most popular application, other VCO-based architectures have gained popularity as well. Five main classes can be distinguished:

- *Open-loop VCO-based ADCs*: The analog input is connected directly to the VCO and no further analog processing is performed. These designs have the highest scaling potential because they require almost no analog circuitry,

M. Baert and W. Dehaene are with the Department of Electrical Engineering, ESAT-MICAS, Katholieke Universiteit Leuven, 3001 Leuven, Belgium (e-mail: maarten.baert@esat.kuleuven.be, wim.dehaene@esat.kuleuven.be).

M. Baert is an SB PhD fellow at Research Foundation – Flanders (FWO).

however the resolution is limited. They are also very sensitive to PVT variations. Digital calibration is usually required to correct the nonlinearity of the VCO tuning curve. [1], [4]–[7]

- *Hybrid VCO-based ADCs*: VCO-based converters are embedded in a SAR or pipeline architecture, where they fulfill a function similar to a flash ADC. These designs are able to overcome the resolution limitations of purely VCO-based designs and still achieve high speed, however the architecture is quite complex and usually contains more analog circuitry than purely VCO-based designs, reducing the scaling potential. [8]–[11]
- *Purely VCO-based noise-shaping ADCs*: One or more VCOs are used as quantizers in a noise-shaping architecture. Simple first-order designs usually rely on the intrinsic noise shaping properties of the VCO, whereas higher-order designs use a feedback loop to create a Delta-Sigma ADC. MASH architectures are also possible. When feedback is used, the resolution and linearity of the VCO are less of a concern, since the feedback loop improves both properties, at the cost of bandwidth. However these higher-order designs contain significant analog circuitry and as a result are less scalable. Since a VCO is able to process continuous-time signals, they are most commonly seen in continuous-time Delta-Sigmas, but discrete-time designs also exist. [2], [12]–[16]
- *Hybrid VCO-based noise-shaping ADCs*: A VCO is added to a traditional ADC to create a noise-shaping ADC, similar to a noise-shaping SAR. The traditional ADC acts as the first stage, while a VCO is added as the second stage to provide first-order noise shaping. These designs combine the qualities of traditional ADCs with the inherent noise-shaping properties of VCO-based ADCs, at the expense of increased complexity. [17], [18]
- *PWM-based ADCs*: The analog input signal is converted to a PWM signal whose duty cycle is measured by a VCO. The PWM signal can be applied directly as the control voltage of the VCO such that the average frequency of the VCO changes linearly with the duty cycle, or the VCO can be used as a time-to-digital converter (TDC) in order to measure the duty cycle directly. This approach avoids distortion due to tuning curve nonlinearity, however it shifts the problem of creating a sufficiently linear voltage-to-time conversion to the PWM generator, which tends to become the limiting factor. [19]–[21]

This paper will focus specifically on open-loop VCO-based ADCs, however many of the techniques that will be described can also be applied to other types.

III. CHALLENGES AND PROPOSED IMPROVEMENTS

Although the basic architecture of a VCO-based ADC is relatively simple, practical implementations face several challenges:

- In order to obtain both high speed and high resolution, a VCO with an extremely high oscillation frequency is required. For example, a 5 GS/s 8 ENOB ADC would require a VCO with a tuning range of 1.28 THz, which is clearly not realistic.

Fig. 2. Multi-phase VCRO-based ADC with multiple counters.

- A high-speed VCO requires an equally high-speed cycle counter, however the maximum frequency of most counters is significantly lower than the maximum oscillation frequency of a ring oscillator.
- Both the VCO and the counter consume significant power.
- Sampling the state of the counter is nontrivial due to timing skew and metastability: straightforward sampling of the counter output may result in corrupted values.
- In order to avoid distortion, the nonlinearity of the VCO tuning curve must be compensated with digital calibration. In practice this limits the usable region of the VCO tuning curve since the distortion compensation also amplifies quantization noise.

In order to improve the resolution without requiring excessively high frequencies, most VCO-based ADCs use voltage-controlled ring oscillators (VCROs) which produce multiple phases. One such implementation [2] is shown in figure 2. A separate counter is connected to each of the 8 output phases of the VCRO, and all counter values are added together. By doing this, the apparent frequency of the VCRO increases by a factor of 8, resulting in a 3-bit resolution improvement. Unfortunately, adding more stages to a ring oscillator in order to improve the resolution doesn't work: if the sample rate is kept constant, doubling the number of stages does double the number of output phases, but it also halves the oscillation frequency of the VCRO. As a result, the effective resolution remains constant (for a constant sample rate). In other words, the resolution is determined not by the oscillation frequency of the VCRO, but by the delay of a single delay element. In practice, the minimum number of stages will be chosen based on two constraints: the minimum number of stages required to obtain a stable oscillation, and the maximum operation frequency of the counters, the latter generally being the limiting factor.

Unfortunately the large number of counters required for a multi-phase VCRO-based ADC drastically increases the power consumption and exacerbates the challenge of sampling the counters without corrupted values. The power can be reduced significantly by switching to an architecture with just one counter [6], as shown in figure 3. By sampling not only the counter but also the state of the VCRO itself, the same resolution is achieved without multiple power-hungry counters. Unfortunately this architecture still requires a lot

Fig. 3. Multi-phase VCRO-based ADC with single counter.

Fig. 4. Overview of one channel of the proposed ADC architecture.

of sampling flip-flops, which also consume significant power. This results in a power tradeoff between the counter and sampling flip-flops: increasing the number of stages reduces the operating frequency of the counter, thus reducing its dynamic power consumption, however it also increases the number of sampling flip-flops, thus increasing the energy required for sampling. The optimal value depends not only on the design of the counter and sampling flip-flops, but also on the sampling frequency. One drawback of this design is that it requires synchronized sampling of the VCRO and counter state, which is challenging due to the delay of the counter.

Although the power is reduced and the resolution is improved compared to the initial single-phase design, the VCRO frequency is still limited by the fact that the counter needs to be able to operate at the maximum VCRO frequency without consuming excessive power. This makes it difficult to create a low-power VCRO-based ADC with sufficient resolution.

IV. PROPOSED ADC ARCHITECTURE

In this work, several improvements over existing designs will be introduced. Figure 4 shows a high-level overview of one channel of the proposed ADC architecture [22]. Each channel consists of a bootstrapped sampling circuit, a pseudodifferential VCO-based core, and a digital calibration circuit. Each core consists of a VCRO, a counter, sampling flip-flops and decoders. Since the maximum frequency of the counter is limiting the performance of the ADC, a high-speed asynchronous counter is used. This enables the use of either a VCRO with faster delay cells to gain resolution, or a VCRO with less stages to reduce the number of sampling flip-flops and save power. However, the propagation delay of the asynchronous counter complicates sampling: straight-forward sampling would produce a mixture of old and new bits. In order to compensate for this propagation delay and

Fig. 5. Voltage controlled ring oscillator with (a) standard cross-coupling and (b) feedforward cross-coupling.

Fig. 6. Current starving with (a) individual current-starving transistors and (b) shared tail transistor.

sample the counter value without risk of metastability, a novel asynchronous double sampling technique is introduced. The nonlinearity of the VCRO tuning curve is compensated with low-power on-chip digital calibration based on lookup tables. The VCO-based core is discussed in the following sections, followed by the clock generation and sampling process, and the on-chip digital calibration.

A. Voltage-controlled Ring Oscillator

Most VCRO-based ADCs use differential ring oscillators, since these can be implemented with an even number of stages. Differential implementations also avoid propagation delay mismatch between the rising and falling edges [13], and can be sampled more easily with differential sampling flip-flops. However differential ring oscillators can be susceptible to common-mode oscillations, which must be avoided by ensuring that the common-mode gain is below unity, or at least significantly smaller than the differential gain. Typically this is done by using OTA-like delay cells with a differential pair as input. However this approach has some downsides. A differential pair suppresses common-mode gain only when biased correctly, i.e. when the tail transistor is in saturation. Large common-mode signals can drive the tail transistor into the linear region, which results in a common-mode signal at the output. In some situations this can lead to large-signal common-mode oscillations even when the small-signal common-mode gain is below unity. In order to avoid this, resistors or diode-connected MOSFETs must be added to the output in order to limit the output swing of the delay cell, such that the next delay cell remains in saturation. This approach increases the power consumption of the oscillator significantly. In addition to this, it is challenging to design a ring oscillator with a wide tuning range and a sufficiently linear tuning curve while simultaneously ensuring that the delay cells remain properly biased to avoid common-mode oscillations.

The proposed design uses a different ring oscillator design based on current-starved inverters. Instead of relying on differential pairs, cross-coupling is used to suppress common-mode

Fig. 7. Transient response of the proposed VCRO design to control voltage changes. For clarity only the first four output phases and the drain voltage of the shared tail transistor are shown.

gain. Figure 5a shows the standard cross-coupling approach for inverter-based differential ring oscillators. Although this design is very simple and reliable, the addition of cross-coupled inverters reduces the oscillation frequency and also increases the power consumption. In the proposed design these disadvantages are avoided by using a feedforward cross-coupling approach as shown in figure 5b, where the inputs of the cross-coupled inverters are connected two stages back rather than within the same stage. When all inverters have the same drive strength, this design suppresses common-mode gain almost completely. Furthermore, the feedforward path actually increases the oscillation frequency rather than reducing it. When used as a VCO-based ADC, this change roughly doubles the resolution for the same power consumption.

In order to control the oscillation frequency, current starving is used. Current starving can be done either separately for each inverter [3] (as shown in figure 6a), or with a single shared tail transistor for the entire VCRO [13] (as shown in figure 6b). The first method has several downsides though. If only the low side or high side is current-starved, then only the fall time or rise time respectively will be affected. This results in significant imbalance which makes it harder to match the drive strength of different inverters in order to suppress common-mode gain. If both sides are current-starved, this problem is avoided and the tuning range is improved as well, however two bias voltages with opposite polarity and a different DC offset are required to achieve this, which is inconvenient. In addition to this, in both cases the VCRO frequency is quite sensitive to the supply voltage. As a result, supply noise can reduce the accuracy of the VCRO-based ADC dramatically. In the proposed design these downsides are avoided by using a single shared tail transistor. Rather than limiting the current of each transistor individually, this approach essentially adjusts the supply voltage seen by the ring oscillator in order to limit the total current. This approach has several advantages. Since the PMOS and NMOS transistors see the same gate-source voltage, the rise time and fall time remain balanced even though only one side is current-starved, so only one bias voltage is required. The shared tail approach can obtain a wide tuning range with a reasonably linear tuning curve. Since the current starving transistor is now outside the critical path, its size can be increased in order to reduce flicker noise and mismatch without affecting the oscillation frequency. Since the

Fig. 8. Tuning curve of the proposed VCRO design across process corners.

Fig. 9. Sensitivity of the VCRO tuning curve to voltage and temperature variations.

tail transistor is used in saturation for most of the input range, it acts as a current source which inherently isolates the VCRO from the supply voltage. As a result, it makes the VCRO far less susceptible to supply noise, which is particularly important in time-interleaved applications with many VCROs.

One downside of the shared tail transistor is that it introduces a low-frequency pole due to the parasitic capacitance at the drain of the tail transistor, which limits how quickly the VCRO will respond to a change in bias voltage. The size of the tail transistor must be chosen with this in mind. Figure 7 shows the transient response of the proposed VCRO design. Note that due to the design of the sampling circuit, the VCRO control voltage is reset to V_{CM} between samples. The response is sufficiently fast to allow the VCRO to settle during this reset phase, which is important to avoid memory effects and dynamic nonlinearity. For this reason no pole mitigation techniques [23] are required.

In order to minimize the power consumption, the transistors of the delay cells are designed close to their minimum size. As a result the device parasitics are very small and the wiring capacitance is dominant. In contrast, the shared tail transistor is given a large width and length in order to minimize the mismatch and flicker noise. The resulting tuning curve is shown in figure 8. The nominal usable tuning range is 2 GHz to 12.3 GHz, which corresponds to an input range of 0.3 V. Within this range the VCRO has good linearity and an average power consumption of just 0.5 mW. The tuning curve may vary significantly across process corners, however even under worst-case conditions the resolution loss is only 10%, which corresponds to an SNDR reduction of just 1 dB. Figure 9

Fig. 10. Simulated noise contribution due to quantization noise, delay cell mismatch and thermal noise before and after calibration.

shows the simulated sensitivity of the VCRO tuning curve to supply voltage and temperature variations. The effect of the supply voltage and temperature is significant, however it is smaller than the effect of process variations. These variations must be corrected with digital calibration.

Figure 10 shows the relative noise contribution due to quantization noise, delay cell mismatch and thermal noise. It is clear that quantization noise is the dominant noise source. Before calibration, the noise level is relatively constant across the entire input range, however after calibration the noise level is higher at the edges. This is because the slope of the VCRO tuning curve is lower in those regions, such that the signal (and noise) must be amplified by digital calibration in order to obtain a linear ADC. The optimal input range is determined by the tradeoff between this noise amplification effect and the ADC resolution.

B. High-Speed Asynchronous Counter

The oscillation frequency of the VCRO is limited by the maximum operating frequency of the counter. Increasing this frequency enables the use of faster delay cells in the VCRO in order to improve the resolution of the ADC, or reduce the number of stages and corresponding sampling flip-flops to reduce the power. For this reason, the choice of counter architecture is quite important for the overall performance of the ADC. Most VCO-based ADCs use simple synchronous counters (e.g. [2], [3], [6], [24]), which can only operate at a limited speed due to the long critical path created by carry propagation. The proposed design uses an asynchronous counter in order to overcome this limitation and operate at a much higher frequency.

The asynchronous counter, shown in figure 11, is an 8-bit ripple counter formed by four cascaded divide-by-four blocks. The state of each divider produces two bits of information, for a total of 8 bits. Compared to a standard ripple counter, this requires the same number of logic gates, but halves the propagation delay. Each divider consists of four cascaded latches, which are implemented as differential tri-state buffers with a shared clock transistor in order to reduce the load on the previous divider. Minimum sized cross-coupled inverters are used to preserve the latched value. This circuit can run at a much higher frequency than a synchronous counter, and since only the first divider runs at the full clock speed, the overall power consumption remains low. The counter has an

Fig. 11. High-speed asynchronous counter based on cascaded dividers, implemented using differential latches.

input frequency range from DC to 15 GHz, which is sufficient to cover the full frequency range of the VCRO. It consumes an average power of just 0.06 mW.

C. Asynchronous Double Sampling

Since the counter value changes asynchronously with respect to the sampling clock, it is important to consider metastability and signal skew, which can easily corrupt the sampled value. This is especially true for asynchronous counters. In order to minimize metastability, a fast differential sampling flip-flop similar to the ‘double tail’ circuit [25] is used. This allows reliable sampling of the VCRO state, however this alone is insufficient to reliably sample the counter state. At the maximum operating frequency, the propagation delay of the entire counter is much longer than one clock cycle. As a result, a straightforward sampling circuit will produce a mixture of old, new, and potentially metastable bits.

For synchronous counters, this problem can be solved with double sampling [1], which is illustrated in figure 12. The counter value is delayed by a half VCRO cycle, and both the original and the delayed value are sampled. At any point in time, at least one of these two values will be uncorrupted by transitions. By also sampling the VCRO output itself, the sampled VCRO output value can be used to select either the original or the delayed counter value. This approach always works: When the VCRO value is zero, there can’t be any transitions in the original counter signal, so its sampled value is always valid. Likewise, when the VCRO value is one, there can’t be any transitions in the delayed counter value, so the sampled delayed value is always valid. Of course, the VCRO output itself may be transitioning during sampling, so metastability is still possible in the flip-flop that samples the VCRO output. However in practice this risk is negligible: The internal feedback loop of the sampling flip-flop has a time constant of approximately 10 ps, so the probability that metastability will persist until the end of the 1.4 ns sampling period is in the order of 10^{-60} . For this reason it can be assumed that the metastable flip-flop will resolve to either zero or one arbitrarily. This will result in either the original or the delayed value being selected, both of which are valid. Note

Fig. 12. Synchronous double sampling timing diagram.

Fig. 13. Asynchronous double sampling timing diagram.

that the goal is not to eliminate metastability entirely, which is fundamentally impossible, but to make the probability of metastability sufficiently low and to ensure that the worst-case output error is less than one LSB. This method also makes the counter insensitive to propagation delay mismatch and jitter, since the impact of these effects is significantly smaller than the signal skew produced by the propagation delay itself.

In this work, the same principle is extended to asynchronous counters, as illustrated in figure 13. Since the output of each divide-by-four block is used as the counter for the next block, each divider value should be delayed by a half cycle relative to its own clock. The latch-based dividers shown in figure 11 naturally produce these delayed outputs, so this requires no extra hardware. Both the original and delayed value of each divider are sampled, as well as the VCRO output itself. The sampled VCRO output value is used to select either the original or delayed value of the first divider, as was the case with synchronous double sampling. However since the output of each divider acts as the clock for the next divider, the decoded value of each divider is used to select the correct value for the next divider. This ensures that the decoded bits are consistent with each other, even when the propagation delay of the entire counter is longer than one clock cycle. As a result, this asynchronous counter behaves just like a synchronous counter operating at frequencies up to 15 GHz.

D. Decoders

The sampled state of the ring oscillator produces a cyclic thermometer code, which is converted to a binary value using a binary search thermometer decoder. The decoder for the counter, as shown in figure 14, consists of three layers. The first layer implements the asynchronous double sampling scheme as described earlier. The second layer consists of XOR gates which convert the state of each divider to a binary value, in order to obtain a binary output. If the counter was connected directly to the VCRO, this would be sufficient, however in practice two buffer stages must be added between the VCRO and the counter in order to create sufficient drive strength. These buffers introduce some latency, so the counter value may

Fig. 14. Counter with sampling flip-flops and decoder for asynchronous double sampling and buffer latency correction.

lag one cycle behind the VCRO value if the falling edge of the VCRO output is still propagating through the buffers. In order to correct this, the state of the buffers is sampled as well, and a falling edge detector is added to detect when a transition is propagating through the buffers but hasn't reached the counter yet. The third layer of the counter decoder then increments the counter value by one when a falling edge is detected, and produces a counter value which is always consistent with the value from the VCRO.

E. Time Interleaving

In order to relax the speed requirements of the VCO-based ADC, 8 independent ADC channels are time-interleaved. When multiple VCO-based ADC channels are interleaved, the intrinsic noise shaping properties of the VCROs will create a periodic noise spectrum as discussed in [24]. The periodic noise spectrum may be useful for some niche applications, but for a general-purpose wideband ADC it is not particularly beneficial. Since a sampled input is provided to each channel, they essentially behave as independent first-order incremental Delta-Sigma ADCs. Therefore it seems attractive to reset the state of the VCRO and counter between samples, which would result in a flat noise spectrum and improve the SNR by 3 dB. Unfortunately it is quite challenging to reset and restart the VCRO sufficiently quickly without causing excessive supply noise that would cause interference with the other ADC channels. For this reason such a reset was not implemented in this work.

F. Clock Generation

An overview of the clock distribution network is shown in figure 15. The sampling clocks for the 8 time-interleaved channels are generated on-chip from the 5 GHz full-speed clock. The full-speed clock drives the sequence generator, which generates a pulse every 8 cycles that acts as a sequencing signal. This signal controls when the first channel will activate its sampling switch. Each channel delays the sequencing signal by one cycle before sending it to the next channel.

The schematic of the sequence generator is shown in figure 16. The implementation is fully differential to avoid skew between the positive and negative version of the signals. The full-speed clock is divided by 8 and then delayed by one

Fig. 15. Overview of the clock distribution network.

Fig. 16. Sequence generator.

cycle. The original and delayed version are combined to create a 1-cycle pulse every 8 cycles, which is then re-synchronized with the full-speed clock using another flip-flop.

Figure 17 shows the schematic of the channel clock generator. Each channel requires two clocks: the main sampling clock, and the ADC core clock which is delayed by a half cycle relative to the sampling clock. Both clocks can be disabled for testing purposes. Since the sampling process is very sensitive to jitter, each channel re-synchronizes the sequence signal with the full-speed clock using a low-jitter latch. The sampling clock path also contains a tunable capacitor for sampling time offset compensation.

G. Sampling Circuit

The sampling approach is similar to bottom-plate sampling in SAR ADCs. A timing diagram is shown in figure 18. During the tracking phase, the input side of the sampling capacitor is connected to the input using a bootstrapped switch, while the output side is connected to V_{CM} . During the hold phase, the input side is connected to V_{CM} instead, and the output side is connected to the VCRO-based ADC core. The clock signal to the input switch is delayed slightly such that the V_{CM} switch always switches off before the input switch does. Compared to traditional top-plate sampling, this approach is more linear because the sampling instant is determined by the V_{CM} switch rather than the input switch. Since the V_{CM} switch is not directly exposed to the input voltage, there are less signal-dependent effects that can affect the switch-off time

Fig. 17. Channel clock generator.

Fig. 18. Timing diagram of the sampling process.

and injected charge of the switch. Once the V_{CM} switch is open, the switch-off process of the input switch can no longer affect the charge on the sampling capacitor.

The input switch is bootstrapped in order to maintain a constant R_{on} . A standard bootstrapping circuit [26] is used with one modification: the bulk of the main switch is connected to the bottom plate of the bootstrap capacitor, so that the bulk voltage is equal to the input voltage during the tracking phase. This eliminates the bulk effect of the main switch and improves the linearity of the sampling circuit significantly. The resulting circuit is shown in figure 19. The V_{CM} switch is a simple PMOS since this is sufficient to meet the linearity requirements.

Thanks to the bottom-plate sampling approach, the VCRO-based core sees only V_{CM} during the tracking phase, so it runs at a fixed frequency, rather than being exposed to the transients produced by the sampling circuit. However in practice, the R_{on} of the V_{CM} switch still produces some ripple at the input of the VCRO-based core at high input frequencies, which is not desired. In order to eliminate this ripple, an additional set of isolation switches is introduced, as shown in figure 20. With these extra switches, the voltage at the input of the VCRO-based core is exactly V_{CM} . In addition to this, an extra switch is added in parallel to the two V_{CM} switches to provide a direct path from the positive to the negative side. This reduces the combined differential resistance of the switches by a factor 3, while only increasing the clock load by a factor 1.5. The original V_{CM} switches could theoretically be removed if the signal was perfectly differential, however in practice this often will not be the case. The VCRO-based core is sampled with a half-cycle delay compared to the sampling switch, such that the core is sampled in the middle of the tracking phase. This ensures that the VCRO-based core is not affected by transitions of the analog input during sampling. Furthermore, by resetting the VCRO control voltage to a fixed voltage between samples, dynamic nonlinear effects are avoided and

Fig. 19. Bootstrapped sampling switch.

Fig. 20. Sampling circuit with third V_{CM} switch and extra isolation switches.

the calibration process is greatly simplified.

H. Digital Calibration

Since the VCRO tuning curve is inherently nonlinear, the output value must be corrected in order to obtain a linear analog-to-digital conversion. Traditionally this is done using polynomial fitting. A 7th order polynomial is sufficient to meet the linearity requirements, however this would require 7 multiplications per sample, and the intermediate results would have to be represented with significantly more bits than the ADC itself in order to avoid excessive accumulation of rounding errors. The resulting calibration circuit would consume too much power.

In order to reduce the power consumption, a 32-segment piecewise linear fitting is used instead. The calibration circuit is shown in figure 21. A configurable barrel shifter is added in order to rescale the input value such that it matches the range of the lookup table. The 5 most significant input bits are used as the address for the lookup table. In order to perform linear interpolation, two consecutive points must be read out of the lookup table each cycle. Rather than implementing a two-port lookup table, we take advantage of the fact that exactly one even and one odd address is read each cycle. By splitting the lookup table into two separate smaller lookup tables for the even and odd addresses respectively, one read port for each lookup table is sufficient. Once the two endpoints of the current segment of the piecewise linear function have been read from the lookup table, they are linearly interpolated using the remaining 7 least significant input bits. This requires just one subtraction, one multiplication and one addition. Since the difference between two consecutive points is relatively small, the size of the intermediate result can be reduced from 12 to 10 bits in order to reduce the power consumption of the multiplier.

Fig. 21. Digital calibration circuit.

Fig. 22. Die photo with the most important functional blocks highlighted.

The total power consumed by on-chip digital calibration is only 2.7 mW.

V. MEASUREMENT RESULTS

In order to verify the performance of the proposed ADC architecture, a prototype was implemented in 28nm CMOS. Figure 22 shows a die photo with the most important functional blocks highlighted. The ADC consists of 8 time-interleaved channels, each of which contains a separate sampling circuit, ADC core and digital calibration circuit. A 1.0 V analog supply is used in order to improve the linearity and tuning range of the VCRO, while a 0.85 V digital and clock supply is used to minimize the power consumption. The total power consumption (including clock management and calibration) is 22.7 mW: 8.24 mW analog, 7.05 mW digital, 6.92 mW clock and 0.52 mW common-mode bias.

Figure 23 shows the measured transfer function of 32 ADC channels (across 4 dies). The mismatch between the channels is quite limited and can be easily calibrated. In terms of ADC resolution, these measurements show a standard deviation of 0.4% between channels on the same die, and an inter-die standard deviation of 1.2%. However the transfer function is quite sensitive to the supply voltage, which demonstrates the importance of a stable, low-noise supply. Figure 24 shows the change in SNDR and analog power consumption due to supply voltage changes.

The ADC is calibrated in two steps. During the first step, the on-chip digital calibration is configured to correct the nonlinearity and mismatch of each ADC channel. A full-scale 5 MHz sine wave is applied as a test signal, and the uncorrected ADC output is captured. The ADC output is then corrected using a polynomial that represents the inverse of the VCRO tuning curve. A numerical optimization algorithm is used to calculate the polynomial coefficients that minimize the difference between the corrected output signal and the

Fig. 23. Measured transfer function mismatch between 32 channels and transfer function sensitivity to supply voltage changes.

Fig. 24. Measured change in SNDR and analog power consumption due to supply voltage changes.

known test signal. These coefficients may be different for each channel due to mismatch. The resulting polynomials are converted to lookup table values and then loaded into the chip. From this point forward, the on-chip calibration circuit corrects the nonlinearity as well as the offset and gain error of each channel. Note that since the VCRO tuning curve is sensitive to the supply voltage and temperature, it is necessary to re-run the calibration procedure if either of these parameters changes.

During the second calibration step, the tunable delays in the sampling circuit of each channel are configured in order to correct the sampling time error. For this step a full-scale 2.4 GHz sine wave is applied to the input, and the ADC output is captured. When the output samples of each ADC channel are observed separately, a 100 MHz tone is seen in each channel. The phase of this tone is compared to the ideal value in order to determine the sampling time error. The tunable delays are then adjusted to correct this error. This procedure is repeated until the error is minimized.

Figure 25 shows the measured spectrum for a full-scale 400 MHz sine wave before and after calibration. Before calibration, the distortion is dominated by the third harmonic, which has an amplitude of -24.55 dBc. The fifth harmonic and various interleaving spurs caused by channel mismatch also degrade the performance. After calibration, the third harmonic is reduced to -66.55 dBc, and other harmonics and interleaving spurs are similarly suppressed, such that the SNDR improves from 24.37 dB to 46.42 dB, and the SFDR improves from 24.55 dB to 64.25 dB, however the noise floor is raised slightly. This demonstrates that the presented on-chip digital calibration is able to efficiently correct the mismatch and distortion inherent to VCO-based ADCs.

Fig. 25. Measured spectrum of a 400 MHz tone before and after calibration.

Fig. 26. Measured two-tone spectrum of a 1.7 GHz and 1.8 GHz tone with amplitude -6 dBFS before and after calibration.

Figure 26 shows the measured two-tone spectrum for a 1.7 GHz and 1.8 GHz tone with amplitude -6 dBFS before and after calibration. Before calibration, the third-order distortion results in high IM3 tones, but after calibration the IM3 tones are reduced greatly and the SFDR improves from 33.10 dB to 69.05 dB.

Figure 27 shows the SNDR and SFDR as a function of the input frequency for a 5 GHz sampling rate, and as a function of the sampling rate for a 330 MHz input frequency. The SNDR peaks at 46.7 dB and drops to 45.2 dB at the Nyquist frequency. The SFDR peaks at 60.3 dB and drops to 57.1 dB at the Nyquist frequency. Beyond the Nyquist frequency, the distortion of the sampling switch becomes the limiting factor. As predicted by theory, the ADC resolution improves as the sampling rate is reduced. However, thermal noise also becomes more dominant. Below 2 GHz, the bootstrapped sampling switch no longer works well due to charge leakage from the bootstrap capacitor, resulting in degraded linearity. Below 1 GHz, the counter inside the VCO-based core overflows due to the increased resolution and the output becomes invalid.

Figure 28 shows the SNDR as a function of the input amplitude for a 600 MHz input frequency and a 5 GHz sampling rate, before and after calibration. Before calibration, the performance is limited by channel mismatch at low amplitudes and by nonlinear distortion at higher amplitudes. After calibration, channel mismatch is almost completely eliminated and nonlinear distortion is reduced significantly. However at high amplitudes the SNDR still degrades slightly, since the calibration algorithm amplifies the quantization noise. Due to the S-shape of the VCRO tuning curve, the ADC inherently has better resolution in the center of the input range, and worse resolution near the edges. As a result, the dynamic range of 50.3 dB is noticeably higher than the peak SNDR of

Fig. 27. Measured SNDR and SFDR as a function of input frequency and sampling rate (average of 3 samples).

Fig. 28. Measured SNDR as a function of input amplitude for 3 samples.

TABLE I
COMPARISON TO STATE-OF-THE-ART ADCs

	This Work	[28]	[29]	[30]	[31]
Published	ISSCC 2019	ISSCC 2017	ISSCC 2015	CICC 2015	VLSI 2013
Architecture	TI VCO	TI SAR	TI Pipe/SAR	TI SAR	TI Pipe
Technology	28nm CMOS	28nm FDSOI	28nm CMOS	28nm CMOS	28nm CMOS
Supply (V)	1.0/0.85	1.8/1.0/0.8	1.0	1.0	1.8/0.9
Sample Rate (GS/s)	5.0	5.0	5.0	5.0	5.4
Interleaving	8x	16x	2x/8x	12x	2x
Power (mW)	22.7	150	150	76	500
SNDR (dB)	45.2	51.7	46.1	42	50
SFDR (dB)	57.1	59.2	58.5	54.4	65
Walden FOM (fJ/cs)	30.5	95.5	181.9	165	358.4
Schreier FOM (dB)	155.6	153.9	148.3	147.2	147.3
Core Area (mm ²)	0.023	0.184	0.45	0.57	0.4

46.2 dB. Although this is a disadvantage for single-tone tests, it is actually beneficial for multi-tone applications where most of the samples are located in the center of the input range.

Table I shows a comparison with several state-of-the-art ADCs. These ADCs all use a similar technology (28nm), and have similar sampling rates and resolution. The ADC design presented in this work consumes significantly less power, resulting in a Walden FOM that is more than three times lower than the next best design. The Schreier FOM is also the best in this comparison. Furthermore, the core area is 8 times smaller than the next smallest design, as a result of the compactness of the VCO-based approach. Figure 29 compares this work to all ADCs published at ISSCC and VLSI which achieve at least 6 ENOB, demonstrating a significant improvement over the state of the art.

Fig. 29. This work compared to other published designs. [27]

VI. CONCLUSION

A complete time-interleaved VCO-based ADC architecture with on-chip digital calibration has been implemented in 28nm CMOS. It uses a high-speed voltage-controlled ring oscillator with feedforward cross-coupling and a shared tail transistor to improve the resolution and linearity. A high-speed asynchronous counter is used to minimize the power consumption. Asynchronous double sampling ensures reliable sampling of the VCRO and asynchronous counter state. The resulting ADC achieves an SNDR of 45.2 dB near Nyquist with a power consumption of just 22.7 mW, resulting in a Walden FOM of 30.5 fJ/cs, which is three times lower than the best reported result achieving similar speed and resolution. Thanks to the compactness of the VCO-based approach, the area is 8 times smaller than comparable ADCs. Due to its highly digital nature, the performance of this architecture can be expected to further improve in smaller technology nodes. These results demonstrate that VCO-based ADCs are a viable alternative to traditional ADC architectures for high-speed, medium resolution applications such as next-generation Ethernet and high-speed wireless communication.

REFERENCES

- [1] J. Daniels, W. Dehaene, and M. Steyaert, "All-digital differential VCO-based A/D conversion," in *Proceedings of 2010 IEEE International Symposium on Circuits and Systems*, May 2010, pp. 1085–1088.
- [2] J. Kim, T. Jang, Y. Yoon, and S. Cho, "Analysis and design of voltage-controlled oscillator based analog-to-digital converter," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 57, no. 1, pp. 18–30, Jan 2010.
- [3] A. Iwata, N. Sakimura, M. Nagata, and T. Morie, "The architecture of delta sigma analog-to-digital converters using a voltage-controlled oscillator as a multibit quantizer," *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, vol. 46, no. 7, pp. 941–945, July 1999.
- [4] Y. Yoon and S. Cho, "A 1.5-GHz 63dB SNR 20mW direct RF sampling bandpass VCO-based ADC in 65nm CMOS," in *2009 Symposium on VLSI Circuits*, June 2009, pp. 270–271.
- [5] S. Rao, K. Reddy, B. Young, and P. K. Hanumolu, "A 4.1mW, 12-bit ENOB, 5MHz BW, VCO-based ADC with on-chip deterministic digital background calibration in 90nm CMOS," in *2013 Symposium on VLSI Circuits*, June 2013, pp. C68–C69.
- [6] P. K. Sharma and M. S. Chen, "A 6b 800MS/s 3.62mW Nyquist AC-coupled VCO-based ADC in 65nm CMOS," in *Proceedings of the IEEE 2013 Custom Integrated Circuits Conference*, Sep. 2013, pp. 1–4.
- [7] T. Wu and M. S. Chen, "A noise-shaped VCO-based nonuniform sampling ADC with phase-domain level crossing," *IEEE Journal of Solid-State Circuits*, vol. 54, no. 3, pp. 623–635, March 2019.

- [8] A. K. Gupta, K. Nagaraj, and T. R. Viswanathan, "A 1.1mW, 63.7dB-SNDR, 10MHz-BW hybrid voltage-time domain ADC," in *2014 IEEE 57th International Midwest Symposium on Circuits and Systems (MWSCAS)*, Aug 2014, pp. 1053–1056.
- [9] P. Zhu, X. Xing, and G. Gielen, "A 40MHz-BW 35fJ/step-FoM nonlinearity-cancelling two-step ADC with dual-input VCO-based quantizer," in *ESSCIRC 2014 - 40th European Solid State Circuits Conference (ESSCIRC)*, Sep. 2014, pp. 63–66.
- [10] W. El-Halwagy, P. Mousavi, and M. Hossain, "A 79dB SNDR, 10MHz BW, 675MS/s open-loop time-based ADC employing a 1.15ps SAR-TDC," in *2016 IEEE Asian Solid-State Circuits Conference (A-SSCC)*, Nov 2016, pp. 321–324.
- [11] Z. Ding, X. Zhou, and Q. Li, "A 0.5-1.1V 10B adaptive bypassing SAR ADC utilizing oscillation cycle information of VCO-based comparator," in *2018 IEEE Symposium on VLSI Circuits*, June 2018, pp. 93–94.
- [12] M. Park and M. Perrott, "A 0.13 μ m CMOS 78dB SNDR 87mW 20MHz BW CT $\Delta\Sigma$ ADC with VCO-based integrator and quantizer," in *2009 IEEE International Solid-State Circuits Conference - Digest of Technical Papers*, Feb 2009, pp. 170–171, 171a.
- [13] G. Taylor and I. Galton, "A mostly-digital variable-rate continuous-time delta-sigma modulator ADC," *IEEE Journal of Solid-State Circuits*, vol. 45, no. 12, pp. 2634–2646, Dec 2010.
- [14] —, "A reconfigurable mostly-digital delta-sigma ADC with a worst-case FOM of 160 dB," *IEEE Journal of Solid-State Circuits*, vol. 48, no. 4, pp. 983–995, April 2013.
- [15] S. Huang, N. Egan, D. Kesharwani, F. Opteynde, and M. Ashburn, "A 125MHz-BW 71.9dB-SNDR VCO-based CT $\Delta\Sigma$ ADC with segmented phase-domain ELD compensation in 16nm CMOS," in *2017 IEEE International Solid-State Circuits Conference (ISSCC)*, Feb 2017, pp. 470–471.
- [16] H. Maghami, P. Payandehnia, H. Mirzaie, R. Zangbaghi, S. Dey, J. Goins, K. Mayaram, and T. S. Fiez, "0.9V, 79.7dB SNDR, 2MHz-BW, highly linear OTA-less 1-1 MASH VCO-based $\Delta\Sigma$ with a novel phase quantization noise extraction technique," in *2019 IEEE Custom Integrated Circuits Conference (CICC)*, April 2019, pp. 1–4.
- [17] A. Sanyal and N. Sun, "A 18.5-fJ/step VCO-based 0-1 MASH $\Delta\Sigma$ ADC with digital background calibration," in *2016 IEEE Symposium on VLSI Circuits (VLSI-Circuits)*, June 2016, pp. 1–2.
- [18] K. Ragab and N. Sun, "A 12-b ENOB 2.5-MHz BW VCO-based 0-1 MASH ADC with direct digital background calibration," *IEEE Journal of Solid-State Circuits*, vol. 52, no. 2, pp. 433–447, Feb 2017.
- [19] S. Rao, B. Young, A. Elshazly, W. Yin, N. Sasidhar, and P. K. Hanumolu, "A 71dB SFDR open loop VCO-based ADC using 2-level PWM modulation," in *2011 Symposium on VLSI Circuits - Digest of Technical Papers*, June 2011, pp. 270–271.
- [20] P. Gao, X. Xing, J. Craninckx, and G. Gielen, "Design of an intrinsically-linear double-VCO-based ADC with 2nd-order noise shaping," in *2012 Design, Automation Test in Europe Conference Exhibition (DATE)*, March 2012, pp. 1215–1220.
- [21] E. Prefasi, E. Gutierrez, L. Hernandez, S. Paton, S. Walter, and U. Gaijer, "A 0.03mm², 40nm CMOS 1.5GS/s all-digital complementary PWM-GRO," in *2014 21st IEEE International Conference on Electronics, Circuits and Systems (ICECS)*, Dec 2014, pp. 116–119.
- [22] M. Baert and W. Dehaene, "A 5GS/s 7.2 ENOB time-interleaved VCO-based ADC achieving 30.5fJ/conv-step," in *2019 IEEE International Solid-State Circuits Conference - (ISSCC)*, Feb 2019, pp. 328–330.
- [23] A. Mukherjee, M. Gandara, B. Xu, S. Li, L. Shen, X. Tang, D. Pan, and N. Sun, "A 1-GS/s 20 MHz-BW capacitive-input continuous-time $\Delta\Sigma$ ADC using a novel parasitic pole-mitigated fully differential VCO," *IEEE Solid-State Circuits Letters*, vol. 2, no. 1, pp. 1–4, Jan 2019.
- [24] Y. Yoon, J. Kim, T. Jang, and S. Cho, "A time-based bandpass ADC using time-interleaved voltage-controlled oscillators," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 55, no. 11, pp. 3571–3581, Dec 2008.
- [25] D. Schinkel, E. Mensink, E. Klumperink, E. van Tuijl, and B. Nauta, "A double-tail latch-type voltage sense amplifier with 18ps setup+hold time," in *2007 IEEE International Solid-State Circuits Conference. Digest of Technical Papers*, Feb 2007, pp. 314–605.
- [26] M. Dessouky and A. Kaiser, "Input switch configuration suitable for rail-to-rail operation of switched op amp circuits," *Electronics Letters*, vol. 35, no. 1, pp. 8–10, Jan 1999.
- [27] B. Murmann, "ADC performance survey 1997-2018," [Online], available: <http://web.stanford.edu/~murmann/adcsurvey.html>.
- [28] J. P. Keane, N. J. Guilar, D. Stepanovic, B. Wuppermann, C. Wu, C. W. Tsang, R. Neff, and K. Nishimura, "An 8GS/s time-interleaved SAR ADC with unresolved decision detection achieving -58dBFS noise and 4GHz bandwidth in 28nm CMOS," in *2017 IEEE International Solid-State Circuits Conference (ISSCC)*, Feb 2017, pp. 284–285.
- [29] M. Brandolini, Y. Shin, K. Raviprakash, T. Wang, R. Wu, H. M. Geddada, Y. Ko, Y. Ding, C. Huang, W. Shin, M. Hsieh, W. Chou, T. Li, A. Shrivastava, Y. Chen, J. Hung, G. Cusmai, J. Wu, M. M. Zhang, G. Unruh, A. Venes, H. S. Huang, and C. Chen, "A 5GS/s 150mw 10b SHA-less pipelined/SAR hybrid ADC in 28nm CMOS," in *2015 IEEE International Solid-State Circuits Conference - (ISSCC) Digest of Technical Papers*, Feb 2015, pp. 1–3.
- [30] J. Fang, S. Thirunakkarasu, X. Yu, F. Silva-Rivas, K. Y. Kim, C. Zhang, and F. Singor, "A 5GS/s 10b 76mw time-interleaved SAR ADC in 28nm CMOS," in *2015 IEEE Custom Integrated Circuits Conference (CICC)*, Sep. 2015, pp. 1–4.
- [31] J. Wu, A. Chou, C. Yang, Y. Ding, Y. Ko, S. Lin, W. Liu, C. Hsiao, M. Hsieh, C. Huang, J. Hung, K. Y. Kim, M. Le, T. Li, W. Shih, A. Shrivastava, Y. Yang, C. Chen, and H. Huang, "A 5.4GS/s 12b 500mw pipeline ADC in 28nm CMOS," in *2013 Symposium on VLSI Circuits*, June 2013, pp. C92–C93.

Maarten Baert received the M.Sc. degree in electrical engineering (integrated circuits) from KU Leuven, Leuven, Belgium, in 2014. In the same year he joined the ESAT-MICAS research group at KU Leuven, where he currently works as a research assistant pursuing the Ph.D. degree. As of 2016 he is also a SB PhD fellow at FWO (Research Foundation - Flanders). His research interests include high-speed ADCs, time-based processing, high-speed asynchronous digital circuits and digital calibration techniques.

Wim Dehaene was born in Nijmegen, The Netherlands, in 1967. He received the M. Sc. degree in electrical and mechanical engineering in 1991 from the Katholieke Universiteit Leuven. In November 1996 he received the Ph.D. degree at the Katholieke Universiteit Leuven. After receiving the M.Sc. Degree Wim Dehaene was a research assistant at the ESAT-MICAS Laboratory of the Katholieke Universiteit Leuven. His research involved the design of novel CMOS building blocks for hard disk systems. The research was first sponsored by the IWONL (Belgian

Institute for Science and Research in Industry and agriculture) and later by the IWT (the Flemish institute for Scientific Research in the Industry). In November 1996 Wim Dehaene joined Alcatel Microelectronics, Belgium. There he was a senior project leader for the feasibility, design and development of mixed mode Systems on Chip. The application domains were telephony, xDSL and high speed wireless LAN. In July 2002 Wim Dehaene joined the staff of the ESAT-MICAS laboratory of the Katholieke Universiteit Leuven where he is now a full professor and head of the MICAS division. His research domain is circuit level design of digital circuits. The current focus is on ultra low power signal processing and memories in advanced CMOS technologies. Part of this research is performed in cooperation with IMEC, Belgium, where he is also a part time principal scientist. Wim Dehaene is teaching several classes on electrical engineering and digital circuit and system design. He is also very interested in the didactics of engineering and STEM. As such he is guiding several projects aiming to bring STEM to youngsters in secondary education and he is a teacher in the teacher education program of the KULeuven. Wim Dehaene is a senior member of the IEEE. Wim Dehaene was the technical program-chair for ESSCIRC 2017. He is a member of the ESSCIRC/ESSDERC steering committee and the ESSCIRC technical program committee. He has also served for several years on the ISSCC program committee.