

Exploring Speed Maximization of Frequency-to-Digital Conversion for Ultra-Low-Voltage VCO-Based ADCs

Viet Nguyen^{ID}, *Graduate Student Member, IEEE*, Filippo Schembari^{ID}, *Member, IEEE*,
and Robert Bogdan Staszewski^{ID}, *Fellow, IEEE*

Abstract—A frequency-to-digital converter (FDC) performs the role of precise frequency digitization within a voltage-controlled oscillator (VCO)-based ADC. To be compatible with energy-harvesting (EH) Internet-of-Things (IoT) devices, the development of ultra-low-voltage (ULV) FDCs is crucial, where the primary focus must be directed towards the maximization of data throughput under dramatic constraints of reliability and timing variability associated with deep-subthreshold operation. This article investigates the speed maximization of a 0.2 V full-custom ULV FDC design, consisting of an array of several parallel XOR-based FDC units, and the multi-rate decimation-filtering digital back-end. At the core of this broad exploration is a high-speed sense-amplify phase sampler (PS) featuring hardware redundancy, capable of sampling the phase of low-voltage-swing inputs. Particular focus is placed on the yield-based reliability-driven design methodology for the sense-amplify phase-sampling circuits running up to 40 MS/s and practical variability-mitigation strategies. To overcome the speed bottleneck in the digital back-end, a fully parallel bitstream-processing architectural composition of the computations for summation and decimation are proposed. Experimental verification through measurements of the FDC integrated within a 10-bit 160 kHz bandwidth (BW) open-loop VCO-based ADC across clock frequency with supply variations demonstrate robust operation of the first 0.2 V multi-phase FDC in the advanced 28 nm CMOS process.

Index Terms—0.2 V, bitstream, decimation, deep-subthreshold, energy-harvesting, frequency-to-digital converter (FDC), hardware redundancy, Internet-of-Things (IoT), multi-rate, phase-sampling, sense-amplifier, ultra-low voltage (ULV), variability, voltage-controlled oscillator (VCO)-based ADC.

I. INTRODUCTION

SUPPLY voltage (V_{DD}) scaling is commonly known to be the most effective technique in reducing the energy

Manuscript received 15 July 2022; revised 10 October 2022 and 15 November 2022; accepted 16 November 2022. Date of publication 28 November 2022; date of current version 27 February 2023. This work was supported in part by the Science Foundation Ireland under Grant 14/RP/I2921, in part by the European Union's Horizon 2020 Programme through Marie Skłodowska-Curie under Grant 747585, and in part by the Microelectronic Circuits Centre Ireland (MCCI) through Enterprise Ireland under Grant TC-2015-0019. This article was recommended by Associate Editor S. Zhu. (Corresponding author: Robert Bogdan Staszewski.)

Viet Nguyen and Robert Bogdan Staszewski are with the School of Electrical and Electronic Engineering, University College Dublin, Dublin 4, Ireland (e-mail: viet.nguyen@ucdconnect.ie).

Filippo Schembari was with the School of Electrical and Electronic Engineering, University College Dublin, Dublin 4, Ireland. He is now with Huawei Technologies, 20090 Segrade, Italy.

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TCSI.2022.3223468>.

Digital Object Identifier 10.1109/TCSI.2022.3223468




Fig. 1. Propagation delay (t_{inv}) of a digital inverter in super-threshold (0.8 V), near-threshold (0.4 V) and sub-threshold (0.2 V) regions of operation.

consumption of digital circuits [1], [2], [3]. The demand for ultra-low-power (ULP) systems has pushed the digital circuits to run from an ultra-low-voltage (ULV) supply below 0.3 V [4], [5], [6], [7], [8], where the transistors enter the subthreshold region. While its operational speed was exponentially slowed down, the aforementioned subthreshold circuits provided a compelling solution for energy-constrained applications, such as biomedical and environmental monitoring, because the signals of interest hardly exceeded a few kHz. On the other hand, transformational advances in energy harvesting (EH) have produced a resurgence of general interest in expanding the ULV design space beyond ULP digital processors [9], [10], [11], [12], [13], [14], [15], as ambient energy harvesters can produce output voltages ranging from 50 mV to 200 mV [16]. To enable the widespread adoption of EH for a broader range of IoT applications, the achievable data throughput should therefore be improved by several orders-of-magnitude. This means that conventional ULV design methodologies for circuit *energy minimization* must now be concurrently accompanied with their *speed maximization*.

The digital inverter gate in 28 nm LP CMOS technology experiences a dramatic increase in propagation delay (t_{inv}) of 200 $\times$ as V_{DD} scales from 0.8 V down to 0.2 V, as shown in Fig. 1. Despite this voltage scaling offering a 16 $\times$ switching-energy saving (CV_{DD}^2), the t_{inv} of around 4 ns indicate that ULV digital circuits can function efficiently with a clock rate pushed to only a few MHz. Moreover, a caveat of exploiting the shorter transistor channel length afforded with a nanometer-scaled CMOS process [i.e., by migrating to the 28 nm node to increase the transistor's transit frequency (f_T)] is that local intra-die mismatches become more prominent and prohibitively difficult to manage, especially in weak-inversion [17]. To cover the random delay variability

(i.e., $\sigma_{\text{inv}}/\mu_{\text{inv}}$) of 20% at 0.2 V compared to less than 5% at 0.8 V, extensive timing margins must be allocated such that the circuit will rarely operate at its fastest intended speed.

The class of hardware used for frequency-to-digital conversion seems particularly affected by the deep-subthreshold operation, yet, it has remained relatively unexplored. Frequency-to-digital converters (FDCs) are most known for their use in the voltage-controlled-oscillator (VCO)-based analog-to-digital converter (ADC), where they provide a precise digitization of the frequency-modulated (FM) VCO output [18], [19], [20], [21], [22], [23], [24], [25], [26], [27]. FDCs have also found importance in all-digital phase-locked-loops (ADPLLs) requiring high-speed phase accumulators [28], [29], [30]. Despite using the same foundational building blocks (i.e., latches and combinational logic gates) as digital processors and memories, their functionality and requirements are fundamentally and uniquely different. A basic digital processor performs arithmetic computations synchronous to the system clock. The processing throughput can be enhanced by dividing operations into multiple pipeline stages, while budgeting timing margins for the critical delay paths of the pipeline stage prevent timing violations from occurring. Timing variations are addressed with techniques for dynamic error monitoring, detection and resilience [31], [32]. FDCs, on the contrary, operate *asynchronously* to the system clock to sample and synchronize the continuous-time frequency information, presented as VCO phase outputs with sluggish transition edges and highly amplitude-modulated voltage swings, resembling more of an analog rather than strictly digital waveform. To digitize the frequency information with sufficient precision, FDCs employ *parallelism*, *oversampling* and *noise-shaping* principles, taking the form of a multitude of $\Delta\Sigma$ modulators operating in the multi-phase configuration. In addition to the high-speed digital processing for summation and decimation, the need for a fixed and fast clock rate (i.e., a constant and short sampling period which defines the time-base and timing resolution of the frequency measurement respectively) at the asynchronous phase-sampling interface raise serious concerns regarding reliability.

In this article, we introduce the deep-subthreshold multi-phase FDC (consisting of parallel XOR-based FDC units with embedded sense-amplify phase-sampling) of the 0.2-V open-loop VCO-based ADC shown in Fig. 2, whose front-end (consisting of the VCO core and the analog phase-processing circuits) was described in [1]; we also cover the downstream decimation filtering (DEC block). We wish to emphasize that the common theme for the proposed solutions throughout this work is *not* to propose novel FDC architectures, *nor* to over-exploit circuit sizing optimization strategies. Rather, the aim is to provide an in-depth study of the FDC operation in the deep-subthreshold regime and, moreover, to smartly and efficiently re-engineer the FDC and DEC circuit blocks for high-speed PVT-tolerant operation at a V_{DD} of 0.2 V.

Framed into context more familiar to the designer, 40 MS/s at 0.2 V translates to upwards of 10 GS/s operation around 1 V for 28 nm LP CMOS technology (a ballpark approximation based on the inverter propagation delay characteristics in Fig. 1) but exhibiting more than $4\times$ delay variability, impacting the yield dramatically. To address these critical concerns, the contributions of this article are as follows: 1) the ‘bottom-up’ exploration of the high-speed sense-amplify asynchronous phase-sampling interface; 2) the incorporation of hardware




Fig. 2. Frequency-to-digital conversion within a VCO-based ADC.

redundancy techniques for the mitigation of circuit variability; 3) the integration of the multi-phase FDC array for large BW frequency digitization; and 4) the architectural composition for the multi-rate decimation filtering and summation logic of the digital back-end. We begin in Section II by describing the main classes of open-loop FDCs and their challenges encountered at 0.2 V. Sections III and IV investigate the sense-amplifier phase-sampling FDC at, respectively, the *circuit* and *system* levels of abstraction. Section V describes the digital processing of the FDC bit-stream outputs, namely to provide digital filtering and downsampling necessary for output decimation to the Nyquist data-rate. Section VI offers experimental characterizations of FDC and DEC blocks around 0.2 V, at different clock frequencies.

II. FREQUENCY-TO-DIGITAL CONVERTERS

A. XOR-Based FDC

The simple FDC structure illustrated in Fig. 3(a) observes whether the VCO output oscillating at frequency f_{VCO} transitions within a clock period t_{CLK} of the synchronous sampling clock CLK (of frequency f_{CLK}) by sampling the phase state of the VCO and differentiating two consecutive samples using an XOR gate. With the multi-phase VCO outputs (e.g., ϕ_1 to ϕ_4), the expected digital output is $D_{\text{out}} = \mathcal{E}\{2N_{\text{FDC}}f_{\text{VCO}}/f_{\text{CLK}}\}$, where N_{FDC} represents the number of parallel FDC slices working in tandem. In the $N_{\text{FDC}} = 4$ configuration of Fig. 3(a), the toggling between counts 3 and 4 maps f_{VCO} to an intermediate frequency between $3f_{\text{CLK}}/8$ and $4f_{\text{CLK}}/8$. When the clock oversampling rate ($\text{OSR} = f_{\text{CLK}}/2\text{BW}$) is beyond 100, a single bit-stream can encode frequency information with an effective resolution of more than 10 bits after the decimation to the Nyquist rate [26]. The maximum detectable frequency is upper-bounded by $f_{\text{CLK}}/2$ but, interestingly, sub-sampled operation [33] is also permitted $[nf_{\text{CLK}}/2 < f_{\text{VCO}} < (n + 1)f_{\text{CLK}}/2]$ for integer values of $n > 0$.

B. Counter-Based FDC

The structure in Fig. 3(b) accumulates the number of VCO periods in the consecutive cycles of the much slower CLK [thus, allowing $f_{\text{VCO}} < (2^{N_{\text{CNT}}} - 1)f_{\text{CLK}}$], where N_{CNT} is the word length of the digital incrementer (CNT). The digital subtractor differentiates the consecutive stored counter values to determine a measure of the frequency information in its digital format ($D_{\text{out}} = \mathcal{E}\{f_{\text{VCO}}/f_{\text{CLK}}\}$) [28], [34].

The need to synchronize the asynchronously incremented binary values of CNT to the sampling clock leads to regular occurrences of metastability. Furthermore, mismatch-induced propagation delay skews across the incrementer bits can result in severe timing misalignment and cause unrecoverable hard-failures of the FDC even if the sampling process itself




Fig. 3. Conceptual architecture and timing diagrams for the (a) XOR-based FDC, (b) counter-based FDC and (c) coarse-fine FDC topologies processing the phase/frequency information of the upstream VCO core.

is free of metastability. Several solutions for high-voltage designs successfully combat this erroneous synchronization. The Gray coding scheme [35] limits metastability-induced sampling errors to within a single LSB, but sacrifices the speed of operation. Double sampling introduces redundancy to avoid timing windows where the count transitions, but at the cost of more than $2\times$ power consumption. The valid sampling window is less than half of the VCO period, demanding the critical path delay of the counting circuitry (e.g., the frequency divider [36]) to be significantly shorter. From the perspective of maximizing the clock frequency to increase OSR, the counter-based topologies become nonviable at ULV due to the need to run both sequential and combinational logic of the timing-sensitive multi-bit digital incrementer at the VCO frequency, a much higher speed than the sampling registers already running at the outlined upper limit of f_{CLK} .

The coarse-fine architecture shown in Fig. 3(c) extends the counter-based structure (integer-count value) to detect the phase transitions in a power-efficient manner (fractional-count value). The power efficiency of this architecture is similar to the XOR-based FDC (with decimation, it actually becomes superior) [24]. However, all of the issues encountered in the counter-based FDC (and their remedies) are inherited here. In fact, the coarse-fine partition brings additional timing issues due to the delay skew between the coarse and fine quantizers, along with phase delay mismatches within the fine-quantizer itself. Architectures utilizing phase reordering [37], scrambling [38] or linear-state feedback registers [39] can be employed, at the cost of an increased system complexity. It remains to be seen whether such solutions can be efficiently ported to the deep-subthreshold environment without functionally limiting its operation to a very slow, sub-optimal clock speed, where constraints such as leakage power, mismatches, variability and exponential decrease in maximal speeds place significant boundaries on what can be practically implemented.

C. Proposed XOR-Based FDC Unit Implementation

The bit-stream processing nature of the XOR-based FDC allows us to simultaneously maximize the “raw” f_{CLK} speed and increase the N_{FDC} parallelism while eliminating potential

sources of errors that can arise from the unpredictable timing variability. Therefore, the XOR-based FDC is utilized in our work as a base to explore *high-speed, multi-phase* frequency-to-digital conversion. For completeness, we briefly describe our circuit implementation of the single FDC unit presented in Fig. 4. The sense-amplifier flip-flop (SAFF) phase sampler (PS) reads the state of the asynchronous differential VCO waveforms (ϕ_p and ϕ_n). The amplification of the FDC differential inputs ϕ_p and ϕ_n upon assertion of $\overline{CLK_i}$ leads to a reduction in the FDC metastability window when sampling slow, commutating transition edges of the VCO phase outputs [40]. Negative-edge triggered master-slave D-flip-flops (DFF) store the regenerated signal Q0. The digital differentiator XOR uses Q1 and Q2 to isolate its transient dynamics from the input-dependent CLK-to-Q delay of the SAFF stage. The clock is routed in the opposite direction of data flow, and the negative clock skew is introduced to prevent hold violations [41]. $D_{out,d}$ is obtained by passing the FDC unit bitstream D_{out} through a 2^{nd} -order, downsample-by-4 decimation filter.

III. SENSE-AMPLIFY PHASE SAMPLING

A. StrongARM Sense-Amplifier

An intriguing aspect yet to be explored in deep-subthreshold is the underlying structure of the *high-speed analog* phase-sampling SAFF circuit. For the design of this PS, we concern mostly with its speed and thus metastability behavior. The FDC’s first-order noise-shaping of phase sampling non-idealities through its differentiation ($1 - z^{-1}$) means it is not additionally burdened by the strict constraints on the input-referred offset, noise and kickback encountered with high-speed voltage comparators. A baseline SAFF to first consider is the StrongARM-style sense-amplifier [40] shown in Fig. 5(a). The latch embeds the back-to-back ‘inverters’ M_{4-7} , whereby outputs Q_i and $\overline{Q_i}$ are preset to V_{DD} during the pre-charge phase ($\overline{CLK_i}$ is low). The back-end SR-latch to hold the state of $Q, \overline{Q}$ must be of the NAND-type. When $\overline{CLK_i}$ is high, M_{4-7} form a positive feedback network. The regeneration time constant (τ) equals to $C_{latch}/G_{m,latch}$, the large-signal transconductance discharging the effective capacitance seen at one of Q_i or $\overline{Q_i}$ to zero, dependent on the polarity of the differential input voltage ($\Delta V_\phi = \phi_p - \phi_n$). Detriment to




Fig. 4. Circuit implementation of the standard XOR-based FDC unit. The front-end asynchronous sampling interface and back-end decimation filter (DEC) are particularly difficult to design at 0.2V for speed-maximized conditions.

the latch regeneration, and thus the speed of the overall SAFF is the method whereby the analog inputs (ϕ_p, ϕ_n), through the input pair $M_{2,3}$, unbalance the latch inverters.

Observe that for a small ΔV_ϕ , the input common-mode sets the gate-source voltage V_{GS} of $M_{2,3}$ to be just $V_{DD}/2$, (i.e., 0.1V). The latch inverters become severely current-starved. The integration of a differential voltage across $Q_i, \bar{Q}_i$ is therefore sluggish, impeding the kick-start of the positive feedback action by M_{4-7} . Even for rail-to-rail inputs (ΔV_ϕ of ± 200 mV), the finite on-resistance R_{on} of $M_{2,3}$ within the series stacking of four transistors results in less voltage headroom allocated for M_{4-7} , effectively degenerating $G_{m,latch}$. The only solution that can remedy the detrimental effects of the $M_{2,3}$ input devices is to dramatically up-size their widths. To compensate for the subsequent increase in the capacitive loads seen at ϕ_p and ϕ_n , power-consuming input buffers would be required.

B. Modified Sense-Amplifier

An alternative approach taken in this work is to employ a modified SA structure, as shown in Fig. 5(b). Here, the input pair is connected in parallel with the latch to form a series stack of only three transistors. An obvious advantage of this configuration is that M_{4-7} have the maximum available voltage headroom (largely independent from the sizing of $M_{2,3}$), to achieve fast regeneration. To ensure that only dynamic power is consumed, a clock-gating pMOS device M_8 is inserted so that both outputs Q_i and $\bar{Q}_i$ are pulled to ground when CLK_i is high (during the pre-charge phase). The back-end SR-latch is modified to the NOR-type. To unbalance the latch inverters, signals ϕ_p, ϕ_n modulate the respective impedance seen looking into the drain of input devices M_2/M_3 . This offsets the latch away from the bi-stable state during the pre-charge phase, unlike in the StrongARM SA where the unbalancing of the latch only occurs at the start of the regeneration phase. Upon the activation of M_8 at the negative edge of CLK_i , M_{4-7} immediately initialize and direct the latch regeneration.

Figure 6 plots the regeneration time (t_{latch}) for both the StrongARM and modified SA structures (in red and blue color, respectively), lending further understanding to the intuitive explanations of their operation. To investigate the effects of ΔV_ϕ and input device $M_{2,3}$ width sizing on t_{latch} , the devices M_4, M_5 ($3\mu\text{m}/30\text{nm}$) and M_6, M_7 ($6\mu\text{m}/30\text{nm}$) are identically sized for both SA structures. The main clock devices (M_1 and M_8 in Fig. 5(a) and (b), respectively) need to exhibit low R_{on} so are sized to be $15\mu\text{m}/30\text{nm}$.




Fig. 5. (a) The commonly used StrongARM sense-amplifier flip-flop (SAFF) with 4 stacked transistors, and (b) the modified SAFF structure with 3 stacked transistors. The dashed lines in the timing waveforms depict the differential nature of $Q_i, \bar{Q}_i$.

Evidently, even with the maximum rail-to-rail input offset ($\Delta V_\phi = \pm 200$ mV), the StrongARM SA requires a $M_{2,3}$ width of $16\mu\text{m}$ to match the t_{latch} of ~ 6 ns exhibited by the modified SA with a corresponding $M_{2,3}$ width in the range of only $1\mu\text{m}$ to $4\mu\text{m}$. The StrongARM SA necessitates a large input device sizing, where for a ΔV_ϕ of ± 10 mV and ± 200 mV, the respective t_{latch} improves from 55 ns to 10 ns and from 9 ns to 5 ns as $M_{2,3}$ is swept from $1\mu\text{m}$ to $16\mu\text{m}$. The speed of the modified SA actually deteriorates when its input pair is inappropriately oversized (e.g., for $M_{2,3}$ width of $16\mu\text{m}$) as the direct coupling of the input branches to Q_i and $\bar{Q}_i$ adds unnecessary capacitance to C_{latch} . In that case, for a ΔV_ϕ of ± 10 mV and ± 200 mV, t_{latch} is slowed down from 11 ns to 17 ns and 6 ns to 8 ns, respectively as $M_{2,3}$ of the modified SA is swept from $1\mu\text{m}$ to $16\mu\text{m}$.

C. Discrete-Time Pre-Amplification

For low-voltage-swing signals, or in the occurrence of a small ΔV_ϕ , t_{latch} is > 12 ns, which may not be sufficient to fully regenerate and trigger the SR-latch in time. To further enhance the sampling capabilities of the SAFF, an upstream pre-amplification (*pre-amp*) stage [see Fig. 7(a)], identical to the modified SA of Fig. 5(b), is inserted to form the high-speed PS circuit shown in Fig. 7(b). The timing coordination between this latch stage and the modified SAFF is demonstrated in Fig. 7(c). When encountering a metastable state at L, $\bar{L}$ [enclosed within the red circle in Fig. 7(c)] due to the sampling of a small ΔV_ϕ , this on-going regenerating signal is further amplified to obtain a rail-to-rail waveform at $Q_i, \bar{Q}_i$ (SA output of the SAFF stage, enclosed within the green circle), thus improving the functionality of the SR-latch and the following flip-flop re-sampling stage (respectively the Q_0 and Q_1 outputs of Fig. 4). In terms of metastability, the effect of the pre-amp sampling latch is shown in Fig. 8 where t_{latch} , defined as the time from the negative edge of CLK_i to the fully regenerated $Q_i, \bar{Q}_i$ within the SAFF stage, now exhibits a much lower metastability window and appears to remain almost constant across ΔV_ϕ .

While this may give the impression that VCO phase outputs (ϕ_p, ϕ_n) with *minuscule* voltage swing levels




Fig. 6. Sense-amplifier latch regeneration time for the StrongARM (red) and modified (blue) configurations with various $M_{2,3}$ input transistor widths.




Fig. 7. High-speed phase-sampling interface incorporating (a) the discrete-time pre-amp sampling latch, (b) its integration with SAFF and (c) the associated timing diagram.

($\Delta V_{\phi, \max} \ll V_{DD, FDC}$) can be sampled correctly, the front-end pre-amp's internal mismatches must also be considered. As visualized with the 0-to-1 output transition threshold of the pre-amp sampling latch in Fig. 9 for 1000 Monte-Carlo runs, local mismatches shift the ideal output transition point away from $\Delta V_{\phi} = 0$. If $\Delta V_{\phi, \max}$ is too small, the SA output may remain 'stuck' to either 0 or 1 regardless of the VCO phase outputs' sampled differential voltage ($|\Delta V_{\phi}| < \Delta V_{\phi, \max}$). The histogram plot shows an input-referred voltage offset standard deviation σ of 26 mV, predominantly arising from the back-to-back regeneration inverters. This sets a 3σ lower bound of 80 mV on $\Delta V_{\phi, \max}$, such that $\Delta V_{\phi, \max} / V_{DD, FDC} > 0.4$.

D. Phase-Sampling FDC Dynamic Performance

Consider the three PS configurations shown in Fig. 10, namely (#1) D-flip-flop (DFF), (#2) SAFF, consisting of the




Fig. 8. Simulated regeneration time of the modified sense-amplifier latch without/with pre-amp sampling stage.




Fig. 9. Pre-amp's 0-to-1 output transition threshold and (b) simulated input-referred voltage offset characteristics with local mismatches.




Fig. 10. Signal flow of the various phase sampler configurations: (#1) DFF PS, (#2) SAFF PS and (#3) Pre-amp+SAFF PS.

modified SA followed by the SR-latch, shown in Fig. 5(b) and (#3) the pre-amp stage followed by the SAFF, shown in Fig. 7(b). Note that the DFF is sized to have a similar power consumption to the modified SA. To gain a deeper insight into the advantages of these structures, the PS configurations are empirically investigated by computing the VCO-based ADC's signal-to-quantization-noise ratio (SNR_Q) with a full-scale 150 kHz sinusoidal input, covering around 70% of the XOR-based FDC quantization range, where SNR_Q is dependent on the VCO tuning range, number of FDC readout phases (N_{FDC}), f_{CLK} and signal BW [1]. For a BW of 160 kHz and $N_{FDC} = 1$, the ideal SNR_Q is approximately 60 dB with f_{CLK} set to 40 MHz. A yield-based reliability-driven design methodology is imperative for deep-subthreshold operation, with the yield of the FDC embedding the PS disturbed by local mismatch-induced effects, defined as:

$$\text{Yield} = [1 - P(SNR_Q < SNR_{\text{Threshold}})] \cdot 100\% \quad (1)$$

where $P(SNR_Q < SNR_{\text{Threshold}})$ is the probability that the FDC digitizes the FM signal with an SNR_Q below the defined pass/fail threshold (taking a lower bound of 58 dB here).

The yield contour plots of Fig. 11 sweep $V_{DD, FDC}$ and $\Delta V_{\phi, \max} / V_{DD, FDC}$. For a $\Delta V_{\phi, \max} / V_{DD, FDC}$ ratio of 1, the (#1) DFF PS outperforms both the (#2) SAFF and (#3) Pre-amp+SAFF configurations (at $V_{DD, FDC}$ of 210 mV, its yield is 100% instead of 96% for the SAFF-based configuration). While the equivalent t_{latch} regeneration time for the DFF (its CLK-to-Q delay) is also around 6 ns, the SAFF must additionally incur the propagation delay through its back-end SR-latch, which may result in the setup-time violation of the following DFF in the FDC processing chain.




Fig. 11. Simulated FDC yield (%) contours versus $V_{DD,FDC}$ and $\Delta V_{\phi,max}/V_{DD,FDC}$ for the PS configurations: (#1) DFF PS, (#2) SAFF PS and (#3) Pre-amp+SAFF PS.

The advantages of the SA-based PS become evidently obvious for lower voltage-swing inputs. Even for the high $V_{DD,FDC}$ of 210mV, as $\Delta V_{\phi,max}/V_{DD,FDC}$ goes below 0.8, the yield quickly deteriorates from 95% to below 20% for the DFF PS. On the contrary, the (#2) SAFF experiences a yield drop from around 95% to 85% and the superior (#3) Pre-amp+SAFF configuration sees its yield drop to only 90%. Note also that in these simulations, the common-mode (CM) voltage of ϕ_p, ϕ_n is set to $V_{DD,FDC}/2$. Due to the single-ended nature of DFF, any slight shift in this CM voltage would prove further detrimental for small $\Delta V_{\phi,max}$ conditions.

In the context of the VCO-based ADC integration, these simulations imply that with a heavily modulated VCO output swing [26], or when the VCO core circuitry operates from a lower supply voltage relative to the FDC domain, the introduction of the sense-amplification phase-sampling mechanisms aid to relax the requirements on any explicit level-shifters (LS) and/or VCO buffers that may be necessary (facing issues with static power consumption, loading on the VCO outputs, phase mismatches and harmonic distortion induced by group delay dispersion [1]) in order to functionally interface with the FDC domain.

E. Second-Order Effects and PVT Variations

The following dynamic effects of the SAFF PS, with/without the presence of the pre-amp sampling latch, are further presented:

- 1) *VCO output edge transition time*: Reducing the VCO edge transition time (tt_{VCO} , quantified between $0.1V_{DD}$ to $0.9V_{DD}$) relative to its period (T_{VCO}), narrows the width of the SAFF metastability window. Consequently,




Fig. 12. Simulated (#2 and #3) PS yield versus (a) VCO output transition time (tt_{VCO}) relative to its period (T_{VCO}), (b) pre-amp stage clock delay (t_{pre}) and (c) clock transition time (tt_{CLK}).

the yield can be improved from 50% to above 65% when tt_{VCO}/T_{VCO} reduces from 0.2 to 0.1 [see Fig. 12(a)]. The addition of the pre-amp sampling latch not only improves the equivalent PS yield to above 80%, it also becomes relatively independent of tt_{VCO} .

- 2) *Pre-amp sampling stage delay*: The negative edge of CLK_i must be asserted in advance of the positive edge of CLK_{ii} to ensure $L, \bar{L}$ are sampled before the front-end pre-amp sampling latch resets, otherwise a hold violation can occur. If t_{pre} is below 1 ns, the pre-amp sampling latch actually worsens the yield of the overall PS. The optimal delay is 3 ns [see Fig. 12(b)]. When t_{pre} is longer than 5 ns, the yield reduces by 10% since less time is dedicated for the regeneration phase of the pre-amp sampling stage. A simple inverter element is inserted to accomplish the delay of t_{pre} , obviating the need for complicated non-overlapping clock generation. Although, such sensitivity of t_{pre} on the cascaded PS structure means it would be more desirable to insert a coarsely programmable delay generator.
- 3) *Clock edge transition time*: This transition time (tt_{CLK} , quantified between $0.1V_{DD}$ to $0.9V_{DD}$) affects mainly transistor M_8 of the sense amplifier. Ideally, at the negative edge of the clock, this pMOS device should switch to a low R_{on} state instantaneously, giving maximal voltage headroom to the latching transistors M_{4-7} for the entirety of the regeneration phase. Since the local clock drivers only realistically produce tt_{CLK} above 3 to 4 ns, M_8 takes longer to turn on, thus prolongs t_{latch} and, in turn, lowers the yield [see Fig. 12(c)]. The addition of the pre-amp stage thus relaxes the rise-fall times requirements of the high-speed clock signal.

Figure 13(a) plots the (#3) PS simulated yield (i.e., including mismatch effects) versus $V_{DD,FDC}$ at the temperatures of 0°C, 20°C and 40°C. An upward temperature shift of 20°C (e.g., 0 to 20°C) moves the minimum functional $V_{DD,FDC}$ down by around 20–25 mV. In other words, there is almost a 1 mV-to-1°C correspondence between supply and temperature (VT) variations at 0.2 V. These PS yield simulations are repeated for the slow and fast corners (SS and FF, respectively) in Fig. 13(b) to demonstrate the extreme effects of global process corners. An SS process requires a minimum $V_{DD,FDC}$ of 300 mV, equivalent to a 10× reduction in speed. On the other hand, an FDC operating in the FF process corner can tolerate $V_{DD,FDC}$ as low as 150 mV.




Fig. 13. Simulated (#3) Pre-amp + SAFF PS yield versus (a) temperature variations and (b) process corners.

IV. SYSTEM TECHNIQUES AND CONSIDERATIONS

A. Standby-Hardware Redundancy

Local variations may cause a fault in the operation of the phase sampling mechanism and, therefore, cannot guarantee sufficient yield (i.e., in the order of parts per million). One solution could be to up-size every device, as mismatch is inversely proportional to the square root of its gate area, at the cost of large loading capacitance. Alternatively, digital calibration can be employed to tune the switching threshold (i.e., imbalances between pMOS and nMOS devices in both differential paths), but the source of imbalances, random in nature, can arise from anywhere in the SR-latch, the sense-amplifier within the SAFF or the pre-amp stage. The hardware overhead, which would then be required to enable the re-configurability of nearly every MOS device in the PS makes such a solution highly impractical. Instead of relying on the absolute robustness of a single PS, consider an array of N_{red} redundant PSs. Now, by selecting one of N_{red} phase samplers that is verified as functional, the joint probability that all of the statistically uncorrelated phase samplers fail is the product of their individual failure rate. The *composite* yield therefore becomes:

$$\text{Yield}_{\text{red}} = [1 - P(\text{SNR}_Q < \text{SNR}_{\text{Threshold}})]^{N_{\text{red}}} \cdot 100\% \quad (2)$$

Such “standby hardware redundancy” techniques are hallmarks of fault-tolerant system design, and have been exploited to improve mismatch resilience in the implementation of flash ADCs [42] and SRAM memories [43].

The improvement in PS yield is verified as illustrated in Fig. 14(a). Note that there are actually two benefits of employing the hardware redundancy. At a relatively high supply voltage of 210 mV, the yield of a single phase sampler is 94%, while the composite yield reaches 99.7%, 99.98% and 99.999% respectively for N_{red} of 2, 3, and 4. Moreover, the large random variability associated with deep-subthreshold operation means that, in some cases, the phase sampler could be in reality, much *faster* than its mean performance. Exploiting this phenomenon essentially reduces the minimum supply voltage, as the yield curve shifts to the left of the $V_{DD,FDC}$ axis. In this work, we implement the 2-choose-1 (2C1) configuration. This choice is governed by three main factors. First, leakage energy consumption grows linearly with the amount of hardware redundancy (from 4% for 2C1 to 8% for 4C1, see Section IV.B). Although, this additional leakage would still be rather small, with further circuit-level leakage-minimization optimizations possible. Second, yield maximization in conjunction with V_{DD} minimization is not aggressively pursued. For example, at $V_{DD,FDC}$ of 190 mV, the original yield is improved from 80% to 96% (2C1) and $> 99\%$ (4C1). An industry-oriented product would certainly





Fig. 14. Phase sampling: (a) hardware redundancy leading to improvement in the yield of functioning FDC units, (b) two-choose-one (2C1) hardware redundant PS with clock and data path gating.

demand significantly tighter margins on yield, hence 4C1 (and beyond) might be preferred. Third, the fault tolerance mechanism of the asynchronous PS interface is not applicable in the downstream flip-flops, XOR gates, decimation and digital recombination blocks (designed instead to meet the timing margins of the synchronous clock domain). Without the protection of hardware redundancy, further reduction in V_{DD} may lead these blocks to fail the timing slack requirements and become the new speed bottleneck. This could negate any yield/speed improvements reaped with a high N_{red} -choose-1 PS configuration.

The implementation of the 2C1 hardware-redundant PS (i.e., $N_{\text{red}} = 2$) is shown in Fig. 14(b). Multiplexers in the clock and output data paths incur minimal digital hardware overhead. For the PS that is not selected (i.e., in standby), its clock signal is replaced with a constant voltage bias of V_{STBY} . With reference to Fig. 5(b), the clock device M_8 morphs into a sleep transistor to minimize the leakage current through the ‘standby’ PS. Furthermore, AND gates are placed at the output, so that in the very rare scenario where neither redundant phase samplers work, its output is disabled by the static control signal EN to not affect the multi-bit digital output of the entire FDC array, composed of several FDC units.

B. FDC Power-Leakage Trade-Off

Figure 15 shows the power consumption breakdown of the single FDC unit at $V_{DD,FDC}$ of 210 mV, consisting of the clock buffering (CLK), standard XOR-based FDC cell (2 DFFs and 1 XOR) and the phase sampler [standby leakage, pre-amp, sense-amplifier (SA) and SR-latch (SR)]. As shown in Fig. 15(a), the total power consumption is 490 nW when f_{CLK} is 40 MHz and V_{STBY} (the bias voltage applied to the M_8 clock




Fig. 15. Simulated power consumption of the FDC unit for f_{CLK} of 40 MHz, with V_{STBY} at (a) 210 mV and (b) 420 mV, and for f_{CLK} of 10 MHz at the corresponding V_{STBY} values, in (c) and (d).

devices of the modified SA circuit when it is in standby) is 210 mV. The phase sampler (158 nW) allocates 58 nW each for the pre-amp and the SAFF's sense-amplifier, along with 42 nW for its SR-latch. The local clock buffering for the FDC unit (partitioned roughly equal between driving the PS circuits and the 2DFF+1XOR FDC block) consumes a further 143 nW, with the downstream standard XOR-based FDC cell (2 DFFs and 1 XOR gate) consuming 150 nW.

The standby leakage of the redundant PS within the 2C1 arrangement dissipates 40 nW (10 nW each for the pre-amp and SA, 20 nW for the SR-latch). Despite the M_8 sleep transistor cutting off the power supplied to the standby-redundant PS blocks, the leakage power (10 nW per SA) is still significant. To mitigate the issue of leakage for standby circuits, an on-chip switched-capacitor voltage doubler [13] for *high-impedance loads* is implemented in our ADC prototype to straightforwardly boost the V_{STBY} bias to 420 mV, with the consequent power breakdown shown in Fig. 15(b). The leakage current of both SAs combine to less than 1 nA (i.e., virtually non-existent). Since this technique was not applied to the SR-latches, its leakage power remains at 20 nW.

It is interesting to show the effect of leakage currents at f_{CLK} of 10 MHz with f_{VCO} scaled accordingly [see Fig. 15(c),(d)]. The standby leakage component is static and thus remains unchanged, dissipating a bigger portion of the total power budget. Leakage is also concerning for the actively switching components of the FDC unit. The total power reduces from 490 nW to 275 nW as f_{CLK} is scaled from 40 MHz to 10 MHz. This *sub-linear* power-frequency scaling indicates that the deep-subthreshold energy efficiency is much degraded for slower speeds of operation.

Therefore, it is recommended for ULV standby circuits in a large-scale system to employ sleep transistors reverse-biased (for a pMOS device) through a shared, simple switched-capacitor voltage doubler, while actively switching circuits must operate at their fastest possible speeds to remain maximally energy-efficient. This insight implies that systems

which promote lower parallelism coupled with high-speed processing slices (e.g., the XOR-based FDC), are preferred over low-clock-rate multi-bit processing (typical of counter-based/coarse-fine architectures) for efficient deep-subthreshold operation.

C. In-Situ FDC Performance Monitoring

To determine the selection of one functional PS within an array of N_{red} redundant PSs (and moreover evaluate the functionality of the overall FDC structures against severe PVT variations), we must first provide a method to characterize the individual FDC units. Figure 16(a) shows the SNR_Q (BW = 160 kHz) of a *single FDC unit* for a frequency-modulated VCO waveform when a full-scale (FS) voltage sinusoidal input is applied, across 50 Monte-Carlo runs. Of course, observing the achieved SNR_Q makes it trivial to determine whether the FDC functions or not. However, this computationally intensive fast Fourier transform (FFT)-based testing method is only feasible in a laboratory environment. It becomes apparent that only very simple methods should be used to facilitate a built-in self-test (BIST) of a microwatt-level ULV design.

Due to the asynchronous nature of the phase sampling process, it is not strictly necessary to provide a dynamic input stimulus (e.g., a sinusoid). A VCO with constant frequency (f_{VCO}), *asynchronous* to f_{CLK} , generates an intrinsic periodic phase-ramp at the input of PS, which exercises the full 0-to- 2π phase of the clock period. This greatly simplifies the test stimulus, as the ADC input can now be tied to either the supply or ground power rails. Another factor which makes it difficult to diagnose faults in the single XOR-based FDC unit is that it outputs an oversampled $\Delta\Sigma$ bit-stream which toggles between just two levels (the high and low levels being V_{DD} and ground, respectively) with a very high toggle density. In other words, the switching activity induced by high-frequency noise-shaped quantization components may make the FDC appear functional but in fact, its faulty operation has corrupted the low-frequency signal of interest. To alleviate this issue, we pass the output of the individual FDC unit through a 2nd-order, decimate-by-4 digital filter to obtain a better time-domain behavior of the FDC output. The implementation of the high-speed decimation filtering stages is discussed in Section V.

With both transient waveforms of the sinusoid and DC input cases at respectively Run 1 and 30 for the *decimated* digital output ($D_{out,d}$), see Fig. 16(c) and Fig. 16(d), we can now observe that FDC is most prone to faulty frequency digitization when f_{VCO} is close to its full-scale value of $f_{CLK}/2$, where metastability events cause catastrophic glitches in $D_{out,d}$. This glitching is quantified with $\Delta D_{out,d}$, the jump in consecutive digital output codes, otherwise known as the derivative of $D_{out,d}$. For Run 1, the PS/FDC slice consistently fails to sample correctly the VCO phase information along its processing chain, leading to a largely unusable and corrupt digital output. For Run 30, the glitches, less frequent in time, occur when the PS/FDC samples metastable states which result in the occurrence of a bitstream string of three or more consecutive 1's or 0's, leading to the incorrect overflow of the XOR-based FDC structure.

Consequently, a functioning PS results in the decimated digital output [see Fig. 16(e)] to toggle by 1 only




Fig. 16. 50 Monte-Carlo runs of a single phase-sampling XOR-based FDC unit processing (a) SNR_Q (BW = 160 kHz) derived from the frequency modulated VCO when a full-scale (FS) voltage sinusoidal input is applied, (b) $\max(\Delta D_{out,d})$ at the maximum VCO frequency ($f_{VCO,max}$) when a DC input voltage is applied to the VCO. The corresponding transient waveforms are shown in (c), (d) and (e) for Runs 1, 30 and 41 respectively.

[i.e., exhibiting a maximum derivative $\max(\Delta D_{out,d})$ of only 1]. Note that this assumes the VCO (i.e., FDC input) is noiseless, but in the realistic case of the VCO exhibiting phase noise (PN), being equivalent to a noisy DC input source, $\max(\Delta D_{out,d})$ is shown to be less than 4 for our design. This sets the threshold by which the FDC under test is designated as functional ($\leq \text{Threshold}=4$) as opposed to faulty ($> \text{Threshold}=4$). Figure 16(b) demonstrates this direct *one-to-one* correspondence between quantifying $\max(\Delta D_{out,d})$ for a DC input, which places f_{VCO} near $f_{CLK}/2$, compared to the SNR_Q measurement for a full-scale input sinewave visualized in Fig. 16(a). The proposed FDC testing protocol is visualized with the flowchart in Fig. 17. For each available FDC unit, its internal hardware redundant PSs ('A' and 'B' in this case) are individually tested with $\max(\Delta D_{out,d})$ computed. Only the first instance of a functional PS needs to be detected to determine the PS selection bits (SEL_A or SEL_B). The calibration procedure ends after the iteration through all FDC units under test.

V. DIGITAL DECIMATION AND FILTERING

The decimation filter, by virtue of being in the digital domain, is often neglected in the literature and in most prototypes not considered for on-chip implementation. Here, the need for a large OSR dictates the first stages of the




Fig. 17. Flowchart for the proposed in-situ testing/monitoring of faulty hardware redundant phase-sampling FDCs.

digital filter to operate as fast as the speed-maximized phase sampling interface. Moreover, utilizing the multi-phase outputs of the ring-VCO results in a significant portion of the ADC's hardware dedicated to the output summation logic to perform digital recombination of the parallel FDC $\Delta\Sigma$ streams [24]. Aside from the necessary additional power consumption, the combination of high-speed operation and large mismatch-induced delay variability may compromise the functionality of these *seemingly trivial* digital blocks in deep-subthreshold.

A. High-Speed Design Challenges

We first describe the timing characterizations for D-flip-flop (DFF) and full-adder (FA) circuits (implemented with extremely-low V_t "elvt" devices in 28 nm LP CMOS), being the backbone of all digital processing hardware, to demonstrate the high-speed design difficulties. Pertaining to the clocked storage elements, the non-ratioed master-slave DFF topology has been proven to provide relatively robust functionality in subthreshold [4]. Suppose at 0.2 V, we may operate the FDCs at a realistic clock rate of 40 MS/s (i.e., a t_{CLK} of 25 ns). We estimate the setup time to be around 4 ns for a CLK-to-Q delay t_{CQ} of 5 ns. In other words, the 'delay' of the flip-flop within a pipeline stage accounts for nearly 40% of the timing budget. With local mismatch-induced delay variations, a setup time of 4 ns translates statistically into t_{CQ} below 5 ns for only 20% of DFFs. To guarantee the reliability of every flip-flop in the design, 13 ns (more than 50% of the cycle timing budget) is sacrificed just to allow for pipelining, since this precious time and associated energy costs are not expended for useful computations.

Consequently, the remaining allocation of 12 ns is permitted for the inter-stage combinational propagation delay (t_{comb}) in the best-case scenario. Our evaluation of an FA cell sees a simulated compute time of around 11 ns (from input transition to carry output generation) and a standard deviation (σ_{comb}) of 1 ns for a worst-case delay approaching 15 ns. Therefore, a single pipeline stage can accommodate at most 1 or 2 FA stages. The registers required for storing the intermediate values along the pipeline stages end up consuming significant power and furthermore, takes up half of the timing budget. Latch-based super-pipelines [45], [46] mitigate this to an extent through time borrowing, but are affected by hold violations, necessitating strict coordination between all clock and data paths of the digital signal processing chain.

A simple circuit optimization technique may be to aggressively up-size the width of every transistor and use alternative




Fig. 18. Signal flow diagram for the conventional FDC and DEC configuration requiring high speed summation logic and multi-bit decimation filtering.




Fig. 19. Signal flow diagram for FDC and DEC blocks in this work, with bit-stream polyphase decimation and the summation logic shifted to the downsampled rate ($M = 4$).

flip-flop variants [44]. This improves the timing situation, but the performance gains remain rather incremental as the decreased gate delays cannot sufficiently cover slower PVT conditions. The bulky transistors are hampered by their increased capacitive loading and leakage currents, leading to a diminishing throughput return, all the while incurring a decrease in both static and dynamic energy efficiency.

B. $\Delta\Sigma$ Bit-Stream Processing

To minimize the dependence on optimization strategies such as circuit sizing as a means to overcome the encountered timing bottleneck, we must first revisit, at the architecture level, how the digital summation and decimation processes are carried out. In the conventional processing chain of Fig. 18, the parallel outputs of the FDCs are digitally recombined at the full-rate clock (CLK) with the output summation logic (e.g., implemented as a Wallace adder tree). The multi-bit output is then processed by the classical cascaded-integrator-comb (CIC) Hogenauer filter [48] to provide 2nd-order low-pass filtering (necessary to filter out the 1st-order noise-shaped FDC output) before the alias-free downsampling by a factor of $M = 4$. This CIC topology moves the combing differentiators to the quarter-rate clock (CLK/4), but high-speed integrators, built from multi-bit digital adders, are still required. The computations this digital back-end entails make such digital processing impractical under the aforementioned timing constraints. Interestingly, it is neither necessary to perform the digital summation nor filtering operations at the full-rate clock. We suggest two modifications to the traditional back-end implementation in Fig. 19 in order to restructure the digital signal processing. Using the Noble identities [49], the order of downsampling and anti-aliasing filtering can be commuted, so as to move the filter coefficient multiplication and summation operations to the low-rate clock domain. Furthermore, applying linear superposition, it is possible to perform digital recombination after the downsampling, rather than immediately following the FDC digitization.

The polyphase composition, through the expansion of the CIC's filter response, reveals in its non-recursive FIR form a triangular-window $\text{sinc}^2()$ filter impulse response. The implementation of the polyphase 2nd-order, decimate-by-4 filter in Fig. 19 thus allows us to process the individual $\Delta\Sigma$ bit-streams, such that *all* combinational adders are conveniently relocated to the timing-relaxed quarter-rate clock domain. It is important to note that the modular nature of the proposed parallel decimation filter paths naturally allows the built-in testing protocol and characterization of the individual FDC units discussed earlier in Section IV. This limits the digital output glitching behavior, quantified by $\Delta D_{\text{out},d}$, to be caused by the single PS unit under test, and not the combined faults of the entire FDC array within the VCO-based ADC.

C. Back-End Implementation

The circuit implementation of the decimation filtering back-end for our pseudo-differential VCO-based ADC prototype is shown in Fig. 20(a). Four FDC $\Delta\Sigma$ streams ($N_{\text{FDC}} = 4$) in both positive and negative complementary halves (a total of 8 FDC units outputting $D_{\text{out},A\pm}$ to $D_{\text{out},D\pm}$) are individually fed to their own polyphase decimation filters to produce the respective 4-bit outputs $D_{\text{out},dA\pm}$ to $D_{\text{out},dD\pm}$, organized as in Fig. 20(b). Only the shift registers are clocked at CLK1 (40 MHz). The flip-flops operating at CLK4 (10 MHz) perform the downsample-by-4 operation. As the input values are bit-streams, the filter coefficient multiplication simply uses logical bit shift operations (thus, costless and delay-free). Three pipeline stages, for a total processing time of approximately 300 ns, are used. The first pipeline stage adds the filter taps with 2-bit to 4-bit carry-ripple adders [see Fig. 20(c)], with an arithmetic depth of 9 FA cells. The subsequent pipeline stages perform the digital recombination of the decimated FDC outputs within the single-ended ADC halves (producing $D_{\text{tot},d\pm}$) and the 2's complement subtraction between the pseudo-differential ADC halves (an arithmetic depth of 9 and 8 FA cells, respectively) to obtain the final ADC output ($D_{\text{tot},d}$). This partitioning of digital processing at CLK4 rate results in the critical path delay of each pipeline stage to be fairly similar. The respective circuit implementation of the constituent DFF and FA blocks are shown in Fig. 20(d) and (e). The 4 $\times$ longer clock period (100 ns) allows for smaller transistor sizes. The large arithmetic depth per pipeline stage averages out the accumulating propagation delay mismatches of a long cascade of FAs [47], resulting in a more robust operation. Pipelining is now straightforward, as it contributes to only around 10% of the available timing cycle budget. It is important to note that the data movement interface between the multi-rate clock domains (i.e., the re-sampling of data clocked at CLK1 by CLK4 within the polyphase decimation filter) represents a potential point-of-failure. It is imperative that CLK4 leads CLK1 to prevent any hold violations.

Such fully parallel signal composition in both space and time (for recombination and decimation, respectively) allows for the greatest achievable throughput as the operational speed is limited by only the delay of a single DFF. Of course, the complete removal of any computational speed bottleneck does come at a cost to power and area consumption, which must be carefully balanced with increased parallelism beyond




Fig. 20. Circuit implementation of (a) the multi-rate, multi-phase decimation filtering digital back-end, (b) the polyphase decimation filter, (c) the quarter-rate combinational logic within the digital filter. The backbone of all digital processing hardware is constructed with (d) the D flip-flop (DFF) and (e) the full-adder (FA) cell.

$N_{FDC}=4$ implemented in this work. The cost of decimation increases linearly with N_{FDC} . Moreover, as decimation increases the word length (from 1-bit to 4-bit), the benefit of the quarter-rate processing may be outweighed by more multi-bit computations required in general for the digital adders. Another consideration is that while this decimation filtering hardware demonstrates a down-sampling factor of 4, further decimation stages down to the Nyquist data-rate are still required since the OSR of our VCO-based ADC design is >100 . These decimation stages face much more benign timing considerations (e.g., an additional $4\times$ decimator stage would operate with a clock rate of $f_{CLK}/16$), thus omitted from the prototype implementation.

VI. EXPERIMENTAL RESULTS

Figure 21 shows the chip micrograph and core structure of FDC and DEC digital back-end blocks integrated within the 0.2-V VCO-based ADC prototype, whose front-end was described in [1]. The IC is fabricated in TSMC 28-nm LP CMOS and the ADC occupies an active area of 0.12mm^2 , of which the FDC and DEC array take up 0.03mm^2 . The serial-to-parallel interface (SPI) and design-for-test (DFT) digital read-out enable the outputs of FDC units and DEC processing slices to be multiplexed out for individual characterization.

The spectrum at the output of the individual FDC+DEC processing slice (e.g., $D_{out,dA+}$), sampled at 45MS/s (digital data stream at the decimated 11.25MS/s rate) with the VCO core modulated by a 20kHz , $0.2V_{pp}$ single-ended input sinewave, is shown in Fig. 22(a). For this intermediate single FDC unit, the SNR for ‘A’ (blue) and ‘B’ (red) PS configurations are 40.7dB and 52.9dB , respectively. Clearly, PS ‘A’ malfunctions, destroying the noise-shaping property of phase quantization errors, thus leading to an elevated noise floor. PS ‘B’ operates as intended, so the dominant in-band noise contribution originates from the VCO phase noise. This measured 52.9dB value is therefore less than the ideal $\text{SNR}_Q > 60\text{dB}$ with only the quantization noise of the FDC unit taken into account. The distinction in functionality




Fig. 21. Implemented VCO-based ADC: (a) chip micrograph embedding the FDC and DEC blocks, and (b) the signal flow of the FDC and DEC digital back-end and the measured SNR along the processing chain.




Fig. 22. (a) Measured spectrum at node $D_{out,dA+}$ (i.e., single FDC+DEC slice), with a $0.2V_{pp}$ single-ended 20kHz sinusoidal input. The measured SNR (SNR_{Q+PN}) includes both FDC Q-noise and VCO phase-noise. (b) Decimated digital data stream at $D_{out,dA+}$ with a DC voltage input, in the ‘A’ and ‘B’ PS configurations of the single FDC unit.

between both phase samplers can be easily demonstrated with the digitization of a DC input as shown in Fig. 22(b).




Fig. 23. Measured SNR for all 8 individual FDC units (FDC₁ to FDC₈) incorporating 2C1 standby redundancy ('A' and 'B' phase samplers) versus $V_{DD,FDC}$ for f_{CLK} equal to (a) 30 MHz, (b) 40 MHz and (c) 50 MHz. The SNR threshold is set to 50 dB, where PSs with SNR < 50 dB are classified as 'faulty', while PSs with SNR > 50 dB are classified as 'functional'.

The decimated digital output $D_{out,d}$ toggling between codes 10 and 11 places f_{VCO} from 14 MHz to 15.5 MHz. However, the FDC output with the PS 'A' selected experiences large glitches in the output code. For the digitization of a DC input, this may only occur if the FDC's operation is faulty.

Figure 23 extends the SNR characterization to all available FDC units (8 in total), labeled from FDC₁ to FDC₈, across FDC supply voltage $V_{DD,FDC}$ at sampling rates of 30 MHz, 40 MHz and 50 MHz, summarizing the results of a total of 192 unique data measurement spectra. The SNR of the individual FDC units is computed by multiplexing the internal 4-bit signals $D_{out,dA\pm}$ to $D_{out,dD\pm}$ off-chip one-by-one through the DFT circuitry. In order for the ADC to function correctly, the entire FDC array (4 FDC units per positive and negative complementary halves) must not fail; in this case, the complete SNR measured at node $D_{tot,d}$ reaches 62 dB [1]. By introducing the 2C1 hardware-redundant PS, only one of two PSs within a FDC unit needs to work. Take for example Fig. 23(a), PS 'A' of FDC₆ does not function at $V_{DD,FDC}$ of 205 mV. By switching to PS 'B', the introduced hardware redundancy ensures the entire FDC array is not compromised in spite of internal faulty components. Similarly, PS 'B' is stronger for FDC₄ at 40 MHz [see Fig. 23(b)] while PS 'A' is stronger for FDC₇ at 50 MHz [see Fig. 23(c)].

The Shmoo plot of Fig. 24(a) maps out the minimum $V_{DD,FDC}$ of the entire FDC array (i.e., where all 8 FDC+DEC processing slices function properly) at different f_{CLK} . At 30, 40 and 50 MHz, $V_{DD,FDC}$ must be respectively above 200, 215 and 235 mV. The total power consumed by the FDC array




Fig. 24. Measured XOR-based FDC array (total of 8 FDC units): (a) pass/fail Shmoo plot, and (b) power consumption versus f_{CLK} and $V_{DD,FDC}$. Measured decimation-filtering digital back-end: (c) pass/fail Shmoo plot, and (d) power consumption versus f_{CLK} and $V_{DD,DEC}$.

is visualized with the contour plot of Fig. 24(b). At $V_{DD,FDC}$ and f_{CLK} of 215 mV and 40 MHz, the FDCs consume $4.4 \mu W$ (i.e., 550 nW per FDC unit).

The functionality and power consumption of the decimation-filtering digital back-end are characterized respectively in Fig. 24(c) and (d). This digital back-end consists of 8 polyphase decimator slices, the quarter-rate digital recombination logic, an 8-bit subtractor and local clock drivers. At 30, 40 and 50 MHz, $V_{DD,DEC}$, the supply voltage of this power domain, must be respectively above 205, 225 and 235 mV to meet the necessary timing requirements. At $V_{DD,DEC}$ and f_{CLK} of 225 mV and 40 MHz, the total power consumed by the digital back-end is 8.6 μ W.

To the best of our knowledge, this work represents the first full demonstration of a deep-subthreshold 0.2 V multi-phase FDC architecture running at $\sim$ 40 MHz in advanced CMOS, such as a 28-nm node. Additionally, it is worth mentioning that further experimental verification of the entire 0.2 V open-loop VCO-based ADC, such as SNR performance across several ICs, dynamic measurements versus input amplitude and input frequency (i.e., including the VCO analog front-end core) and its comparison to state-of-the-art VCO-based ADC implementations are available in [1].

VII. CONCLUSION

Energy harvesting IoT network solutions call for the development of ultra-low-voltage (ULV) circuits operating deep into the subthreshold regime of the CMOS devices. The expositions in this article explore the speed maximization of the multi-phase frequency-to-digital converter (FDC) architecture integrated within an open-loop VCO-based ADC prototype, operating with the supply voltage approaching 0.2 V at clock frequencies between 30 MHz and 50 MHz. Such a feat is made possible by the high-speed sense-amplify asynchronous phase-sampling interface, capable of sampling small-voltage-swing signals. Hardware redundancy is incorporated to mitigate large circuit variability associated with deep-subthreshold operation for significantly improved fault tolerance and yield. The FDC array consisting of several parallel XOR-based FDC units are decimated and combined with polyphase decimation filters and quarter-rate digital logic. Consequent speed bottlenecks due to the operation with timing margins of merely 10's of nanoseconds imposed at the clock-sampling limit are overcome. Experimental characterizations verify that at 40 MS/s, the FDCs and decimation-filtering back-end are functional at the minimum supply voltage of 215 mV and 225 mV, consuming 4.4 μ W and 8.6 μ W respectively.

ACKNOWLEDGMENT

The authors would like to thank the TSMC University Shuttle for the chip fabrication; Dr. Hsieh-Hung Hsieh (TSMC) and Dr. Teerachot Siriburanon for help with the tapeout; the Microelectronic Circuits Centre Ireland (MCCI) for technical and administrative support; Hieu Minh Nguyen and Suoping Hu for technical discussions, Erik Staszewski and Sunisa Staszewski for measurement and lab assistance. They would also like to acknowledge the anonymous reviewers for their invaluable comments and suggestions.

REFERENCES

[1] V. Nguyen, F. Schembari, and R. B. Staszewski, "A deep-subthreshold variation-aware 0.2-V open-loop VCO-based ADC," *IEEE J. Solid-State Circuits*, vol. 57, no. 6, pp. 1684–1699, Jun. 2022.

[2] A. P. Chandrakasan and R. W. Brodersen, "Minimizing power consumption in digital CMOS circuits," *Proc. IEEE*, vol. 83, no. 4, pp. 498–523, Apr. 1995.

[3] M. Alioto, "Ultra-low power VLSI circuit design demystified and explained: A tutorial," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 59, no. 1, pp. 3–29, Jan. 2012.

[4] A. Wang and A. Chandrakasan, "A 180-mV subthreshold FFT processor using a minimum energy design methodology," *IEEE J. Solid-State Circuits*, vol. 40, no. 1, pp. 310–319, Jan. 2005.

[5] B. H. Calhoun, A. Wang, and A. Chandrakasan, "Modeling and sizing for minimum energy operation in subthreshold circuits," *IEEE J. Solid-State Circuits*, vol. 40, no. 9, pp. 1778–1786, Sep. 2005.

[6] S. Hanson et al., "Exploring variability and performance in a sub-200-mV processor," *IEEE J. Solid-State Circuits*, vol. 43, no. 4, pp. 881–891, Apr. 2008.

[7] B. H. Calhoun and A. P. Chandrakasan, "A 256-kb 65-nm sub-threshold SRAM design for ultra-low-voltage operation," *IEEE J. Solid-State Circuits*, vol. 42, no. 3, pp. 680–688, Mar. 2007.

[8] T.-H. Kim, J. Liu, J. Keane, and C. H. Kim, "A 0.2 V, 480 kb subthreshold SRAM with 1k cells per bitline for ultra-low-voltage computing," *IEEE J. Solid-State Circuits*, vol. 43, no. 2, pp. 518–529, Feb. 2008.

[9] S. Yang, J. Yin, H. Yi, W.-H. Yu, P.-I. Mak, and R. P. Martins, "A 0.2-V energy-harvesting BLE transmitter with a micropower manager achieving 25% system efficiency at 0-dBm output and 5.2-nW sleep power in 28-nm CMOS," *IEEE J. Solid-State Circuits*, vol. 54, no. 5, pp. 1351–1362, May 2019.

[10] H. Yi et al., "A 0.18-V 382- μ W Bluetooth low-energy receiver front-end with 1.33-nW sleep power for energy-harvesting applications in 28-nm CMOS," *IEEE J. Solid-State Circuits*, vol. 53, no. 6, pp. 1618–1627, Jun. 2018.

[11] M. H. Ghaed et al., "Circuits for a cubic-millimeter energy-autonomous wireless intraocular pressure monitor," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 60, no. 12, pp. 3152–3162, Dec. 2013.

[12] C.-C. Li et al., "All-digital PLL for Bluetooth low energy using 32.768-kHz reference clock and $\leq$ 0.45-V supply," *IEEE J. Solid-State Circuits*, vol. 53, no. 12, pp. 3660–3671, Dec. 2018.

[13] C.-C. Li, M.-S. Yuan, Y.-T. Lin, C.-C. Liao, C.-H. Chang, and R. B. Staszewski, "A 0.2-V three-winding transformer-based DCO in 16-nm FinFET CMOS," *IEEE Trans. Circuits Syst. II, Exp. Briefs*, vol. 67, no. 12, pp. 2878–2882, Dec. 2020.

[14] A. Esmailiyan, F. Schembari, and R. B. Staszewski, "A 0.36-V 5-MS/s time-mode flash ADC with dickson-charge-pump-based comparators in 28-nm CMOS," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 67, no. 6, pp. 1789–1802, Jun. 2020.

[15] S. Hu et al., "A type-II phase-tracking receiver," *IEEE J. Solid-State Circuits*, vol. 56, no. 2, pp. 427–439, Feb. 2021.

[16] Y. Lin et al., "A 180 mV 81.2%-efficient switched-capacitor voltage doubler for IoT using self-biasing deep N-well in 16-nm CMOS FinFET," *IEEE Solid-State Circuits Lett.*, vol. 1, no. 7, pp. 158–161, Jul. 2018.

[17] D. Bol, R. Ambroise, D. Flandre, and J. D. Legat, "Interests and limitations of technology scaling for subthreshold logic," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 17, no. 10, pp. 1508–1519, Oct. 2009.

[18] G. G. E. Gielen, L. Hernandez, and P. Rombouts, "Time-encoding analog-to-digital converters: Bridging the analog gap to advanced digital CMOS—Part 1: Basic principles," *IEEE Solid-State Circuits Mag.*, vol. 12, no. 2, pp. 47–55, Spring 2020.

[19] M. Hovén, A. Olsen, T. S. Lande, and C. Toumazou, "Delta-sigma modulators using frequency-modulated intermediate values," *IEEE J. Solid-State Circuits*, vol. 32, no. 1, pp. 13–22, Jan. 1997.

[20] M. Z. Straayer and M. H. Perrott, "A 12-bit, 10-MHz bandwidth, continuous-time $\Sigma\Delta$ ADC with a 5-bit, 950-MS/s VCO-based quantizer," *IEEE J. Solid-State Circuits*, vol. 43, no. 4, pp. 805–814, Apr. 2008.

[21] J. Kim, T. K. Jang, Y. G. Yoon, and S. Cho, "Analysis and design of voltage-controlled oscillator based analog-to-digital converter," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 57, no. 1, pp. 18–30, Jan. 2010.

[22] A. Babaie-Fishani and P. Rombouts, "A mostly digital VCO-based CT-SDM with third-order noise shaping," *IEEE J. Solid-State Circuits*, vol. 52, no. 8, pp. 2141–2153, Aug. 2017.

[23] J. Borgmans, R. Riem, and P. Rombouts, "The analog behavior of pseudo digital ring oscillators used in VCO ADCs," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 68, no. 7, pp. 2827–2840, Jul. 2021.

- [24] J. Borgmans, E. Sacco, P. Rombouts, and G. Gielen, "Methodology for readout and ring oscillator optimization toward energy-efficient VCO-based ADCs," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 69, no. 3, pp. 985–998, Mar. 2022.
- [25] G. Taylor and I. Galton, "A mostly-digital variable-rate continuous-time delta-sigma modulator ADC," *IEEE J. Solid-State Circuits*, vol. 45, no. 12, pp. 2634–2646, Dec. 2010.
- [26] V. Nguyen, F. Schembari, and R. B. Staszewski, "A 0.2-V 30-MS/s 11b-ENOB open-loop VCO-based ADC in 28-nm CMOS," *IEEE Solid-State Circuits Lett.*, vol. 1, no. 9, pp. 190–193, Sep. 2018.
- [27] H. Wang, V. Nguyen, F. Schembari, and R. B. Staszewski, "An adaptive-resolution quasi-level-crossing delta modulator with VCO-based residue quantizer," *IEEE Trans. Circuits Syst. II, Exp. Briefs*, vol. 67, no. 12, pp. 2828–2832, Dec. 2020.
- [28] R. B. Staszewski and P. T. Balsara, "Phase-domain all-digital phase-locked loop," *IEEE Trans. Circuits Syst. II, Exp. Briefs*, vol. 52, no. 3, pp. 159–163, Mar. 2005.
- [29] M. Z. Straayer and M. H. Perrott, "A multi-path gated ring oscillator TDC with first-order noise shaping," *IEEE J. Solid-State Circuits*, vol. 44, no. 4, pp. 1089–1098, Apr. 2009.
- [30] C. Weltin-Wu, G. Zhao, and I. Galton, "A 3.5 GHz digital fractional-PLL frequency synthesizer based on ring oscillator frequency-to-digital conversion," *IEEE J. Solid-State Circuits*, vol. 50, no. 128, pp. 2988–3002, Dec. 2015.
- [31] K. A. Bowman et al., "Energy-efficient and metastability-immune resilient circuits for dynamic variation tolerance," *IEEE J. Solid-State Circuits*, vol. 44, no. 1, pp. 49–63, Jan. 2009.
- [32] I. J. Chang, S. P. Park, and K. Roy, "Exploring asynchronous design techniques for process-tolerant and energy-efficient subthreshold operation," *IEEE J. Solid-State Circuits*, vol. 45, no. 2, pp. 401–410, Feb. 2010.
- [33] U. Wismar, D. Wisland, and P. Andreani, "A 0.2 V 0.44 μ W 20 kHz analog to digital $\Sigma\Delta$ modulator with 57 fJ/conversion FoM," in *Proc. 32nd Eur. Solid-State Circuits Conf.*, Montreux, Switzerland, Sep. 2006, pp. 187–190.
- [34] P. Prabha et al., "A highly digital VCO-based ADC architecture for current sensing applications," *IEEE J. Solid-State Circuits*, vol. 50, no. 8, pp. 1785–1795, Aug. 2015.
- [35] J. Huang and P. P. Mercier, "A 112-dB SFDR 89-dB SNDR VCO-based sensor front-end enabled by background-calibrated differential pulse code modulation," *IEEE J. Solid-State Circuits*, vol. 56, no. 4, pp. 1046–1057, Apr. 2021.
- [36] M. Baert and W. Dehaene, "A 5-GS/s 7.2-ENOB time-interleaved VCO-based ADC achieving 30.5 fJ/cs," *IEEE J. Solid-State Circuits*, vol. 55, no. 6, pp. 1577–1587, Jun. 2020.
- [37] V. Unnikrishnan and M. Vesterbacka, "Mitigation of sampling errors in VCO-based ADCs," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 64, no. 7, pp. 1730–1739, Jul. 2017.
- [38] A. Quintero et al., "A coarse-fine VCO-ADC for MEMS microphones with sampling synchronization by data scrambling," *IEEE Solid-State Circuits Lett.*, vol. 3, pp. 29–32, 2020.
- [39] C. Perez, A. Quintero, P. Amaral, A. Wiesbauer, and L. Hernandez, "A 73 dB-A audio VCO-ADC based on a maximum length sequence generator in 130 nm CMOS," *IEEE Trans. Circuits Syst. II, Exp. Briefs*, vol. 68, no. 10, pp. 3194–3198, Oct. 2021.
- [40] M. Hassanpourghadi, P. K. Sharma, and M. S.-W. Chen, "A 6-b, 800-MS/s, 3.62-mW Nyquist rate AC-coupled VCO-based ADC in 65-nm CMOS," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 64, no. 6, pp. 1354–1367, Jun. 2017.
- [41] J. M. Rabaey, A. Chandrakasan, and B. Nikolic, *Digital Integrated Circuits: A Design Perspective*. Englewood Cliffs, NJ, USA: Prentice-Hall, 2003.
- [42] M. P. Flynn, C. Donovan, and L. Sattler, "Digital calibration incorporating redundancy of flash ADCs," *IEEE Trans. Circuits Syst. II, Analog Digit. Signal Process.*, vol. 50, no. 5, pp. 205–213, May 2003.
- [43] N. Verma and A. P. Chandrakasan, "A 256 kb 65 nm 8T subthreshold SRAM employing sense-amplifier redundancy," *IEEE J. Solid-State Circuits*, vol. 43, no. 1, pp. 141–149, Jan. 2008.
- [44] M. Alioto, E. Consoli, and G. Palumbo, "Variations in nanometer CMOS flip-flops: Part I—Impact of process variations on timing," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 62, no. 8, pp. 2035–2043, Aug. 2015.
- [45] R. B. Staszewski, K. Muhammad, and P. Balsara, "A 550-MSample/s 8-tap FIR digital filter for magnetic recording read channels," *IEEE J. Solid-State Circuits*, vol. 35, no. 8, pp. 1205–1210, Aug. 2000.
- [46] D. Jeon, M. Seok, C. Chakrabarti, D. Blaauw, and D. Sylvester, "A super-pipelined energy efficient subthreshold 240 MS/s FFT core in 65 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 47, no. 1, pp. 23–34, Jan. 2012.

- [47] B. Zhai, S. Hanson, D. Blaauw, and D. Sylvester, "Analysis and mitigation of variability in subthreshold design," in *Proc. Int. Symp. Low Power Electron. Design (ISLPED)*, 2005, pp. 20–25.
- [48] E. Hogenauer, "An economical class of digital filters for decimation and interpolation," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. ASSP-29, no. 2, pp. 155–162, Apr. 1981.
- [49] P. P. Vaidyanathan, "Multirate digital filters, filter banks, polyphase networks, and applications: A tutorial," *Proc. IEEE*, vol. 78, no. 1, pp. 56–93, Jan. 1990.



Viet Nguyen (Graduate Student Member, IEEE) was born in Hanoi, Vietnam, in 1994. He received the B.Sc., M.E., and Ph.D. degrees in electronic engineering from University College Dublin (UCD), Dublin, Ireland, in 2016, 2017 and 2022, respectively, where he is currently pursuing a Post-Doctoral Researcher at UCD. In 2016, he was an IC Design Intern at Xilinx, Dublin for nine months. His current research interests include ultra-low-voltage circuit design and time-mode data conversion.



Filippo Schembari (Member, IEEE) was born in Codogno, Italy, in 1988. He received the B.Sc. degree in biomedical engineering in 2010, and the M.Sc. and Ph.D. degrees in electrical engineering from the Politecnico di Milano, Milan, Italy, in 2012 and 2016, respectively. During his M.Sc. and Ph.D., he worked on low-noise multi-channel readout ASICs for X- and γ -ray spectroscopy and imaging applications. From 2016 to 2019, he worked as a Post-Doctoral Researcher at the University College Dublin (UCD), Ireland, focusing on deep-subthreshold time-mode ADCs, level-crossing-sampling ADCs, mismatch-calibrated SAR ADCs, and SAR TDCs. During this period he also worked for six months as an IC Design Intern at Xilinx, Dublin, Ireland. In 2019, he joined Huawei Technologies, Milan, Italy, as a RFIC Designer. He was a recipient of the 2018 IEEE Emilio Gatti and Franco Manfredi Best Ph.D. Thesis Award in Radiation Instrumentation and the Marie Skłodowska-Curie European Individual Fellowship (EU-IF) in 2017.



Robert Bogdan Staszewski (Fellow, IEEE) was born in Białystok, Poland. He received the B.Sc. (*summa cum laude*), M.Sc., and Ph.D. degrees in electrical engineering from the University of Texas at Dallas, Richardson, TX, USA, in 1991, 1992, and 2002, respectively. From 1991 to 1995, he was at Alcatel Network Systems, Richardson, involved in SONET cross-connect systems for fiber optics communications. He joined Texas Instruments Incorporated, Dallas, TX, USA, in 1995, where he was elected Distinguished Member of Technical Staff (limited to 2% of technical staff). From 1995 to 1999, he was engaged in advanced CMOS read channel development for hard disk drives. In 1999, he co-started the Digital RF Processor (DRP) Group within Texas Instruments with a mission to invent new digitally intensive approaches to traditional RF functions for integrated radios in deeply scaled CMOS technology. He was appointed as a CTO of the DRP Group, from 2007 to 2009. In 2009, he joined the Delft University of Technology, Delft, The Netherlands, where currently he holds a guest appointment of a Full Professor (*Antoni van Leeuwenhoek Hoogleraar*). Since 2014, he has been a Full Professor with the University College Dublin (UCD), Dublin, Ireland. He has authored or coauthored seven books, 11 book chapters, 160 journals, and 220 conference publications. He holds 210 issued U.S. patents. His research interests include nanoscale CMOS architectures and circuits for frequency synthesizers, transmitters and receivers, and quantum computers. He was a recipient of the 2012 IEEE Circuits and Systems Industrial Pioneer Award.