

A 0.65-pJ/bit 3.6-TB/s/mm I/O Interface with XTalk Minimizing Affine Signaling for Next-Generation HBM with High Interconnect Density

Hyunjun Park, *Graduate Student Member, IEEE*, Jiwon Shin, *Member, IEEE*,
Hanseok Kim, *Graduate Student Member, IEEE*, Jihee Kim, *Graduate Student Member, IEEE*,
Haengbeom Shin, *Graduate Student Member, IEEE*, Taehoon Kim, *Graduate Student Member, IEEE*,
Jung-Hun Park, *Member, IEEE*, and Woo-Seok Choi, *Member, IEEE*

Abstract—This paper presents an I/O interface with Xtalk Minimizing Affine Signaling (XMAS), which is designed to support high-speed data transmission in die-to-die communication over silicon interposers or similar high-density interconnects susceptible to crosstalk. The operating principles of XMAS are elucidated through rigorous analyses, and its advantages over existing signaling are validated through numerical experiments. XMAS not only demonstrates exceptional crosstalk removing capabilities but also exhibits robustness against noise, especially simultaneous switching noise. Fabricated in a 28-nm CMOS process, the prototype XMAS transceiver achieves an edge density of 3.6 TB/s/mm and an energy efficiency of 0.65 pJ/b. Compared to the single-ended signaling, the crosstalk-induced peak-to-peak jitter of the received eye with XMAS is reduced by 75 % at 10 GS/s/pin data rate, and the horizontal eye opening extends to 0.2 UI at a bit error rate $< 10^{-12}$.

Index Terms—High bandwidth memory (HBM), edge density, crosstalk cancellation, simultaneous switching noise, ultra-short-reach (USR) link

I. INTRODUCTION

HIGH-performance computing (HPC) catalyzes transformative developments across diverse domains, encompassing artificial intelligence, cloud computing, natural sciences such as astronomy and physics, and humanities including economics and sociology. In light of its significance, there is a pressing need to enhance HPC's capabilities and address its technological challenges. Nevertheless, as the CMOS technology scaling slows down, the effort for enhancement faces obstacles. Acknowledging these challenges, the emergence of advanced packaging technologies offers promising avenues to sustain technological progress and prolong Moore's Law [1].

High Bandwidth Memory (HBM), where data between a host and memory are transmitted over thousands of silicon interposer channels, can offer high bandwidth suitable for HPC applications, but the demands for even higher bandwidth are rapidly growing to support emerging applications. To maintain overall power budget, higher bandwidth demand should be accompanied by I/O energy efficiency scaling in HBM. Pursuing high bandwidth with low area and energy consumption of the I/O interface leads to the adoption of single-ended

(SE) signaling, which exhibits $2\times$ higher pin efficiency than differential signaling for data transmission, over unterminated channels [2], [3]. However, SE brings several disadvantages, particularly its vulnerability to various noise. In systems like HBM having extremely large number of I/O's, the signal deterioration due to data-dependent simultaneous switching noise (SSN) becomes especially pronounced. Various strategies have been explored and developed to mitigate these challenges [4]–[11].

Another key obstacle to higher bandwidth is the crosstalk (XTalk) between the channels. Higher edge density requires small spacing between the channels, increasing crosstalk between neighboring channels [12], [13], which becomes a major threat to signal integrity. To meet the ever-growing demand for higher bandwidth, we should maximize throughput and ensure robustness against noise and crosstalk, while not sacrificing pin efficiency. Addressing this pivotal question forms the core objective of this paper, introducing XTalk Minimizing Affine Signaling (XMAS) as a novel solution to this multifaceted challenge.

Our approach begins with a comprehensive mathematical modeling of XMAS to capture the system performance such as eye width and eye height at the receiver in the presence of crosstalk. This modeling allows design space exploration and strategic co-optimization of the channel and XMAS design to achieve the highest edge density without compromising signaling integrity.

Compared to prior art employing coding or circuit-level techniques for crosstalk cancellation (XTC), the proposed XMAS shows better performance as follows. Conventional bus encoding techniques [14]–[17] add redundant bits to reduce crosstalk, which significantly compromises pin efficiency. Furthermore, as it inherently adopts SE signaling, it is susceptible to noise, making it less robust in systems like HBM. Another prevalent approach involves direct compensation of distortion caused by crosstalk using equalizers [18]–[25]. Despite its effectiveness, this method introduces significant hardware complexity and overhead, worsening the overall I/O energy efficiency. Unlike these methods, XMAS assigns optimized correlation across multiple wires, achieving a remarkable pin efficiency of 87.5 %. Moreover, XMAS ensures robustness against noise and crosstalk without incurring circuit overhead. This novel approach not only addresses the limitations of the previous methods but also offers a balance between pin efficiency and robustness against noise and crosstalk.

This work was supported by the Samsung Electronics Company, Ltd., Hwaseong, Korea.

H. Park, J.-H. Park, J. Shin, H. Kim, J. Kim, H. Shin, T. Kim, and W.-S. Choi are with the Department of Electrical and Computer Engineering and the Inter-University Semiconductor Research Center, Seoul National University, Seoul 08826, South Korea (e-mail: spp098@snu.ac.kr, wooseok-choi@snu.ac.kr).




Fig. 1. Multi-input multi-output linear time invariant channel.

The remainder of this paper is organized as follows. Section II presents a mathematical model for XMAS, which serves as the foundation for the co-optimization techniques with the channel detailed in Section III. Section IV introduces the XMAS transmitter and receiver implementations, and Section V presents the measurement results of the prototype transceiver, followed by a conclusion of this work in Section VI.

II. MATHEMATICAL MODELING

In XMAS, the transmitter (TX) transmits voltage levels after applying an affine transformation to the incoming parallel binary data, and the receiver (RX) recovers the binary data by linearly transforming the received voltage levels. In the following, we present an analytical model for XMAS that enables cooptimization of signaling and interconnect design presented in Section III.

A. Channel Modeling

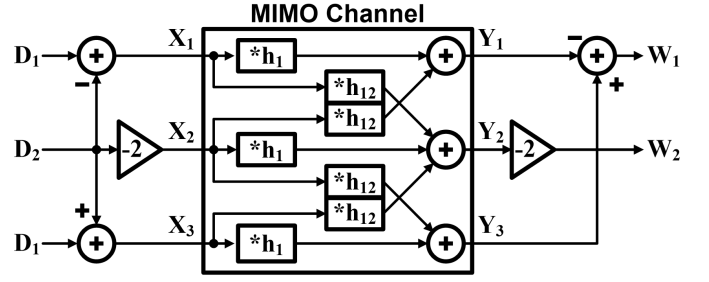
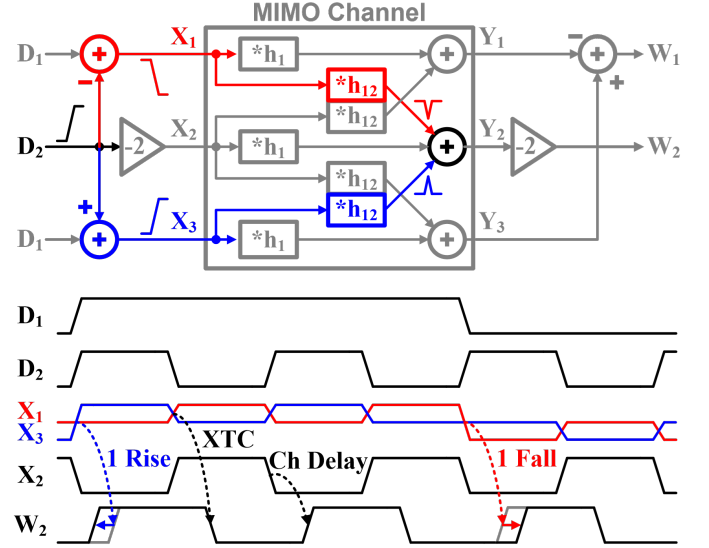
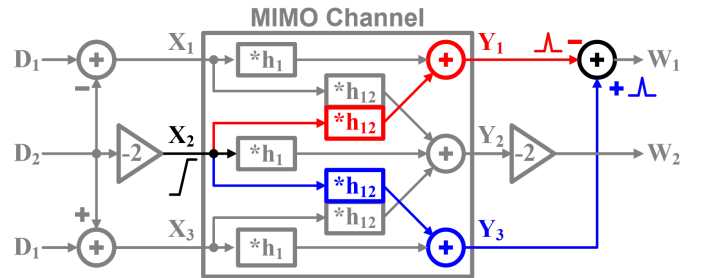
Parallel channels with crosstalk can be considered a multi-input multi-output (MIMO) linear time-invariant (LTI) system, which can be described using a set of channel impulse response as shown in Fig. 1. (h_{ij} denotes the channel output Y_j when the input X_i is given as an impulse.) Specifically, when the channel receives the pulse-amplitude-modulated signals, i.e., $X_i(t) = \sum_{i=-\infty}^{\infty} a_{l,i} \Pi(t - iT)$ where $\Pi(t)$ is 1 for $0 \leq t < T$ and 0 otherwise, the output can be represented as follows.

$$Y_j = \sum_{i=1}^m X_i(t) * h_{ij}(t) = \sum_{i=1}^m \sum_{k=-\infty}^{\infty} a_{i,k} E_{ij}(t - kT). \quad (1)$$

$E_{ij}(t)$ denotes the response of the j -th channel to a single-bit-pulse input originating from the i -th channel. If the channel loss is sufficiently small (i.e., intersymbol interference is negligible), the k -th output of j -th channel depends only on the k -th m -parallel inputs $a_{i,k}$ ($i \in [m]$) and $E_{ij}(t)$. For a MIMO channel, we define \mathbf{H}_k as an $m \times m$ matrix whose elements are $E_{ij}(t - kT)$.

B. Encoding/Decoding with Affine/Linear Transformation

In XMAS, an affine transformation is applied to the parallel input data, which are then transmitted by TX using pulse-amplitude modulation. Specifically, an $n \times m$ integer matrix \mathbf{T} is used to encode m -parallel incoming binary data to an n -dimensional integer vector. Then, the encoded elements are

Fig. 2. XMAS example ($n = 3, m = 2$).Fig. 3. XTC for W_2 .Fig. 4. XTC for W_1 .

mapped to the voltage levels \mathbf{a}_i between 0 and V_{DDQ} (supply voltage of the TX output driver), which can be expressed as:

$$\mathbf{a}_i = [a_{1,i} \dots a_{n,i}]^T = 0.5 \cdot V_{DDQ} (\mathbf{T}_{\text{eff}} \mathbf{d}_i + [1 \dots 1]^T) \quad (2)$$

where $\mathbf{d}_i \in \{-1, 1\}^m$ denotes a vector representing m -parallel binary input to be transmitted, and \mathbf{T}_{eff} is the normalized \mathbf{T} such that the ℓ_1 -norm of each row vector in \mathbf{T}_{eff} becomes 1 to guarantee that the voltage levels \mathbf{a}_i are within 0 and V_{DDQ} . Then the channel outputs can be represented as $\mathbf{H}_k \mathbf{a}_i$, which undergoes a linear transformation at the RX front end. If the linear transformation applied by RX is an $m \times n$ integer matrix




Fig. 5. Channel design parameter (S, W) optimization: (a) Physical structure of channels. (b) Coupling capacitance ratio (C_1/C_2) between adjacent channels as channel width (W) and spacing (S) vary. (c) Channel transfer function with $(S, W, L) = (0.126 \mu\text{m}, 0.36 \mu\text{m}, 1.26 \text{mm})$.

\mathbf{R} , then the decoded outputs are¹

$$\mathbf{W} = \sum_{i=-\infty}^{\infty} \mathbf{R} \mathbf{H}_i \mathbf{a}_i = \sum_{i=-\infty}^{\infty} 0.5 V_{DDQ} (\mathbf{R} \mathbf{H}_i \mathbf{T}_{\text{eff}} \mathbf{d}_i). \quad (3)$$

As a toy example, Fig. 2 illustrates XMAS with below matrices, where two input data are encoded over three channels.

$$\mathbf{T} = \begin{bmatrix} 1 & -1 \\ 0 & -2 \\ 1 & 1 \end{bmatrix} \quad \mathbf{R} = \begin{bmatrix} -1 & 0 & 1 \\ 0 & -2 & 0 \end{bmatrix}$$

$$\mathbf{H} = \begin{bmatrix} h_{11} & h_{12} & 0 \\ h_{12} & h_{11} & h_{12} \\ 0 & h_{12} & h_{11} \end{bmatrix}$$

Inputs D_1 and D_2 are encoded by \mathbf{T} into a set of signals (X_1, X_2, X_3) transmitted through the channels. The channel characteristics are defined by \mathbf{H} , where each channel has identical pulse response h_{11} , and the symmetric coupling between adjacent channels is denoted by h_{12} . Channel outputs (Y_1, Y_2, Y_3) are then decoded by matrix \mathbf{R} into symbols $W_1 = \hat{D}_1$ and $W_2 = \hat{D}_2$. Due to the channel structure, the channel inputs X_1 and X_3 influence Y_2 , causing crosstalk-induced jitter (CIJ). However, as depicted in Fig. 3, the input D_2 is added with the opposite signs into the red and blue paths, effectively canceling the crosstalk at Y_2 . In more detail, when input data in Fig. 3 are given, X_1 and X_3 transition in the exactly opposite directions, perfectly canceling out the crosstalk, or when one signal (X_1) undergoes a full swing transition, the transition in the other signal (X_3) is always prevented, thereby reducing CIJ. Another source of crosstalk is the influence of X_2 on Y_1 and Y_3 . As shown in Fig. 4, X_2 causes the same amount of distortion in Y_1 and Y_3 , which is perfectly canceled during decoding. Such properly designed XMAS matrices thus can hold significant potential for XTC. Hence, in the following, we show how to carefully design the XMAS matrices with the channels to maximize the interface edge density by taking advantage of excellent XTC with XMAS.

III. XMAS DESIGN

This section focuses on determining the XMAS and channel design parameters, aiming to maximize edge density in dense

channel environments. Edge density is affected by various design parameters such as channel dimensions, per-pin data rate, and XMAS matrices. The intricate interplay of diverse parameters significantly influences edge density, making it complex and challenging to find optimal parameters achieving the best performance. Moreover, since the number of possible encoding and decoding matrices in XMAS is vast and simulating each case is computationally complex and time-consuming, finding an optimal parameter set without a theoretical model is practically infeasible. Thus, the analytical model described in Section II is leveraged to efficiently navigate the expansive parameter space and to find the optimal parameters for the maximum edge density. Specifically, the optimization problem for the prototype XMAS transceiver is defined as follows:

$$\begin{aligned} \max_{S, W, L, \mathbf{T}_{n \times m}, \mathbf{R}_{m \times n}, B} \quad & \text{Edge Density} \\ \text{subject to} \quad & \text{Eye Width} \geq 0.7 \text{ UI, Height} \geq 100 \text{ mV} \\ & \text{Channel Loss} \leq 10 \text{ dB, } 0.9 \leq \frac{C_1}{C_2} \leq 1.1 \end{aligned} \quad (4)$$

where S, W, L, C_1 , and C_2 denote the channel spacing, width, length, capacitance between the adjacent channels in the same layer and different layers, respectively (see Fig. 5(a)), and B represents the symbol rate (Baud). Note that the last two constraints in (4) are added so the designed channels have low loss and symmetric capacitance between adjacent channels.

A. Codesign of Interconnects with Signaling

Fig. 5(a) depicts the adopted channel layout, following the densely structured approach of [19] to maximize the interconnect density. The number of wires grouped for encoded data transmission, determined by the number of rows (n) in the XMAS encoding matrix \mathbf{T} , can be adjusted to alter the channel configuration, and the width (W) and spacing (S) of the channels determine their characteristics such as the channel resistance and capacitance as well as the coupling capacitance between adjacent channels. Fig. 5(b) illustrates the variation of the coupling capacitance ratio C_1/C_2 as W and S change. A ratio close to 1 indicates symmetric crosstalk from adjacent channels, represented as the blue dots in Fig. 5(b)². Among them, the configuration with the lowest insertion loss (IL),

¹In (3), the bias term due to $[1 \dots 1]^T$ in (2) is omitted since \mathbf{R} will be chosen to make the bias term become zero.

²Although not mandatory, symmetric crosstalk yields simpler XMAS design, so the ratio close to 1 is chosen for the prototype implementation.




Fig. 6. Impact of crosstalk-induced jitter on SBR: (a) XMAS with optimized matrices. (b) XMAS with degenerate matrices. (c) Single-ended signaling.




Fig. 7. XMAS parameter optimization (n, m, B): (a) Representing whether required eye mask is satisfied or not as symbol rate (B) and number of channels (n) vary. (b) Maximum achievable edge density for different number of channels. (c) Comparison of edge density, maximum symbol rate, and energy efficiency between different signaling schemes.

R								T							
4	-4	0	0	0	0	0	0	4	0	-3	0	0	0	0	-2
0	0	-4	4	0	0	0	0	-4	0	-3	0	0	0	0	-2
-2	-2	2	2	0	0	0	0	0	-4	3	0	0	0	0	-2
0	0	0	0	-4	4	0	0	0	4	3	0	0	0	0	-2
0	0	0	0	0	0	0	-4	4	0	0	-4	0	-3	2	2
0	0	0	0	-2	-2	2	2	0	0	0	4	0	-3	2	2
-1	-1	-1	-1	1	1	1	1	0	0	0	0	-4	3	2	2
								0	0	0	0	4	3	2	2

Fig. 8. Proposed XMAS matrices.

which corresponds to $S = 0.126\mu\text{m}$ and $W = 0.36\mu\text{m}$, is marked with the green square and chosen for the prototype. When the eight wires are designed with $(S, W, L) = (0.126\mu\text{m}, 0.36\mu\text{m}, 1.26\text{mm})$, the channel characteristics are obtained as Fig. 5(c). At a frequency of 5 GHz, the insertion loss is around 10 dB, and the far-end crosstalk (FEXT) reaches -24dB , primarily emanating from the four adjacent channels. These channels emerge as the principal sources of interference, whereas more distant channels have a negligible impact on the overall crosstalk.

B. XMAS Matrix Design

For the described channel configuration, the XMAS matrices \mathbf{T} and \mathbf{R} are designed to possess the following properties that significantly improve signal integrity.

1) *Binary Decision*: Although pulse-amplitude-modulated signals are transmitted through channels, RX makes a binary decision with XMAS, which minimizes the sensitivity to intersymbol interference due to channel loss, similar to chord

signaling [10]. To this end, the XMAS encoding and decoding matrices, \mathbf{T} and \mathbf{R} , are designed such that every row vector in \mathbf{R} and every column vector in \mathbf{T} are orthogonal. In other words, for some diagonal matrix $\mathbf{\Lambda}$,

$$\mathbf{RT} = \mathbf{\Lambda}. \quad (5)$$

With this condition, due to (3), the decoded outputs will have binary levels.

2) *Minimal Crosstalk-Induced Jitter (CIJ)*: For the designed channels, Fig. 6 shows the simulation result demonstrating the impact of CIJ on single-bit response (SBR), where the zero-crossing times are either delayed or advanced depending on the data pattern. This distorted SBR can be accurately captured by substituting appropriate patterns into \mathbf{d}_i as each element of \mathbf{W} in (3) represents the output waveform at the RX. For instance, CIJ-induced SBR for the channel #4 (W_4) can be calculated by evaluating (3) after setting \mathbf{d}_i such that a single bit pulse is given to d_4 (i.e., $d_{i,4} = 1$ and $d_{j,4} = -1$ for $j \neq i$), and all the possible combinations of data patterns are provided to the other data. Then, CIJ_4 , which is the largest difference between the zero-crossing times for W_4 , can be readily calculated based on the computed SBRs. Note that, since the SBR waveforms are determined by \mathbf{T} and \mathbf{R} , proper matrix selection for given channels (\mathbf{H}_i) helps minimize CIJ. As illustrated in Fig. 6, with the optimal selection of \mathbf{T} and \mathbf{R} , CIJ can be greatly reduced; otherwise, CIJ can become worse than SE.

3) *Minimal SSN*: The encoding matrix in XMAS is constructed so the set of voltage levels formed by the n TX drivers can be always constant, i.e., each driver may transmit different voltage during each unit interval (UI), but as a group

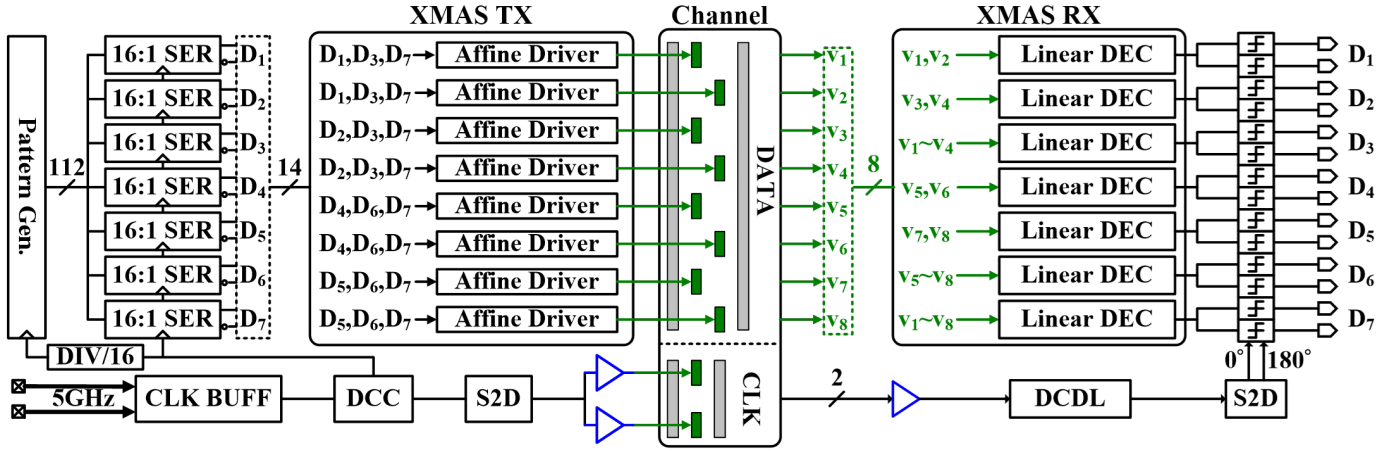


Fig. 9. Overall architecture of the proposed transceiver with XMAS.

of n drivers, they always transmit the identical voltages. For instance, for the 7×8 encoding matrix \mathbf{T} presented in Fig. 8, eight voltages of $V_{DDQ} \cdot [0, 2/9, 3/9, 4/9, 5/9, 6/9, 7/9, 1]$ are invariably used by eight drivers. In other words, regardless of seven input data, the eight encoded data have one of the values in $V_{DDQ} \cdot [0, 2/9, 3/9, 4/9, 5/9, 6/9, 7/9, 1]$, and each driver transmits distinct voltage levels. This ensures that the current provided from the supply to the drivers remains constant irrespective of input data, thereby greatly removing SSN.

Having discussed the desired XMAS properties, now we describe how to design \mathbf{T} and \mathbf{R} for XMAS. Once the channel dimensions S , W , and L are determined, for the fixed values of m and n , integer matrices \mathbf{T} and \mathbf{R} can be determined to satisfy the required properties. To improve pin efficiency and maximize edge density, the number of parallel data to be encoded (m) is fixed at $n - 1$. Fig. 7(a) depicts the symbol rates (B) that satisfy the eye mask constraints in (4), with each set of optimized matrices \mathbf{T} and \mathbf{R} uniquely determined for the fixed channel configuration across various values of n . Note that, for n equal to 7 or greater than 8, the orthogonality condition (5) is unattainable, which is illustrated by a red line in Fig. 7(a). As n decreases, the maximum symbol rate meeting the eye mask constraint tends to increase, suggesting that a smaller n yields better XTC performance of XMAS. However, this also results in lower pin efficiency, thereby not necessarily yielding the highest edge density. Fig. 7(b) illustrates the edge density at the maximum symbol rate for each n value. While the XTC performance of XMAS may decrease as n increases, higher pin efficiency leads to an increase in edge density. Consequently, for maximum edge density, we choose $n = 8$ and the corresponding optimal orthogonal matrices \mathbf{T} and \mathbf{R} are depicted in Fig. 8. For the identical channel configuration and eye-opening conditions, XMAS outperforms SE and differential signaling thanks to its high XTC performance even at a maximum symbol rate of 10 GS/s. Moreover, its high pin efficiency (7/8) allows XMAS to achieve 1.7 \times and 2.25 \times higher edge density compared with SE and differential signaling, respectively, while maintaining comparable energy efficiency, as illustrated in Fig. 7(c).




Fig. 10. Affine driver implementation corresponding to the row vector $[0 \ 4 \ 3 \ 0 \ 0 \ 0 \ -2]$.

IV. XMAS INTERFACE ARCHITECTURE

Fig. 9 illustrates the overall architecture of the proposed interface with XMAS. TX serializes the 16×7 bits of parallel data by 16:1 serializers (SERs) and supplies 7 parallel differential data stream to the 8 affine drivers, which drive the channels with appropriate voltages. The affine drivers encode the parallel incoming data following the row vectors of \mathbf{T} in Fig. 8. Since each row of \mathbf{T} consists of three non-zero elements, each driver takes three out of seven parallel incoming data as shown in Fig. 9. At the RX front end, the linear decoders, implemented following the row vectors of \mathbf{R} , convert the received voltages from 8 channels into 7 binary symbols to restore the data. For the clock path, a limited-swing half-rate differential clock is forwarded to the RX, amplified by a CML-to-CMOS converter [26], and a digitally-controlled delay line (DCDL) [27] is used to compensate for the skew between data and clock.

The TX affine driver consists of multiple N-over-N drivers, where each row vector of \mathbf{T} determines the weight of each driver. For instance, as illustrated in Fig. 10, the driver corresponding to the row vector $[0 \ 4 \ 3 \ 0 \ 0 \ 0 \ -2]$ consists of N-over-N drivers with weights of 4, 3, and 2, respectively. Pull-up inputs for each driver are D_2 , D_3 , and \bar{D}_7 , and pull-down inputs are \bar{D}_2 , \bar{D}_3 , and D_7 . This implementation generates the output voltage level of $V_{CM} + V_R(4D_2 + 3D_3 - 2D_7)$, which is an affine transformation of (D_2, D_3, D_7) . Supply voltage of the driver in the prototype is chosen to be 0.4 V to minimize power. Fig. 11 illustrates the device-mismatch-

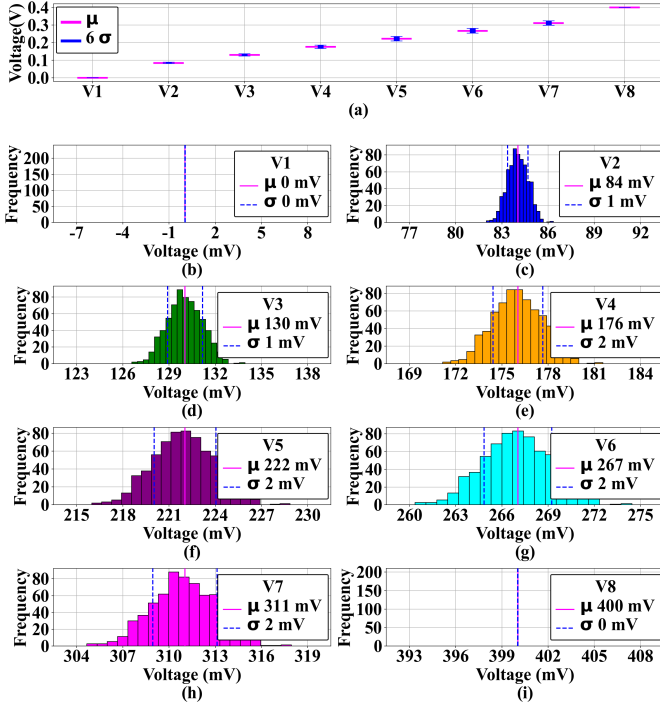



Fig. 11. Simulated affine driver output voltage variation.

Fig. 12. Linear decoder implementation corresponding to the row vector $[-2 -2 2 2 0 0 0]$.

induced distribution of 8 analog voltage levels produced by the designed affine driver. It also presents the histograms of voltage distributions for each output level, which demonstrates that the output voltage level variation due to device mismatch is negligible.

Similar to the TX, the RX decoders at the front end are implemented, following the row vectors of \mathbf{R} , to restore the binary symbol through a linear transformation. For instance, Fig. 12 shows the decoder implementation corresponding to the row vector $[-2 -2 2 2 0 0 0]$ in \mathbf{R} . Differential pairs with capacitive source degeneration allows RX to compensate for channel loss and to perform $v_3 - v_1 + v_4 - v_2$ operation with some gain for recovering D_3 . All the seven decoders employ the identical topology with the appropriately chosen channel inputs based on \mathbf{R} .

We validate the effectiveness of the proposed XMAS interface through simulation results. As shown in Fig. 13(a,b), com-




Fig. 13. RX eye diagrams: (a) SE without SSN, (b) XMAS without SSN, (c) SE with SSN, and (d) XMAS with SSN (5 nH supply inductance).

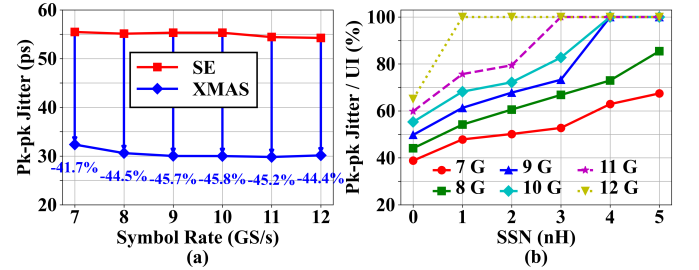


Fig. 14. (a) Peak-to-peak jitter reduction with XMAS across various symbol rates. (b) Peak-to-peak jitter with SE for various supply inductances and data rates.




Fig. 15. (a) CIJ for various bus encoding techniques across different data rates. (b) Performance comparison of bus encoding techniques and XMAS.

pared to SE, thanks to its XTC performance XMAS reduces CIJ from 55ps to 30ps. This timing margin improvement becomes even more pronounced in the presence of SSN. With a supply inductance of 5 nH, as shown in Fig. 13(c,d), while the SE eye completely closes, the XMAS eye is not degraded at all. In case of SE, even with small supply inductance, SE experiences huge peak-to-peak jitter caused by both SSN and CIJ, which leads to significantly degraded signal integrity especially at higher data rates (see Fig. 14(b)). On the other hand, XMAS does not suffer from SSN at all, and as shown in Fig. 14(a), XMAS consistently demonstrates around 45 % jitter reduction on average across various symbol rates compared to SE even in SSN-free environments.

Prior bus coding techniques typically encode the data to




Fig. 16. Chip photomicrograph and power breakdown.

avoid transmitting symbols sensitive to crosstalk by adding redundant bits, which limits the achievable pin efficiency. Comparison between prior art and XMAS is depicted in Fig. 15. For a fair comparison, all methods are compared at a fixed pin efficiency of 75 %³ with the identical channels. Among prior art, [17] achieves the highest XTC performance but uses a differential code, and [14] employs Fibonacci coding for XTC, where the encoders add logical distance between symbols. These works use digital logic to perform encoding/decoding, which incurs a long end-to-end latency. On the other hand, since XMAS encodes/decodes the data without any digital logic, it does not incur any coding latency, while providing superior XTC performance. In the case of chord signaling [10], the XTC performance can vary significantly and may even boost the crosstalk depending on the MIMO channel characteristics. In other words, chord signaling [10] does not necessarily guarantee XTC performance. Fig. 15(b) compares the pin efficiency and peak-to-peak jitter of XMAS with those of prior art with the highest XTC performance. In conclusion, XMAS shows the highest pin efficiency, while achieving the lowest CIJ without incurring additional coding latency.

V. MEASUREMENT RESULTS

The prototype transceiver was fabricated in a 28 nm CMOS technology. Fig. 16 shows the die photo and the power breakdown of the chip. 7-bit-parallel data are encoded and transmitted at 10 GS/s/pin through eight channels (i.e., 70 Gb/s aggregate bandwidth). The overall transceiver occupies an active area of 0.0074 mm² and consumes 44.53 mW.

Fig. 17 shows the measured bathtub curves of XMAS and SE at 10 Gb/s. As shown in Fig. 17(a), the proposed XMAS allows all seven data to have a timing margin at least 0.2 UI at a BER of 10⁻¹². Since the test chip does not have any per-pin deskew circuit, D_7 has a timing skew of about 0.175 UI, but the proposed XMAS achieves error-free operation for all the data. However, the timing margin with SE is severely degraded with crosstalk, and BER lower than 10⁻⁵ cannot be achieved with four aggressors (see Fig. 17(b)). The voltage margin is also significantly improved with XMAS as demonstrated in Fig. 18. For SE, the eye width and height




Fig. 17. Measured bathtub curves at 10 Gb/s: (a) XMAS tested with PRBS15, and (b) SE tested with PRBS7.

decrease rapidly as the number of aggressors increases. Even with three aggressors, SE cannot achieve BER less than 10⁻⁹, while XMAS has a margin of 0.24 UI width and 32 mV height with four aggressors. XMAS provides larger voltage and timing margin compared to SE with just one aggressor (see Fig. 18(f)). Peak-to-peak jitter for SE without any crosstalk is measured to be 68 ps, and Fig. 18(f) shows that SE suffers from CIJ even with one aggressor, increasing jitter to 82 ps, while XMAS is capable of greatly suppressing CIJ and jitter is only 76 ps with four aggressors.

Compared with other state-of-the-art XTC schemes, XMAS shows the best XTC performance and achieves the highest edge density with comparable energy efficiency and area. One popular XTC method [20] uses circuit-level techniques, where the receiver integrates a continuous-time compensation equalizer with a decision feedback equalizer (DFE). The DFE is further augmented with logic that removes post-taps due to crosstalk. However, this approach suffers from significant power and area overhead due to its complex equalizer configuration. Another method [28] mitigates crosstalk by dividing the transmission bandwidth and splitting the phase in high-frequency bands. Though different from adjusting transmission delay for XTC [29], it shares similarities with phase domain equalization as it separates noise sources in the phase domain and applies to the filter. Coding-based XTC schemes, on the other hand, eliminates the need for complex equalizers and exhibit good area and power efficiency, but they generally face limitations such as low pin efficiency and large coding latency.

Fig. 19(a) illustrates the area, energy, and pin efficiency of various XTC schemes, where XMAS, despite being coding-based, shows almost the best energy efficiency and minimal hardware overhead with pin efficiency close to 1. Additionally,

³The XMAS matrices are modified to transmit 3 data over 4 channels for comparison.




Fig. 18. Measured eye diagrams: (a) SE with 1 aggressor, (b) SE with 2 aggressors, (c) SE with 3 aggressors, (d) SE with 4 aggressors, (e) XMAS with 4 aggressors, and (f) eye height and peak-to-peak jitter across for different cases (BER at 10^{-9}).

TABLE I
COMPARISON WITH STATE-OF-THE-ART XTC SCHEMES AND ON-CHIP INTERFACES.

	[20]	[28]	[19]	[25]	[14]	[32]	This work
Technology	32 nm	28 nm	65 nm	28 nm	28nm	5 nm	28 nm
Signaling	SE	SE	SE	SE	SE	SE	XMAS
Data Rate/pin (Gb/s)	7	6	4	28	10	25.2	10
Pin Efficiency(%)	100	100	100	100	75	100	87.5
Edge Density (TB/s/mm)	0.0032	N/A	0.75	N/A	0.0354	0.725	3.6
XTC scheme	CTXC+DFXC	PXC	FFE	CTXC	Coding	N/A	Coding
FEXT@Nyquist(dB)	-29.5	-12	-31.8	+12	-7.8	-42.2	-24
Jitter Reduction	N/A	N/A	78%	N/A	45%	N/A	75 %
Eye Opening (BER)	0.13 UI (10^{-12})	0.2UI (10^{-12})	0.4UI (10^{-9})	0.3UI (10^{-12})	0.58UI (10^{-12})	0.66UI (10^{-12})	0.2UI (10^{-12})
Area (mm²/lane)	0.012	0.004	0.008	0.035	0.004	0.0042	0.001
Energy Efficiency (pJ/bit)	8 (RX)	0.6 (RX)	1.5 (TX+RX)	0.85 (RX)	1.29 (TX+RX)	0.19 (TX+RX)	0.65 (TX+RX)




Fig. 19. Performance comparison: (a) area, power and pin efficiency with other XTC schemes, and (b) edge density with recent on-chip interfaces.

XMAS achieves the highest edge density compared to the studies over the past three years (see Fig. 19(b)) [19], [30]–[33]). Table I compares state-of-the-art interfaces over the past five years. Notably, the XMAS transceiver excels by achieving a great edge density of 3.6 TB/s/mm, an energy efficiency of

0.65 pJ/bit, and an area efficiency of 0.0012 mm²/lane.

VI. CONCLUSION

This paper introduces an I/O interface with the novel signaling XMAS, which is designed to support high-speed data transmission over channels with high crosstalk. XMAS not only demonstrates exceptional crosstalk removing capabilities but also exhibits robustness against noise, especially simultaneous switching noise. To maximize the edge density, co-optimizing the design of channels and signaling is performed using an analytical model of XMAS. The prototype XMAS transceiver is fabricated in a 28-nm CMOS process and achieves an edge density of 3.6 TB/s/mm with an energy efficiency of 0.65 pJ/b. Compared to SE, CIJ of the received eye with XMAS is reduced by 75 % at 10 GS/s/pin data rate, and the horizontal eye opening extends to 0.2 UI at a bit error rate less than 10^{-12} .

ACKNOWLEDGMENTS

The EDA tool was supported by the IC Design Education Center (IDEC), Korea.

REFERENCES

- [1] S. Mirabbasi, L. C. Fujino, and K. C. Smith, "Through the Looking Glass—The 2022 Edition: Trends in solid-state circuits from ISSCC," *IEEE Solid-State Circuits Magazine*, vol. 14, no. 1, pp. 54–72, 2022.
- [2] W.-S. Choi, G. Shu, M. Talegaonkar, Y. Liu, D. Wei, L. Benini, and P. K. Hanumolu, "A 0.45–0.7V 1–6 Gb/s 0.29–0.58 pJ/b source-synchronous transceiver using near-threshold operation," *IEEE Journal of Solid-State Circuits*, vol. 53, no. 3, pp. 884–895, 2018.
- [3] J. C. Lee, J. Kim, K. W. Kim, Y. J. Ku, D. S. Kim, C. Jeong, T. S. Yun, H. Kim, H. S. Cho, Y. O. Kim, J. H. Kim, J. H. Kim, S. Oh, H. S. Lee, K. H. Kwon, D. B. Lee, Y. J. Choi, J. Lee, H. G. Kim, J. H. Chun, J. Oh, and S. H. Lee, "18.3 a 1.2V 64Gb 8-channel 256GB/s HBM DRAM with peripheral-base-die architecture and small-swing technique on heavy load interface," in *2016 IEEE International Solid-State Circuits Conference (ISSCC)*, 2016, pp. 318–319.
- [4] J. W. Poulton, W. J. Dally, X. Chen, J. G. Eyles, T. H. Greer, S. G. Tell, J. M. Wilson, and C. T. Gray, "A 0.54 pJ/b 20 Gb/s ground-referenced single-ended short-reach serial link in 28 nm CMOS for advanced packaging applications," *IEEE Journal of Solid-State Circuits*, vol. 48, no. 12, pp. 3206–3218, 2013.
- [5] J. W. Poulton, J. M. Wilson, W. J. Turner, B. Zimmer, X. Chen, S. S. Kudva, S. Song, S. G. Tell, N. Nedovic, W. Zhao *et al.*, "A 1.17-pJ/b, 25-Gb/s/pin ground-referenced single-ended serial link for off-and on-package communication using a process-and temperature-adaptive voltage regulator," *IEEE Journal of Solid-State Circuits*, vol. 54, no. 1, pp. 43–54, 2018.
- [6] Y. Kwon, H. Park, Y. Choi, J. Sim, J. Choi, S. Park, K.-M. Kim, C. Choi, H.-K. Jung, and C. Kim, "A 33-Gb/s/Pin 1.09-pJ/Bit Single-Ended PAM-3 Transceiver With Ground-Referenced Signaling and Time-Domain Decision Technique for Multi-Chip Module Memory Interfaces," *IEEE Journal of Solid-State Circuits*, 2023.
- [7] A. Shokrollahi, D. Carnelli, J. Fox, K. Hofstra, B. Holden, A. Hormati, P. Hunt, M. Johnston, J. Keay, S. Pesenti *et al.*, "10.1 A pin-efficient 20.83 Gb/s/wire 0.94 pJ/bit forwarded clock CNRZ-5-coded SerDes up to 12mm for MCM packages in 28nm CMOS," in *2016 IEEE International Solid-State Circuits Conference (ISSCC)*. IEEE, 2016, pp. 182–183.
- [8] A. Tajalli, M. Bastani, D. Carnelli, C. Cao, J. Fox, K. Gharibdoust, D. Gorret, A. Gupta, C. Hall, A. Hassanin *et al.*, "A 1.02 pJ/b 417Gb/s/mm USR link in 16nm FinFET,"
- [9] A. Tajalli, M. B. Parizi, D. A. Carnelli, C. Cao, K. Gharibdoust, A. Gupta, A. Hassanin, K. Hofstra, B. Holden, A. Hormati *et al.*, "Short-reach and pin-efficient interfaces using correlated NRZ," in *2020 IEEE Custom Integrated Circuits Conference (CICC)*. IEEE, 2020, pp. 1–8.
- [10] A. Tajalli, M. B. Parizi, D. A. Carnelli, C. Cao, K. Gharibdoust, D. Gorret, A. Gupta, C. Hall, A. Hassanin, K. L. Hofstra *et al.*, "A 1.02-pJ/b 20.83-Gb/s/wire USR transceiver using CNRZ-5 in 16-nm FinFET," *IEEE Journal of Solid-State Circuits*, vol. 55, no. 4, pp. 1108–1123, 2020.
- [11] H. Cronie and A. Shokrollahi, "Orthogonal differential vector signaling," Tech. Rep., 2011.
- [12] S.-K. Lee, B. Kim, H.-J. Park, and J.-Y. Sim, "A 5 Gb/s single-ended parallel receiver with adaptive crosstalk-induced jitter cancellation," *IEEE journal of solid-state circuits*, vol. 48, no. 9, pp. 2118–2127, 2013.
- [13] A. Rainal, "Transmission properties of various styles of printed wiring boards," *Bell System Technical Journal*, vol. 58, no. 5, pp. 995–1025, 1979.
- [14] Q. Liu, L. Du, and Y. Du, "A 0.90-Tb/s/in 1.29-pJ/b Wireline Transceiver With Single-Ended Crosstalk Cancellation Coding Scheme for High-Density Interconnects," *IEEE Journal of Solid-State Circuits*, 2023.
- [15] C. Duan, A. Tirumala, and S. P. Khatir, "Analysis and avoidance of cross-talk in on-chip buses," in *HOT 9 Interconnects. Symposium on High Performance Interconnects*. IEEE, 2001, pp. 133–138.
- [16] B. Victor and K. Keutzer, "Bus encoding to prevent crosstalk delay," in *IEEE/ACM International Conference on Computer Aided Design. ICCAD 2001. IEEE/ACM Digest of Technical Papers (Cat. No. 01CH37281)*. IEEE, 2001, pp. 57–63.
- [17] P. Subrahmanya, R. Manimegalai, V. Kamakoti, and M. Mutyam, "A bus encoding technique for power and cross-talk minimization," in *17th International Conference on VLSI Design. Proceedings*. IEEE, 2004, pp. 443–448.
- [18] S.-Y. Kao and S.-I. Liu, "A 7.5-Gb/s one-tap-FFE transmitter with adaptive far-end crosstalk cancellation using duty cycle detection," *IEEE journal of solid-state circuits*, vol. 48, no. 2, pp. 391–404, 2012.
- [19] H.-G. Ko, S. Shin, J. Oh, K. Park, and D.-K. Jeong, "6.7 An 8Gb/s/μm FFE-Combined Crosstalk-Cancellation Scheme for HBM on Silicon Interposer with 3D-Staggered Channels," in *2020 IEEE International Solid-State Circuits Conference (ISSCC)*. IEEE, 2020, pp. 128–130.
- [20] C. Aprile, A. Cevrero, P. A. Francesse, C. Menolfi, M. Braendli, M. Kessel, T. Morf, L. Kull, I. Oezkaya, Y. Leblebici *et al.*, "An eight-lane 7-Gb/s/pin source synchronous single-ended RX with equalization and far-end crosstalk cancellation for backplane channels," *IEEE Journal of Solid-State Circuits*, vol. 53, no. 3, pp. 861–872, 2018.
- [21] T. Oh and R. Harjani, "A 12-Gb/s multichannel I/O using MIMO crosstalk cancellation and signal reutilization in 65-nm CMOS," *IEEE journal of solid-state circuits*, vol. 48, no. 6, pp. 1383–1397, 2013.
- [22] T. Oh and R. Harjani, "A 6-Gb/s MIMO crosstalk cancellation scheme for high-speed I/Os," *IEEE Journal of Solid-State Circuits*, vol. 46, no. 8, pp. 1843–1856, 2011.
- [23] M. H. Nazari and A. Emami-Neyestanek, "A 15Gb/s 0.5 mW/Gb/s 2-tap DFE receiver with far-end crosstalk cancellation," in *2011 IEEE international solid-state circuits conference*. IEEE, 2011, pp. 446–448.
- [24] S.-Y. Kao and S.-I. Liu, "A 10-Gb/s adaptive parallel receiver with joint XTC and DFE using power detection," *IEEE journal of solid-state circuits*, vol. 48, no. 11, pp. 2815–2826, 2013.
- [25] L. Zhong, H. Wu, W. Wu, W. Xiao, X. Luo, D. Xu, X. Cheng, Z. Li, T. Fan, and Q. Pan, "A 2 × 50 Gb/s Single-Ended MIMO PAM-4 Crosstalk Cancellation and Signal Reutilization Receiver in 28 nm CMOS," in *ESSCIRC 2022-IEEE 48th European Solid State Circuits Conference (ESSCIRC)*. IEEE, 2022, pp. 501–504.
- [26] W.-S. Choi, T. Anand, G. Shu, A. Elshazly, and P. K. Hanumolu, "A burst-mode digital receiver with programmable input jitter filtering for energy proportional links," *IEEE Journal of Solid-State Circuits*, vol. 50, no. 3, pp. 737–748, 2015.
- [27] A. Elmallah, M. G. Ahmed, A. Elkholy, W.-S. Choi, and P. K. Hanumolu, "A 1.6ps peak-INL 5.3ns range two-step digital-to-time converter in 65nm CMOS," in *2018 IEEE Custom Integrated Circuits Conference (CICC)*, 2018, pp. 1–4.
- [28] J. Du, J. Zhou, X. S. Wang, C.-H. Wong, H.-N. Chen, C.-P. Jou, and M.-C. F. Chang, "A Compact Single-Ended Dual-band Receiver with Crosstalk and ISI Reductions for High-density I/O Interfaces," in *2019 IEEE Radio Frequency Integrated Circuits Symposium (RFIC)*. IEEE, 2019, pp. 231–234.
- [29] H. Muljono, K. Peng, L. Sun, I. Abraham, C. Lin, Y. Zhu, and C. Song, "A 2.666 GT/s 128GB/s 14nm Memory I/O with Jitter and Crosstalk Cancellation," in *2019 IEEE Asian Solid-State Circuits Conference (A-SSCC)*. IEEE, 2019, pp. 21–24.
- [30] J.-H. Park, K.-H. Lee, Y. Lee, J.-W. Sull, Y. Song, S. Lee, H. Lee, H. Cho, J. Oh, H.-G. Ko *et al.*, "A 68.7-fJ/b/mm 375-GB/s/mm Single-Ended PAM-4 Interface with Per-Pin Training Sequence for the Next-Generation HBM Controller," in *2022 IEEE Symposium on VLSI Technology and Circuits (VLSI Technology and Circuits)*. IEEE, 2022, pp. 150–151.
- [31] H. Park, Y. Choi, J. Sim, J. Choi, Y. Kwon, J. Song, and C. Kim, "A 0.385-pJ/bit 10-Gb/s TIA-Terminated Di-Code Transceiver with Edge-Delayed Equalization, ECC, and Mismatch Calibration for HBM Interfaces," in *2022 IEEE International Solid-State Circuits Conference (ISSCC)*, vol. 65. IEEE, 2022, pp. 1–3.
- [32] Y. Nishi, J. W. Poulton, W. J. Turner, X. Chen, S. Song, B. Zimmer, S. G. Tell, N. Nedovic, J. M. Wilson, W. J. Dally *et al.*, "A 0.297-pJ/Bit 50.4-Gb/s/Wire Inverter-Based Short-Reach Simultaneous Bi-Directional Transceiver for Die-to-Die Interface in 5-nm CMOS," *IEEE Journal of Solid-State Circuits*, vol. 58, no. 4, pp. 1062–1073, 2023.
- [33] Y. Nishi, J. W. Poulton, X. Chen, S. Song, B. Zimmer, W. J. Turner, S. G. Tell, N. Nedovic, J. M. Wilson, W. J. Dally *et al.*, "A 0.190-pJ/bit 25.2-Gb/s/wire Inverter-Based AC-Coupled Transceiver for Short-Reach Die-to-Die Interfaces in 5-nm CMOS," in *2023 IEEE Symposium on VLSI Technology and Circuits (VLSI Technology and Circuits)*. IEEE, 2023, pp. 1–2.