Electrode-to-Antenna Activity-Adaptive Cortical Monitoring via Stat-Informed DQ-LC-ADC and an Event-Driven Compressive UWB Transmitter

Mina Sayedi*, Maxim Korosteliov*, Soheila Khakneshin, Hossein Kassiri
Electrical Engineering and Computer Science Department, York University, Toronto, Canada
* Equally-Credited Authors (ECAs)

***Abstract*—This paper presents an end-to-end activity-adaptive architecture for scalable wireless cortical monitoring, in which the full recording chain, including amplification, quantization, and data transmission all inherently scale with input signal activity. The core of the recording path is the first reported direct-quantization level-crossing analog-to-digital converter (DQ-LC-ADC), which merges the analog front end and quantizer so that the entire acquisition chain becomes input-adaptive. To make the DQ-LC-ADC feasible, rather than sizing its bandwidth to a conventional worst-case slope, we exploit intracranial EEG slope statistics to set realistic design specifications, relaxing the comparator preamplifier bandwidth by 78× and reducing its power by 37×, while achieving full cortical dynamic range and lowering the average acquisition data rate by nearly an order of magnitude. The 180-nm CMOS prototype achieves 3.15 μV_{rms} integrated IRN, 87.7 dB CMRR, and >100 MΩ input impedance. The DQ-LC-ADC pulse stream is further losslessly compressed using a greatest common divisor (GCD)-based gap encoding with a stride-based extension, and transmitted directly by a fully-digital impulse-radio UWB transmitter whose power scales with pulse rate. Measurements show a 40 Mb/s link at 5 cm with −11.5 dBm output power and 3.3×10^{-5} BER. When compression is accounted for, this corresponds to an effective throughput of 3.4 Gb/s and 0.21 pJ/b effective energy efficiency, an ∼84× improvement over a Nyquist baseline and more than an order-of-magnitude reduction in effective energy per bit compared to state-of-the-art wireless transmitters for brain implants, all while preserving diagnostic accuracy.**

I. Introduction

Higher channel counts are the key to richer neurological insight, but scaling cortical implants raises power consumption against a budget capped by implant size and safety limits. In today's multi-channel systems, an average of 80% of that budget is consumed by data transmission (Fig. 1(a)), making lossless data-rate reduction essential for sustainable scaling. Full reconstruction of raw recordings must be preserved at the receiver, as clinicians rely on it for both real-time interpretation and offline analysis.

Fortunately, cortical signals themselves offer a path to this reduction. Their highly nonuniform spectra, with amplitudes decaying steeply with frequency ($|V(f)| \propto 1/f^2$–$1/f^4$ [1]) as shown in Fig. 1(b), imply limited instantaneous slope and an effective information rate far below Nyquist. In Landau terms, the occupied bandwidth, the portion of the spectrum carrying significant energy, is typically one to two orders of magnitude smaller than the formal Nyquist bandwidth. We define the ratio between these as the overdesign factor (ODF), quantifying how much sampling or transmission is inflated under worst-case assumptions relative to the signal's statistical requirements.

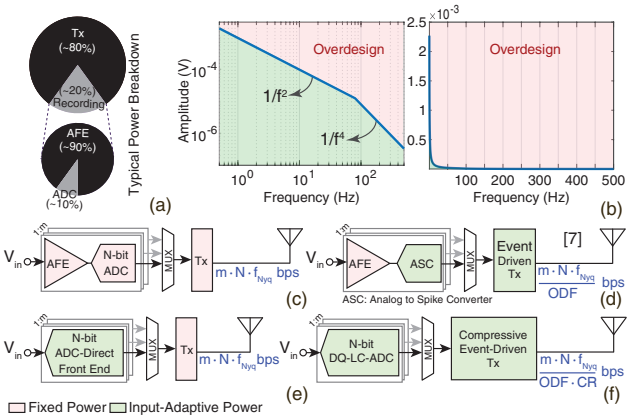

Fig. 1. Motivation and prior art: (a) system-level and recording power breakdown in cortical implants, (b) cortical spectrum overdesign (log and linear scale), and (c-e) representative conventional non-adaptive or partially-adaptive architectures and (f) proposed DQ-LC-ADC-based architecture.

This motivates adaptive data acquisition that scales power with the signal's information rate. LC-ADCs are well suited for this, naturally adapting to signal dynamics while preserving delta-based spectral shaping. However, prior work (Fig. 1(c-e) [4]–[7]) left the analog front end non-adaptive despite it accounting for most of the acquisition power (Fig. 1(a)), and sized LC-ADC bandwidth for rare worst-case slopes, causing significant overdesign. Effective adaptivity must span the entire chain from electrode to antenna.

This work introduces an end-to-end activity-adaptive architecture (Fig. 1(f)) featuring: (i) a DQ-LC-ADC that merges the AFE and quantizer, making the entire recording path power-adaptive; (ii) a statistically informed design of the DQ-LC-ADC that relaxes its bandwidth and power by 78× without loss of diagnostic fidelity; (iii) an IR-UWB transmitter that directly sends DQ-LC-ADC pulses (i.e., no multi-bit encoding or serialization), with pulse-rate-dependent power that naturally tracks signal activity; and (iv) a lossless compression method that leverages inactivity gaps in the DQ-LC-ADC pulse stream to further reduce transmitted data rate. These techniques collectively lead to an overall effective energy efficiency of 0.21 pJ/b (normalized 4.2 pJ/b/m), an 84× improvement compared to Nyquist baseline.

II. Design Methodology

A. Stat-Informed DQ-LC-ADC Design

Fig. 2 shows the implemented DQ-LC-ADC, which employs a bandwidth-limited differential-difference input stage for DC

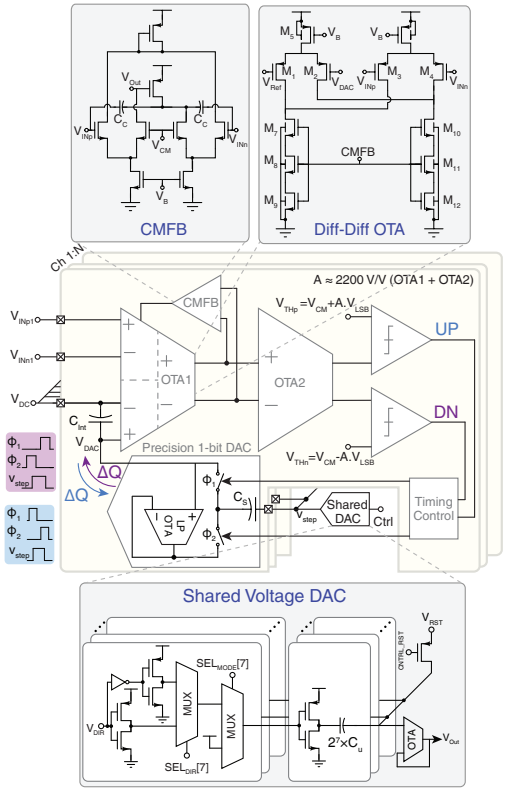

Fig. 2. DQ-LC-ADC architecture: differential-difference input stage, comparators with startup-calibrated thresholds $V_{\mathrm{TH}} = \pm AV_{\mathrm{LSB}}$ for PVT robustness, per-channel precision integrating DAC, and shared binary-weighted capacitive voltage DAC.

offset removal ($\leq$600 mV) while meeting the CMR, Z_{in}, and noise requirements. This is followed by a second g_m stage for additional gain, comparators (threshold trimmed at startup for PVT variations), and a novel 1-bit charge-mode precision DAC. The DAC pushes or pulls charge to or from an integrating capacitor C_{Int} using a 9.5-fF C_S, while a low-power OTA copies the C_{Int} voltage onto the C_S bottom plate before each step to achieve μV-level precision. In parallel, a capacitive DAC, shared among all channels, sets the V_{Step} applied to C_S, refining the up/down step sizes. Transistor-level implementations of key blocks are also shown in Fig. 2.

Unlike conventional LC-ADCs with amplified inputs, the DQ-LC-ADC operates on raw signals, requiring in-loop gain whose delay consumes T_{GRAN} and tightens the timing budget. For cortical signals with a theoretical maximum slope of 3.14 V/s (1 mV at 500 Hz), this would require a continuous-time comparator with impractically high bandwidth and power ($\sim$200 μW). This motivates using discrete-time comparators (DTC) with cross-coupled latches, which exploit positive feedback for high gain and short delay at much lower power. Although discretizing the loop compromises full asynchronicity, a sufficiently high sampling clock prevents noticeable SNR degradation [2]. Using DTC reduces comparator power to $\sim 5\mu$W, which is still excessive.

To reduce power further, we analysed >500 h of raw, unfiltered multi-patient multi-region intracranial EEG [3], including *seizure and artifact segments*, to determine realistic

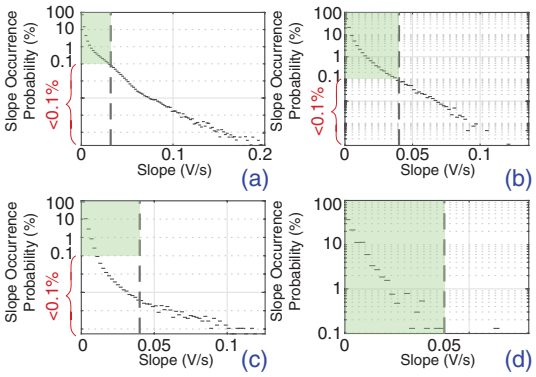

Fig. 3. Representative slope distributions from multiple raw (unfiltered, no denoising, no artifact removal) neural datasets, motivating a statistically grounded and robust 99.9th-percentile slope design target.

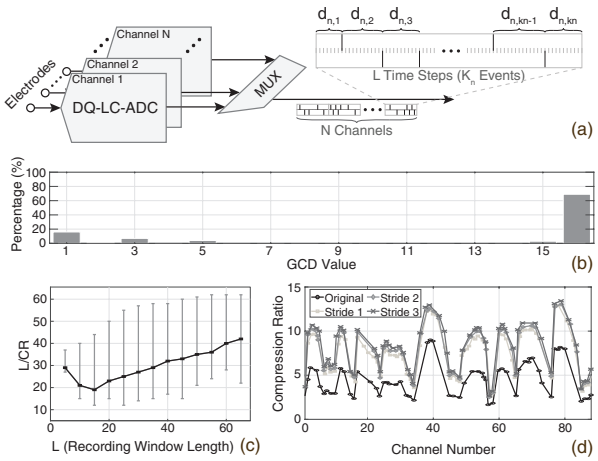

Fig. 4. Stride-extended GCD gap encoding: (a) inter-event gap model, (b) measured GCD distributions across channels, (c) window-length (L) tradeoffs, and (d) stride-based gains validated with <0.5% reconstruction error, showing diminishing returns beyond one stride.

slope limits. Fig. 3(a-b) show the resulting slope distribution (two patients) over 10 h of neural signals including seizure periods, computed directly from raw recordings with no denoising, filtering, or artifact removal, stabilizing at a maximum slope <0.2 V/s. Fig. 3(c-d) isolates seizure-only segments, where the maximum slope decreases to 0.14 V/s. The analysis confirms that the largest slopes occur rarely and are primarily artifacts, while clinically relevant signals occupy much smaller slope ranges. Instead of designing for the theoretical worst case (3.14 V/s), we target a slope capturing 99.9% of observed slopes, excluding only the top 0.1% outliers. This corresponds to 40 mV/s (green area in Fig. 3), representing a 78$\times$ reduction in required input slope and a proportional relaxation in comparator and preamplifier bandwidth and power. Using this method, for a 10h recording, the average data rate drops from 8 kb/s for a Nyquist ADC to 0.9 $\pm$ 0.2 kb/s, with diagnostic fidelity remaining intact, as shown in Section III.

B. Compressive UWB Transmitter

The sparse DQ-LC-ADC output, dominated by long zero runs during low-activity periods, enables further data-rate reduction across channels. As shown in Fig. 4(a), each channel

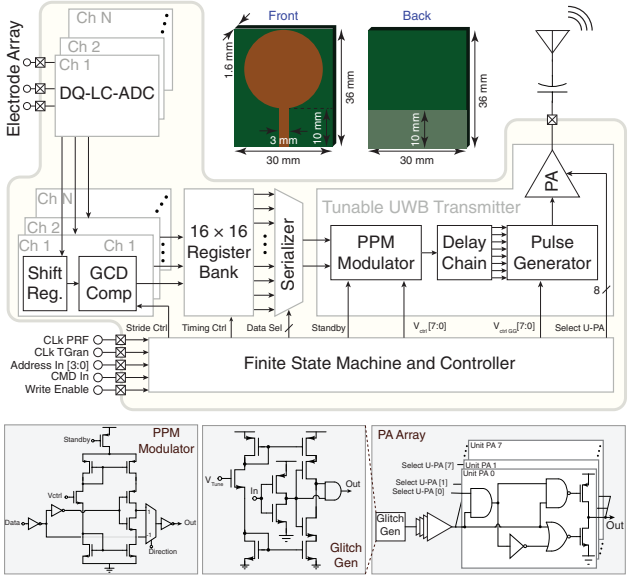

Fig. 5. Top-level block diagram: DQ-LC-ADC array, packetization with GCD compression, pulse-based IR-UWB PPM transmitter, and custom-designed implant-size-representative patch antenna. UWB transmitter circuits: PPM modulator, glitch generator, and fully programmable PA array.

produces inter-event gaps $d_n = \{d_{n,1}, d_{n,2}, \ldots, d_{n,K_n+1}\}$. If these gaps share a common factor $g_n = \mathrm{GCD}(d_n)$, each gap is divided by g_n, giving a compressed packet length $L'_n = K_n + (L - K_n)/g_n$, where L is the window length and K_n is the number of events. The receiver restores timing by multiplying each gap by g_n, yielding a per-channel compression ratio $\mathrm{CR}_n = L/L'_n$.

Selecting L requires balancing latency, preamble overhead, probability of empty packets, likelihood of higher GCD values, and the number of channels that can be grouped per packet while keeping the pulse repetition frequency (PRF) within limits. Shorter windows improve GCD occurrence and allow more channels per packet, while longer ones reduce preamble cost. As shown in Fig. 4(b), which reports the range of L/CR values across 66 iEEG channels, a window length of $L = 16$ achieves the most favorable tradeoff. Using this window length, real iEEG data produce the GCD distribution in Fig. 4(c), where most packets are empty (GCD $= 16$) and fewer than 20% are incompressible (GCD $= 1$), demonstrating strong compression potential.

Additionally, Fig. 4(d) illustrates the stride-based extension, where shifting each event by one tick significantly increases the chances of a higher GCD, hence the effective compression (average: $\sim 2\times$) with negligible reconstruction error ($<$0.5%), thanks to the slow dynamics of iEEG. As shown, beyond the first stride, additional strides offer minimal improvement.

Fig. 5 shows the end-to-end adaptive architecture and the transistor-level implementation of its key blocks. The DQ-LC-ADC array provides per-channel up/down streams that enter shift registers, after which each 16-sample packet is compressed by the GCD block, buffered in a 16×16 register bank, serialized, and delivered to the fully-digital edge-combination UWB transmitter comprising the pulse position

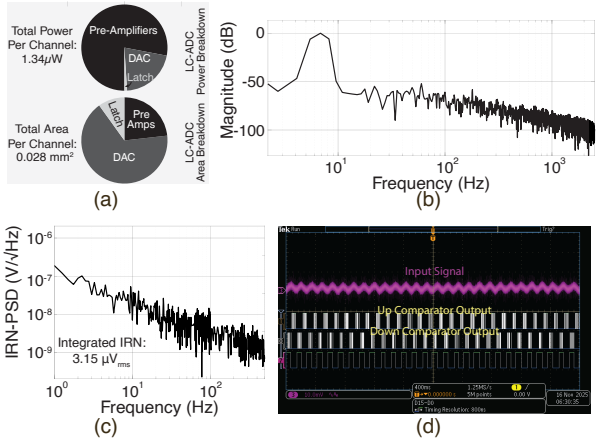

Fig. 6. DQ-LC-ADC measurements: (a) power/area breakdown, (b) output spectrum, (c) integrated IRN, and (d) example oscilloscope waveform and UP/DN bitstreams (input signal: 2mVpp, 7Hz).

modulator (PPM), delay chain, pulse generator, and a power amplifier (PA) array. By design, the transmitter avoids high-order modulation and any PLL because the compression relaxes the required PRF, enabling simple PPM timing control, thus much lower power. All timing parameters, including PPM delay, delay-cell taps, glitch width, edge-combination pattern, and the number of active PA units for gain control, are tunable in silicon. The antenna is a circular patch on FR4, with dimensions matched to a typical implant size.

III. Measurement Results

The 180-nm CMOS prototype was evaluated as both a direct-digitization front end and a fully adaptive wireless system. Measurements in Fig. 6 show that when tested with an input whose slope matches the 40 mV/s statistical maximum observed across the entire neural dataset, the prototype achieves an SNDR covering the full cortical dynamic range, confirming that the stat-informed bandwidth relaxation introduces no loss in recording fidelity. The prototype further achieves 3.15 $\mu\mathrm{V_{rms}}$ input-referred noise (1–100 Hz), Z_{In} >100 MΩ, and 87.7 dB CMRR (measurement setup shown in Fig. 7 (c)), while comparison with a conventional LC-ADC ((Fig. 7 (a)) demonstrates a 37× reduction in required preamplifier bandwidth and proportional power savings.

To verify that the statistically informed approach preserves diagnostic fidelity, seizure-detection accuracy was evaluated using labelled neural signals from epileptic patients (16 seizures over 6 h), where a random-forest classifier achieved identical performance (Fig. 7 (b)) on raw and LC-ADC–quantized data, with rare motion-induced saturation events recovering immediately without affecting reconstruction or classification.

Fig. 8(a–b) shows representative UWB output pulses in time and frequency domains. Fig. 8(c) shows the measurement setup with the chip, power/timing board, custom antennas, and a commercial LNA (ZX60-63GLN+) and band-pass filter (VBFZ-4000-S+) on the receive path. The measured BPF response and antenna return loss appear in Fig. 8(d–e). Because the DQ-LC-ADC and GCD compression greatly reduce data

DQ-LC-ADC Specifications		
	Conventional	Proposed
Target BW (Hz)	500	500
Target Amp_{max} (V)	1m	1m
Required $Slope_{max}$(V/s)	3.14	0.04
Res. (bit)	8	8
T_{Gran} (µs)	2.48	195
Req. Amp BW (Hz)	192	5.2

(a)

Seizure Detection using Raw vs. Proposed DQ-LC-ADC Output		
Input	Raw iEEG	Proposed DQ-LC-ADC Output
True Pos	16	17
False Pos	0	1
False Neg	10	9
True Neg	694	693
Accuracy	98%	98%
Specificity	100%	99.8%

(b)

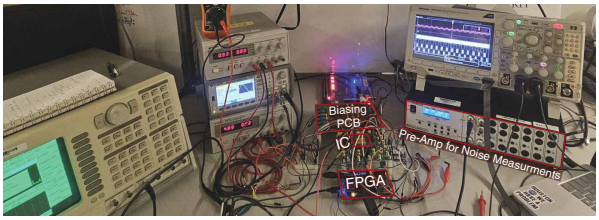

(c)

Fig. 7. (a) Performance benefits in speed and power from statistically informed slope targeting, (b) diagnostic-fidelity results comparing seizure detection (raw vs. DQ-LC-ADC), and (c) ADC tests experimental setup.

rate, the modulation order, and thus required SNR, can be lowered, enabling a simpler transmitter without PLLs or tight timing and reducing circuit complexity and power. Exploiting this, we limit the transmit output to –11.5 dBm, well below FCC limits, resulting in a total TX power of 0.73 mW while achieving a 40-Mbps link at 5 cm with a BER of 3.3×10^{-5}.

With an average GCD-based compression (one stride) ratio of 9.5 and a DQ-LC-ADC data-rate reduction of $\sim$8.9$\times$ for an 8-bit equivalent resolution, the 40-Mbps physical rate corresponds to an effective throughput of 3.4 Gbps. Consequently, the measured 17.85 pJ/b physical energy efficiency maps to 0.21 pJ/b effective, which is more than an order-of-magnitude improvement over the state of the art. Given the large margin between the measured spectrum and the FCC mask, the transmit power (and thus PA power) can be increased if longer-range operation or lower BER is required. Fig. 9 shows the annotated chip micrograph, layout breakdown, and a comparison table against recent UWB transmitter ICs.

IV. Conclusions

This work presented an end-to-end activity-adaptive architecture for scalable wireless cortical monitoring. The first reported DQ-LC-ADC merges the analog front end and quantizer, making the full recording chain, including analog amplification, event-driven and power-adaptive. Its bandwidth is sized from intracranial EEG slope statistics rather than worst-case assumptions, relaxing design specifications by 78$\times$ without compromising noise, input impedance, or diagnostic accuracy. The sparse DQ-LC-ADC output enables lossless GCD-based gap compression with stride extension, and maps directly onto an IR-UWB link whose power tracks neural activity. Together, these yield an effective energy efficiency of 0.21 pJ/b, an $\sim$84$\times$ improvement over a Nyquist baseline and more than an order-of-magnitude reduction in effective energy per bit over the state of the art, demonstrating that exploiting statistical overdesign and pushing adaptivity across the entire chain can unlock higher-channel-count cortical implants under tight power budgets.

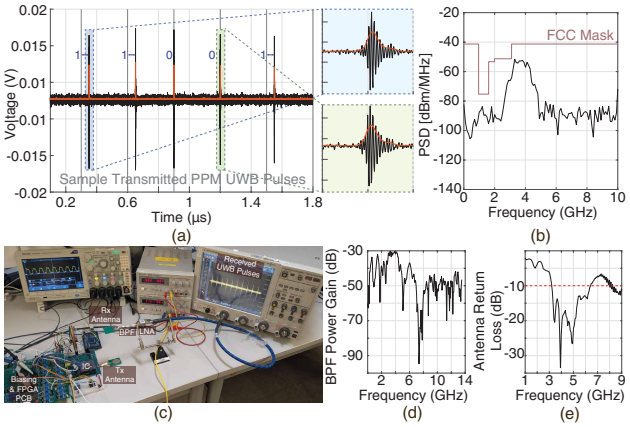

Fig. 8. UWB measurements: transmitted pulses, output spectrum vs FCC mask, full wireless system test setup, BPF response, and antenna return loss.

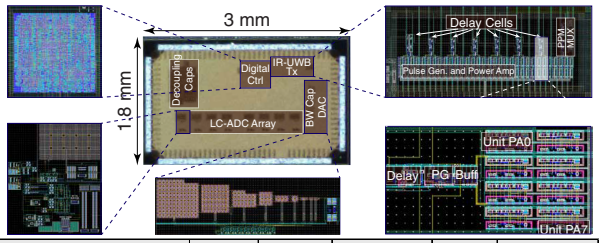

	[8] ISSCC 22	[9] ISSCC 23	[10] CICC24	[11] CICC 25	This work	
					Raw	Effective
CMOS Process (nm)	28	40	65	65	180	
Frequency Band	6 - 9	3.1 - 5	3.1 - 6	6.7 - 9.3	3.1 - 4.4	
Architecture	Up Conversion	Edge Combination	Up Conversion	Up Conversion	Edge Combination	
Modulation Scheme	4PPM+8PSK +4PAM	D16 PPM + PWM+DBPSK	OOK + Envelope Shaping	EPWM+PSK	PPM	
Max Tested Datarate (bps)	1.66G	1.8G	100-800M	1.8G	40M	3.4G
Bit Error Rate	1.00E-04	1.00E-04	1.00E-04	1.00E-03	3.30E-05	
Max output power (dBm)	-8	-7.1	1.4	-	-11.5	
Tx power (mW)	9.69	4.09*	13.2	28.2	0.73	
Energy efficiency (pJ/bit)	5.8	2.3*	16.5	15.7	17.85	0.21
Tx antenna gain (dB)	-8.5	2.6	-19.2	-	3.5	
Rx anteanna gain (dB)	5	3.5	8.5	-	3.5	
Norm. energy efficiency (pJ/b/m) [8]	290	15.3*	26.4	N/R	357	4.2

*Tx Only

Fig. 9. Annotated chip micrograph and block layouts; and comparison table benchmarking the presented work versus recent UWB transmitters in raw/physical and effective energy efficiency.

References

[1] W. Smith *et al.*, IEEE ESSCIRC, 2014, pp. 1–4. doi: 10.1109/ESSCIRC.2014.6942033.

[2] J. Van Assche *et al.*, IEEE TBioCAS, Aug. 2020, pp. 746-756. doi: 10.1109/TBCAS.2020.3009027.

[3] A. Burello *et al.*, IEEE BioCAS, 2018, pp. 1–4. doi: 10.1109/BIOCAS.2018.8584751.

[4] H. Kassiri et al., IEEE J. Solid-State Circuits, vol. 51, no. 5, pp. 1274-1289, 2016. doi: 10.1109/JSSC.2016.2528999.

[5] T. Moeinfard et al., IEEE CICC, 2022, pp. 1-2. doi: 10.1109/CICC53496.2022.9772782

[6] H. Kassiri et al., IEEE JSSC, vol. 52, no. 11, pp. 2793–2810, Nov. 2017. doi: 10.1109/JSSC.2017.2749426

[7] Y. He et al., IEEE JSSC, Oct. 2022, pp. 3058–3070. doi: 10.1109/JSSC.2022.3193846.

[8] M. Song *et al.*, IEEE ISSCC, 2022, vol. 65, pp. 394–396. doi: 10.1109/ISSCC42614.2022.9731608.

[9] J. Lei *et al.*, IEEE ISSCC, 2023, vol. 66, pp. 464–466. doi: 10.1109/ISSCC42615.2023.10067667.

[10] C. Ding *et al.*, IEEE CICC, 2024, pp. 1–2. doi: 10.1109/CICC60959.2024.10528987.

[11] L. Lin et al., IEEE CICC, 2025, pp. 1–3. doi: 10.1109/CICC63670.2025.10983449.