

12.3 A Scalable and Instantaneously Wideband 5GS/s RF Correlator Based on Charge Thresholding Achieving 8-bit ENOB and 152 TOPS/W Compute Efficiency

Kareem Rashed¹, Aswin Undavalli², Shantanu Chakrabarty², Aravind Nagulu², Arun Natarajan¹

¹Oregon State University, Corvallis, OR

²Washington University in St. Louis, St. Louis, MO

Correlators are fundamental building blocks in radar/communication signal processing and analog-to-information (A-to-I) applications such as spectrum sensing [1]. Typically, correlation, which is equivalent to an inner (dot) product, is performed using digital multiply-and-accumulate (MAC) operations (Fig. 12.3.1), with power consumption scaling with frequency, compute, and ADC power. Analog correlators eliminate ADCs but suffer from high power/area and/or small correlation lengths (~10 samples) [2,3]. While matrix multiplication in compute-in-memory (CiM) cores can be used for correlation, such blocks typically operate with lower speeds for multi-bit inputs, require new materials/memory IP and DACs and ADCs at the input/output [4,5].

Importantly, the correlation between two signals can be viewed as a pattern-matching computation where ensemble redundancy provides robustness to approximation errors. For example, an L1 distance-based digital-domain approximation for the dot-product is used in [4], achieving ~100TeraOps/s/W (TOPS/W), targeting applications where inputs are available as multi-bit words. Notably, practically viable analog-domain approximation in correlators must (i) support long sequences (>1000), (ii) exhibit improved accuracy with higher sequence length (similar to MAC), and (iii) achieve high compute energy efficiency comparable to digital MAC while eliminating input ADCs. This paper presents a direct-RF, wideband correlator based on the margin-propagation (MP) paradigm in standard 65nm CMOS that achieves: (i) correlation of instantaneously wideband RF inputs (DC-2.5GHz, i.e., 2x better than prior works), (ii) large correlation length of 1024 in analog domain (>100x better than prior work), (iii) 8-bit computing accuracy (a.k.a. hardware-dynamic range of 50.3dB), while also providing (iv) high compute efficiency of 152TOPS/W traditionally provided only by digital-intensive compute schemes. We also demonstrate system-level measurements/applications such as radar signal detection, code-domain processing, and spectrum sensing using compressive sampling.

As described in Fig. 12.3.1, the true cross-correlation (R_{∞}) between two random input sequences, $\mathbf{X}$ and $\mathbf{Y}$, is only measurable over long sequence lengths. For a given sequence pair, even MAC cross-correlators only converge to R_{∞} as sequence length is increased, with estimation errors for finite lengths depending upon sequence periodicity and probability distributions. As shown in Fig. 12.3.1, correlation can also be estimated using MP-functions that operate on additive and subtractive operands such as $\mathbf{X} \pm \mathbf{Y}$ [6]. For random sequences, the correlation predicted by MP-approximation (R_{MP}) follows a similar error distribution to MAC schemes (R_{MAC}) (i.e., $\epsilon_{MP}^2 = |R_{MP} - R_{\infty}|^2 \sim \epsilon_{MAC}^2 = |R_{MAC} - R_{\infty}|^2$) with both R_{MAC} and R_{MP} converging to R_{∞} as sequence length, $N \rightarrow \infty$ (Fig. 12.3.1). While various possibilities exist for MP functions (Fig. 12.3.1), thresholding-based ReLU function is well-suited for low-power implementations, where $R_{MP} = G^{-1}(z^+ - z^-)$, with $z^{\pm}$ satisfying $\sum_{i=1}^N \text{ReLU}(|x_i \pm y_i| - z^{\pm}) = \gamma$, where x_i and y_i are elements of $\mathbf{X}$ and $\mathbf{Y}$, γ is a hyper-parameter, and G is a monotonic one-to-one mapping function (5^{th} -order polynomial in this work).

As shown in Fig. 12.3.2, the MP approximation is analogous to a reverse water-filling problem of finding an output voltage for a given set of inputs when the total charge across all compute capacitors (C_c), Q_{total} , is constrained. Prior current-domain MP-compute circuits require bias currents that lead to high power for long correlation lengths [6]. The charge-domain scheme to compute MP-based correlation is shown conceptually in Fig. 12.3.2. The thresholding-MP function can be realized using charge-coupled diode-capacitor circuits and the hyperparameter $\gamma = Q_{\text{total}}/C_c$ (Fig. 12.3.2). Four-quadrant operands ($\pm x_i \pm y_i$) are applied to the diodes, as shown in Fig. 12.3.2. At reset, the output voltages $V_{\text{out}}^+ = V_{\text{out}}^- = V_{DD}$ and total charge $Q_{\text{total}} = 0$. During the compute phase, the total charge $Q_{\text{total}} = I_0 t$, the diode outputs $V_i^{\pm} = \max(|x_i \pm y_i|, V_{\text{out}}^{\pm})$, and the thresholding-MP condition $\sum_{i=1}^N C_c \times \text{ReLU}(|x_i \pm y_i| - V_{\text{out}}^{\pm}) = Q_{\text{total}}$ is satisfied. Thus, the differential voltage $V_{\text{out},d} = V_{\text{out}}^+ - V_{\text{out}}^-$ estimates correlation, A . Non-zero diode knee-voltage translates to a constant DC shift that does not impact $V_{\text{out},d}$. The differential implementation is also insensitive to parasitic capacitances on the output node, enabling scalability to large correlation lengths. The predicted correlation R_{MP} , correlation-error ϵ_{MP} follow similar behavior as a MAC-correlator and the $\epsilon_{MP} \rightarrow 0$ as N increases (Fig. 12.3.2). In this MP-compute scheme, energy of $E_{\text{core}} = 2Q_{\text{total}}V_{DD}$ is drawn per N -length correlation with a precision of ENOB=8 and E_{sampler} is consumed for driving the MP-core, resulting in a high compute-efficiency of $N(\text{ENOB}^2 + \text{ENOB})/(E_{\text{core}} + E_{\text{sampler}}) \sim 150$ TOPS/W.

Figure 12.3.3 shows the implemented low-power 5GS/s, 1024-sample analog correlator using the thresholding-MP core shown in Fig. 12.3.2. The RF correlator can be divided into three sub-blocks – a sampler that sequentially stores input samples on 1024 capacitors, an operand generator, and an MP correlation compute engine. The sampler

supports up to 5GS/s with two-layer sampling to reduce the parasitic capacitance at the RF input node. Following sampling, in the operand generation phase, the sampling capacitors, C_{sj} , are stacked to generate the four quadrant operands required by the MP-correlation estimator. The sampling capacitors, C_{sj} (50fF) are conservatively sized to ensure charge leakage during sampling and operand generation does not impact overall compute error. Finally, the operands are applied to the MP-core in the compute phase to calculate the cross-correlation. In the conceptual MP compute cell in Fig. 12.3.2, the compute charge Q_{total} flowing through the diodes is sourced from the sampling capacitors, thus changing the operand value and resulting in computing errors. In the CMOS implementation, the diode-capacitor configuration is replaced with common-drain transistor ($M_{1,2}$) and capacitor to separate the signal and Q_{total} transfer paths as shown in Fig. 12.3.3. The computation speed is determined by the charging rate of the compute cap, C_c (25fF) which is selected conservatively at 10x the lower limit for 8-bit ENOB based on process mismatches and noise simulations. A 50 μ A cascode current source imposes the MP constraint, balancing speed and power trade-offs. The analog output of the MP-core is driven off-chip using a high input-impedance amplifier with 1GHz bandwidth. At 5GS/s, input sampling requires ~200ns, the operand generation settles in < 2ns, and a compute time of ~100ns is used to sample the output.

The IC is implemented in 65nm CMOS with core power consumption of 1.2mW at 1.2V (300 μ W in sampler switch drivers, 680 μ W in operand generator and 220 μ W in MP-compute cell) at 5GS/s. The multi-phase LO generation consumes 4.2mW and LO input buffers consume 22.2mW. Each unit compute cell includes the sampling and operand generation circuits with overall 1024 correlator area of 0.97mm². Measured MP correlator performance is shown in Fig. 12.3.4. Errors in correlation computations arise from (i) errors due to finite input length of sequences $\epsilon_{\text{len}} \sim 1/\sqrt{N}$ with $\epsilon_{\text{len}} \rightarrow 0$ for larger sequence lengths, (ii) errors arising from the hardware MP implementation, ϵ_{HW} , that includes the MP-approximation and noise/mismatch (hardware-dynamic-range, HDR = $20\log(1/\epsilon_{HW})$). The errors from finite sequence lengths can be minimized with periodic inputs (i.e., $\epsilon_{\text{len}} = 0$ and $R_{N,\text{finite}} = R_{\infty}$). Thus, an HDR=50.3dB (i.e., ENOB = 8.06 bits) is measured when correlating two phase-shifted sinusoidal inputs. Measurements using random 5GS/s input sequences with known correlations, show the measured correlation tracking R_{∞} . In this case, the errors are ~30.7dB and are dominated by ϵ_{len} , with measured performance that follows simulations in Fig. 12.3.2. Correlator scalability measurements are shown in Fig. 12.3.4 where longer inputs are portioned into 1024-sample subsequences with correlation for longer sequence lengths (up to $8 \times 1024 = 8192$) computed by summing the outputs for each subsequence. The measured higher accuracy with increasing sequence length demonstrates the feasibility of the MP-based approach for longer correlations. Measurements across $\pm 20\%$ supply voltage demonstrate robustness to voltage variations. System-level measurements for typical high-speed correlation functions are shown in Fig. 12.3.5. The correlation for a noisy 1.25GHz BW radar pulse (SNR = 0dB) sampled at 5GS/s demonstrates analog-domain performance comparable to a MAC correlator across time-shifted pulse templates. Similarly, an input PN code at 2.5GHz chip rate with SNR= 0dB is correlated with delayed versions of the PN code and shows the expected impulse response for zero lag. The correlator is also used to perform compressive sensing measurements where a spectrally-sparse input 1024 sequence with four 5MHz frequency bins occupied is correlated with K (=128) sequences that represent a 128×1024 sensing matrix. The outputs are fed to a CoSAMP recovery in Matlab, with an estimated spectrum close to the input spectrum. These measurements show how the MP correlator can be applied across signal processing applications, with correlation occurring in the analog domain at RF. The MP correlator performance is summarized in Fig. 12.3.6 and compared to the prior art, including digital-intensive compute and analog/RF correlators. The IC demonstrates high input frequency, higher correlator energy efficiency and longer correlation lengths while operating in the analog domain. The die photo is shown in Fig. 12.3.7.

Acknowledgement:

This work was supported by the NSF (#2128535) and the DARPA Massive Cross-Correlation (DARPA MAX) program. The authors also thank Dr. James Wilson (DARPA) for valuable feedback.

References:

- [1] R. T. Yazicigil et al., "Wideband Rapid Interferer Detector Exploiting Compressed Sampling With a Quadrature Analog-to-Information Converter," *IEEE JSSC*, Dec. 2015.
- [2] V. Mangal and P. R. Kinget, "Clockless, Continuous-Time Analog Correlator Using Time-Encoded Signal Processing Demonstrating Asynchronous CDMA for Wake-Up Receivers," *IEEE JSSC*, Aug. 2020.
- [3] Q. Wu et al., "Monolithically Integrated Directly Pumped Time-Varying Transmission Lines," *IEEE TMTT*, vol. 71, no. 4, pp. 1509-1517, April 2023.
- [4] Y. He et al., "A 28nm 38-to-102-TOPS/W 8b Multiply-Less Approximate Digital SRAM Compute-In-Memory Macro for Neural-Network Inference," *ISSCC*, pp. 130-131, Feb. 2023.
- [5] D. Wang et al., "DIMC: 2219TOPS/W 2569F2/b Digital In-Memory Computing Macro in 28nm Based on Approximate Arithmetic Hardware," *ISSCC*, pp. 266-267, Feb. 2022.
- [6] M. Gu and S. Chakrabarty, "Bias-scalable Inner-Product Approximation Circuit Using Analog Margin Propagation," *IEEE MWSCAS*, pp. 525-528, Aug. 2013.




Figure 12.3.1: Analog-friendly threshold-based margin-propagation function provides correlation estimate with precision that is comparable to traditional digital multiply-and-accumulate (MAC) correlators. Similar to MAC correlators, the precision of the MP-correlators improves for longer sequence lengths.




Figure 12.3.2: Thresholding MP can be understood as satisfying reverse water-filling with a diode and capacitor-based circuit constraining total charge. Simulations across random inputs show that output voltage estimates correlation, R_{MP} , with error performance comparable to MAC and > 150 TOPS/W efficiency.




Figure 12.3.3: Schematic of the ReLU MP-based correlator implemented in 65nm CMOS. Operation is divided into sampling phase, operand generation phase and correlation compute phase.




Figure 12.3.4: With periodic sinusoid inputs, measured MP correlation (R_{MP}) demonstrates 8-bit performance. Measurements across random sequences with different sequence lengths show measured performance matching theory/simulations. Charge-domain MP compute is robust to supply voltage variations.




Figure 12.3.5: Application demonstrations using MP-compute IC. Radar measurement for input 1.25GHz BW pulse correlated with pulse template. Measurement for 2.5Gb/s PN code-modulated signal correlated with delayed PN code. Compressive spectrum sensing uses signal correlation with basis sequences to detect sparse frequency spectrum.




Figure 12.3.6: The MP correlator IC supports correlation for long sequences and is equivalent to a digital inner-product computation. IC performance is compared to state-of-the-art digital/CIM vector multipliers, analog correlators and spectrum sensors. The analog-domain IC demonstrates high energy-efficiency while operating at RF frequencies.




Figure 12.3.7: Die photo of implemented correlator IC with 1024 samples for each input, operand generation and MP compute. The IC occupies $2\text{mm} \times 1.1\text{mm}$ (limited by MPW requirements) with MP correlator occupying 0.97mm^2 .