

Energy and Performance Benchmarking of a Domain Wall-Magnetic Tunnel Junction Multibit Adder

T. PATRICK XIAO¹, CHRISTOPHER H. BENNETT¹ (Member, IEEE),
XUAN HU² (Student Member, IEEE), BEN FEINBERG¹ (Member, IEEE),
ROBIN JACOBS-GEDRIM¹, SAPAN AGARWAL¹ (Member, IEEE),
JOHN S. BRUNHAVER³ (Member, IEEE), JOSEPH S. FRIEDMAN² (Senior Member, IEEE),
JEAN ANNE C. INCORVIA⁴ (Member, IEEE), and
MATTHEW J. MARINELLA¹ (Senior Member, IEEE)

¹Sandia National Laboratories, Albuquerque, NM 87185-1084 USA

²Department of Electrical and Computer Engineering, The University of Texas at Dallas, Richardson, TX 75080 USA

³School of Electrical, Computer, and Energy Engineering, Arizona State University, Tempe, AZ 85281 USA

⁴Department of Electrical and Computer Engineering, The University of Texas at Austin, Austin, TX 78712 USA

CORRESPONDING AUTHOR: T. P. XIAO (txiao@sandia.gov)

The work of T. P. Xiao, C. H. Bennett, B. Feinberg, R. Jacobs-Gedrim, S. Agarwal, and M. J. Marinella was supported by Sandia's Laboratory-Directed Research and Development Program. This work was supported by the National Science Foundation, CCF, under Award 1910800 and Award 1910997.

ABSTRACT The domain-wall (DW)-magnetic tunnel junction (MTJ) device implements universal Boolean logic in a manner that is naturally compact and cascadable. However, an evaluation of the energy efficiency of this emerging technology for standard logic applications is still lacking. In this article, we use a previously developed compact model to construct and benchmark a 32-bit adder entirely from DW-MTJ devices that communicates with DW-MTJ registers. The results of this large-scale design and simulation indicate that while the energy cost of systems driven by spin-transfer torque (STT) DW motion is significantly higher than previously predicted, the same concept using spin-orbit torque (SOT) switching benefits from an improvement in the energy per operation by multiple orders of magnitude, attaining competitive energy values relative to a comparable CMOS subprocessor component. This result clarifies the path toward practical implementations of an all-magnetic processor system.

INDEX TERMS Benchmarking, domain wall (DW), magnetic logic, magnetic tunnel junction (MTJ), post-CMOS logic, spintronics.

I. INTRODUCTION

SPINTRONIC devices, which exploit transformations between electron spin and electronic charge at the nanosecond or subnanosecond time scale, allow for new frontiers in emerging electronics in terms of speed, energy efficiency, and durability [1]. Due to their nonvolatility, fast switching speed, low-energy barrier for typical energy per bit writing, small feature size, and back-end-of-the-line (BEOL) CMOS compatibility, spintronic devices are, in general, a leading low-power emerging memory candidate [2]. Spintronic device candidates include two-terminal switching devices, such as spin-transfer torque magnetic tunnel junctions (STT-MTJ), which use spin-polarized current to manipulate the state of a thin magnetic layer; three-terminal spin-orbit torque magnetic tunnel junctions (SOT-MTJ), which additionally use spin-orbit coupling or spin-Hall effect

physics [3]; and domain-wall (DW) style devices, which rely upon the movement of a magnetic DW in a ferromagnetic thin film using spin-polarized current and typically require three terminals, or in some cases, a specially structured nanotrack [4], [5].

Presently, emerging spintronic devices are being considered to replace or augment some components of the modern memory hierarchy [6], in more exotic memory applications, e.g., neuromorphic computing [7], and lastly, to implement new styles of energy-efficient Boolean logic [8]. In this article, we consider the last case and significantly advance the analysis of DW logic devices for the next-generation logic systems relative to previous models.

At present, many proposed spintronic logic designs, e.g., hybrid magnetic device and CMOS flip-flops, heavily rely on CMOS devices, while only a small portion of the

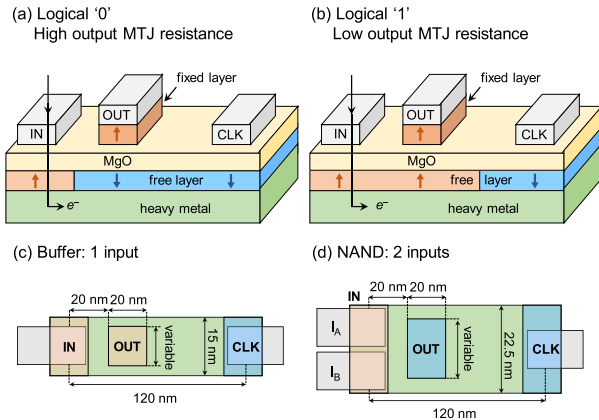


FIGURE 1. (a) Diagram of a DW-MTJ buffer device that uses the spin-orbit torque effect induced by a heavy metal to translate a DW along a PMA ferromagnetic track (free layer). The output MTJ is in the high-resistance antiparallel state and produces a low current (“0”) when the DW in the underlying track lies to its left. (b) When the DW moves to the right-hand side of the MTJ, the MTJ is in the low-resistance parallel state and produces a high current (“1”). (c) Top-down view of a buffer gate. The MTJ width is adjusted to provide different fan-outs. The fixed layer of the MTJ is aligned with the magnetization on the far left side of the track. (d) Top-down view of the NAND gate, which has a wider track and thus a larger threshold current for DW movement. The fixed layer of the MTJ is aligned with the magnetization on the far right-hand side of the track.

functionality is implemented with magnetic devices [9], [10]. This is, in part, due to the difficulty of achieving all necessary logic functions with STT-MTJ or SOT-MTJ devices, which do not contain sufficient information capacity or dynamics to implement intrinsic logic gates. While proposals for all-spin-logic exist, these results are largely based on modeling [11], [12]. In contrast, DW magnetic memory devices possess key properties for intrinsic logic due to spatial and temporal manipulation, and these device properties have been experimentally demonstrated. DW logic allows for flexibility, and thus, energy efficiency in allowing multiple varieties (SOT and STT) of spin-polarized current to drive DWs in ferromagnetic materials [13], allows for higher density and better cascading ability than standard MRAM (e.g., STT-MTJ) logic circuits [14], and allows for varieties of chiral DW motion via nanotrack engineering [15], [16]. While a proposal for better-than-CMOS DW-logic called mLogic has been made and simulated [17], this article lacked a realistic device model and additionally only considered STT-modulated DW motion.

In contrast, this article models a three-terminal DW magnetic tunnel junction (DW-MJT) device from which we have extrapolated realistic values. This device contains input-output and clocking ports and relies upon the movement of a DW along a ferromagnetic strip under the junctions to switch states (see Fig. 1). It has already been fabricated and used to realize small proof-of-concept logic circuits [18] and is being engineered to utilize both SOT and STT style DW movements.

This report significantly extends early simulated work, which demonstrated a 1-bit full adder [19] by utilizing a recent SPICE model [20] to implement a complex, multi-bit adder system. The SPICE model has been benchmarked

against and reproduces the DW behavior predicted by micro-magnetic simulations. Our results take into account multiple realistic DW-MTJ device parameters, imperfect circuit effects, and register and communication (interconnect) components, and are clocked and pipelined in a way which allows for intrinsic logic performance. Although incorporation of many of these realistic effects provides less favorable benchmark numbers than previously reported for STT DW logic [21], [22], our results suggest that there is nonetheless a potential for SOT-based DW-MTJ logic systems. By providing a realistic device and architecture-level benchmarks, we suggest possible optimization routes for better energy efficiency than that possible with CMOS logic systems.

II. DW-MTJ DEVICE AND ADDER DESIGN

A. DW-DEVICE AND ITS VARIANTS

The operation of the three-terminal DW-MTJ device is shown in Fig. 1. The device state is encoded in the position of a DW along a soft ferromagnetic track, whose magnetization at the left and right ends is pinned. The DW separates the track into two regions of opposing magnetization. An injection current through the IN terminal with the Clock (CLK) terminal grounded can translate the DW along the length of the track by the STT or SOT mechanism. By injecting a current through the input (left) terminal of the device, the DW can be translated along the length of the track. This occurs either by STT applied on the DW by a spin-polarized current through the ferromagnetic track or by a strictly current-induced SOT [23] arising from the spin Hall effect [3], [24] in a heavy metal layer that lies below the track. Fig. 1(a) and (b) shows the device structure for the SOT case.

The ferromagnetic track forms the free layer of the output (center) MTJ, which switches between a low-resistance parallel state R_p and a high-resistance antiparallel state R_{ap} as the DW moves from one side of the track to the other. The two resistances are related by the tunneling magnetoresistance of the MTJ: $TMR = 100\% \times (R_{ap} - R_p) / R_p$. A subsequent injection of current through the Clock (right) terminal of the device moves the DW back to the left side of the track and resets the resistance state of the MTJ. When cascaded, part of this reset current passes through the MTJ and communicates the device state to the next gate. The remainder is sunk into the Clock terminal of the previous gate, which has already been reset without affecting its state.

Motion of the DW is produced only above a threshold current I_{th} . Thus, the device functions as a buffer gate if a sufficiently high input current moves the DW such that it subsequently produces a high output current. Conversely, if a high input current is followed by a low output current—accomplished by reversing the fixed layer magnetization in Fig. 1(a)—the device functions as an inverter. By adding a second input terminal and setting its threshold so that the DW moves only when both input currents are high, we obtain a NAND gate. The use of a current signal at both the input and output allows DW-MTJ devices to be readily cascaded to implement any logic functionality. For design and process simplicity, we construct logic circuits using only the NAND and buffer gates, although a single DW-MTJ device can additionally implement the AND, OR, and NOR gates [19].

The switching current requirements of our device, important for energy efficiency, are notably set by the style of

TABLE 1. Parameters used in circuit simulation at 300 K.

Parameter	Value
Clock voltage V_{CLK}	125 mV (STT), 12.5 mV (SOT)
Clock period τ	15.0 ns
Clock pulse width	2.0 ns
Track width w	15 nm (buffer), 22.5 nm (NAND)
Track length L	120 nm
Track + HM thickness d	2.5 nm
Threshold current density J_{th}	2.4×10^{11} A/m ² (STT) 2.4×10^{10} A/m ² (SOT)
Output MTJ RA product	$1.0 \Omega \mu\text{m}^2$
Output MTJ length L_{MTJ}	20 nm
Output MTJ width w_{MTJ} and ON-resistance R_p	FO1: $w_{\text{MTJ}} = 5.95$ nm, $R_p = 8.4$ k Ω FO2: $w_{\text{MTJ}} = 13.15$ nm, $R_p = 3.8$ k Ω FO3: $w_{\text{MTJ}} = 20.83$ nm, $R_p = 2.4$ k Ω FO4: $w_{\text{MTJ}} = 22.5$ nm, $R_p = 2.2$ k Ω Register: $w_{\text{MTJ}} = 8.33$ nm, $R_p = 6.0$ k Ω Adder out: $w_{\text{MTJ}} = 10$ nm, $R_p = 5.0$ k Ω
Free layer / HM resistivity ρ	$4.0 \times 10^{-7} \Omega \text{m}$
Track + HM resistance R_{write}	1.3 k Ω (buffer), 0.87 k Ω (NAND)
Series resistance R_{series}	1.37 k Ω (buffer), 1.8 k Ω (NAND)
Saturation magnetization M_{sat}	0.8 MA/m
Injected spin polarization P	0.7
Domain wall width	2.5 nm
Edge repulsion factor K_{repe}^*	0.25
Inertia parameter α^*	0.01

FO = Fanout, HM = heavy metal. * See [20] for definitions.

anisotropy in our output MTJ. For all the following simulations, we have assumed that the device has perpendicular magnetic anisotropy (PMA), which typically has a lower energy barrier and thus a reduced switching current compared with in-plane magnetic anisotropy (IMA) [25].

Table 1 lists the parameters used for the comparison of DW-MTJ logic with the 2018 CMOS technology node of the International Technology Roadmap for Semiconductors [26], which assumes a metal-1 half pitch of 15 nm. We use a track width of $w = 15$ nm for the one-input buffer gates and $w = 22.5$ nm for the two-input NAND gates, as shown in Fig. 1(c) and (d), respectively. The length and thickness of the ferromagnetic track are fixed at $L = 120$ nm and $d = 2.5$ nm, respectively. In the SOT case, we assume that the combined thickness of the free layer and the heavy metal is 2.5 nm, with a majority of the current passing through the thicker metal, similar to the geometry in [27] and [28]. In STT and SOT with some material combinations (such as Ta/CoFeB), the DW moves in the direction of electron flow, as shown in Fig. 1.

We assume the threshold current density for DW depinning to be $J_{\text{th}} = 2.4 \times 10^{11}$ A/m² for STT-driven DW motion in PMA layers [29]. We have designed the parallel-state resistance R_p of the MTJ to drive each device with a current density slightly larger than threshold when the input is high. With the parameters listed in Table 1, the DW moves with a velocity in the range of $v = 10$ to 15 m/s [30]. Using SOT, the threshold current density can be reduced by approximately an order of magnitude to the 10^{10} -A/m² range, while still maintaining similar or even lower values for the DW velocity [31]. We, therefore, use a value of $J_{\text{th}} = 2.4 \times 10^{10}$ A/m² for SOT-driven DW motion and assume, in our model, that the same DW velocity can be achieved at the reduced current density. We will later evaluate the technological implications of the critical parameter J_{th} . For the chosen values, the threshold current of the device is $I_{\text{th}} = 9.0 \mu\text{A}$ ($0.9 \mu\text{A}$) in the buffers and $I_{\text{th}} = 13.5 \mu\text{A}$ ($1.35 \mu\text{A}$) in the NAND gates for STT (SOT).

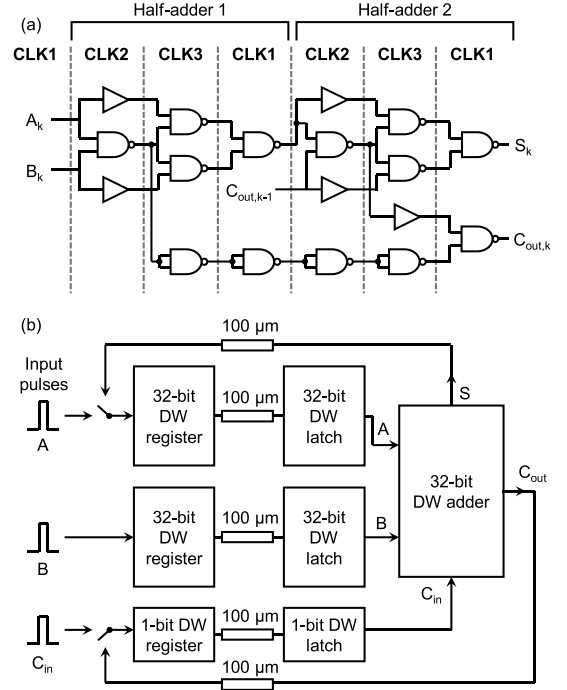


FIGURE 2. (a) Gate-level circuit diagram of the k th 1-bit full adder implemented using DW NAND gates and buffers. (b) DW datapath used for energy benchmarking. The registers are connected to the latches and adder by 100- μm metal interconnects. The registers initially receive their input from CMOS switches that generate the input pulses, and thereafter, from the output of the adder.

For the output MTJ of all devices, we assume a resistance-area product of $RA = 1.0 \Omega \cdot \mu\text{m}^2$ in the parallel state and a fixed length of $L_{\text{MTJ}} = 20$ nm along the ferromagnetic track. Different values of the parallel-state resistance R_p —as necessary to obtain devices of different fan-outs in the logic circuits—are achieved by varying the width of the junction, as shown in Table 1.

As described in [19], the operation of DW-MTJ devices for logic presently relies on the presence of a three-phase clock. We assume that a clock with $V_{\text{CLK}} = 125$ mV is supplied by external CMOS switches with a period of 15 ns and a duration of 2 ns for each clock pulse. A wait time of 3 ns is provided between successive pulses CLK1, CLK2, and CLK3 of the three-phase clock to allow adequate time for the DWs, which possess inertia [20], to settle to their final positions.

B. DW-DEVICE LOGIC CORE

Fig. 2(a) shows the circuit diagram of a DW-MTJ 1-bit full adder implemented using DW NAND and buffer devices of varying fan-out (up to FO4). The full adder circuit is the same as that used in [19] and [20], but with a chain of four buffer elements replaced by four NAND gates operated as inverters. This modification allows the clock terminal of the device to always have at least one low-resistance path to ground, ensuring that each device is reset reliably. A 32-bit full adder is constructed from 1-bit full adders cascaded in a ripple carry scheme: operation of the first half adder for the $(k + 1)$ th bit commences in parallel with the second half adder for the k th bit.

To ensure that a multi-fan-out device distributes current equally among devices of different track widths, we set the sum of the resistance R_{write} (the resistance of the track in the case of STT, and the combined track and HM resistance in the case of SOT) and the tunable series resistance R_{series} to be equal for both the buffer and the NAND gates. This reduces the circuit's sensitivity to the specific values of R_p listed in Table 1.

C. DW-DEVICE DATAPATH

Our energy benchmarking calculations are based on the simulations of the rudimentary processor shown in Fig. 2(b), in which the DW-MTJ adder communicates with DW-MTJ local memory. This system is similar to the 32-bit arithmetic logic unit (ALU) considered in [22] and used for analytical estimations of energy and delay for various emerging logic devices. The primary difference is that we implement only the most time- and energy-expensive arithmetic operation: the 32-bit addition.

The 32-bit adder accesses its inputs by reading from two 32-bit DW-MTJ registers containing the addends A and B, and a 1-bit DW-MTJ register containing the carry input C_{in} . Each register is implemented using a single DW-MTJ buffer that is reset only when its stored bit is accessed in a read or erase operation. The three inputs are first written to the registers via current pulses generated by CMOS switches. Being intrinsically nonvolatile memory elements, the registers can hold their states indefinitely without consuming power.

Upon application of a read-out pulse to its CLK terminal, the register transfers its input across a 100- μm interconnect to a latch implemented using a single fan-out-2 DW-MTJ buffer device, which holds the bit until the rising edge of the CLK1 pulse, at which point it is released to the adder input. The sum output (S) of the adder is written to one of the 32-bit registers and the carry output (C_{out}) is written to the 1-bit register. We set the buffer devices in the registers and the output gates of the adder to be slightly more conductive than fan-out-1 in order to drive sufficient current across the long interconnects. Control signals for memory access are generated by CMOS switches.

D. ENERGY AND DELAY MODELING

Circuit-level simulations of DW-MTJ logic are enabled by SPICE models of the STT DW-MTJ devices [20] and validated by micromagnetic simulations [19] using OOMMF [32]. The SPICE models are implemented in Verilog-A and circuit simulations are performed using the Cadence Virtuoso Spectre simulator [33]. Table 1 lists the model parameters used.

To estimate the energy cost of using SOT-driven DW motion in a strip of PMA material, we use the same SPICE model with a $10\times$ lower threshold current density relative to the STT device but with the same DW velocity (~ 10 m/s) at the reduced current density. We lower the clock voltage to $V_{\text{CLK}} = 12.5$ mV to supply the reduced currents to the SOT devices, with all resistances unchanged from the STT case, listed in Table 1.

We perform all circuit simulations using the same clock timing, leading to the same total delay of 512 ns for the full operation (32-bit read, add, and write). Fig. 3 shows the inputs that are read from the registers and the outputs that are




FIGURE 3. Input and output waveforms for a 32-bit addition obtained from a Cadence Spectre circuit simulation. We use an SOT-driven DW-MTJ with TMR = 200% in the output MTJ. (a) Current that is read out from each of the 32-bit 1-bit DW registers holding the input A. The DW states are read out sequentially starting from the least significant bit. Current that exceeds the threshold of $0.9 \mu\text{A}$ (dashed line) for the buffer element is a logical “1.” (b) Current from the 32-bit 1-bit DW registers holding the input B. (c) Sum output of the 32-bit adder. Individual bits are written to the registers immediately as they are computed. (d) Carry input and carry output bits, which are read out from and written to the 1-bit carry register at the beginning and end of the calculation, respectively. The 32-bit inputs are randomly generated.

written to the registers for the SOT case—the STT case has $10\times$ larger currents. The total delay is measured starting from the rising edge of the input pulse (CLK2) that writes all input bits to the registers. The first input bits (A_0 and B_0) are read out from the registers on the next clock phase (CLK3) after the input write pulse. Subsequently, the bits of the output S are written to the register as they are computed by the adder; a new bit is written every clock cycle, as shown in Fig. 3(c). The computation ends at the falling edge of the pulse (CLK1) that writes the final output bits, S_{31} and C_{out} , back to the registers. The delay of the adder alone is two clock cycles for bit 0 and one clock cycle for every subsequent bit or 495 ns for all 32 bits. Multiple pipelined additions are feasible with a new 33-bit output (32-bit S and 1-bit C_{out}) every clock cycle. We validate the functional operation of the adder using randomly generated 32-bit numbers for the operands, as shown in Fig. 3.

The total energy cost of the operation is calculated from the circuit simulation by integrating the power dissipated in the DW devices, the interconnects, and the input drivers. We do not include in our model the energy consumed in generating and distributing the clock signals to the spintronic devices. Nonetheless, this is potentially a significant energy overhead for spintronic logic circuits that should be investigated in the future work.

III. EVALUATION AND ENERGY PERFORMANCE

A. EFFECT OF TMR and R_p ON ENERGY EFFICIENCY

One way to improve the energy efficiency of the system is to reduce the device OFF-state current by improving the TMR.

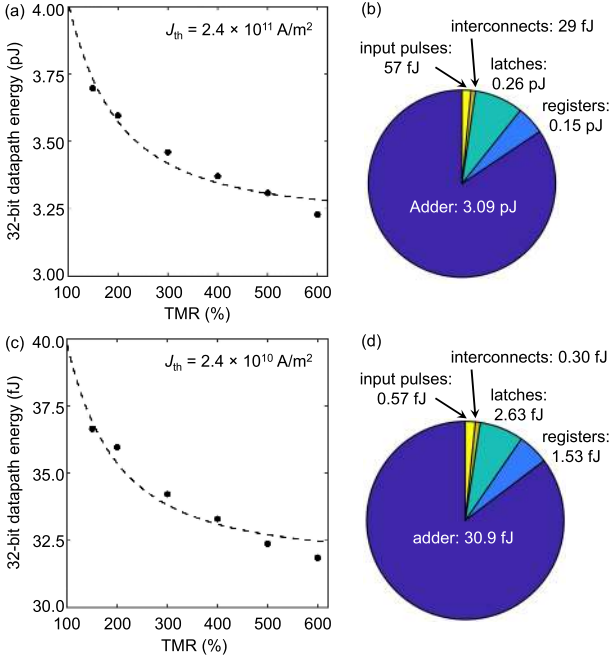


FIGURE 4. (a) Energy cost of the 32-bit STT-driven DW-MTJ datapath as a function of the TMR of the MTJ for the same total delay. The energy savings with increasing TMR follow the dependence given by (1) (dashed line). (b) Share of the energy consumption by the different components of the 32-bit datapath for TMR = 200%. (c) and (d) Corresponding values for the case of SOT-driven DW motion with a 10 \times reduced threshold current density.

Notably, the value of TMR is determined by the style of anisotropy (PMA) and the quality of the growth and fabrication processes, especially at the interfaces of the output MTJ. In principle, because we have set the threshold current of the NAND gate to be 50% larger than that of the buffer, logical operations are possible with a TMR of at least 100%. In practice, because of the lower current swing induced by series resistances (R_{write} and R_{series}), a somewhat larger TMR is necessary. The state of the art for this variety of device ranges between 130% and 200% [34], [35]. However, additional fabrication steps, such as high-temperature annealing, have been used to achieve $\sim 600\%$ TMR in IMA materials [36], showing that there is a path to increasing the PMA TMR. As such, we consider how values in the range from 150% to 600% affect the performance of the system. As we will see in the later section, a larger TMR—which leads to a larger current swing $I_{\text{ON}}/I_{\text{OFF}}$ —provides greater robustness to variability in the threshold current and MTJ resistances.

The total energy consumption of the system using STT devices is shown in Fig. 4(a) as a function of the TMR, and Fig. 4(b) shows the share in energy cost of the different system components for the TMR = 200% case. For this value of TMR, the energy cost of the full datapath is 3.60 pJ, with the dominant component being the adder, which consumes 3.09 pJ. We note that we obtain a substantially larger energy cost for the 32-bit adder than the analytical estimate of 17.3 fJ for STT DW-MTJ devices in [22]. We attribute this in significant part to the much larger clock voltage used in our simulations, which is necessary to accommodate the values of threshold current and MTJ resistances that we have assumed for feature sizes close to 15 nm. The operation of our circuit

TABLE 2. Energy and delay results for DW-MTJ devices.

	STT DW (TMR = 200%)	SOT DW (TMR = 200%)	CMOS High Performance [22]	CMOS Low Voltage [22]
Energy (32-bit)				
Adder:	3.09 pJ	30.9 fJ	19.9 fJ	2.3 fJ
Datapath:	3.60 pJ	36.0 fJ	72.1 fJ*	9.6 fJ*
Delay (32-bit)				
Adder:	495 ns	495 ns	428 ps	4.21 ns
Datapath:	512 ns	512 ns		

*We make a CMOS energy estimate for the datapath in Fig. 2(b) using the expressions and parameters in [22].

also requires additional current pulses than those included in the energy estimates in [22].

For the SOT adder circuit, the 10 \times reduction in current and supply voltage leads to an approximately 100 \times energy reduction, shown in Fig. 4(c) and (d). Table 2 summarizes the results of the SOT circuit simulation and provides a comparison to the STT devices.

The benefit of a larger TMR to the energy cost is modest. This can be readily seen from a simple model of the power consumption, assuming that, on average, half of the current pulses are large (I_{ON}) and the other half are small (I_{OFF})

$$P \sim I_{\text{ON}}^2 + I_{\text{OFF}}^2 \sim 1 + \frac{1}{(1 + \text{TMR})^2}. \quad (1)$$

This relation fits well with the simulation results in Fig. 4(a) and (c) and suggests that the benefit of improving the TMR lies more in providing robustness to interdevice variability effects than in improving energy efficiency.

A different perspective on reducing energy might focus on design choices related to the critical dimensions of the output MTJ, such as to reduce (increase) R_p so that a higher (lower) current flows through the circuit. This creates a direct tradeoff between energy savings and speed, since a higher current moves the DWs more rapidly. Additionally, either smaller critical dimensions [34] or resistance-area product engineering of the oxide layer [37] might equally be used to achieve ultralow switching energy, opening multiple engineering pathways for ultralow power but slower DW logic systems.

B. EFFECT OF INTERCONNECTS

To estimate the energy and latency cost of communication using DW-MTJ logic, we assume that the registers and the adder in our datapath are connected by 100- μm metal interconnects characteristic of the 14-nm CMOS technology node. As shown in Fig. 4(b), we find that the interconnects are responsible for a very small portion of the dissipated energy, largely because the MTJs in the system are more resistive. For devices that operate at the same currents but at a reduced voltage (necessitating lower MTJ resistances), the interconnect energy may play a larger role. In the future, this budget may also be reduced by integrating PMA magnetic nanostraps as low-current interconnect replacements [38]. The additional latency incurred by the interconnect capacitance and the MTJ capacitance is negligible in comparison to the delay associated with translating the DW.

C. BENCHMARK RELATIVE TO CMOS

We compare our results for DW-MTJ logic with analytically estimated values for CMOS devices. Table 2 gives the energy and delay estimates from [22] for high-performance (HP) CMOS in the 2018 technology node (15-nm metal 1 half-pitch). STT-driven DW-MTJ logic is not competitive




FIGURE 5. Projected energy consumption of the 32-bit DW-MTJ adder as a function of the threshold current density. The black dots represent circuit simulation results, assuming the parameter values in Table 1, with the exception of the clock voltage V_{clk} , which we scale linearly with J_{th} . We also use a TMR value of 150%. The CMOS HP and LV values are obtained from [22]. We predict that the DW-MTJ device must attain a threshold current density of $1.9 \times 10^{10} \text{ A/m}^2$ or $6.4 \times 10^9 \text{ A/m}^2$ to achieve energy parity with CMOS HP and LV, respectively.

with CMOS in either energy consumption or latency. However, promisingly, the DW-MTJ logic system using SOT-driven DW motion now lies between HP and low-voltage (LV) CMOS in the datapath energy cost. If the DW velocity can indeed be made faster in these SOT devices than in STT devices at low current densities, as suggested in [31], DW-MTJ logic would become more competitive with CMOS in energy as well as speed.

Based on circuit simulation results at several values of J_{th} , we infer that the energy consumption of the 32-bit adder follows the square of the current density: $E \sim J_{\text{th}}^2$, as shown in Fig. 5. As described previously, the reduction in threshold current is concomitant with a reduction in the clock voltage V_{CLK} by the same proportion, assuming fixed resistances for the magnetic tracks and MTJs. From this model, we predict that a threshold current density for SOT of approximately $J_{\text{th}} = 1.9 \times 10^{10} \text{ A/m}^2$ is needed to achieve energy parity with HP CMOS ($V_{\text{dd}} = 0.73 \text{ V}$), while $J_{\text{th}} = 6.4 \times 10^9 \text{ A/m}^2$ is needed to match LV CMOS ($V_{\text{dd}} = 0.3 \text{ V}$) for the 32-bit adder. In prior numerical studies, a DW velocity of $v = 20 \text{ m/s}$ was predicted with SOT current densities as low as $2.0 \times 10^9 \text{ A/m}^2$ [31], which suggests that additional energy savings are feasible.

D. SENSITIVITY TO DEVICE VARIABILITY

Variation in the fabrication process directly leads to variability in critical device parameters, notably the threshold current I_{th} and the MTJ resistances, which are functions of the device dimensions. Additionally, device parameters, such as the TMR, may be variably reduced on device-by-device basis due to defects, in particular the migration of B atoms to the CoFe/MgO interface [39]. Fig. 6(a) shows the tolerance of the SOT DW-MTJ adder to these fluctuations, which is necessary for reliable logic operation. For these results, we evaluated the accuracy of the 1-bit full adder in Fig. 2(a) over a large number of randomly generated inputs presented sequentially. For each value of mean intrinsic TMR in Fig. 6(a), we first

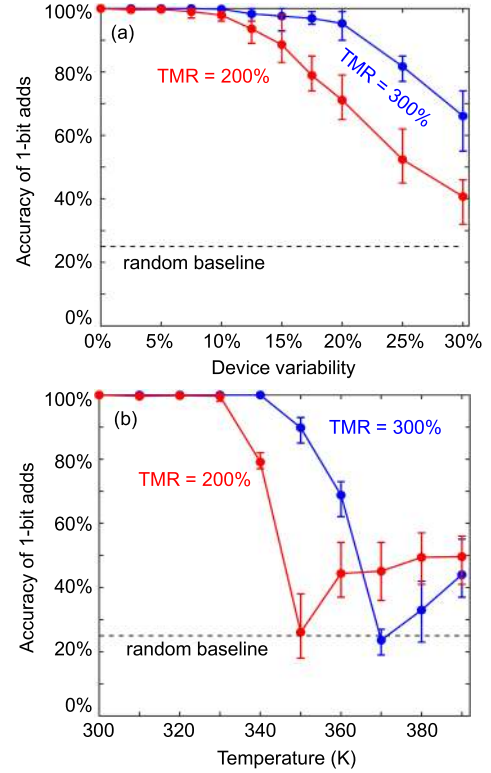


FIGURE 6. (a) Sensitivity of the SOT DW-MTJ adder to variations in threshold current and tunneling magnetoresistance. (b) Sensitivity of the adder to operating temperature, including device-to-device variability of 1%. The mean accuracy is evaluated over 1000 1-bit additions and the error bars indicate the range of the mean accuracies over 100 constituent 1-bit additions.

fine-tuned the values of the desired MTJ resistances R_p in Table 2 to lie at the center of the tolerable range of values without incurring bit errors. Uniform random variation was then introduced to the values of I_{th} and R_p in every device, and a new random value was generated on every clock cycle. Since the accuracy is evaluated on the correctness of two output bits (S and C_{out}), the baseline for random guesses is 25% accuracy.

With a mean TMR of 200% in the MTJs, the circuit can tolerate a device variability of up to 7.5% while maintaining >99.5% accuracy. With a mean TMR of 300%, the tolerable amount of device variability increases to 10% for >99.5% accuracy. A larger intrinsic TMR improves robustness to variability by allowing the circuit to operate at ON-currents well above threshold, while still holding the OFF-currents well below threshold. Since random variations in R_p introduce random fluctuations in the device output currents, these results also suggest some degree of robustness to circuit noise.

E. SENSITIVITY TO TEMPERATURE

The critical DW-MTJ device parameters are known to be functions of temperature. In particular, the current density threshold for DW motion induced by the spin Hall effect has been predicted to decrease with the temperature T as

$$J_{\text{th}}(T) \sim \frac{M_{\text{sat}}(T)H_k(T)}{M_{\text{sat}}(T_0)H_k(T_0)} \quad (2)$$

where M_{sat} is the saturation magnetization, H_k is the perpendicular anisotropy field, and $T_0 = 300$ K [28]. To find both $M_{\text{sat}}(T)$ and $H_k(T)$, we use the temperature model presented in [40]. We also model the temperature dependence of the MTJ tunneling resistance R_p using the model in [40], assuming an oxide thickness of 1.1 nm and an energy barrier of 0.39 eV. A decrease in M_{sat} also influences the velocity of the DWs [20].

Fig. 6(b) shows the accuracy of the DW-MTJ adder when operated at elevated temperatures. The sensitivity to temperature arises from the fact that while the threshold current decreases, the device current increases because of the lower tunneling resistance—this combination leads to bit errors when the device is sufficiently heated. These temperature-induced perturbations to the device truth tables are not random. The decrease and then increase in accuracy at the higher temperatures in Fig. 6(b) correspond first to the conversion of the buffer to an always-high gate, then the NAND gate to the NOR gate.

To cover the maximum temperature range of operation, we set the value of R_p at 300 K to be $\sim 6\%$ higher than the fine-tuned values in Fig. 6(a) for TMR = 200%, and $\sim 13\%$ higher for TMR = 300%. For the same reason as above for variability, a higher TMR leads to greater temperature insensitivity. With TMR = 300%, $>99.5\%$ accuracy is maintained from 300 to 340 K (67 °C), while with TMR = 200%, this is true up to 330 K (57 °C). We note that although we have treated the TMR as a fixed parameter, it is also prone to degrade with temperature, with an absolute decrease by $\sim 20\%$ reported in [41] and [42] from room temperature to 85 °C.

The operational temperature range of DW-MTJ logic can be extended with device and materials engineering. A larger energy barrier in the MTJ, which is also important for state retention in magnetic memory, can reduce the temperature sensitivity of the MTJ resistance [40], though this may come at the expense of a larger resistance-area product. Materials that become demagnetized more slowly with temperature may also provide greater temperature stability to the critical switching current for the spin Hall effect.

IV. DISCUSSION AND FUTURE WORK

For the DW-MTJ datapath in Fig. 2(b), the energy costs are dominated by the adder core rather than by the registers [see Fig. 4(b) and (d)]. On the other hand, based on the methodology in [21] and [22], we estimate that the registers and state elements comprise more than half of the total energy cost of the equivalent CMOS datapath. This leads to an overall advantage in switching energy for the SOT DW-MTJ, as shown in Table 2. Furthermore, the need to retain the states of the SRAM devices in a CMOS processor requires significant standby power not included in Table 2; this significant energy cost is absent in the DW-MTJ registers due to their nonvolatility.

Even though a good representative of combinational logic circuits, the adder is just one component of a modern processor. The fidelity of the CMOS benchmark presented can be improved through accounting for a more general-purpose datapath, which may provide a greater energy advantage. We plan to extend this datapath analysis in the future with the hope of discovering a true crossover point

between HP CMOS and low-energy DW-MTJ logic. Another important area for future analysis is the energy benchmarking of the clocking circuitry and exploration of alternative schemes to a three-phase pulsed clock that can be generated and distributed with greater energy efficiency.

The benefits of DW-MTJ devices may be expanded through logic architectures that take better advantage of device properties. The ability of these devices to store the result of computations without additional memory may allow more efficient implementations of spatial architectures and systolic arrays' architectures that have recently attracted significant interest for neural network and machine-learning workloads [43], [44].

Spatial architectures exploit the data reuse in the matrix-multiply operations at the core of machine-learning inference tasks. The computation is performed within an array of processing elements (PEs), each of which perform a multiply and accumulate (MAC) operation on data that enters the PE and transmits the data to the next PE. The exact energy breakdown of spatial architectures depends on the workload and specific dataflow; however, prior work has found that in dataflows optimized for convolutional neural networks, more than half of the energy is consumed by local storage within the PE [44].

The application of DW-MTJ logic to these architectures may be able to improve the efficiency of these architectures by combining the MAC and result buffer within a PE into a single unit. Additionally, for PEs with larger local storage, DW-MTJ logic can further improve efficiency by reducing reads and writes to this storage. In either case, an even more complicated clocking structure would likely be required; however, efficiency gains may still be realizable as a function of reduced local storage and data movement energy.

Finally, while the present design employs a standard nanotrack, new varieties of DW-MTJ species have since been proposed that utilize either gradients in anisotropy [45] or a shape-based modulation of the nanotrack [46]. Although these effects have so far been used in a neuromorphic context, similar effects could be used to modulate DW motion or increase the efficiency of the STT/SOT current injection via interface engineering. The aim of this engineering would be to reduce critical/switching current, iso-DW speed.

V. CONCLUSION

We have designed and simulated a realistic logic system composed almost entirely of DW-MTJ devices and used the results to evaluate this technology as a candidate for post-CMOS Boolean logic. We find that while STT-driven DW motion in these devices is unlikely to produce systems that are competitive with scaled CMOS in energy efficiency and speed, further advances in these devices can feasibly make them competitive. In particular, since SOT-driven DW motion offers energy-efficiency savings of two orders of magnitude and lower voltage operation relative to STT, DW-MJT logic systems using this approach are within striking distance of optimized CMOS competitor circuits from the perspective of core logic (adder) costs. This result is an important stepping stone toward our goal of benchmarking an in-house optimized CMOS processor against an optimized all-magnetic logic processor and highlights the importance of codesign between device and logic application moving forward.

ACKNOWLEDGMENT

This article describes objective technical results and analysis. Any subjective views or opinions that might be expressed in this article do not necessarily represent the views of the U.S. Department of Energy or the United States Government. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under Contract DE-NA0003525.

REFERENCES

- [1] A. Hirohata and K. Takanashi, "Future perspectives for spintronic devices," *J. Phys. D, Appl. Phys.*, vol. 47, no. 19, May 2014, Art. no. 193001.
- [2] K. L. Wang, J. G. Alzate, and P. K. Amiri, "Low-power non-volatile spintronic memory: STT-RAM and beyond," *J. Phys. D, Appl. Phys.*, vol. 46, no. 7, 2013, Art. no. 074003.
- [3] A. Brataas and K. M. D. Hals, "Spin-orbit torques in action," *Nature Nanotechnol.*, vol. 9, no. 2, pp. 86–88, 2014.
- [4] S. Fukami et al., "20-nm magnetic domain wall motion memory with ultralow-power operation," presented at the IEDM Tech. Dig., Dec. 2013, doi: [10.1109/iedm.2013.6724553](https://doi.org/10.1109/iedm.2013.6724553).
- [5] D. A. Allwood, G. Xiong, C. C. Faulkner, D. Atkinson, D. Petit, and R. P. Cowburn, "Magnetic domain-wall logic," *Science*, vol. 309, pp. 1688–1692, Sep. 2005.
- [6] K. C. Chun, H. Zhao, J. D. Harms, T.-H. Kim, J.-P. Wang, and C. H. Kim, "A scaling roadmap and performance evaluation of in-plane and perpendicular MTJ based STT-MRAMs for high-density cache memory," *IEEE J. Solid-State Circuits*, vol. 48, no. 2, pp. 598–610, Feb. 2013.
- [7] N. Hassan et al., "Magnetic domain wall neuron with lateral inhibition," *J. Appl. Phys.*, vol. 124, no. 15, 2018, Art. no. 152127.
- [8] X. Fong et al., "Spin-transfer torque devices for logic and memory: Prospects and perspectives," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 35, no. 1, pp. 1–22, Jan. 2016.
- [9] K. Ali, F. Li, S. Y. H. Lua, and C.-H. Heng, "Energy- and area-efficient spin-orbit torque nonvolatile flip-flop for power gating architecture," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 26, no. 4, pp. 630–638, Apr. 2018.
- [10] M. Kazemi, E. Ipek, and E. G. Friedman, "Energy-efficient nonvolatile flip-flop with subnanosecond data backup time for fine-grain power gating," *IEEE Trans. Circuits Syst. II, Exp. Briefs*, vol. 62, no. 12, pp. 1154–1158, Dec. 2015.
- [11] B. Behin-Aein, D. Datta, S. Salahuddin, and S. Datta, "Proposal for an all-spin logic device with built-in memory," *Nature Nanotechnol.*, vol. 5, no. 4, pp. 266–270, Apr. 2010.
- [12] B. Behin-Aein, A. Sarkar, S. Srinivasan, and S. Datta, "Switching energy-delay of all spin logic devices," *Appl. Phys. Lett.*, vol. 98, no. 12, Mar. 2011, Art. no. 123510.
- [13] M. Hayashi, L. Thomas, R. Moriya, C. Rettner, and S. S. P. Parkin, "Current-controlled magnetic domain-wall nanowire shift register," *Science*, vol. 320, no. 5873, pp. 209–211, Apr. 2008.
- [14] R. Venkatesan, V. Kozhikkottu, C. Augustine, A. Raychowdhury, K. Roy, and A. Raghunathan, "TapeCache: A high density, energy efficient cache based on domain wall memory," presented at the ACM/IEEE Int. Symp. Low Power Electron. Design (ISLPED), 2012, doi: [10.1145/2333660.2333707](https://doi.org/10.1145/2333660.2333707).
- [15] S. S. P. Parkin, M. Hayashi, and L. Thomas, "Magnetic domain-wall racetrack memory," *Science*, vol. 320, no. 5873, pp. 190–194, 2008.
- [16] K. A. Omari and T. J. Hayward, "Chirality-based vortex domain-wall logic gates," *Phys. Rev. Appl.*, vol. 2, no. 4, 2014, Art. no. 044001.
- [17] D. Morris, D. Bromberg, J.-G. Zhu, and L. Pileggi, "mLogic: Ultra-low voltage non-volatile logic circuits using STT-MTJ devices," presented at the 49th Annu. Design Automat. Conf. (DAC), 2012, doi: [10.1145/2228360.2228446](https://doi.org/10.1145/2228360.2228446).
- [18] J. A. Currivan-Incorvia et al., "Logic circuit prototypes for three-terminal magnetic tunnel junctions with mobile domain walls," *Nature Commun.*, vol. 7, no. 1, 2016, Art. no. 10275.
- [19] J. A. Currivan, Y. Jang, M. D. Mascaró, M. A. Baldo, and C. A. Ross, "Low energy magnetic domain wall logic in short, narrow, ferromagnetic wires," *IEEE Magn. Lett.*, vol. 3, 2012, Art. no. 3000104.
- [20] X. Hu, A. Timm, W. H. Brigner, J. A. C. Incorvia, and J. S. Friedman, "SPICE-only model for spin-transfer torque domain wall MTJ logic," *IEEE Trans. Electron Devices*, vol. 66, no. 6, pp. 2817–2821, Jun. 2019.
- [21] D. E. Nikonov and I. A. Young, "Uniform methodology for benchmarking beyond-CMOS logic devices," presented at the IEDM Tech. Dig., Dec. 2012, doi: [10.1109/iedm.2012.6479102](https://doi.org/10.1109/iedm.2012.6479102).
- [22] D. E. Nikonov and I. A. Young, "Benchmarking of beyond-CMOS exploratory devices for logic integrated circuits," *IEEE J. Explor. Solid-State Comput. Devices Circuits*, vol. 1, no. 1, pp. 3–11, Dec. 2015.
- [23] N. Murray et al., "Field-free spin-orbit torque switching through domain wall motion," *Phys. Rev. B, Condens. Matter*, vol. 100, no. 10, Sep. 2019, Art. no. 104441.
- [24] S. Emori, U. Bauer, S.-M. Ahn, E. Martinez, and G. S. D. Beach, "Current-driven dynamics of chiral ferromagnetic domain walls," *Nature Mater.*, vol. 12, pp. 611–616, Jun. 2013.
- [25] M. Wang, Y. Zhang, X. Zhao, and W. Zhao, "Tunnel junction with perpendicular magnetic anisotropy: Status and challenges," *Micromachines*, vol. 6, pp. 1023–1045, Jun. 2015.
- [26] *International Technology Roadmap for Semiconductors*. Accessed: Sep. 9, 2019. [Online]. Available: <http://www.itrs.net>
- [27] J. Torrejon et al., "Interface control of the magnetic chirality in CoFeB/MgO heterostructures with heavy-metal underlayers," *Nature Commun.*, vol. 5, Aug. 2014, Art. no. 4655.
- [28] L. Liu, O. J. Lee, T. J. Gudmundsen, D. C. Ralph, and R. A. Buhrman, "Current-induced switching of perpendicularly magnetized magnetic layers using spin torque from the spin Hall effect," *Phys. Rev. Lett.*, vol. 109, no. 9, Aug. 2012, Art. no. 096602.
- [29] S. Fukami et al., "Micromagnetic analysis of current driven domain wall motion in nanostrips with perpendicular magnetic anisotropy," *J. Appl. Phys.*, vol. 103, no. 7, 2008, Art. no. 07E718.
- [30] G. S. D. Beach, M. Tsoi, and J. L. Erskine, "Current-induced domain wall motion," *J. Magn. Magn. Mater.*, vol. 320, no. 7, pp. 1272–1281, Apr. 2008.
- [31] A. V. Khvalkovskiy et al., "Matching domain-wall configuration and spin-orbit torques for efficient domain-wall motion," *Phys. Rev. B, Condens. Matter*, vol. 87, no. 2, Jan. 2013, Art. no. 020402(R).
- [32] M. J. Donahue and D. G. Porter, *OOMMF User's Guide, Version 1.0*. Gaithersburg, MD, USA: NISTIR, 1999.
- [33] *Cadence Virtuoso Spectre Circuit Simulator*, Cadence, San Jose, CA, USA, 2009.
- [34] L. Xue, A. Kontos, C. Lazik, S. Liang, and M. Pakala, "Scalability of magnetic tunnel junctions patterned by a novel plasma ribbon beam etching process on 300 mm wafers," *IEEE Trans. Magn.*, vol. 51, no. 12, Dec. 2015, Art. no. 4401503.
- [35] L. Xue et al., "Process optimization of perpendicular magnetic tunnel junction arrays for last-level cache beyond 7 nm node," presented at the IEEE Symp. VLSI Technol., Jun. 2018, doi: [10.1109/vlsit.2018.8510642](https://doi.org/10.1109/vlsit.2018.8510642).
- [36] S. Ikeda et al., "Tunnel magnetoresistance of 604% at 300K by suppression of Ta diffusion in CoFeB/MgO/CoFeB pseudo-spin-valves annealed at high temperature," *Appl. Phys. Lett.*, vol. 93, no. 8, 2008, Art. no. 082508.
- [37] C. Grezes et al., "Ultra-low switching energy and scaling in electric-field-controlled nanoscale magnetic tunnel junctions with high resistance-area product," *Appl. Phys. Lett.*, vol. 108, no. 1, 2016, Art. no. 012403.
- [38] M. Sharad and K. Roy, "Spintronic switches for ultralow energy on-chip and interchip current-mode interconnects," *IEEE Electron Device Lett.*, vol. 34, no. 8, pp. 1068–1070, Aug. 2013.
- [39] A. P. Chen, J. D. Burton, E. Y. Tsymal, Y. P. Feng, and J. Chen, "Effects of B and C doping on tunneling magnetoresistance in CoFe/MgO magnetic tunnel junctions," *Phys. Rev. B, Condens. Matter*, vol. 98, no. 4, 2018, Art. no. 045129.
- [40] M. Kazemi, G. E. Rowlands, E. Ipek, R. A. Buhrman, and E. G. Friedman, "Compact model for spin-orbit magnetic tunnel junctions," *IEEE Trans. Electron Devices*, vol. 63, no. 2, pp. 848–855, Feb. 2016.
- [41] C. Park et al., "Temperature dependence of critical device parameters in 1 Gb perpendicular magnetic tunnel junction arrays for STT-MRAM," *IEEE Trans. Magn.*, vol. 53, no. 2, Feb. 2017, Art. no. 3400104.
- [42] S. G. Wang, R. C. C. Ward, G. X. Du, X. F. Han, C. Wang, and A. Kohn, "Temperature dependence of giant tunnel magnetoresistance in epitaxial Fe/MgO/Fe magnetic tunnel junctions," *Phys. Rev. B, Condens. Matter*, vol. 78, no. 18, Nov. 2008, Art. no. 180411.
- [43] N. P. Jouppi et al., "In-datacenter performance analysis of a tensor processing unit," presented at the 44th Annu. Int. Symp. Comput. Archit., Toronto, ON, Canada, Jun. 2017.
- [44] Y.-H. Chen, J. Emer, and V. Sze, "Eyeriss: A spatial architecture for energy-efficient dataflow for convolutional neural networks," in *Proc. ACM/IEEE 43rd Annu. Int. Symp. Comput. Archit. (ISCA)*, Jun. 2016, pp. 367–379.

- [45] W. H. Brigner *et al.*, “Graded-anisotropy-induced magnetic domain wall drift for an artificial spintronic leaky integrate-and-fire neuron,” *IEEE J. Explor. Solid-State Computat. Devices Circuits*, vol. 5, no. 1, pp. 19–24, Jun. 2019.
- [46] W. H. Brigner *et al.*, “Shape-based magnetic domain wall drift for an artificial spintronic leaky integrate-and-fire neuron,” 2019, *arXiv:1905.05485*. [Online]. Available: <https://arxiv.org/abs/1905.05485>

T. PATRICK XIAO received the B.A. degree in physics and the Ph.D. degree in electrical engineering and computer sciences from the University of California at Berkeley, Berkeley, CA, USA, in 2014 and 2019, respectively. His thesis work explored the emergent properties of ultraefficient optoelectronics and the design of non-von Neumann hardware accelerators for combinatorial optimization problems.

He is currently a Postdoctoral Researcher with Sandia National Laboratories, Albuquerque, NM, USA, where he investigates emerging compute-in-memory devices and their use in both conventional and novel architectures for computing.

CHRISTOPHER H. BENNETT (M’14) received the B.Sc. and M.Sc. degrees from Stanford University, Stanford, CA, USA, in 2011, a joint M.Sc. degree from KU Leuven, Leuven, Belgium and the Chalmers University of Technology, Gothenburg, Sweden, in 2014, and the Ph.D. degree from Université Paris-Saclay, Saint-Aubin, France, in 2018. During his thesis work at the Centre de Nanosciences et Nanotechnologies (C2N), he built hardware learning systems with analog polymeric nanodevices and designed nanoarchitectures for hardware learning.

At Sandia National Laboratories, Albuquerque, NM, USA, he contributes to preindustrial designs of ReRAM accelerators and explores the integration of magnetic memories for logic and online learning applications.

XUAN HU (S’16) received the B.S. degree in electrical and information engineering from Huaqiao University, Xiamen, China, in 2013, and the M.S. degree in electrical engineering from Arizona State University, Tempe, AZ, USA, in 2015. He is currently pursuing the Ph.D. degree in electrical engineering with the Erik Jonsson School of Engineering and Computer Science, The University of Texas at Dallas, Richardson, TX, USA.

His current research interest includes circuit design and modeling of efficient memristive, spintronic, and carbon nanotube logic circuits.

BEN FEINBERG (S’15–M’19) received the B.S. degree in electrical and computer engineering and the M.S. and Ph.D. degrees in electrical engineering from the University of Rochester, Rochester, NY, USA, in 2012, 2014, and 2019, respectively.

He is currently a Postdoctoral Appointee with Sandia National Laboratories, Albuquerque, NM, USA. His research is in computer architecture with an emphasis on memory-centric accelerators, heterogeneous architectures, and energy-efficient data representations.

ROBIN JACOBS-GEDRIM received the B.A. degree in physics from the New College of Florida, Sarasota, Florida, in 2010, and the M.Sc. and Ph.D. degrees in nanoscale engineering from the State University of New York at Albany, Albany, NY, USA, in 2013 and 2015, respectively.

He is currently a Research Associate with Sandia National Laboratories, Albuquerque, NM, USA, with a focus on hardware acceleration of neural computing and beyond CMOS radiation-hardened memory technologies.

SAPAN AGARWAL (M’06) received the B.S. degree in electrical engineering from the University of Illinois at Urbana–Champaign, Champaign, IL, USA, in 2007, and the Ph.D. degree in electrical engineering from the University of California at Berkeley, Berkeley, CA, USA, in 2012.

He is currently a Senior Member of Technical Staff at Sandia National Laboratories, Albuquerque, NM, USA. He is interested in everything from explainable machine learning and neuromorphic hardware to semiconductor and photonic devices.

JOHN S. BRUNHAVER (S’04–GS’08–M’13) received the bachelor’s degree in electrical and computer engineering from Northeastern University, Boston, MA, USA, in 2008, and the master’s and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, USA, in 2011 and 2015, respectively, where his thesis, written as part of his doctoral work, is titled “The Design and Optimization of a Stencil Engine” and examines the procedural generation of hardware for image processing and image understanding.

In 2015, he joined the Arizona State University faculty, Tempe, AZ, USA, as an Assistant Professor in electrical computer and energy engineering with the School of Electrical, Computer and Energy Engineering. His current research focuses on the design of energy-efficient computer architectures and the design automation techniques for implementing them.

JOSEPH S. FRIEDMAN (S’09–M’14–SM’19) received the A.B. and B.E. degrees from Dartmouth College, Hanover, NH, USA, in 2009, and the M.S. and Ph.D. degrees in electrical and computer engineering from Northwestern University, Evanston, IL, USA, in 2010 and 2014, respectively.

He joined The University of Texas at Dallas, Richardson, TX, USA, in 2016, where he is currently an Assistant Professor of electrical and computer engineering and the Director of the NeuroSpinCompute Laboratory. From 2014 to 2016, he was a Research Associate with the Centre national de la recherche scientifique, Institut d’Electronique Fondamentale, Université Paris-Sud, Orsay, France. He has also been a Summer Faculty Fellow at the U.S. Air Force Research Laboratory, Rome, NY, USA, a Visiting Professor with the Politecnico di Torino, Turin, Italy, and a Guest Scientist with RWTH Aachen University, Aachen, Germany. He worked on logic design automation as an Intern at Intel Corporation, Santa Clara, CA, USA. His current research interests include the invention and design of novel logical and neuromorphic computing paradigms based on nanoscale and quantum mechanical phenomena with particular emphasis on spintronics.

Dr. Friedman is a member of the Editorial Board of the *Microelectronics Journal*, the technical program committees of the Design Automation Conference (DAC), Design, Automation and Test in Europe Conference (DATE), SPIE Spintronics, NANOARCH, the Great Lakes Symposium on VLSI (GLSVLSI), and the IEEE International Conference on Electronics Circuits and Systems (ICECS), the Review Committee of ISCAS, and the Nanoelectronics and Gigascale Systems Technical Committee of the IEEE Circuits and Systems Society. He has been a member of the Organizing Committee of NANOARCH 2019 and DCAS 2018. He was a recipient of the Fulbright Postdoctoral Fellowship.

JEAN ANNE C. INCORVIA (M’11) received the B.A. degree in physics from the University of California at Berkeley, Berkeley, CA, USA, in 2008, and the M.A. and Ph.D. degrees in physics from Harvard University, Cambridge, MA, USA, in 2012 and 2015, respectively.

She was a Postdoctoral Research Associate with Stanford University, Stanford, CA, USA, from 2015 to 2017. In 2017, she joined the Department of Electrical and Computer Engineering, The University of Texas at Austin, Austin, TX, USA, as an Assistant Professor. Her technical contributions include magnetic compute-in-memory devices and circuits, and device and materials research in novel materials, including low-dimensional materials and emerging memories.

MATTHEW J. MARINELLA (SM’17) received the Ph.D. degree in electrical engineering from Arizona State University, Tempe, AZ, USA, in 2008, under the supervision of D. Schroder.

He is currently a Principal Member of the Technical Staff with Sandia National Laboratories, Albuquerque, NM, USA. He is also a Principal Investigator for Sandia’s Nonvolatile Memory Program and numerous neuromorphic and low-power computing projects.

Dr. Marinella is the Chair of the Emerging Memory Devices Section for the IRDS Roadmap Beyond CMOS Chapter and serves on various technical program committees, including the IEEE International Conference on Rebooting Computing.