Bit-Sliced, Ultra-Compact Slow-Light Mach–Zehnder Modulator for Low-Power Operation at 6-Bit Resolution

NICHOLAS W. GANGI[1], MENG ZHANG[1], BOWEN LIU[1], BRETT BAMFO[1], AMIR BEGOVIĆ[2], JIAQI GU[3], AND RENA HUANG[1,*]

[1] Department of Electrical and Computer Engineering, Rensselaer Polytechnic Institute, Troy, NY 12204, USA
[2] American Institute of Manufacturing (AIM) Photonics, Albany, NY 12203, USA
[3] School of Electrical, Computer and Energy Engineering, Arizona State University, Tempe, AZ 85281, USA
[*] huangz3@rpi.edu

Compiled July 3, 2026

This paper presents a compact, slow-light-enhanced segmented Mach-Zehnder modulator (MZM) designed for power-efficient photonic computing systems. By leveraging an apodized Bragg grating to induce slow-light effects, we achieve a total phase-shifter length of only 200 µm. The device operates as an electro-optic digital-to-analog converter (eoDAC) using two cascaded 3-bit segments to achieve 6-bit resolution, shifting bit weighting into the optical domain via the segment length ratio to bypass the exponential power scaling of high-resolution electronic DACs. The modulator achieves a normalized mean square error (NMSE) of $< 10^{-2}$ at clock rates up to 1 GHz and power reduction of approximately 15%. This architecture provides a scalable, high-precision solution for dense, high-throughput silicon photonic accelerators in machine learning applications.

http://dx.doi.org/10.1364/ao.XX.XXXXXX

1. INTRODUCTION

Photonic computing requires high-throughput, high-bit-width electro-optic digital-to-analog converters (eoDACs). Unlike optical communications that push for extreme single-channel baud rates at low bit-depths, photonic processors scale throughput by operating at energy-efficient computing clock speeds while demanding stringent multi-bit precision. Limited eoDAC resolution would introduce algorithmic quantization errors that directly degrade computing accuracy. To date, silicon electro-optic modulators have been reported to achieve 6-to-7-bit resolutions [1, 2]. One key consideration is the power consumption of these high bitwidth eoDACs. Data conversion, both DACs and analog-to-digital converters (ADCs), frequently dominates the system energy budget. For instance, prior work reported that DAC operations alone can account for over 70% of total system power [3]. The power consumption of conventional binary-weighted DAC architectures generally increases approximately exponentially with bit resolution because the total weighted capacitance and switching complexity scale with ~2^{N_b} [4].

Scaling photonic computing systems also requires compact EO modulators to increase on-chip integration density. Conventional silicon Mach–Zehnder modulators (MZMs) typically require millimeter-scale phase shifters. Slow-light MZMs provide an attractive alternative to achieve compact size in the range of 100 µm to several hundred micrometers [5]. While photonic crystals can achieve exceptional modulation efficiency [5], their sub-100 nm features may limit the yield on standard Si photonic foundry utilizing 193 nm lithography. Therefore, this work utilizes Bragg gratings, a simpler structure to demonstrate the bit-slicing concept on a slow-light Si modulator.

Segmented MZMs have previously been explored for increasing data rates in optical communication systems. In these designs, a long traveling-wave electrode is divided into a series of shorter, equal-length phase shifters to suppress microwave attenuation due to transmission-line effect [6–8]. There is also reported work utilizing unequal segments to generate PAM 4 signals, equivalent to 2 bit modulation for optical communication [9, 10]. For example, Jafari et. al. [10] reported a binary weighted two-segment MZM with length ratio of 1.5:1 between the most significant bit (MSB) and least significant bit (LSB) electrodes to generate DAC-less PAM-4 optical signals. Similarly, Hojo et al. [11] recently demonstrated segmented slow-light modulators explicitly optimized for resolving high-speed RF eye diagrams at 10 and 15 Gbaud for PAM-4 transmission. While this prior telecom-centric literature excels at raw transmission baud rates, a critical challenge remains in scaling these segmented architectures in computing application where it requires the eoDAC to achieve higher bit-widths to perform accurate matrix math with emphasis on multi-level static precision, linearity, and aggressive energy efficiency over single-channel speed.

In contrast, bit-slicing utilizes segmentation as a specific data-mapping technique, dividing a larger digital word to drive physically independent segments with individual bits for digital-to-analog conversion. To simultaneously achieve high resolution, low power, and a compact footprint for photonic computing, this work utilizes a bit-slicing technique to implement a multi-level segmented slow-light silicon MZM with 6-bit resolution. In this architecture, individual modulator segments operate at a reduced bitwidth (3 bits), while their optical outputs are weighted and combined to realize a higher effective bit precision (6 bits). A detailed theoretical framework is reported here. A total phase-

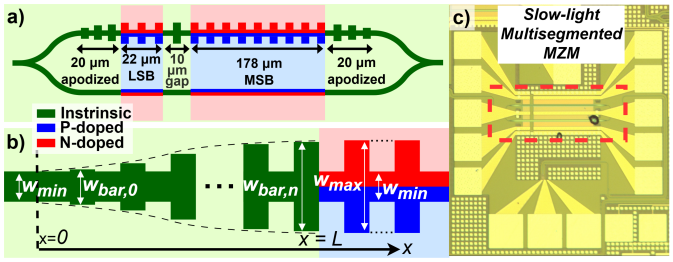

Fig. 1. (a) Schematic of the two-segment SL-MZM. (b) Detailed diagram of the apodized region. (c) An optical micrograph of the fabricated two segment SL-MZM chip.

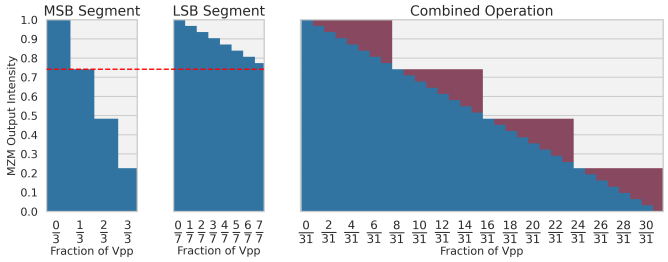

Fig. 2. Diagram of the expected modulator intensity for a 5 bit illustrative example output applying varying ratios of V_{pp} to (a) the long segment (2 bits) and (b) the short segment (3 bits). Operating both segments in tandem results in the intensity distribution in (c), where the major and minor steps originate from the MSB segment and LSB segment respectively.

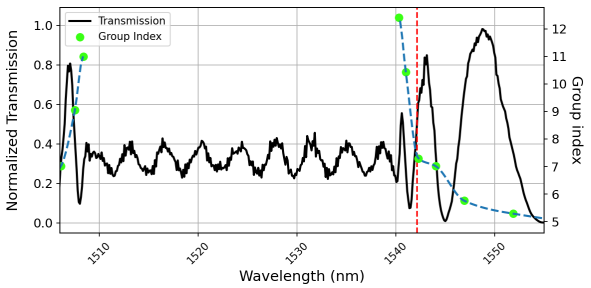

Fig. 3. Multisegment MZM wavelength sweep and calculated n_g. The stop-band ranges from approximately 1509 nm to 1540 nm and the vertical dashed red line shows the wavelength of operation.

shifter length of only 200 µm is achieved in this work, representing a 2.8-fold reduction in physical length compared with [10]. This work demonstrates a LSB segment length of only 22.3 µm for 3-bit operation at over GHz clock rate. To the best of our knowledge, this is the shortest phase-shifter length reported for multi-bit silicon eoDAC operation.

2. DEVICE DESIGN

Bragg grating incorporated slow light MZM (SL-MZM) is employed for the device design. A schematic of the MZM is shown in Fig. 1 where one arm of the modulator has Bragg grating structures while the other arm utilizes a Si/SiO_2 rib waveguide. The grating structure over the Si waveguide is outward apodized to suppress right side band oscillations, with a waveguide core width of 400 nm and a maximum grating bar width of 1.5 µm, and a grating period of 282 nm. The width of the apodized grating bars follows

$$W_{\mathrm{bar},i} = \frac{W_{\max} - W_{\min}}{1 + (\frac{L-x_i}{x_i})^{1.5}} + W_{\min}, \tag{1}$$

where $W_{\mathrm{bar},i}$ is the width of the grating bar at a coordinate x_i along the apodization length L, where $L = 20$ µm. $W_{\max}$ denotes the maximum grating bar width and $W_{\min}$ is the waveguide core width. Grating bar length variation in Eq. 1 was chosen to emulate the well-established super-Gaussian profile [12], while maintaining a uniform grating region with $W_{\mathrm{bar}} = W_{\max}$ in the doped electro-optic modulation active regions. In contrast, the two apodized regions and the 10 µm gap between those two segments remain undoped.

Within the two active regions, $W_{\max}$ is held constant at 1.5 µm for a total of 200 µm. A PPN doping regime was applied to the core of the grating. Detailed doping designs can be found in [2]. The two sets of RF electrodes are separated by a 10 µm undoped gap to ensure electrical isolation. The length ratio of the two segments is a design choice that must balance desired modulation strength at a chosen voltage, desired bit significance, and power usage.

Assuming linear operation near the quadrature point of the MZM, the segment length and voltages can be examined with

$$\Delta P_{\mathrm{res}} = \alpha L \frac{V_{\mathrm{pp}}}{2^{N_b} - 1}, \tag{2}$$

where ΔP_{res} is the change in output intensity between adjacent bit levels when the selected peak-to-peak voltage V_{pp} is divided by the number of bits N_b, α is an efficiency constant specific to the modulator, and L is the electrode length. For the multisegmented case where two sets of electrodes (a and b) are cascaded and when bit significance must be weighted between the electrodes, the desired relationship can be described as:

$$\Delta P_{\mathrm{res},2} = 2^{N_{b,1}} \Delta P_{\mathrm{res},1}. \tag{3}$$

Eq. 3 guarantees there will be continuous, evenly spaced representation of bit levels. Fig. 2 illustrates this concept using an example where MSB is 2 bits while LSB is 3 bits, i.e. $N_{b,2} = 2$ and $N_{b,1} = 3$ (note the actual fabricated devices utilizes $N_{b,2} = 3$ and $N_{b,1} = 3$). Eqs. 2 and 3 can be combined to find:

$$\frac{L_2}{L_1} = (2^{N_{b,2}} - 1) \frac{2^{N_{b,1}}}{2^{N_{b,1}} - 1} \frac{\alpha_1}{\alpha_2} \frac{V_{\mathrm{pp},1}}{V_{\mathrm{pp},2}}, \tag{4}$$

where $N_{b,2}$, L_2, and $V_{\mathrm{pp},2}$ are the number of bits, length, and voltage applied to the MSB segment, respectively, and $N_{b,1}$, L_1, and $V_{\mathrm{pp},1}$ are the number of bits, length, and voltage applied to the LSB segment, respectively. A generalized length ratio derivation for arbitrary cascaded segments is provided in Supplement 1. To establish the design framework, we initially assume an idealized linear and lossless regime where the modulation efficiency is uniform across all segments ($\alpha_1 = \alpha_2 = \alpha$). Physical deviations from this idealized uniform assumption are discussed in detail alongside the experimental results in Section 3C. In this work, electrodes are designed to maximize power efficiency according to [13]. Electrode lengths are chosen such that bit significance is weighted passively for a signal with $N_{b,1} = N_{b,2} = 3$ bits for a net total of 6 bits, assuming the driving voltages $V_{\mathrm{pp},1}$ and $V_{\mathrm{pp},2}$ are equivalent. The optimal length ratio of the LSB and MSB segments is 1:8 from Eq. 3. With an effective total electrode

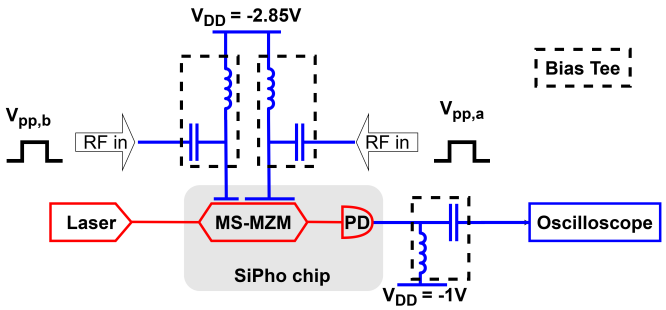

Fig. 4. Testing setup diagram for the multisegment MZM.

length of 200 µm, the LSB and MSB segments were 22.3 µm and 178.2 µm respectively with RF electrodes driven separately.

3. RESULTS AND DISCUSSION

The device was fabricated at AIM Photonics foundry with active multiproject wafer (MPW) runs. The fabricated chip subsequently underwent optical and electrical characterization.

A. Optical Characterization

The optical response of the slow-light segmented MZM is shown in Fig. 3. A photonic stop band from 1509 nm to 1540 nm was identified. Using the interferometric method [14], the maximum measured group index n_g is 12 near the band edge. A gap (10 µm) between those two segments could form a cavity if there is substantial reflection at the interfaces. We were able to examine the effective refractive index of the doped waveguide versus the undoped region in the grating. Carrier dynamics were simulated at varying biases and compared to the undoped regions to determine undesired reflectivities. Calculated reflectivities are quite low (less than 10^{-3}), indicating the cavity effect is weak and the gap region length (10 µm) is a flexible parameter that can be adjusted depending on application.

Due to limited chip space, in-arm thermal phase shifters for tuning to the quadrature point were excluded. Instead, the wavelength was scanned near the photonic band edge while balancing linearity and dynamic range. A wavelength of 1542.1 nm with an n_g of 7.42 was identified near the quadrature point and chosen as the operation wavelength for the NMSE testing.

B. Testing Setup

The EO segmented MZM was driven in single-drive mode along the grating arm. The LSB and MSB segments were driven with a 2-channel Tektronix 70002A Arbitrary Waveform Generator (AWG). The AWG output with 500 mV peak-to-peak was connected to a Picosecond Pulse Labs 5867 low noise amplifier (LNA) to achieve V_{pp} of several volts. Each segment of the MZM is under single-ended drive at a DC bias of 2.85 V reverse bias to the PN junction. The testing setup diagram is shown in Fig. 4.

C. Modulator Characterization for Computing

The device was designed to operate with $V_{pp,1} = V_{pp,2}$ for MSB = LSB = 3 bits; however, in experimental characterization the short segment of the MZM exhibits a 2.42× higher modulation efficiency, α as defined in Eq. 2, than the longer segment. We attribute this difference to optical-loss-induced intensity mismatch between the short and long segments, combined with inherent nonlinearity of the MZM due to the sinusoidal transfer function. Specifically, the longer MSB segment induces a larger phase shift that pushes the optical response off the quadrature point further and operates in the compressed regions of the sinusoidal transfer curve. Concurrently, its longer physical length accumulates more propagation loss, creating an amplitude mismatch that reduces the effective modulation depth. Because the segments are independently driven, we compensated for these physical non-idealities by reducing the LSB driving voltage during operation to restore the proper binary weighting.

Using Keysight ADS simulation, the RF group index was estimated to be 6.04 for the short electrode across the measured frequency range up to 10 GHz, compared with approximately 3.74 for the long segment. The higher index of the LSB segment

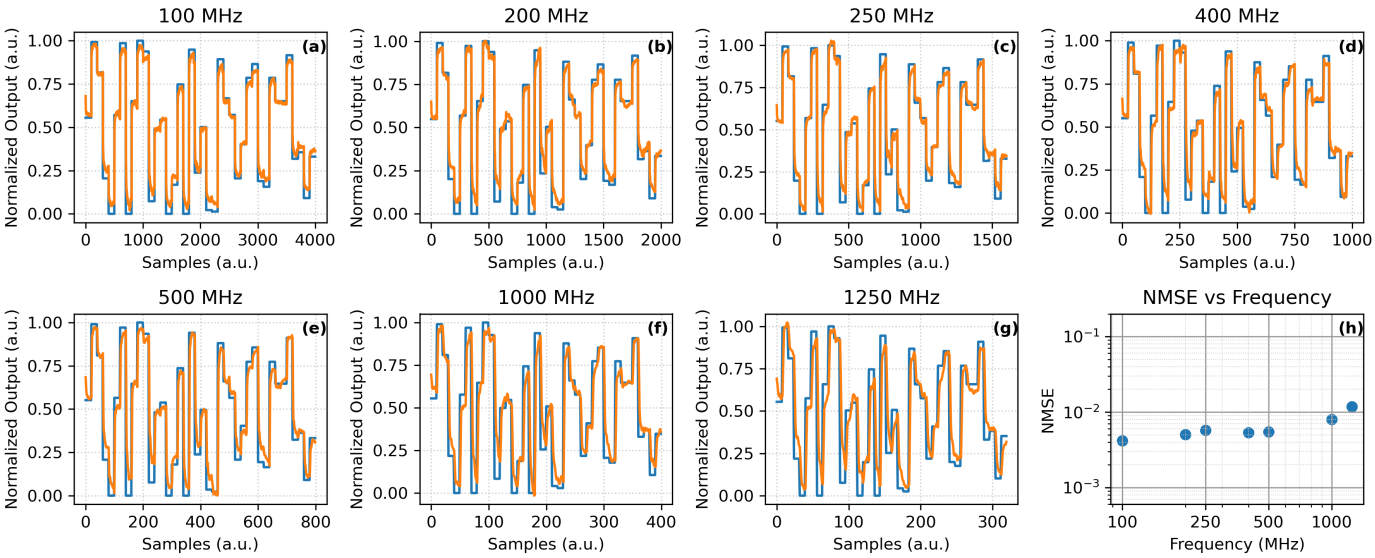

Fig. 5. a-g) Random sequences with 3 MSB and 3 LSB at clock rates from 100 MHz to 1.25 GHz. h) Measured NMSE at varying clock rates.

represents the effective RF group index of the full short-electrode path, rather than the intrinsic group index of only the uniform phase-shifter section. The RF pads, taper, and transition regions contribute a larger fraction of the total group delay, leading to a larger normalized delay and thus a higher extracted RF group index than the MSB electrode. Given the optical group index of 7.42, the short segment provides better RF/optical velocity matching and therefore more effective RF drive delivery to the PN junction than the long segment of the MZM. To compensate for this imbalance, we empirically adjust the drive amplitudes to $V_{pp,2} = 2.4\,V$ and $V_{pp,1} = 1.2\,V$, thereby reconstructing the MSB and LSB optical-weight matching. Due to the relatively low driver voltage and ultra-short phase shifter length, the measured dynamic extinction ratio (ER) is 1.22 dB, comparable to previously reported values [2]. The device is characterized by an estimated insertion loss of approximately 4 dB and an electrical bandwidth of 6.5 GHz [2, 12].

Random sequential 3-bit signals were output by the AWG to the LNAs. Nonlinearity in the LNA output was characterized and precompensated. Despite a known 5-bit LNA bottleneck and 6-bit MZM capacity [2] our two-segment design achieves 6-bit total resolution by driving 3 bits per segment, bypassing LNA bit-resolution constraints. The device was tested using random sequences with a length of 1000 symbols. An averaged staircase-shaped waveform was first used to generate a lookup table for normalized mean-square error (NMSE) calculation [2]. The measured random-signal waveforms are shown in Fig. 5 for clock rates up to 1.25 GHz. At this maximum sampling rate, the system exhibited an NMSE of 1.18×10^{-2}, reflecting the residual linearity distortion caused by the MSB segment operating in the compressed region. The highest effective number of bits (ENOB) extracted from NMSE using the method in [15] is 3.66 bits.

For a modulator with $2^{N_b} = M$ intensity levels and M^2 transition permutations, the average energy consumption per switch is:

$$E_{\text{switch}} = \frac{E_T}{M^2}, \tag{5}$$

where, for a p-n junction modulator of capacitance C, the total rising transition switching energy (E_T) is given in [16] as:

$$E_T = CV_{\text{pp}}^2 \sum_{i=1}^{M-1} (M-i)\left(\frac{i}{M-1}\right)^2. \tag{6}$$

Eq. 5 and Eq. 6 can be reduced to:

$$E_{\text{switch}} = CV_{\text{pp}}^2 \frac{M+1}{12(M-1)}. \tag{7}$$

For the segmented MZM, $M = 8$ (3 bits) for both the LSB and MSB segments for combined 6 bit computing. The measured capacitance of the short and long phase-shifter segments was 260 fF and 355 fF respectively, corresponding to estimated energy consumptions of 40 fJ/switch and 219 fJ/switch. For the combined two-segment 6-bit eoDAC operation, the total estimated energy is 259 fJ/switch, remaining below the pJ/bit level. Simulations also predicted crosstalk between the two segments to be below -40 dB, which was experimentally verified to remain below -65.9 dB over the relevant frequency range.

For comparison, we estimate the switching energy of a single-segment MZM with a length of 200 µm at $M = 64$ (6 bits), with capacitance estimated as the sum of those two segments (615 fF). Assuming this single-segment MZM is driven with $V_{\text{pp,single}} = 2.4$ V, its capacitive switching energy per symbol according to Eq. 7 is 304 fJ/switch. Therefore, the two-segmented MZM has a power reduction of 15% compared with the single-segment implementation if only the MZM capacitive load is considered. Total system power is dominated by the DAC/driver, which scales with unit elements ($2^{N_b} - 1$). Replacing a single 6-bit DAC (63 units) with two 3-bit DACs (14 units) reduces driver energy by approximately 78%. By offloading bit-weighting optically, this bit-sliced architecture reduces overall dynamic power by approximately 74% (see calculations in Supplement 1).

4. CONCLUSION

We demonstrated a compact, slow-light silicon MZM utilizing bit-slicing techniques for energy-efficient, 6-bit computational tasks with phase shifter lengths of only 22 µm and 178 µm. This approach allows multiple low-power, low-resolution eDACs to be aggregated effectively, maintaining NMSE below 10^{-2} at clock rates up to 1 GHz with 6-bit resolution while reducing the modulator's capacitive switching power by approximately 15%. Notably, we demonstrate NMSE of 4.15×10^{-3} at a 100 MHz clock rate and 7.95×10^{-3} at a 1 GHz clock rate. Moving forward, our research will focus on the co-optimization of optical segments and electrode parameters to enhance bandwidth and balance loss across the segments. Ultimately, these segmented MZMs serve as a scalable foundation for dense photonic computing cores, providing a viable path to overcome the power-density bottlenecks currently facing AI and machine learning hardware.

Funding. This work was supported by the Naval Surface Warfare Center Crane and by the NEEC program (Grant N00178-25-0001).

Acknowledgment. The work of Nicholas Gangi was supported by the DoD SMART Scholarship Program. Bowen Liu and Meng Zhang acknowledge the support from the New York State Focus center.

Disclosures. The authors declare no conflicts of interest.

Data availability. Data underlying the results presented in this paper are not publicly available at this time but may be obtained from the authors upon reasonable request.

Supplemental document. See Supplement 1 for supporting content.

REFERENCES

1. Y. Yuan, S. Cheung, T. Van Vaerenbergh, *et al.*, 2023 Opto-Electronics Commun. Conf. (OECC) pp. 1–3 (2023).
2. A. Begović, M. Zhang, D. Yin, *et al.*, Opt. Express **32**, 42016 (2024).
3. M. Zhang, D. Yin, N. Gangi, *et al.*, J. Appl. Phys. **135**, 223105 (2024).
4. Y. Kim, H. Kim, D. Ahn, and J.-J. Kim, Proc. Int. Symp. on Low Power Electron. Des. (2018).
5. Y. Terada, K. Kondo, H. Abe, and T. Baba, Opt. Lett. **42**, 5110 (2017).
6. M. Jacques, A. Samani, E. El-Fiky, and D. V. Plant, 2021 Opt. Fiber Commun. Conf. Exhib. (OFC) pp. 1–3 (2021).
7. Q. Liao, Y. Zhang, S. Ma, *et al.*, IEEE J. Solid-State Circuits **57**, 767 (2022).
8. B. G. Lee, N. Dupuis, R. Rimolo-Donadio, *et al.*, 2017 Opt. Fiber Commun. Conf. Exhib. (OFC) pp. 1–3 (2017).
9. D. Patel, A. Samani, V. Veerasubramanian, *et al.*, IEEE Photonics Technol. Lett. **27**, 2433 (2015).
10. O. Jafari, S. Zhalehpour, W. Shi *et al.*, J. Light. Technol. **39**, 5074 (2021).
11. K. Hojo, Y. Terada, N. Yazawa, *et al.*, IEEE Photonics Technol. Lett. **28**, 1438 (2016).
12. S. R. Anderson, A. Begović, H. Jiang, and Z. R. Huang, IEEE Photonics Technol. Lett. **35**, 697 (2023).
13. Z. Yin, N. Gangi, M. Zhang, *et al.*, 2024 ACM/IEEE Int. Conf. On Comput. Aided Des. (ICCAD) pp. 1–9 (2024).
14. E. Dulkeith, F. Xia, L. Schares, *et al.*, Opt. Express **14**, 3853 (2006).
15. M. Tayyib, M. Amir, U. Javed, *et al.*, PLOS ONE **15**, e0225397 (2020).
16. R. Dubé-Demers, S. LaRochelle, and W. Shi, Optica **3**, 622 (2016).

FULL REFERENCES

1. Y. Yuan, S. Cheung, T. Van Vaerenbergh, *et al.*, "A 7-bit precision linearized mach-zehnder interferometer for high accuracy optical neural networks," 2023 Opto-Electronics Commun. Conf. (OECC) pp. 1–3 (2023).
2. A. Begović, M. Zhang, D. Yin, *et al.*, "Foundry fabricated compact slow-light mach-zehnder modulator and photodetector for on-chip analog photonic computing," Opt. Express **32**, 42016–42030 (2024).
3. M. Zhang, D. Yin, N. Gangi, *et al.*, "Tempo: Efficient time-multiplexed dynamic photonic tensor core for edge ai with compact slow-light electro-optic modulator," J. Appl. Phys. **135**, 223105 (2024).
4. Y. Kim, H. Kim, D. Ahn, and J.-J. Kim, "Input-splitting of large neural networks for power-efficient accelerator with resistive crossbar memory array," Proc. Int. Symp. on Low Power Electron. Des. (2018).
5. Y. Terada, K. Kondo, H. Abe, and T. Baba, "Full c-band si photonic crystal waveguide modulator," Opt. Lett. **42**, 5110–5113 (2017).
6. M. Jacques, A. Samani, E. El-Fiky, and D. V. Plant, "Silicon photonic mzm architectures for 200g per lambda im/dd transmission," 2021 Opt. Fiber Commun. Conf. Exhib. (OFC) pp. 1–3 (2021).
7. Q. Liao, Y. Zhang, S. Ma, *et al.*, "A 50-gb/s pam-4 silicon-photonic transmitter incorporating lumped-segment mzm, distributed cmos driver, and integrated cdr," IEEE J. Solid-State Circuits **57**, 767–780 (2022).
8. B. G. Lee, N. Dupuis, R. Rimolo-Donadio, *et al.*, "Driver-integrated 56-gb/s segmented electrode silicon mach zehnder modulator using optical-domain equalization," 2017 Opt. Fiber Commun. Conf. Exhib. (OFC) pp. 1–3 (2017).
9. D. Patel, A. Samani, V. Veerasubramanian, *et al.*, "Silicon photonic segmented modulator-based electro-optic dac for 100 gb/s pam-4 generation," IEEE Photonics Technol. Lett. **27**, 2433–2436 (2015).
10. O. Jafari, S. Zhalehpour, W. Shi, and S. LaRochelle, "Dac-less pam-4 slow-light silicon photonic modulator providing high efficiency and stability," J. Light. Technol. **39**, 5074–5082 (2021).
11. K. Hojo, Y. Terada, N. Yazawa, *et al.*, "Compact qpsk and pam modulators with si photonic crystal slow-light phase shifters," IEEE Photonics Technol. Lett. **28**, 1438–1441 (2016).
12. S. R. Anderson, A. Begović, H. Jiang, and Z. R. Huang, "Compact Slow-Light Integrated Silicon Electro-Optic Modulators With Low Driving Voltage," IEEE Photonics Technol. Lett. **35**, 697–700 (2023).
13. Z. Yin, N. Gangi, M. Zhang, *et al.*, "Scatter: Algorithm-circuit co-sparse photonic accelerator with thermal-tolerant, power-efficient in-situ light redistribution," 2024 ACM/IEEE Int. Conf. On Comput. Aided Des. (ICCAD) pp. 1–9 (2024).
14. E. Dulkeith, F. Xia, L. Schares, *et al.*, "Group index and group velocity dispersion in silicon-on-insulator photonic wires," Opt. Express **14**, 3853–3863 (2006).
15. M. Tayyib, M. Amir, U. Javed, *et al.*, "Accelerated sparsity based reconstruction of compressively sensed multichannel EEG signals," PLOS ONE **15**, e0225397 (2020).
16. R. Dubé-Demers, S. LaRochelle, and W. Shi, "Ultrafast pulse-amplitude modulation with a femtojoule silicon photonic modulator," Optica **3**, 622–627 (2016).

Bit-Sliced, Ultra-Compact Slow-Light Mach–Zehnder Modulator for Low-Power Operation at 6-Bit Resolution: supplemental document

This supplement provides supporting mathematical models and power analyses for the reported bit-sliced slow-light Mach-Zehnder Modulator. It includes generalized analytical derivations for determining exact physical electrode length ratios in arbitrary multi-segment architectures. Furthermore, it presents a comprehensive dynamic power calculation demonstrating that a two-segment (3+3 bit) design reduces total system-level dynamic power by roughly 74% compared to a single-segment 6-bit implementation.

1. SEGMENTED MZM LENGTH RATIO ANALYSIS

To establish a continuous, evenly spaced representation of bit levels across a multisegment Mach-Zehnder Modulator (MZM), the resolution of each segment must be carefully weighted. In a two-segment architecture, where the segments are physically and logically ordered such that segment 1 represents the Least Significant Bits (LSB) and segment 2 represents the Most Significant Bits (MSB), the resolution step size of the MSB segment must scale by the total number of discrete states in the LSB segment:

$$\Delta P_{res,2} = 2^{N_{b,1}} \Delta P_{res,1}, \tag{S1}$$

where $N_{b,1}$ is the number of bits in the first (LSB) segment. Substituting the linear operation relationship $\Delta P_{res} = \alpha L \frac{V_{\text{pp}}}{2^{N_b}-1}$ for both segments yields the physical length ratio:

$$\frac{L_2}{L_1} = 2^{N_{b,1}} \frac{2^{N_{b,2}}-1}{2^{N_{b,1}}-1} \frac{\alpha_1}{\alpha_2} \frac{V_{\text{pp},1}}{V_{\text{pp},2}}. \tag{S2}$$

Assuming equal phase modulation efficiencies ($\alpha_1 = \alpha_2$) and uniform peak-to-peak driving voltages ($V_{\text{pp},1} = V_{\text{pp},2}$) across both segments, the scaling relationship is dictated purely by the bit allocation $(N_{b,1}, N_{b,2})$. Table S1 illustrates this simplified ratio for various bit configurations, as well as the absolute segment lengths assuming a fixed total modulator active length of 200 µm (where $L_{total} = L_1 + L_2 = 200$ µm). While equal modulation efficiency was assumed for illustrative simplicity, physical devices experience both microwave attenuation and optical propagation losses that vary by segment length and placement. The inclusion of the segment-specific efficiency parameter, α, ensures that these localized electrical and optical signal degradations are fully accounted for in the final length ratio calculations.

Table S1. Two-Segment Length Ratios and Physical Lengths ($L_{total} = 200$ µm)

$N_{b,1}$ **(LSB)**	$N_{b,2}$ **(MSB)**	**Total Bits**	L_2/L_1	L_1 **(µm)**	L_2 **(µm)**
3	3	6	8	22.22	177.78
2	4	6	20	9.52	190.48
4	4	8	16	11.76	188.24
6	6	12	64	3.08	196.92

This concept can be abstracted to a higher level for an arbitrary number of M cascaded segments. By arranging the segments sequentially from the smallest LSB segment (index $i = 1$) up to the longest MSB segment (index $i = M$), the generalized step-size requirement for any two adjacent segments i and $i+1$ becomes:

$$\Delta P_{res,i+1} = 2^{N_{b,i}} \Delta P_{res,i}. \tag{S3}$$

To calculate the characteristics of any arbitrary segment relative to the smallest baseline segment ($i = 1$), the weighting must account for the cumulative number of bits in all preceding, lower-significance segments. The generalized resolution scaling compared to the initial LSB segment is expressed as:

$$\Delta P_{res,i+1} = 2^{\left(\sum_{k=1}^{i} N_{b,k}\right)} \Delta P_{res,1}. \tag{S4}$$

By substituting the physical modulator parameter equation into Eq. S3, we can derive the length ratio between any two adjacent segments in the array as:

$$\frac{L_{i+1}}{L_i} = 2^{N_{b,i}} \frac{2^{N_{b,(i+1)}} - 1}{2^{N_{b,i}} - 1} \frac{\alpha_i}{\alpha_{i+1}} \frac{V_{\mathrm{pp},i}}{V_{\mathrm{pp},i+1}}. \tag{S5}$$

Finally, for layout and fabrication purposes, it is highly advantageous to design the entire electrode cascade relative to the smallest physical unit. Substituting the physical parameters into Eq. S4 provides the generalized length ratio of any arbitrary segment $i + 1$ directly compared to the first LSB segment:

$$\frac{L_{i+1}}{L_1} = 2^{\left(\sum_{k=1}^{i} N_{b,k}\right)} \frac{2^{N_{b,(i+1)}} - 1}{2^{N_{b,1}} - 1} \frac{\alpha_1}{\alpha_{i+1}} \frac{V_{\mathrm{pp},1}}{V_{\mathrm{pp},i+1}}. \tag{S6}$$

To demonstrate the practical application of Eq. S6, Table S2 presents scaling ratios and absolute physical lengths for a three-segment architecture ($M = 3$). Similar to the two-segment analysis, these calculations assume uniform phase modulation efficiencies and equal peak-to-peak driving voltages across all segments, with the total cascaded electrode length constrained to $L_{total} = L_1 + L_2 + L_3 = 1000$ µm.

Table S2. Three-Segment Length Ratios and Physical Lengths ($L_{total} = 1000$ µm)

$N_{b,1}$	$N_{b,2}$	$N_{b,3}$	**Total Bits**	L_2/L_1	L_3/L_1	L_1 **(µm)**	L_2 **(µm)**	L_3 **(µm)**
3	3	3	9	8.00	64.00	13.70	109.60	876.70
4	3	2	9	7.47	25.60	29.35	219.20	751.45
6	2	1	9	3.05	4.06	123.30	375.75	500.95

It is important to note that while these equations provide exact mathematical length ratios, practical implementations must carefully evaluate whether the resulting segment lengths are physically feasible, particularly as bit resolution increases. A baseline segment length configuration that works reliably for a low bit-depth design may become entirely impractical at higher bit depths. Due to the exponential scaling required by the equations, accommodating a higher bit depth either shrinks the LSB segment to dimensions that are impossible to fabricate and dominated by parasitic fringing fields, or forces the MSB segment to exceed practical chip dimensions and RF attenuation limits.

2. FULL DAC POWER CONSUMPTION ANALYSIS

The total electrical power consumption is estimated by summing the intrinsic DAC/driver power and the capacitive switching power of the silicon modulator:

$$P_{\mathrm{total}}(b) = P_{\mathrm{driver}}(b) + P_{\mathrm{mod}}(b) + P_{\mathrm{static}}, \tag{S7}$$

where b is the DAC bitwidth and P_{static} accounts for static bias and leakage power. For a b-bit DAC, the number of quantization levels is

$$M = 2^b. \tag{S8}$$

A. Si Modulator capacitive Switching Power Analysis

The capacitive switching energy of the silicon modulator can then be written as

$$E_{\mathrm{switch}} = C_{\mathrm{mod}} V_{\mathrm{pp}}^2 \frac{M+1}{12(M-1)} = C_{\mathrm{mod}} V_{\mathrm{pp}}^2 \frac{2^b+1}{12(2^b-1)}, \tag{S9}$$

where C_{mod} is the modulator capacitance and V_{pp} is the peak-to-peak drive voltage. This is Eq. (7) in the paper. Consider f_s is the DAC switch rate or symbol rate, the corresponding modulator switching power is

$$P_{\text{mod}} = f_s E_{\text{switch}}(b). \tag{S10}$$

In this work, the average modulator switching energy is calculated to be 259 fJ/switch and the equivalent single segment modulator switching energy is estimated to be 304 fJ/switch.

B. DAC Driver Power Consumption Estimation

Assuming that the DAC/driver dynamic power scales approximately with the number of DAC unit elements, the DAC/driver energy can be modeled as

$$E_{\text{driver}} = E_{\text{unit}}(2^b - 1), \tag{S11}$$

and the corresponding DAC/driver dynamic power is

$$P_{\text{driver,dyn}} = f_s E_{\text{unit}}(2^b - 1), \tag{S12}$$

where the unit DAC/driver switching energy, E_{unit} refers to a unit-cell switching energy that can be extracted from a reference DAC/driver implementation. This extracted E_{unit} is used to scale the DAC/driver contribution with bitwidth, while the modulator contribution is calculated separately from the capacitive switching energy of the silicon modulator.

For a single-segment MZM with 6-bit precision, the DAC/driver energy can be estimated on the number of effective unit element, i.e. $2^6 - 1 = 63$ unit elements. the DAC/driver energy per update can be estimated as

$$E_{\text{driver}} = 63 E_{\text{unit}}. \tag{S13}$$

In comparison, the two segment MZM with 3-bit control in each segment, the the number of effective unit elements in the two segmented drivers would be

$$N_{\text{unit,total}} = 2(2^3 - 1) = 14. \tag{S14}$$

Therefore, the DAC/driver energy per update for the two segment MZM can be estimated as

$$E_{\text{driver}} \approx 2(2^3 - 1) E_{\text{unit}} = 14 E_{\text{unit}}. \tag{S15}$$

The unit DAC/driver switching energy, E_{unit}, is introduced as a first-order normalization of the intrinsic DAC/driver energy per effective unit element per signal update. This value depends on the DAC driver architecture, CMOS technology node, supply voltage, output swing, sampling rate and other detailed driver design. In this paper, we have utilized arbitrary waveform generator (AWG) to drive the Si modulator, so we will utilize literature reported DAC/driver values for system level energy consumption estimation.

We normalize the literature reported power in GHz range of current-steering DACs by the switch rate (i.e. signal refresh rate) and the number of effective DAC unit elements. For an N-bit DAC, this gives

$$E_{\text{unit}} \approx \frac{P_{\text{DAC}}}{f_s\,(2^N - 1)}, \tag{S16}$$

where P_{DAC} is the reported DAC/driver power and f_s is the sampling rate.

Li et al. [1] reported an 8-bit, 6-GS/s current-steering DAC with a reported power consumption of 95.44 mW corresponds to

$$E_{\text{unit}} \approx \frac{95.44\ \text{mW}}{6\ \text{GS/s} \times 255} \approx 62\ \text{fJ/unit/switch}. \tag{S17}$$

Similarly, Kim et al [2] reported a 6-bit, 3.1-GS/s current-steering DAC with a reported power consumption of 17.7 mW gives

$$E_{\text{unit}} \approx \frac{17.7\ \text{mW}}{3.1\ \text{GS/s} \times 63} \approx 91\ \text{fJ/unit/switch}. \tag{S18}$$

Based on these published values [1–3], we assume E_{unit} = 80 fJ/unit/switch as a representative nominal value. The calculated E_{switch} for a single segment 6-bit DAC driver is 304 fJ/switch while the DAC driver for the 2-segment configuration would be 259 fJ/switch. Table S3 summarized

the power consumption of the DAC driver, Si modulator and the total eoDAC system. In Table S3. only the dynamic power is considered following

$$P_{\text{total,dyn}} = f_s \left[E_{\text{unit}}(2^b - 1) + C_{\text{mod}} V_{\text{pp}}^2 \frac{2^b + 1}{12(2^b - 1)} \right]. \tag{S19}$$

if the static power is also included, the total power becomes

$$P_{\text{total}} = P_{\text{static}} + f_s \left[E_{\text{unit}}(2^b - 1) + C_{\text{mod}} V_{\text{pp}}^2 \frac{2^b + 1}{12(2^b - 1)} \right]. \tag{S20}$$

Table S3. Dynamic Power Consumption Comparison at 1.25 GS/s

System Component	Single-Segment (6-bit)	2-Segment (3+3 bit)
DAC Driver Power (P_{driver})	6.300 mW	1.400 mW
Si Modulator Power (P_{mod})	0.380 mW	0.324 mW
Total Dynamic Power	**6.680 mW**	**1.724 mW**

Compared with a single-segment 6-bit implementation requiring $2^6 - 1 = 63$ DAC unit elements, the two-segment $3 + 3$-bit architecture requires only $2(2^3 - 1) = 14$ effective unit elements. Assuming the same unit DAC/driver switching energy, this reduces the theoretical dynamic DAC/driver energy by approximately:

$$1 - \frac{14}{63} \approx 78\%. \tag{S21}$$

This theoretical savings closely aligns with the complete system-level dynamic power reduction of 74% shown in Table S3. The slight difference is due to the inclusion of the silicon modulator's capacitive charging power in the table. Since the capacitive switching energy of the 200 µm slow-light Si modulator is much smaller than the CMOS DAC/driver energy, the overall system energy is heavily dominated by the driver circuit rather than the intrinsic modulator capacitance.

REFERENCES

1. W. Li, J. Zhang, and C. Sun, "A 6GS/S 8-Bit Current Steering DAC with Optimized Current Switch Drive," in *Proceedings of the 7th International Joint Conference on Pervasive and Embedded Computing and Communication Systems (PECCS 2017)*, (2017), pp. 71–75.
2. S. N. Kim, M. R. Kim, B. R. S. Sung, *et al.*, "A SUC-Based Full-Binary 6-bit 3.1-GS/s 17.7-mW Current-Steering DAC in 0.038 mm^2," IEEE Transactions on Very Large Scale Integration (VLSI) Syst. **24**, 794–798 (2016).
3. G. I. Radulov, J. A. Quinn, and A. H. M. van Roermund, "A 28-nm CMOS 1 V 3.5 GS/s 6-bit DAC with Signal-Independent Delta-I Noise DfT Scheme," IEEE Transactions on Very Large Scale Integration (VLSI) Syst. **23**, 44–53 (2015).