

Toward Comprehensive Shifting Fault Tolerance for Domain-Wall Memories with PIETT

Sebastien Ollivier, *Student Member, IEEE*, Stephen Longofono, *Student Member, IEEE*, Prayash Dutta, *Student Member, IEEE*, Jingtong Hu, *Senior Member, IEEE*, Sanjukta Bhanja, *Senior Member, IEEE*, Alex K. Jones, *Senior Member, IEEE*,

Abstract—Spintronic domain-wall memories (DWMs) offer improved memory density and energy compared to conventional memories but are susceptible to shifting faults. We propose PIETT (Pinning, Insertion, Erasure, and Translation-fault Tolerance) for improved misalignment correction versus the state of the art. PIETT proposes a derived error correction combined with multi-domain access approach to detect and correct a minimum of three misalignment faults after an arbitrary shift distance. Moreover, we characterize the rate of both misalignment and pinning faults in DWM nanowires and demonstrate that pinning faults are a significant concern to DWM. As such, PIETT is the first method combine correction of misalignment and pinning faults in random access DWMs. It also introduces novel PIETT Transverse Access Points (TAPs) that utilize a novel write access mode which can set/reset multiple domains in a single intrinsic operation and can store shift distance detection codes. By allowing checks between shifts of the intrinsic shift distance (*e.g.*, 3 domains), using a single TAP per nanowire expands misalignment protection and determines the needed corrective shifts to correct faults in all nanowires. Two TAPs expands misalignment protection to correct misalignment by more than one position and detects pinning by detecting different shift distances at each extremity of the nanowire. PIETT leverages knowledge of pinned nanowire locations to guide a modified SECDED ECC with one additional parity bit stored in additional parity nanowires. Thus, PIETT in TAP mode can correct unlimited, potentially multi-position, misalignment faults and either up to three pinning faults or up to two pinning faults with up to one bit-flip fault using scrubbing. PIETT provides eight to 21 orders of magnitude improvement in mean-time-to-failure with similar or better area overhead and only a 1% system performance degradation compared to state of the art DWM misalignment correction.

Index Terms—Fault tolerance, spintronic memory, fault modeling, error correction codes



1 INTRODUCTION

Spin-Transfer Torque Magnetic memory (STT-MRAM) has gained traction for on-chip memory deployment due to its near-SRAM performance, CMOS compatibility, low static power, and good endurance [1]. Unfortunately, STT-MRAM has insufficient density for main memory or secondary storage applications. Spintronic domain-wall memory—also referred to as “Racetrack” memory—originally proposed and demonstrated by IBM [2], [3], retains the static energy benefits of STT-MRAM with a 10× density improvement [4]. DWM has a theoretical area per bit as small as $2F^2$, where F is the technology feature size [5]. Moreover, DWM avoids endurance challenges by providing $\geq 10^{16}$ write cycles [6] compared to other emerging memory candidates such as phase-change [7] and resistive [8] memories at $10^8 - 10^9$ and $10^{11} - 10^{12}$ write cycles, respectively [6], [9].

DWM is constructed from ferromagnetic nanowires—also referred to as *tapes* or *racetracks*—separated into domains and connected to one or more access transistor(s) to create access ports. Data is stored by magnetic orientation and accessed by *shifting* the magnetic domains along the nanowire and *aligning* the target domain to a fixed access device [2], [10]. After alignment, data

access is similar to an STT-MRAM Magneto-Tunnel Junction (MTJ). Thus, DWM has been proposed for non-uniform access structures like Non-Uniform Cache Access (NUCA) caches [11].

Unfortunately, slight fluctuations in shifting current can cause shifting faults. These faults include misalignment and pinning faults. Misalignment takes the form of *over-* and *under-shifting*, ranging in frequency from $5 \cdot 10^{-5}$ to 10^{-3} depending on shift distance [12]. Pinning occurs due to imperfections in the domain wall caused by process variation. It can most commonly manifest as an *erasure*¹ where the pinning point functions as a barrier that prevents shifting within the nanowire [13], [14]. Theoretically, an *insertion* may also be possible where the pinning point replicates itself and shifting continues through the whole nanowire. Either pinning fault puts the nanowire in an unrecoverable state.

In memory structures created from DWMs, multiple racetracks are bundled, accessed in parallel, and shifted together [15]. In the bundle, additional racetracks storing Error Correction Codes (ECC) could be added to correct the data perturbed from misalignment or pinning faults. Unfortunately, this form of ECC alone is insufficient to determine when a shifting fault has occurred or to guide its correction. ECC cannot detect faults occurring in part of the nanowire not being read or when the faulty data matches the expected parity value, *e.g.*, when neighboring data contains the same value. Additionally, fault discovery provides no insight into

1. Erasure in this context is a different meaning than in a multi-bit error correction code such as Reed Solomon or low density parity-check “erasure” codes. Instead, erasure is analogous to an erasure (or deletion) in communication theory where a bit is dropped.

- S. Ollivier, S. Longofono, J. Hu, and A. K. Jones are with the Department of Electrical and Computer Engineering, University of Pittsburgh, PA, 15261. E-mail: see <http://www.michaelshell.org/contact.html>
- P. Dutta and S. Bhanja are with the Department of Electrical Engineering, University of South Florida, FL 33620.

This work was partially supported by the laboratory of physical sciences (LPS), NSA, and NSF award 1619027, 1822085.
Manuscript received August 24, 2021

the type of fault, such as misalignment, pinning, or even a bit flip, as each nanowire is only sampled at a single point.

Several recent approaches have been proposed to mitigate misalignment in DWMs. Hi-fi proposes a Johnson code stored in additional synchronization domains to detect alignment [12]. This can result in significant area and performance overheads due to the additional domains and access ports required. Greenflag proposed to correct misalignment using communication theory [16] which was later extended as Foosball to add single bitflip protection [17]. Unfortunately, these approaches require the entire nanowire be accessed in sequence making it unsuitable for implementation of random access memory. Moreover, none of Hi-fi, Greenflag, or Foosball can correct pinning faults.

To provide a more complete solution, we propose PIETT, or **P**inning, **I**nsertion, **E**rasure, and **T**ranslation-fault **T**olerance, to correct faults from misalignment and pinning. PIETT has a high-performance method to correct only misalignment faults during shifting using a Derived Error Correction Coding (DECC) methodology. PIETT-DECC uses a Multi-Domain Reading (MDR) methodology [18], [19], [20] that can determine the number of 1's in multiple adjacent domains in the nanowire. PIETT-DECC uses MDR to access the data *signature*, or number of 1's in the data domains, and stores 1's in the overhead domains to the right of the data domains to record the nanowire position. DECC stores external parity bits to the signature to detect and correct these misalignment faults.

In the presence of both misalignment and pinning faults, PIETT extends the MDR concept to introduce special *Transverse Access Points* (TAPs) deployed in extended padding bits at both extremities of the nanowire and uses them to detect shifting faults. A TAP, conceptually akin to a STT-MRAM Multi-Level Cell (MLC) with t free layers, is constructed with t domains of the nanowire. In one shift operation, all t domains can be preset to '1's or reset to '0's and the number of '1's can be determined with MDR. To detect faults, prior to a shift, both TAPs are reset to a known state and read after the shift. If the shift occurred successfully both TAPs will report the correct alignment state. If there is misalignment, the TAPs will both report the same incorrect alignment state and the nanowire can be correctively shifted. If pinning occurs, it is detected with mismatched TAP alignment.

In this mode, PIETT can independently correct an unlimited number of misalignment faults *including* multiposition misalignment. Using SECDED or Single Error Correction Double Error Detection ECC parity nanowires to protect a group of racetracks, PIETT can correct at least three pinning faults within this group. PIETT supports multi-domain intrinsic shifts and is compatible with bit flip correction, correcting two pinning faults combined with a bit flip fault. To the best of our knowledge, this paper is the first paper first scheme to detect and correct both misalignment and pinning faults in DWM memories.

In particular, we make the following contributions:

- We estimate the shift and pinning fault probability from process variation of domain wall notch width and depth, characterized using micromagnetic device simulation.
- We propose DECC which leverages '1's counting to detect and correct at least three misalignment only faults after arbitrary shift distances in the nanowire.
- We propose TAPs which introduce multi-domain shift-writing and leverage MDR within a DWM.
- We demonstrate how TAPs combined with padding bit encoding can be used to detect alignment or pinning faults

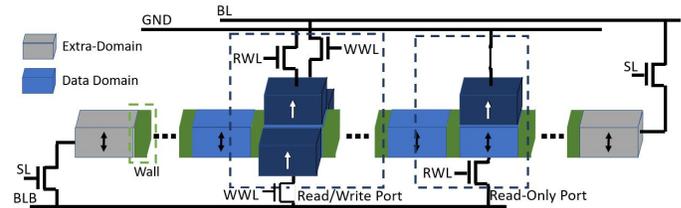


Fig. 1. Anatomy of a DWM nanowire [21].

and directly used to correct misalignment through corrective shifts.

- We demonstrate directed scrubbing based on SECDED ECC guided by TAP-based pinning detection to correct up to three simultaneously pinned nanowires or up to two pinned nanowires in the presence of up to one bit flip fault per data location.
- We provide a detailed analysis of PIETT to evaluate the fault tolerance, performance, energy, and area overheads for a range of incident pinning fault rates.

DECC provides similar fault tolerance to Hi-fi [12] while providing area improvement and more than 50% reduction in dynamic energy. When considering pinning faults, PIETT provides 21 orders of magnitude improvement in mean-time-to-failure based on the 10^{-8} pinning fault rate determined by our model and scales well to higher fault rates, multiposition alignment faults, and longer nanowires. PIETT does increase shift latency, but has only a 1% system performance degradation. PIETT corrects misalignment and pinning with a similar area overhead to fault tolerance schemes with merely misalignment protection.

The remainder of this paper is organized as follows. Section 2 presents more detail on DWM, its shifting challenges, relevant novel access modes, leading solutions for mitigating shift faults, and other related work. The derived error correction mode of PIETT to solve misalignment faults is presented in Section 3. Section 4 explores pinning faults explaining the theory and presenting magnetic simulation results for pinning fault probability. TAPs are described in detail in Section 5. Section 6 demonstrates how PIETT can detect and correct misalignment and pinning with TAPs. The experimental setup and reliability, area, performance, and energy results of PIETT are described in Sections 7 and 8, respectively. Finally, we relate conclusions in Section 9.

2 BACKGROUND AND RELATED WORK

An example of a planar (2D) DWM nanowire with shift write ports is shown in Fig. 1 [21]. The value of each domain is determined by its polarization and illustrated by arrow direction. During read access, a domain is aligned with the fixed layer (dark blue) of the access port. The resistance is detected by a current applied orthogonally through the nanowire across the fixed access port layer. Like STT-MRAM, the resistance is lower if polarization is the same direction as the fixed layer (parallel) and higher if polarization is opposite (antiparallel). Writing uses a much (often an order of magnitude) larger current. Alternatively, shift writing, shown in Fig. 1 at the read/write port, can improve both the speed and energy of writing [21].

An example of DWM data access is shown in Fig. 2. The cross section of the R/W port using shift-based writes is shown in Fig. 2(a) [21]². Presuming the nanowire starts in the center position,

² Note, the design shown in Figs. 1 and 2(a) differ from prior work [21] by adding a second WWL (T3) because, while not needed for reading/writing, it is needed for correct shifting to prevent sneak paths between BLB and BL.

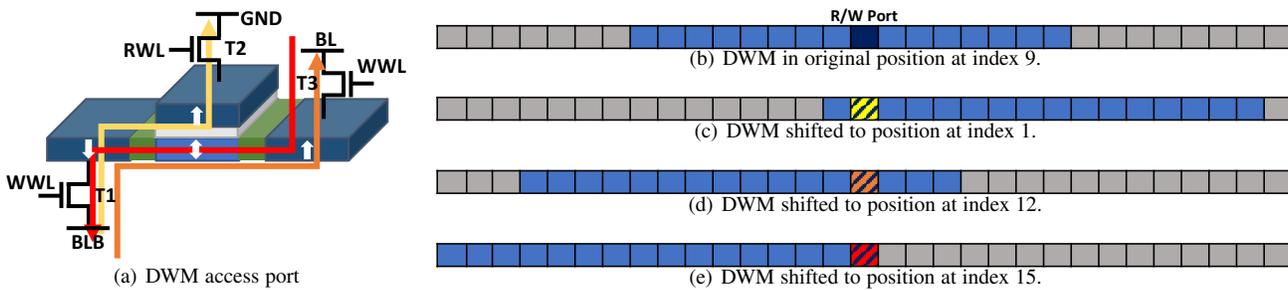


Fig. 2. DWM read and write example with intermediate shifting.

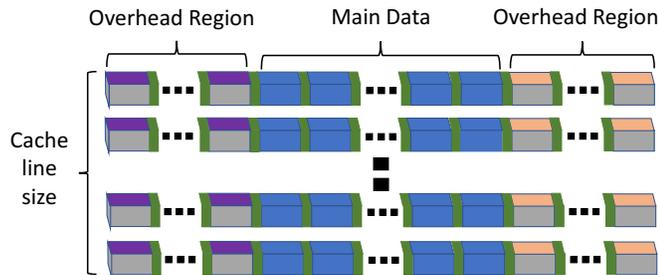


Fig. 3. Domain block cluster example.

index 9 (numbering starting from 0) is aligned to the access point [Fig. 2(b)]. To service a read request to index 1, data must be shifted right by 8 domains which is accomplished by turning on SL allowing current to flow from BLB to BL (see Fig. 1). Once aligned with index 1 [Fig. 2(c)] the example reads by applying a current from BLB to BL traversing WWL (T1) and RWL (T2) shown in yellow. Next to write a '1' to index 12, [Fig. 2(d)] data is shifted left by opening SL and reversing current flow from BL to BLB followed by applying current from BLB to BL through WWL (T1 and T3) which shifts the fixed antiparallel domain into the free layer shown in orange. Finally, a '0' is written into index 15 [Fig. 2(e)] through three more left shifts and reversing the write/shift current as shown in red.

DWM demonstrations of memory array structures [22] and Content Addressable Memories (CAMs) [23] demonstrate fabrication feasibility with great potential for density, performance, and power consumption. Moreover, DWM technology has been proposed for utilization in a variety of positions in the memory hierarchy, including network-on-chips [24], as part of the cache hierarchy representing the last-level cache [11] and multiple cache levels including L1 [25], in GPGPU registers [26] and caches [27], and as a fast main-memory technology [28].

DWM-based memories typically use a traditional hierarchical memory organized into ranks, banks, sub-arrays, tiles, etc. Because a bundle of nanowires contains multiple rows/words of data whose width is determined by the number of nanowires in the bundle, it is treated as a *domain block cluster* [29], [30] or DBC as shown for a cache line granularity in Fig. 3. Thus, data accessed from the memory can directly select the appropriate DBC in the peripheral circuitry, but to access the actual row/word requires shifting all the nanowires for alignment with the access point.

2.1 Shifting Faults

While shifting the DBC, one (or more) nanowires may experience an over- or under-shift misalignment fault and/or a pinning fault.

2.1.1 Misalignment Faults

Misalignment faults, typically due to fluctuations in the shifting current [2], occur due to variation in the operating conditions of

the system. In this case, the entire nanowire over- or under-shifts.

2.1.2 Pinning Faults

Unlike misalignment faults, pinning faults manifest due to operating conditions combined with fabrication imperfections, *i.e.*, where the nanowire is not formed properly due to variations in the process. As discussed in Section 1, pinning can take the form of an erasure where shifting stops in the pinning point of the nanowire [13] or as an insertion where the value is replicated at the pinning point. These behaviors occur when the shifting current is deflected to be near the lower or upper bound of tolerance and a variation defect has impacted the local domain-wall.

When a defect causes an erasure fault, the domain motion stops at the pin point and can be overwritten by the domain that follows. We provide a conceptual example of this fault in Fig. 4(b). When shifting from position (a)(i) and expecting to reach position (a)(ii), *i.e.*, a shift to the left, one bit, d_2 , disappears at the pin point (shown in red) and the remaining domains in the nanowire stop moving.

In the case of an insertion fault, the domain motion for all domains starts at the same speed, however, as they interact with a defect the distance traveled is affected. When sufficiently stretched a replicated (inserted) domain is created. We show this conceptually in Fig. 4(c). The domain at the pin point (d_3) becomes pinned and replicates itself into the adjacent location. Both types of pinning can be detected because the domain motion at the extremities of the nanowire will appear as having different alignments.

2.2 Misalignment Fault Tolerance

Two main techniques have been proposed to detect and correct misalignment, one based on a dedicated code and access points (Hi-fi) [12] and one based on data encoding using Varshamov-Tenegolts (VT) codes (GreenFlag/Foosball) [16], [17]. Hi-fi, like PIETT, targets 2D random access DWM memories with DBCs like Fig. 3.

2.2.1 Hi-fi

Hi-fi presents two techniques, p-ECC and p-ECC-O, which leverage additional access points and encoding techniques for misalignment

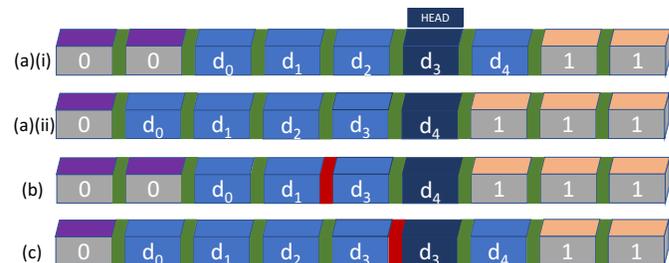


Fig. 4. Pinning example shifting from position (a)(i) expecting to arrive in position (a)(ii) where (b) is an example of erasure and (c) is an example of insertion. Pinned domain-walls shown in red.

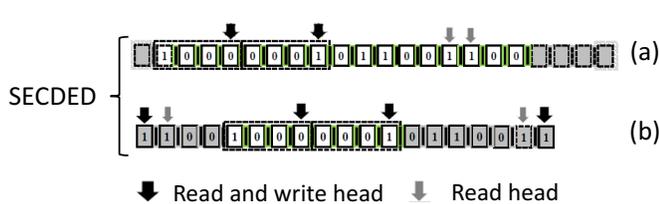


Fig. 5. Hi-fi fault correction (a) p-ECC (b) p-ECC-O [12]. Data bits shown in white with dashed line bounding box. Padding bits shown in gray. Additional encoding bits for p-ECC shown in white.

detection and correction. Fig. 5 shows a SECCED for misalignment example for both approaches. Hi-fi corrects faults by encoding the auxiliary domains with a pattern of alternating groups of two ‘1’s and two ‘0’s. Using the two adjacent read heads, the system can compare two values from the auxiliary bits and compare it against the expected system state. For example, if the system was expecting to read “00” but rather read “01” the tape is misaligned one position too far left. Similarly, reading “10” would signify one position too far right. Reading “11” would indicate misalignment by two, but not which direction.

The main difference between the two Hi-fi techniques is the location where the auxiliary information is stored in the racetrack. In Fig. 5(a), p-ECC adds dedicated domains and two additional associated read-only ports to access the information, but accommodates multiple shifts between checks. In contrast, p-ECC-O, shown in Fig. 5(b), uses the already necessitated extra padding domains for auxiliary information. Unfortunately, one read and one write head are required at each end of the device to maintain and check the pattern, which only allows a single shift between checks.

Both schemes may be scaled to detect bit misalignment by two or more steps by modifying the code and the number of read heads for the auxiliary information. N -domain misalignment correction with $N+1$ -domain misalignment detection requires a total of $N+1$ read ports.

2.2.2 GreenFlag and Foosball

In GreenFlag [16] reading requires the entire nanowire to be read in sequence, requiring a shift and read operation to access each data bit. If an undershift occurs a bit is read twice and if an overshift occurs a bit is lost, similar to what could happen in a communication channel. GreenFlag uses the VT codes and delimiters to recover missing bits and eliminate redundant bits. Thus, when writing much of the nanowire must be rewritten with the new encoding. Foosball [17] extends GreenFlag with a new 8-bit delimiter capable of detecting a misalignment of up to two domains and a bit flip by adding parity nanowires.

2.2.3 Suitability for Pinning Protection

Unfortunately, neither Hi-fi nor Foosball handle pinning faults. Like bit flips and unlike misalignment, pinning is actually destructive as it changes the data stored in the nanowire making it particularly difficult to correct. Foosball does handle bit-flips, but it does not address pinning. Moreover, it requires the assumption that for each the entire nanowire is accessed in sequence. PIETT is designed for 2D planar DBC structures that support parallel access. As Foosball’s access mode is approximately $18\times$ slower and higher energy than these DBCs, we focus on comparisons with more closely related techniques that also target similar DBCs like P-ECC and P-ECC-O.

We were unable to find an obvious way to adapt P-ECC to detect or correct pinning. For P-ECC this can be easily grasped from

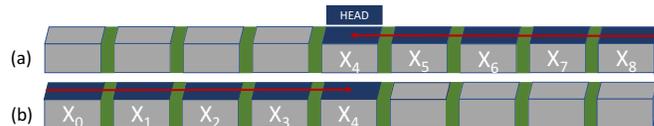


Fig. 6. Shows a transverse read (a) from the right to the access port and (b) from the left to the access port.



Fig. 7. Multi-domain magneto-tunnel junction.

Fig. 5(a) as the alignment domains (shown in white outside of the dotted-line box) are at only one place along the nanowire. Pinning will see different alignments at different parts of the nanowire. P-ECC-O has potential to detect pinning because it adds access points to both ends of the nanowire. Although considerably faster than Foosball, P-ECC-O is slower than all PIETT operation modes and P-ECC because it limits the system to intrinsic shifts by one.

2.3 Multi-Domain Reads

Multi-domain reads determine the number of parallel or anti-parallel domains in a segment of a DWM nanowire. The first technique proposed to implement this function for DWMs is called a transverse read (TR) [19]. TR applies a smaller current in the same direction as the shift current through a portion of the nanowire as shown in Fig. 6. The current is initiated at the end of the nanowire (as shown in the figure) or at an access point and exits through the MTJ of an access point. This allows an access akin to multi-level STT-MRAM cell where multiple free layers are stacked on top of a single fixed layer. Thus, the tunneling magnetoresistance (TMR) of multiple domains impacts the voltage sensed at the access port due to changes of the resistance state. TR has been demonstrated to distinguish the number of parallel or anti-parallel domains within four adjacent domains into different resistance groups [19].

While MDR may also be measurable through the Anomalous Hall Effect (AHE) [31], [32], recently, a multi-domain MTJ was proposed as a scalable alternative to TR for MDR [20]. The multi-domain MTJ creates an access port across multiple domains as shown in Fig. 7. When a read current is applied, each of the domains function as parallel resistors allowing for different resistance levels based on the number of parallel and anti-parallel domains. This work demonstrates resilience to process variation and scalability to seven domains [20].

For an MDR in an arbitrary nanowire segment, $R_{MD} = \sum_{i=0}^{D-1} X_i$. To determine R_{MD} using TR for the system in Fig. 6 the right TR, TR_R , produces $TR_R = \sum_{i=4}^8 X_i$ [Fig. 6(a)]. Similarly, $TR_L = \sum_{i=0}^4 X_i$ [Fig. 6(b)]. X_4 is included in both TR_L and TR_R . Thus, $R_{MD} = TR_R + TR_L - X_4$, requiring two TRs and a standard read. By placing access points appropriately, ‘1’s can be determined through parallel segmented TRs for an arbitrarily long segment of the nanowire in three steps [19]. To determine R_{MD} using multi-domain MTJs, for the configuration in Fig. 7 we can directly obtain the function $R_{MD} = \sum_{i=0}^4 X_i$. By alternating placement of the MTJs on the top and bottom of the nanowire, ‘1’s can be determined in an arbitrarily long segment of the nanowire in two steps.

In Section 4 and beyond we characterize the pinning fault probability for a representative DWM nanowire and discuss a methods to detect and correct pinning faults. However, first we describe PIETT’s advancement in misalignment only shift correction using derived error correction in the next section.

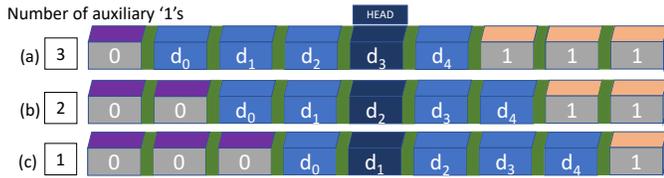


Fig. 8. Auxiliary bit ones for different positions

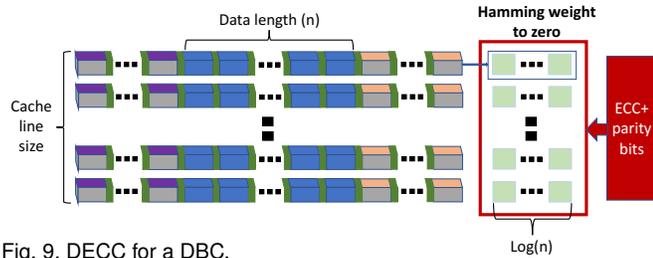


Fig. 9. DECC for a DBC.

3 PIETT WITH DERIVED ERROR CORRECTION

DECC relies on MDR to count the number of ‘1’s in a segment of the nanowire. Encoding of the values stored in the padding bits can report the position of the nanowire. In DECC, each nanowire is constructed with a fixed domain representing a ‘1’ on the right end and another representing a ‘0’ on the left end. Thus, during left and right shifts, appropriate ‘1’s and ‘0’s are shifted into the padding bits on the right and left sides of the nanowire, respectively. The number of ‘1’s indicates the position of the data within the nanowire. As a result, if an under- or over-shift fault occurs, the calculated number of ones will differ from the expected value. Using the difference from the expected value, the fault can be detected and ultimately corrected.

A DECC example is shown in Fig. 8 where the data bits d_i are shown in blue and the data bit aligned with the access port is shown in navy (dark blue). The padding bits on the left side (purple) contain ‘0’s and the right side (beige) contain ‘1’s. The position of the tape corresponds to the number of ‘1’s in the padding bits. DECC uses an MDR to check the number of ‘1’s. It then validates and, if necessary, corrects the alignment.

Consider the case where the racetrack begins in position $\boxed{1}$ [Fig. 8(a)] and attempts to shift to the left by one position to match position $\boxed{2}$ [Fig. 8(b)]. The total number of ‘1’s prior to the shift is $TOT = R_{MD}$. After the shift, the new total TOT' should decrease by one. If an under-shift occurs, $TOT' > TOT - 1$ requiring a left shift to balance the equation. If an over-shift occurs, the tape moves to position $\boxed{3}$ [Fig. 8(c)] and $TOT' < TOT - 1$ and a right shift should correct the misalignment.

3.1 Three Misalignment Correction Guarantee

TOT is the Hamming weight from ‘0’ (H0) of the data bits defined as $\sum_{i=0}^4 d_i$ and the position of the racetrack defined by the ‘1’s in the auxiliary bits. Thus, $TOT - H0$ can be used to verify the racetrack position. We define the H0 as a data *signature*. Rather than storing the signature using $\log_2(n)$ bits for each racetrack, they are created *on demand* after a shifting operation has concluded. We store parity bits and SECDED ECC of the generated signatures in the DBC using STT-MRAM auxiliary bits as shown in red in Fig. 9.

The method to protect three single domain misalignments using this parity information is described in Fig. 10. We use reflected binary Gray codes to represent the signature to ensure that if the shift alignment is off by only one, the signature is only different by one bit. Thus, the parity bits detect misalignment by one position and the ECC is used to repair the signatures of the misaligned

racetracks to guide corrective shifts where where LCL is the length of a cache line [Fig. 10(a)]. One (or more) single misalignment errors with signature deviations in independent columns can all be detected and repaired as shown in Orange in Fig. 10(b). SECDED detects the presence of two errors in the same column and their location is dictated by the parity bit as shown in Purple in Fig. 10(c). This works in the case of a third error in another column shown in Orange. If there are three errors affecting the same SECDED column the ECC correction may point to the wrong location shown in red, but the three parity bits will guide the location of errors for correction as shown in Fig. 10(d).

3.2 Pinning

The signature DECC uses to determine misalignment cannot be guaranteed to change if pinning occurs. For example, consider the pinning examples in Fig. 4. If an erasure occurs, the value d_2 is lost and the signature is expected to be incremented by ‘1’ due to the left shift. However, if d_2 is ‘0’, from an MDR Fig. 4(a)(ii) is indistinguishable from Fig. 4(b) and DECC will not detect a fault. Similarly for insertion faults, in DECC Fig. 4(c) will be indistinguishable from Fig. 4(a)(ii) if d_3 is ‘1’.

We include additional details on DECC including a synthetic uncorrectable fault limit for DECC in a preliminary version of this paper [18]. However, in the next section we demonstrate the presence of runtime pinning faults followed by a discussion of how PIETT improves upon DECC to correct these pinning faults.

4 PINNING FAULT MODELING

To create the domain walls that separate domains in a DWM nanowire, equally spaced fabricated notches are introduced to create pinning sites. The strength or *pinning potential* of a pinning site depends on the geometry of the notch, which can be modeled as described in Eq. 1 where q_{pin} is the pinning site, V_{pin} is the pinning potential at that particular location and M_s is the saturation magnetization of the material used [33], [34], [35]. σ_d is the domain-wall width [36], [37] and E_{pin} is the notch energy density [33] presented in Eq. 2 where A_{ex} , K_u , a , and M are the exchange coefficient, magneto-crystalline anisotropy, material lattice constant, and magnetization amplitude, respectively. A current pulse with adequate amplitude governed by the pinning potential can depin the wall from the notch positions and cause it to travel along the nanowire to the next pinning site. This is governed by the Landau-Lifshitz-Gilbert (LLG) equation [38] in Eq. 3 where H_{eff} , α , γ , and β are the effective field, Gilbert damping constant, gyromagnetic ratio, and non-adiabatic spin-torque coefficients, respectively.

$$V_{pin} = \frac{2M_s E \sigma_d}{q_{pin}(q - q_{pin})^2} \begin{cases} E = E_{pin}, & -\sigma_d \leq q \leq q_{pin} + \sigma_d \\ E = 0, & \text{otherwise} \end{cases} \quad (1)$$

$$\sigma_d = \pi M \sqrt{\frac{2A_{ex}}{K_u a^3}} \quad \text{and} \quad E_{pin} = A_{ex} M^2 \frac{\pi^2}{a \sigma_d} + \frac{\sigma_d K_u}{2} \quad (2)$$

$$\frac{d\vec{M}}{dt} = -\gamma \vec{M} \times \vec{H}_{eff} + \alpha \vec{M} \times \frac{d\vec{M}}{dt} - v_j \frac{\partial \vec{M}}{\partial x} + \beta v_j \vec{M} \times \frac{\partial \vec{M}}{\partial x} \quad (3)$$

To examine the impact of variation, we studied a nanowire with 16 domains where each domain was 200nm long, the full nanowire is 3200nm and the width and thickness were set to 100nm and 4nm, respectively. The material properties are listed in Table 1. We used the most common triangular notches, which are resistant to depinning from thermal perturbation and require a minimized

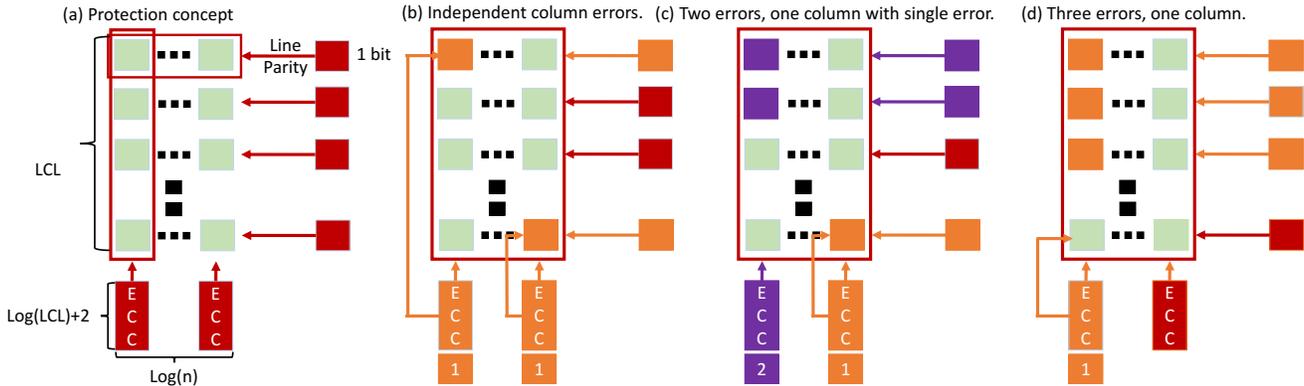


Fig. 10. Signature validation and correction concept as a proxy to detect shift misalignments.

TABLE 1
Material properties used in MuMax simulation.

$A_{ex}(J/m)$	$M_s(A/m)$	α	$K_{ul}(J/m^3)$	current pulse width
2.0×10^{11}	6.5×10^5	0.02	10^6	0.5 ns

shift current. The notches are 50nm wide and 30nm deep. Using Eq. 3 we evaluated the ideal the critical current for a given set of nanowire dimensions and material parameters.

We then modeled the nanowire using the micromagnetic simulation program MuMax3, a widely used GPU accelerated space and time-dependent magnetization dynamics discretized finite-difference solver for nano-sized ferromagnets such as DWM nanowires [39] that has been validated against industry standard simulation such as the Object Oriented MicroMagnetic Framework OOMMF [40]. We characterized the nanowire for changes to the critical shift current density as we varied the notch width and depth by 5%, as described in previous modeling work in the literature [33], [41], at each notch position along the nanowire.

For any given notch, there is lower bound shift current density J_L and an upper bound shift current density J_U to depin and shift one position. For a shift current density J_S in A/m^2 , if $J_S < J_L$ the domain wall will not depin and if $J_S > J_U$ it will travel more than one notch position. The critical shift current density was determined by testing the shifting behavior for different shifting current densities to find the critical shift current density for different variations of width and depth in MuMax. The characterized results showed a monotonically increasing nominal shift current as the notches were farther along the nanowire from the current source as predicted by Eq. 3.

To determine a fault we consider the relationship of J_S to J_L and J_U at all notches in the nanowire using a similar methodology to prior work [42]. Given a notch position i , if $\forall i J_S < J_{i,L}$ or $J_S > J_{i,U}$ then a misalignment fault—undershift or overshift respectively—has occurred. If for a notch k , due to variation in the system, $J_S > J_{k-1,L}$ but $J_S < J_{k,L}$ then domain-wall motion will stop at notch k and pinning (erasure) has occurred. Similarly, insertion can occur in a similar situation near J_U .

To quantify erasure fault probability, we use the total differential method to define the maximum uncertainty of the actual critical shift current density in terms of each of the tested system parameters. Our simulation models determine the partial derivative of J_L with respect to each input parameter determined through characterization. We assume a standard distribution due to process variation on these parameters. J_L is determined by μ centered on the nominal value and σ equal to the overall uncertainty. J_U is calculated in a similar way.

Since a correct shift operation requires all domain walls to shift

in lockstep, for the n th domain wall to shift properly, domain walls $(1, \dots, n-1)$ must also have shifted properly. Counting starts at one, since at zero if the current is under $J_{0,L}$ it is categorized as an under-shift. Thus, the probability of fault free shifting at position n can be defined as $P(n) = \prod_{i=1}^n Q(i)$, where $Q(i)$ is the probability that $J_{i,L} \leq J_S$. A successful full nanowire shift is $P(m)$ where m is the total number of notches in the nanowire. The probability of erasure fault(s) is $1 - P(m)$. Using a similar approach with J_U , we can define the probability of insertion faults.

Using this model, we verified a similar (same order) misalignment probability as prior work [12] and obtained a pinning fault probability reported in Table 2. In the following section, we propose a circuit design for a transverse access point. This TAP forms the foundation for both pinning and misalignment detection in PIETT.

5 TRANSVERSE ACCESS POINTS

To enable PIETT's combined misalignment and pinning detection we propose a TAP circuit as shown in Fig. 11(a). The TAP circuit is related to the shift-write access point [21] but designed along the nanowire to create a segmented, MLC-like device. Our TAP circuit is constructed at the extremity of the nanowire with a fixed domain (in this case aligned right, which we correlate to logic '1') at the very end connected to the shift line (SLB). At the other end of the TAP, we place a fixed left/'0' separated by a standard domain-wall orthogonal to the nanowire and connected to the bit line (BL) through a MOSFET controlled by the V_S signal.

By activating V_S and driving current between SLB and BL (domain-motion happens in the anti-direction of current) and leaving off SL upstream, the free domains between the fixed '1' layer and the out of plane '0' layer can be set to '1's as shown in Fig. 11(b). With sufficient current this can occur in a single intrinsic operation and be slightly overdriven to prevent undershift. Overshift is not a problem because shifting in an extra '1' through the sink results in the same preset configuration. Reversing the polarity of BL and SLB will result in resetting these bits to '0' as shown in Fig. 11(c). Thus, the novel programming concept behind

TABLE 2
Shift error probabilities.

Shifting Distance	Step Fault Rate [12]	Pinning Fault Rate
1	$4.55 \cdot 10^{-5}$	$1.48 \cdot 10^{-8}$
2	$9.95 \cdot 10^{-5}$	$3.23 \cdot 10^{-8}$
3	$2.07 \cdot 10^{-4}$	$6.73 \cdot 10^{-8}$
4	$3.76 \cdot 10^{-4}$	$1.14 \cdot 10^{-7}$
5	$5.94 \cdot 10^{-4}$	$1.80 \cdot 10^{-7}$
6	$8.43 \cdot 10^{-4}$	$2.55 \cdot 10^{-7}$
7	$1.10 \cdot 10^{-3}$	$3.33 \cdot 10^{-7}$

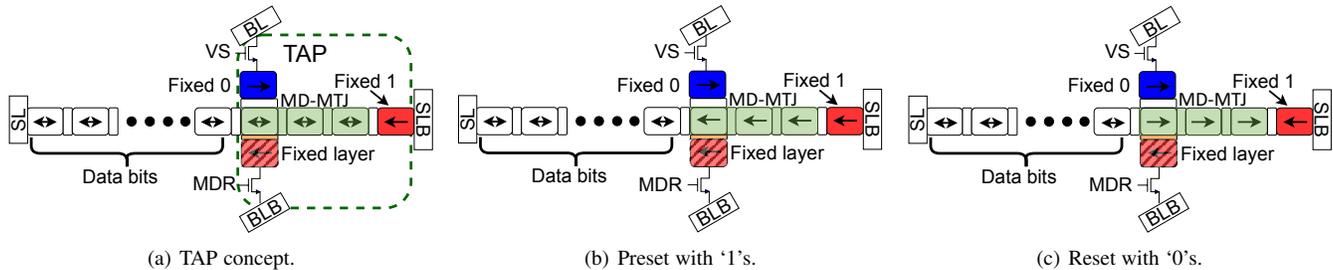


Fig. 11. Transverse access point circuit.

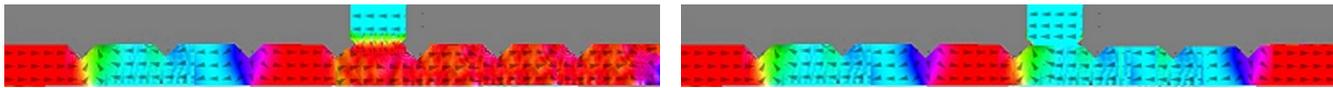


Fig. 12. Magnetic simulation

the TAP is the ability to use a multi-domain shift-based write in a limited subsection of the nanowire.

To verify this capability we conducted a magnetic simulation using the LLG micromagnetic simulator [43] of the TAP circuit from Fig. 11 shown in Fig. 12. In the magnetic simulation we can see the free domains to the left of the TAP moving right to left contain a '1' (red) adjacent to the TAP, followed by two '0's (blue), and a '1' at the far left. Fig. 12(a) shows the alignment after a shift current between \overline{BL} and \overline{SLB} showing that all free domains in the TAP are preset but free domains outside the TAP remained undisturbed. Fig. 12(b) shows resetting to '0's again without disturbing the free domains outside of the TAP.

To conduct an MDR in the TAP, we show the design for two options, TR, and MD-MTJ. For TR, we place another fixed layer orthogonal to the nanowire separated by an insulator (e.g., MgO) shown in hashed red in Fig. 11. This layer is connected to the $\overline{bit\ line}$ (\overline{BLB}) controlled by a MOSFET with a MDR signal. The \overline{VS} MOSFET is turned off and the \overline{MDR} transistor is turned on and a potential is applied between the $\overline{bit\ line}$ (\overline{BLB}) and \overline{SLB} to conduct the MDR. Alternatively, we can directly add an MD-MTJ above or below the nanowire, shown in green in Fig. 11, connected to \overline{BLB} through MDR and GND. MDR is off during preset and set.

Standard domain wall motion through the entire nanowire, including the TAPs, is still possible by turning off both $\overline{VS/MDR}$ MOSFETs and allowing current in the appropriate direction between \overline{SL} and \overline{SLB} . Should the wire shift left, '1's are added to the nanowire similar to the process shown in Fig. 11(b), but they may proceed beyond the fixed '0' domain.

The discussion and simulation are for a single TAP added to the right end of a DWM nanowire. We can build a second mirrored TAP on the left end which can operate entirely independently and in parallel with the other. Moreover, we can swap the alignment to place the fixed '0' at the extremity and the fixed '1' on the internal end of the TAP for either the right or left TAP. The fabrication feasibility of TAPs is similar to fixed magnetic fin-based writes using access transistors connected to \overline{BL} and \overline{BLB} , for which a CMOS layout is demonstrated [21].

6 PIETT WITH TRANSVERSE ACCESS POINTS

Using TAPs from Section 5 PIETT can discover relative position information after conducting a shift of the nanowire. This section describes how TAPs can detect both misalignment and pinning faults. While misalignments can be straightforwardly repaired by

corrective shifts, a technique to correct pinning, or a mixture of pinning and bit-flip faults through scrubbing is described.

6.1 Shift Fault Detection with TAPs

When over-shifting is possible, even with detection, it is necessary to add an additional padding bit at each end of the nanowire so that if over-shifting occurs when attempting to reach the extreme left or right data domain, data is not lost at the other end of the nanowire and corrective shifting is still effective. The TAPs must then be added to each end beyond this additional padding domain. Each TAP must also contain $n + 1$ free domains where n is the length of the maximum intrinsic shift possible in the system.

The TAPs, shown in Fig. 13(a), comprise the outer four padding domains on each side. To detect and distinguish between under-shift, over-shift, and pinning faults, the TAP bits are prepared prior to shifting. Based on their interaction with the other padding bits and the external fixed domain during the shift it is possible to determine whether a fault has occurred. If misalignment occurs, it is reported by both TAPs, simultaneously, both reporting motion that either exceeds or is less than the desired shift amount. Pinning is indicated if part of the nanowire moves a different distance than the other part, indicated by different motion reported by each TAP.

In PIETT, all non-TAP padding bits left of the data are set to '0' and those right of the data are set to '1'. Consider the case that we wish to shift the nanowire in the position from Fig. 13(a) accessing d_2 to be able to access d_3 , requiring a left shift by one domain. Both TAPs are preset to all '1's by shifting both TAPs from left to right by four positions in parallel (see Section 5) as shown in Fig. 13(b). Note, if either TAP was queried at this point with a MDR, the reported value will be 4 '1's as shown in the figure. Upon a successful shift, the nanowire ends up in the position shown in Fig. 13(c). Note that both TAPs, now report "1110" or a read count of 3 '1's. On the left, one of the preset bits was evicted at the left extremity while a '0' padding bit entered the TAP. On the right side, a '0' was inserted into the TAP from fixed '0' domain on the right side.

Fig. 13(d) shows the case where an under-shift occurs, as indicated by both TAPs reporting 4 instead of the expected 3 '1's, requiring a corrective left shift. If over-shifting had occurred, each TAP would read "1100" and report a read count of 2 '1's, requiring a corrective right shift, as shown in Fig. 13(e). Given a TAP contains $n + 1$ free domains, a single TAP allows shifting by n domains in a single step protecting against an over- or under-shift by $k = 1$. In a system free of pinning faults, with both TAPs,

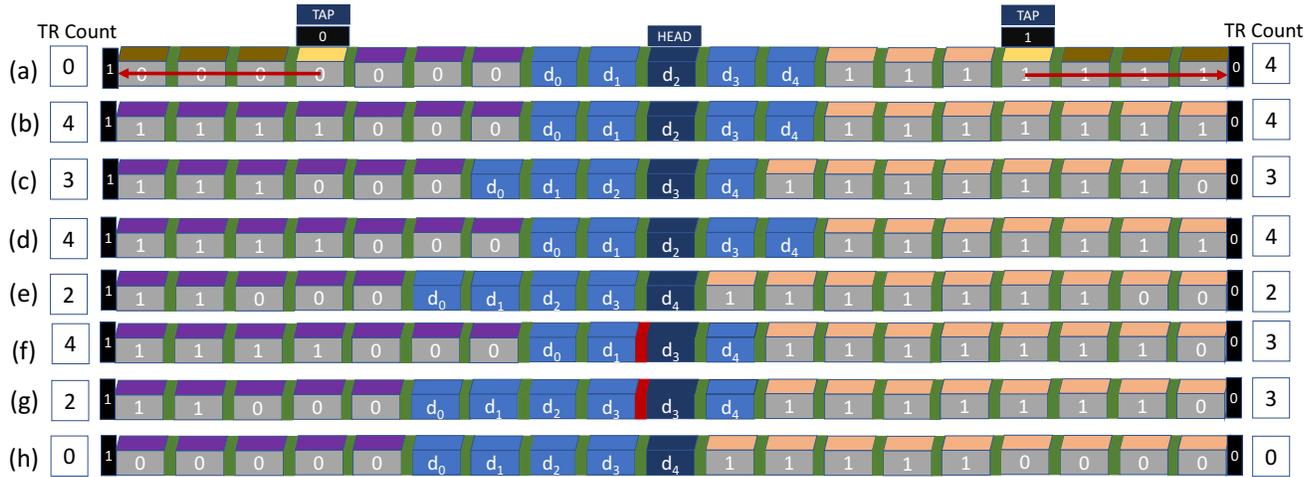


Fig. 13. Example of how TAPs can detect various shifting faults.

PIETT protects against a multiposition over-shift $k > 1$ limited only by potential data loss from exceeding the padding bits. If the over-shift in the example is more than four domains ($k > 3$) the system shifts back by one position until a '1' from the padding bits reenters the right TAP and places the system in a known state. A single corrective shift completes the correction. However, given the probability of misalignment by $k \geq 2$ is $< 10^{-20}$ [12] a pinning fault is more likely to occur.

Fig. 13(f) shows an erasure pinning fault where the pinning point, shown in red, is to the left of the head. Domain motion occurs from left to the pin point and stops. Thus, the right TAP reports "1110" to indicate motion by one position, while the left TAP reports "1111" to indicate no motion. The TAPs mismatched reporting signals pinning has occurred. An insertion pinning fault example shown in Fig. 13(g), occurs when the right portion of the nanowire does not move but domain wall motion starts after the pin point. The left TAP reports "1110" while the right TAP reports "1111" again indicating pinning. In both cases, the pin point cannot be determined and scrubbing is the only remedy.

Right shifting follows the same principle except in this case both TAPs are reset to '0's such that '1's are introduced into the TAPs by domain-wall motion from the leftmost fixed '1' domain or the padding domains right of the data bits but left of the right TAP. We can guarantee that left non-TAP padding bits hold '0's and right non-TAP padding bits hold '1's by examining the behavior of the system at the extremities. For example when shifting left to one data extremity d_4 , all of the '1's preset into the right TAP make their way left into the non-TAP padding bits [Fig. 13(h)], with a similar (complemented) behavior when shifted to d_0 .

The shift steps are: **1** the TAPs are preset for a left shift or reset for a right shift, **2** the shift occurs, **3** if this is a read access and this shift reached the final access location, the read proceeds³, **4** the TAPs are tested and report one of correct shift, misalignment, or pinning, **5** if necessary, correct misalignment through steps **1**⁴, **2**, and **4** until all nanowires are not misaligned, **6** if necessary, correct pinned nanowires. Once these steps are completed a replacement read can be conducted, or a write or subsequent shift are cleared to proceed. Of course,

3. If errors are later detected we assume the system can flush the access and pipeline until the corrected value is determined and returned prior to proceeding. This is standard practice to hide fault-tolerance delay for fault-free accesses.

4. Under-shifts may omit repeating step **1**.

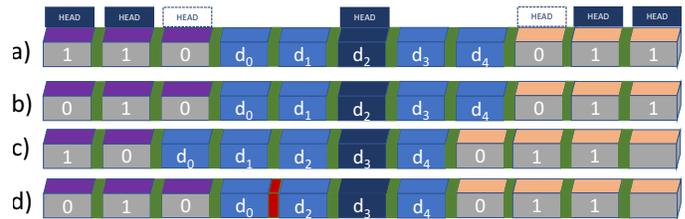


Fig. 14. P-ECC-O against pinning, (a) original position with pattern on the side, (b) changing the last left bit before shifting, (c) Correct shifting operation, (d) pinning during shifting operation

the technique for step **6** was not described. We describe that in Section 6.3. However, in the next section we describe a special case pinning detection for 1-bit TAPs, which can be applied to P-ECC-O from Hi-fi.

6.2 1-bit TAPs and P-ECC-O

Unlike P-ECC and DECC, P-ECC-O writes an alternating pattern into the padding bits. This requires access points at each end of the nanowire as shown in Fig. 5(b). This provides an opportunity to use these access points as a pseudo-TAP for pinning detection. However, to preserve the P-ECC-O misalignment functionality, for pinning detection, the value written must be a function of the bit at the extremity and the penultimate bit. Fortunately, in P-ECC-O there are suitable access points to be able to access these locations in the nanowire after each single domain shift. Recall that P-ECC-O uses the padding bits in a pattern of "110011..." such that it can detect under- or over-shift by one position and misalignment (but not under-/over-shift) by two positions [12].

We show how P-ECC-O can be adapted to address pinning through an example in Fig. 14 for a shift from d_2 to d_3 . In normal P-ECC-O operation, the system, starting in the state from Fig. 14(a) without a fault, would transition directly to Fig. 14(c). Misalignment (over- or under-shift) is checked by testing the two outside bits in the direction of motion against the expected position in the pattern [12]. To add pinning protection we write the complement to the last bit in the direction of domain-wall motion as shown in Fig. 14(b), where the left '1' is replaced by a '0.' If after shifting, the pattern is still overwritten, either pinning or an under-shift fault occurred. An over-shift is detected in the normal way. We can consider an under-shift a pinning fault, however, this conflation may decrease overall fault-tolerance. Instead, an additional head can be added on both sides of the nanowire (white outlined heads). This allows the detection of the alignment on both sides of the nanowire

like a TAP. Thus, Figures 14(b) and (d) can be differentiated. Next, we propose a technique for correcting pinned nanowires.

6.3 Directed Scrubbing

When using nanowires identified by PIETT (or modified P-ECC-O) with pinning and misalignment faults it is relatively simple to fix misalignment through corrective shifts. It is less straightforward to correct pinning. We propose a technique, called *directed scrubbing* which allows the correction of faults from pinning.

Directed scrubbing requires additional nanowires to store parity data based on SECDED ECC for the DBC. First, the DBC is aligned with the farthest left or right data point, whichever is closer, and then read, corrected, and re-written as necessary, moving by single positions until the other extremity is reached. In completing this traversal, in addition to repairing the data domains the encoding domains of the pinned nanowires will all naturally return to the appropriate encoding as described in the discussion of Fig. 13(h).

Of course, during scrubbing, there is a probability of misalignment and additional pinning. Misalignment can be checked and corrected during scrubbing, without need to restart the scrubbing process. If pinning occurs, it can be detected, but scrubbing must start again. Thus, as pinning faults may occur simultaneously during a single intrinsic shift, they may also stack during scrubbing. Single correction ECC may not be sufficient to correct faults in a system with both scrubbing and misalignment protection. Inspired by DECC, which enhances correction by using parity bits to detect the nanowires exhibiting misalignment, by leveraging the location of the nanowires where pinning has occurred we can leverage SECDED with a single parity bit extension to detect and correct as many as three faults.

6.4 Three Pinning-fault Correction Guarantee

TAPs report the nanowires that have experienced pinning. However, even if a nanowire has a pinning fault, it may not report an error during a read while scrubbing. If there are x pinned nanowires, the worst case is that ECC must protect against x errors, but fewer than x errors may also occur. SECDED ECC can correct one error when the location of the fault is unknown. However, if the location of the potential errors is known due to the TAPs, we can correct more errors. We show a variety of error cases during the scrubbing process in Fig. 15, where data bits are shown in blue, Hamming Code parity bits are shown in red, and the Double Error Detection (DED) bit, is shown in gray. Locations of pinned nanowires are illustrated by yellow boxes and actual errors during this particular access are outlined in red. As PIETT reports each possible faulty position by noting the pinned nanowires, with SECDED it is possible to correct faults from two pinned nanowires as follows:

- ECC reports no faults, no re-write is necessary, Fig. 15(a).
- There is one pinned nanowire d_2 , ECC corrects one fault at position d_2 , the corrected bit may be directly re-written [Fig. 15(b)].
- There are two pinned nanowires, d_2 and h_1 and ECC detects two errors (parity bit reports two errors shown in green). The correction is made by flipping the two bits belonging to the pinned nanowires [Fig. 15(c)].

The most interesting case is the last case, where the DED bit, essentially a parity bit for all of the other data and Hamming Code parity bits, reports two errors. The code cannot directly pinpoint which bits are wrong and uses the location of the pinned nanowires to guide correction. However, with the knowledge of potential error location, it is actually possible to correct up to three errors.

In the previous example, either we know the number of errors or there is only one error in the data. In order to correct three errors, we need to reduce the errors to one of these conditions. While the accessed location (cache line) has three simultaneous errors, the DED bit will not report a parity problem as would be the case in two errors. As a consequence, SECDED ECC will report this case as a single error. When combined with location information of the pinned nanowires resolution of all three errors is possible.

In this example, there are three possible faulty locations due to pinning, bits 2 and 4 of the data, and bit 1 of the Hamming Code. Thus for each scrubbing access, the possibilities are as follows:

- ECC reports no errors and no bits are rewritten [Fig. 15(d)].
- ECC reports one error, and it is pointing to a non-pinned nanowire [Fig. 15(e)]. The presumption must be three errors and all three of d'_2, d'_4, h'_1 must be written.
- ECC reports one error, and it is pointing to a pinned nanowire [Fig. 15(f)]. The presumption is that d'_2 must be written. However, if unlucky there may still be three faults. The value is updated with d'_2 and a second ECC check is completed. If the DED bit now indicates two faults, then d'_2, d'_4, h'_1 are written, otherwise only d'_2 is written.
- ECC detects two errors, [Fig. 15(g)]. There are three possibilities, faults in d_2 & d_4 , d_2 & h_1 , or d_4 & h_1 . First we recheck ECC with d'_2, d'_4 , then d'_2, h'_1 , and finally d'_4, h'_1 and write back the pair with an error free code.

This technique is successful as long as the nanowire storing DED bits is not pinned. We can solve this by duplicating the DED bits, requiring one additional nanowire per DBC beyond SECDED ECC. Three fault correction including pinning the DED nanowires(s) is guaranteed as follows:

- Cases with no pinning in either the DED (p_0) or DED₁ (p_1) nanowires—e.g., Fig. 15(h)—resolve to the cases in Fig. 15(b)–(g).
- If the p_0 nanowire is pinned [Fig. 15(i)], p_0 is immediately detected and corrected using p_1 . If there is one other error in either d_4 or h_1 it is corrected using the Hamming code. If SECDED with the corrected DED bit reports two errors both d'_4, h'_1 are written [Fig. 15(i)]. This is similar to the simple case of SECDED ECC described in Fig. 15(c). The case where p_1 is pinned and p_0 is not follows similarly.
- If both p_0 and p_1 are pinned and $p_0=p_1$ [Fig. 15(j)], we cannot know if the DED value is correct. If p_0/p_1 report the incorrect parity we write p'_0 and p'_1 .
- If both p_0 and p_1 are pinned and there is another pinned nanowire (e.g., d_4) [Fig. 15(k)], we use the Hamming code to repair d_4 and then determine the parity and, if necessary, repair the values of p_0 and p_1 .

6.5 Handling Bit Flips

As noted in prior work [17] bit flip faults are possible in DWM due to communication faults over the memory bus when writing, or due to effects like read-disturbance denoted in DWM's spintronic cousin STT-MRAM [44], [45]. Using a similar philosophy to Section 6.4, we can still guarantee three error correction if two of the errors come from pinning and one comes from a bit flip.

Consider in Fig. 15 that d_2 is a bit flip fault, thus we do not know its location. Like any single error, it can be directly corrected by ECC [Fig. 15(b),(f)]. However, if one error is reported, there could be three errors [Fig. 15(e)]. We test again after ECC correction. ECC will report two errors because either one actual

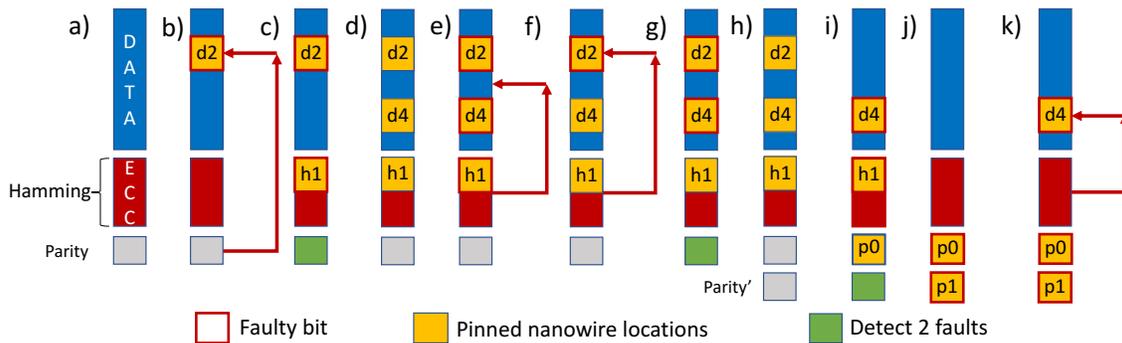


Fig. 15. Example access and fault recovery snapshots during scrubbing. (a) Cache line and parity bits accessed (b) SECDED reports one fault and points in a pinned nanowire (c) SECDED reports two faults due to the DED bit (d) three nanowires are pinned but no bits are actually faulty (e) three fault detection because ECC points to a non-pinned nanowire (f) single fault detected in pinned nanowire while three faults are present (g) two faults detected with three pinned nanowires (h) a duplicate DED is added to protect DED faults (i) three pinned nanowires including a DED bit (j) two DED bits are faulty (k) one data fault and both DED bits are pinned/faulty.

error was corrected, or a new error was added. Either way, the parity will not match signaling that three errors were originally present. Thus, ECC is tested again with both pinned locations corrected (d'_4, h'_1) and now ECC corrects the actual flip at d_2 so that ultimately d'_2, d'_4, h'_1 are written. In the case of two errors [Fig. 15(c),(g)] we flip one pinned location and retest. In the case of (c) ECC will then find bit flip d_2 and h'_1, d'_2 are written. In the case of (g), if we test with d'_4 we are now in the case of (b) and if we test with h'_1 we are now in the case of (e), which are solved.

If there is a bit flip in a DED bit like Fig. 15(i), because $p_0 \neq p_1$ and p_1 reports a parity error the pinned locations are tested. If testing with d'_4 ECC points to h_1 d'_4, h'_1, p'_0 are written, otherwise p'_1 is written. The remaining DED cases [Figs. 15(j) and (k)] follow similarly to Section 6.4.

Thus, $\log_2(\text{data_block_size}) + 3$ additional nanowires per DBC enables repair of either up to three pinned nanowires or up to two pinned nanowires and one bit flip with scrubbing.

7 EXPERIMENTAL SETUP

To evaluate the effectiveness of PIETT, we conducted experiments that study its reliability, area, energy consumption, and performance compared to related schemes. Our DWM memory architecture is based on FusedCache [25], which implements a combination of a set-associative L1 and Last-Level Cache (LLC) in DWM. The domains aligned with the access point belong to L1 and all the other domains logically belong to LLC. When L1 misses, shifting occurs in the DBC in order to access an LLC replacement. Otherwise, FusedCache has a similar organization to TapeCache [11]. To evaluate the latency and energy of shifting we used a modified version of NVSIM designed specifically to model DWM memory [46], [47], [48]. The static energy impact of PIETT is modeled through the inclusion of additional access points for each nanowire and the inclusion of additional nanowires for storing the parity data for each DBC and STT-MRAM elements for DECC.

As PIETT protects against up to three faults in misalignment alone and up to three pinning faults for misalignment with pinning, the size of the data block protected can have a significant impact on reliability. It is the convention to use 64/72 SECDED ECC for a cache line (or memory row) rather than 512/523 where the Length of a Cache Line (LCL) is 512; 64/73 or 512/524 for scrubbing requiring the extra parity bit. We present results for 64/72,73 as it best matches the conventional block size.

To model misalignment and pinning faults during simulation we consider that each DBC contains and shifts R racetracks with n data domains per racetrack, simultaneously. We define the probability

of misalignment after performing a single shift of distance d as $p_{a,d}$. Similarly, we define the probability of pinning faults in one racetrack after performing a single shift of distance d as $p_{p,d}$. We use the values for misalignment and pinning from Table 2, where the pinning probability are obtained through our process discussed Section 4 and the misalignment probability is obtained from the literature [12] and corroborated with the process in Section 4.

Since, fault probability is highly dependent on parameters such as domain size, process variation, shift current, etc., we also consider a sensitivity study of fault probabilities for $p_{p,d}$ from the results in Table 2 (circa 10^{-8}) up to 10^{-4} . Given that correction for misalignment and pinning are corrected orthogonally, we can independently consider $p_{a,d}$ and $p_{p,d}$ as similarly orthogonal. Given the previous treatment of $p_{a,d}$ in previous work that achieves sufficient misalignment protection lifetimes [12], [16], [17], we discuss $p_{a,d}$ alone to evaluate PIETT with DECC in the context of energy and area improvements. Furthermore, given the probability $p_{p,d}$, m is defined as the number of racetracks (out of the R racetracks) which are pinned during an intrinsic shift for the DBC, we can then define the probability of having m racetracks pinned. Using PIETT with TAPs, we focus on $p_{p,d}$ as any number of misalignments can be detected and corrected unless they lead to excessive pinning while conducting corrective shifts.

The memory and fault model were integrated into and simulated using the Sniper multi-core simulator [49]. An architecture with an 8-way 4MB LLC cache and 8-way 32KB L1 cache was studied presuming $n = 32$. Thus, the DBCs are composed of $512 \times 32 = 16384$ bits. Access latencies are as follows: the data read latency is 0.98ns, write latency is 0.65ns, shift latency is 0.32ns, and tag access latency is 0.28ns [25]. The CPU has four out-of-order cores running at a clock speed of 3 GHz. All the benchmarks used to profile the performances are workloads from SPEC-CPU2006 [50].

8 RESULTS

Based on the experimental setup in Section 7 we evaluate the PIETT approach for reliability and examine its impact on energy, performance, and area overheads. In the following sections, P-ECC-O is the version modified to also detect and correct pinning faults.

8.1 Reliability

PIETT-DECC (DECC) exceeds our target 10 year target, achieving a 15 year lifetime. This is the same order of the the 69 year lifetime for SECDED Hi-fi. The tradeoff is that DECC guarantees three misalignment corrections by one with improved area and energy compared to Hi-fi, which corrects all misalignments by one.

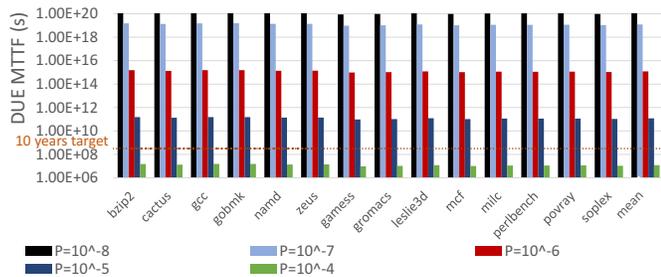


Fig. 16. MTTF: PIETT with TAPs with alignment and pinning faults for different pinning fault rates of 10^{-8} up to 10^{-4} .

For correcting misalignment faults, PIETT-TAPs (PIETT) provides superior fault tolerance as it can natively correct any number of misalignments by at least four making its lifetime essentially unbound, for misalignment fault rates in Table 2. If fault probability increases, Hi-fi and DECC lifetimes would decrease, while PIETT would remain essentially unaffected. As misalignment by two positions is reported as a 10^{-20} fault rate, and misalignment by more than two positions is unmeasurably low [12], the misalignment uncorrectable fault rate of PIETT is better than Hi-fi with double error correction and triple error detection.

PIETT also detects and corrects faults of up to three pinned nanowires. In contrast, P-ECC-O is the only other approach capable (with modification) of detecting pinning faults. We calculated the Mean-Time-To-Failure (MTTF) for pinning incident fault rates ranging from 10^{-8} as obtained from our nanowire model (Table 2) up to 10^{-4} (same order as misalignment fault rates). Without pinning protections, the system MTTF will be between 2 s and 20 μ s for pinning fault rate of 10^{-8} and 10^{-4} , respectively.

Fig. 16 shows the MTTF for PIETT protection of 14 workloads, where the variance is related to frequency of LLC access inducing shifts for the same range of incident fault rates. At 10^{-4} , a particularly high fault rate, PIETT improves MTTF by eight orders of magnitude to 115 days, but still falls short of a 10 year target. As soon as the fault rate is $\leq 10^{-5}$ PIETT improves the MTTF by 14 orders of magnitude to a time of > 385 years, which is well beyond a standard target of 10 years between failures. PIETT improves the MTTF by 21 orders of magnitude for a fault rate of 10^{-8} of the same order as derived from our model. In the following result sections, we consider a pinning probability range of 10^{-8} – 10^{-5} to respect the MTTF target.

8.2 Area Comparison

A standard DWM nanowire consists of data domains, padding domains and an access point. Any additional domains or access points for latency optimization or fault-tolerance, decreases the area efficiency of DWM. P-ECC-O adds four extra heads, two read only and two read/write heads to write their alternating pattern and verify its conformity. In comparison, DECC adds STT-MRAM storage and PIETT adds a fixed number of additional padding domains and logic to provide the transverse write and read capabilities and extra nanowires to store the parity bits for scrubbing. These parity nanowires are also needed for the modified version of P-ECC-O.

Table 3 provides the decomposition of the area (units based on feature size) for the different correction schemes for a nanowire size of $n=32$. The area is broken down by the base DWM area (domains plus heads), the area required to detect and correct misalignment faults, and the overhead to correct pinning faults, when possible. Furthermore, we show two overheads of P-ECC and P-ECC-O for protection against a misalignment of one or two domains, respectively. DECC has the lowest area overhead of all schemes.

TABLE 3

Memory area in $F^2 \times 10^5$ for base array size (base) overhead for misalignment protection (misalign) and for pinning protection (pinning).

Design	DECC	P-ECC	P-ECC-O	P-ECC	P-ECC-O	PIETT
	Misalignment by 1		Misalignment by 2			
Base	8.70	8.70	8.70	8.70	8.70	8.70
Misalign	4.05	5.53	6.91	6.55	7.94	3.94
Pinning	N/A	N/A	1.50	N/A	1.50	1.50
Total	12.75	14.23	17.11	15.25	18.14	14.14

PIETT has comparable overhead to P-ECC while providing pinning protection, and scales better to larger misalignment protection while being 23% less area than pinning modified P-ECC-O.

8.3 Performance

DECC and P-ECC provide similar performance guarantees as both schemes allow shifting to proceed to the final destination prior to misalignment detection/correction. PIETT's improved fault tolerance allows a multi-domain intrinsic shift, but requires a check and write of the TAPs between shift operations. The access latency and system performance in Cycles Per Instruction (CPI) results, shown in Fig. 17 and 18, respectively, are compared to a no-correction baseline. P-ECC has a similar performance to DECC. PIETT and modified P-ECC-O are reported for the fault probabilities from Table 2 with error bars to a pinning probability of 10^{-5} . On average there is a significant latency increase of $1.9\times$ and $2\times$ for PIETT at these pinning probabilities due to the shift and check nature of TAPs. Fortunately, due to the fact that this impacts LLC accesses only, the resulting impact in CPI for the same incident fault rates is only 1% and 2% degradation, respectively. In comparison, modified P-ECC-O, the only other approach that detect pinning has a latency reduction of $5.0\times$ and $5.4\times$ with a more substantial 7% and 9% CPI degradation, respectively.

8.4 Energy Comparison

Fig. 19 shows the energy improvement of DECC over Hi-fi. DECC provides an average of 52% improvement over P-ECC and a 75% reduction over P-ECC-O for misalignment only fault protection.

Fig. 20 shows the energy overhead of PIETT in comparison to P-ECC-O, P-ECC and DECC for the fault probabilities in Table 2 with an error bar that increases the pinning fault probability to 10^{-5} . PIETT is considerably more energy efficient than P-ECC-O, requiring $\frac{1}{3}$ of the energy and reduces energy by more than 35% compared to P-ECC. It does increase energy by about 20% over DECC, but it is important to note that neither P-ECC nor DECC can correct pinning faults and we discuss this comparison further in Section 8.6.

From these results we can observe there is a “fixed” energy overhead (similar to the latency overhead) due to the additional

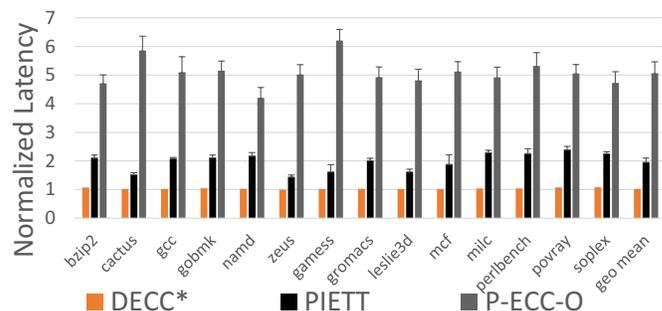


Fig. 17. Latency normalized to no correction with misalignment and pinning fault rates reported in Table 2. Error bars show change in latency if pinning fault rate is increased to 10^{-5} . *DECC reported for misalignment only.

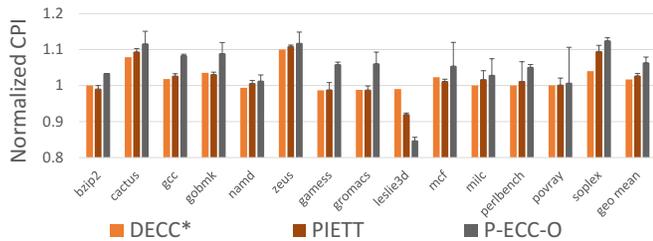


Fig. 18. CPI of DECC, PIETT and P-ECC-O normalized to no correction. Error bars show change in latency if pinning fault rate is increased to 10^{-5} . *DECC provides comparison for misalignment only.

operations to prepare and check amid shifting and the additional parity tapes that shift and consume energy in the DBC, but are necessary when scrubbing is required. There is also a variable cost based on scrubbing the system.

8.5 Bit Flips

In prior work [16], [17] bit flips could be misconstrued as misalignment faults. Bit-flips could also be problematic for DECC (throwing off the signature or encoding bits) unless protected in some other fashion. Prior work has explored how these bit-flip tradeoffs can be considered with correction in STT-MRAM [44].

Due to the TAP concept, bit flips cannot be interpreted as shifting faults in PIETT. In Fig. 21 we show the impact to MTTF of PIETT using the shifting and pinning probabilities from Table 2 with the same range of bit flip probabilities [10^{-9} – 10^{-6}] studied in prior work [17]. PIETT still protects the system well beyond the target 10 years by several orders of magnitude.

8.6 Discussion

PIETT provides two methods for misalignment protection, DECC and PIETT with TAPs. If pinning faults are inconsequential and bit-flips can be managed as is assumed in prior work [12] then DECC provides a reasonable 15 year misalignment guarantee with dramatic savings in energy and area. If pinning is significant then PIETT with TAPs provides significant protection against misalignment, pinning, and even bit flips while maintaining within

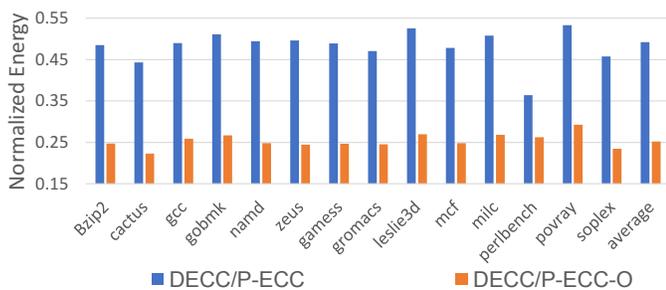


Fig. 19. DECC energy improvement of Hi-fi for misalignment only using fault rates in Table 2.

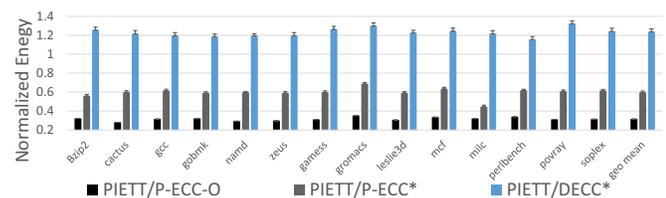


Fig. 20. Energy consumption of PIETT compared to other schemes for fault rates in Table 2. Error bars show change in ratio if pinning fault rate is increased to 10^{-5} . *P-ECC and DECC cannot correct pinning are reported for reference only.

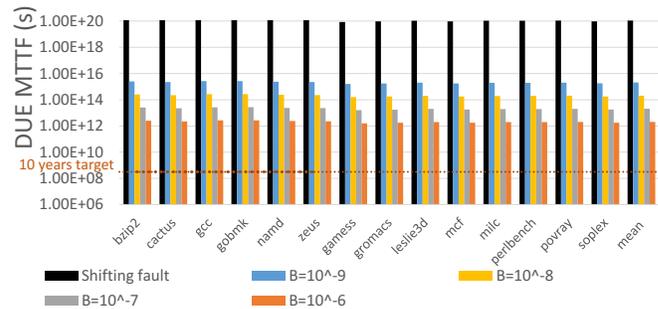


Fig. 21. MTTF: PIETT with bit flips fault rates of 0, 10^{-9} – 10^{-6}

circa 1% performance overhead, dramatically improved energy and similar area overhead of prior work that cannot mitigate pinning or bit flips. Compared to P-ECC-O modified to address pinning, PIETT is considerably better in area, energy, and performance.

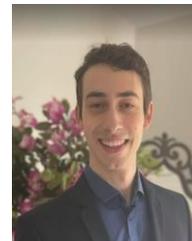
9 CONCLUSION

Manufacturing scaled DWMs will introduce more variation, more defects, and lead to a higher probability for shifting faults to occur. For DWMs to gain traction in real systems, these faults must be efficiently addressed. We propose PIETT that can address misalignment and pinning faults as well as bit flips in random access DWM memories. In PIETT's highest performance and lowest energy mode, DECC, it can provide 15 year reliability guarantees with $>50\%$ energy with area improvements to over the state of the art to correct misalignment-only faults. As pinning fault tolerance is more complex than misalignment because pinning is difficult to detect and harder to correct. PIETT with TAPs is a fault tolerance solution that detects both misalignment and pinning through novel transverse access points placed at the two nanowire extremities. It uses corrective shifts to repair misalignment. PIETT with TAPs leverages knowledge of the location of pinned nanowires to improve the facility of SECDED ECC to repair errors in three pinned nanowires or two pinned nanowires and no more than one bit flip per data element. Without protection from pinning faults, our demonstrated 10^{-8} pinning fault rate indicates DWM devices fail within seconds without pinning protection. In contrast, PIETT can provide effective fault tolerance for pinning fault rates $\geq 10^{-5}$ with MTTF of nearly 400 years. For our modeled fault probabilities (see Table 2), we can guarantee a lifetime over 10^{11} years against pinning faults and a superior protection against misalignment, comparable performance, and an energy reduction of 35% compared to Hi-fi. Important future directions include creating a parameterized fault model for misalignment and pinning of DWM nanowires under different technology nodes, for different amounts of variation, and for different material parameters to further guide fault-tolerant DWM memories. Scaling up multi-domain access for more domains and using MDR and/or TAPs for capabilities beyond fault tolerance are also important future directions.

REFERENCES

- [1] Y. Huai, "Spin-transfer torque mram (stt-mram): Challenges and prospects," *AAPPS bulletin*, vol. 18, no. 6, pp. 33–40, 2008.
- [2] S. S. P. Parkin, M. Hayashi, and L. Thomas, "Magnetic domain-wall racetrack memory," *Science*, vol. 320, no. 5874, pp. 190–194, Apr. 2008.
- [3] S. Parkin and S.-H. Yang, "Memory on the racetrack," *Nature nanotechnology*, vol. 10, no. 3, pp. 195–198, 2015.
- [4] C. Zhang, G. Sun, W. Zhang, F. Mi, H. Li, and W. Zhao, "Quantitative modeling of racetrack memory, a tradeoff among area, performance, and power," *Design Automation Conference*, 2015.

- [5] C. Augustine, A. Raychowdhury, B. Behin-Aein, S. Srinivasan, J. Tschanz, V. K. De, and K. Roy, "Numerical analysis of domain wall propagation for dense memory arrays," in *IEDM*, 2011, pp. 17–6.
- [6] R. Bläsing, A. A. Khan, P. C. Filippou, C. Garg, F. Hameed, J. Castrillon, and S. S. Parkin, "Magnetic racetrack memory: From physics to the cusp of applications within a decade," *Proceedings of the IEEE*, vol. 108, no. 8, pp. 1303–1321, 2020.
- [7] P. Zhou, B. Zhao, J. Yang, and Y. Zhang, "A durable and energy efficient main memory using phase change memory technology," in *ISCA*, 2009, p. 14–23.
- [8] Y.-C. Chen, H. Li, and W. Zhang, "A rram-based memory system and applications," in *The Non-Volatile Memories Workshop*, 2012, (poster).
- [9] J. S. Vetter and S. Mittal, "Opportunities for nonvolatile memory systems in extreme-scale high-performance computing," *Computing in Science Engineering*, vol. 17, no. 2, pp. 73–82, 2015.
- [10] Z. Sun, W. Wu, and H. Li, "Cross-layer racetrack memory design for ultra high density and low power consumption," in *DAC*, 2013.
- [11] 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," in *Proc. of ISLPEd*, 2012, pp. 185–190.
- [12] C. Zhang, G. Sun, X. Zhang, W. Zhang, W. Zhao, T. Wang, Y. Liang, Y. Liu, Y. Wang, and J. Shu, "Hi-fi playback: Tolerating position errors in shift operations of racetrack memory," in *ISCA*, 2015, pp. 694–706.
- [13] S. S. P. Parkin, L. Thomas, and S.-H. Yang, "Method and system for measurement of road profile," April 2014, US Patent 8,687,415 B2.
- [14] M. Al Bahri, B. Borie, T. Jin, R. Sbiaa, M. Kläui, and S. Piramanayagam, "Staggered magnetic nanowire devices for effective domain-wall pinning in racetrack memory," *Phys. Rev. Applied*, vol. 11, p. 024023, Feb 2019.
- [15] Y. Zhang, W. Zhao, D. Ravelosona, J.-O. Klein, J.-V. Kim, and C. Chappert, "Perpendicular-magnetic-anisotropy covef racetrack memory," *Journal of Applied Physics*, vol. 111, no. 9, p. 093925, 2012.
- [16] G. Mappouras, A. Vahid, R. Calderbank, and D. J. Sorin, "Greenflag: Protecting 3d-racetrack memory from shift errors," in *DSN*, 2019.
- [17] S. Archer, G. Mappouras, R. Calderbank, and D. Sorin, "Foosball coding: Correcting shift errors and bit flip errors in 3d racetrack memory," in *DSN*, 2020, pp. 331–342.
- [18] S. Ollivier, D. Kline Jr., R. Kawsher, R. Melhem, S. Bhanja, and A. K. Jones, "Leveraging transverse reads to correct alignment faults in domain wall memories," in *DSN*, 2019.
- [19] K. Roxy, S. Ollivier, A. Hoque, S. Longofono, A. K. Jones, and S. Bhanja, "A novel transverse read technique for domain-wall "racetrack" memories," *IEEE Transactions on Nanotechnology*, vol. 19, pp. 648–652, 2020.
- [20] P. Dutta, A. Lee, K. L. Wang, A. K. Jones, and S. Bhanja, "A multi-domain magneto tunnel junction for racetrack nanowire strips," 2022. [available online] <https://arxiv.org/abs/2205.12494>.
- [21] R. Venkatesan, M. Sharad, K. Roy, and A. Raghunathan, "Dwm-tapestri: an energy efficient all-spin cache using domain wall shift based writes," in *Proc. of DATE*, 2013, pp. 1825–1830.
- [22] A. Annunziata *et al.*, "Racetrack memory cell array with integrated magnetic tunnel junction readout," in *IEDM*, Dec. 2011.
- [23] Y. Zhang, W. Zhao, J.-O. Klein, D. Ravelosona, and C. Chappert, "Ultra-high density content addressable memory based on current induced domain wall motion in magnetic track," *IEEE TMAG*, vol. 48, no. 11, pp. 3219–3222, nov. 2012.
- [24] D. Kline, H. Xu, R. Melhem, and A. K. Jones, "Racetrack queues for extremely low-energy fifos," *IEEE TVLSI*, no. 99, pp. 1–14, 2018.
- [25] H. Xu, Y. Alkabani, R. Melhem, and A. K. Jones, "Fusedcache: A naturally inclusive, racetrack memory, dual-level private cache," *IEEE Transactions on Multi-Scale Computing Systems*, vol. 2, no. 2, pp. 69–82, 2016.
- [26] M. Moeng, H. Xu, R. Melhem, and A. K. Jones, "Contextpref: Enhancing the performance and energy of gpus with nonuniform register access," *IEEE TVLSI*, vol. 24, no. 1, pp. 343–347, 2016.
- [27] R. Venkatesan, S. G. Ramasubramanian, S. Venkataramani, K. Roy, and A. Raghunathan, "Stag: Spintronic-tape architecture for gpgpu cache hierarchies," in *ISCA*, 2014, pp. 253–264.
- [28] Q. Hu, G. Sun, J. Shu, and C. Zhang, "Exploring main memory design based on racetrack memory technology," in *GLSVLSI*, 2016, pp. 397–402.
- [29] A. A. Khan, H. Mewes, T. Grosser, T. Hoefler, and J. Castrillon, "Polyhedral compilation for racetrack memories," *IEEE TCAD*, vol. 39, no. 11, pp. 3968–3980, 2020.
- [30] A. A. Khan, F. Hameed, R. Bläsing, S. S. P. Parkin, and J. Castrillon, "Shiftsreduce: Minimizing shifts in racetrack memory 4.0," *ACM Trans. Archit. Code Optim.*, vol. 16, no. 4, Dec. 2019.
- [31] G. Yu, P. Upadhyaya, Y. Fan, J. Alzate, W. Jiang, K. Wong, S. Takei, S. Bender, L. Chang, Y. Jiang, M. Lang, J. Tang, Y. Wang, Y. Tserkovnyak, P. Amiri, and K. Wang, "Switching of perpendicular magnetization by spin-orbit torques in the absence of external magnetic fields," *Nature Nanotechnology*, vol. 9, no. 7, pp. 548–554, Jul. 2014.
- [32] A. Razavi, H. Wu, Q. Shao, C. Fang, B. Dai, K. Wong, X. Han, G. Yu, and K. L. Wang, "Deterministic spin-orbit torque switching by a light-metal insertion," *Nano letters*, vol. 20, no. 5, pp. 3703–3709, 2020.
- [33] M. Hayashi, *Current driven dynamics of magnetic domain walls in permalloy nanowires*. Stanford University California, 2007.
- [34] L. Thomas, M. Hayashi, X. Jiang, R. Moriya, C. Rettner, and S. S. Parkin, "Oscillator dependence of current-driven magnetic domain wall motion on current pulse length," *Nature*, vol. 443, no. 7108, pp. 197–200, 2006.
- [35] T. Suzuki, S. Fukami, N. Ohshima, K. Nagahara, and N. Ishiwata, "Analysis of current-driven domain wall motion from pinning sites in nanostrips with perpendicular magnetic anisotropy," *Journal of Applied Physics*, vol. 103, no. 11, p. 113913, 2008.
- [36] J. Sampaio, J. Grollier, and P. J. Metaxas, "Domain wall motion in nanostructures," *Magnetism of Surfaces, Interfaces, and Nanoscale Materials*, p. 335, 2015.
- [37] A. Aharoni *et al.*, *Introduction to the Theory of Ferromagnetism*. Clarendon Press, 2000, vol. 109.
- [38] G. Tatara, H. Kohno, and J. Shibata, "Microscopic approach to current-driven domain wall dynamics," *Physics Reports*, vol. 468, no. 6, pp. 213–301, 2008.
- [39] A. Vansteenkiste, J. Leliaert, M. Dvornik, M. Helsen, F. Garcia-Sanchez, and B. Van Waeyenberge, "The design and verification of mumax3," *AIP Advances*, vol. 4, no. 10, p. 107133, 2014.
- [40] M. J. Donahue and D. G. Porter, "Oommf: Object oriented micromagnetic framework," Jan 2016. [Online]. Available: <https://nanohub.org/resources/oommf>
- [41] A. Iyengar and S. Ghosh, "Modeling and analysis of domain wall dynamics for robust and low-power embedded memory," in *DAC*, 2014.
- [42] K. Roxy, S. Longofono, S. Olliver, S. Bhanja, and A. K. Jones, "Pinning fault mode modeling for dwm shifting," *IEEE Transactions on Circuits and Systems II: Express Briefs*, pp. 1–1, 2022.
- [43] M. R. Scheinfein, "Llg micromagnetic simulator," 1997.
- [44] S. Seyedzadeh, R. Maddah, A. Jones, and R. Melhem, "Leveraging ecc to mitigate read disturbance, false reads and write faults in stt-ram," in *DSN*, 2016, pp. 215–226.
- [45] Z. Sun, H. Li, and W. Wu, "A dual-mode architecture for fast-switching stt-ram," in *ISLPEd*, 2012, p. 45–50.
- [46] C. Zhang, G. Sun, W. Zhang, F. Mi, H. Li, and W. Zhao, "Quantitative modeling of racetrack memory, a tradeoff among area, performance, and power," in *ASP-DAC*, 2015, pp. 100–105.
- [47] X. Dong, C. Xu, Y. Xie, and N. P. Jouppi, "Nvsim: A circuit-level performance, energy, and area model for emerging nonvolatile memory," *IEEE TCAD*, vol. 31, no. 7, pp. 994–1007, 2012.
- [48] S. J. Wilton and N. P. Jouppi, "Cacti: An enhanced cache access and cycle time model," *IEEE J. Solid-State Cir.*, vol. 31, no. 5, pp. 677–688, 1996.
- [49] T. E. Carlson, W. Heirman, S. Eyerman, I. Hur, and L. Eeckhout, "An evaluation of high-level mechanistic core models," *ACM TACO*, 2014.
- [50] J. L. Henning, "Spec cpu2006 benchmark descriptions," *ACM SIGARCH Computer Architecture News*, vol. 34, pp. 1–17, 2006.



Sébastien Ollivier graduated from the French engineering school ENSEA (Ecole Nationale Supérieure de l'Electronique et ses Applications) and the university of Pittsburgh with a master's degree in Electrical Engineering in 2018. He completed his PhD candidate at the University of Pittsburgh in electrical and computer engineering. He published several articles focusing on domain-wall memory. He started by working on the novel memory reliability to then use it as computing unit for processing in memory applications.

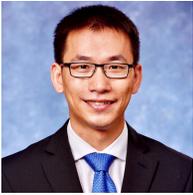


Stephen Longofono received the BS degree in computer engineering from the University of Kansas in 2018, and the MS degree in electrical and computer engineering from the University of Pittsburgh in 2021. He is currently employed as an Avionics Software Engineer at Blue Origin in Seattle, WA. His research interests include computer architecture, memories, information security, and heterogeneous computer systems design. He is a student member of the IEEE and the ACM.



Prayash Dutta completed his Bachelor of Science in Electrical in Electronics Engineering from Bangladesh University of Engineering and Technology (BUET) in 2017. Previously he was employed as the Deputy In charge of Cellphone R&D Department at Walton Digi-Tech Industries Ltd. Now, he is a Teaching Assistant and Research Assistant at the University of South Florida in the Electrical Engineering Department. He is currently a Ph.D. student researching under the supervision of Dr. Sanjukta Bhanja. His research

interest includes Sense Amplifier Design and MTJ properties analysis.



Jingtong Hu is currently an Associate Professor in the Department of Electrical and Computer Engineering at University of Pittsburgh, Pittsburgh, PA, USA. Before that, he was an Assistant Professor at Oklahoma State University from 2013 to 2017. His current research interests include hardware/software co-design for machine learning algorithms, on-device AI, embedded systems. He has served on the technical program committee of many international conferences. He served as a guest editor for IEEE Transactions

on Computers and is currently an associate editor for IEEE Embedded Systems Letters and ACM Transactions on Cyber-Physical Systems.



Sanjukta Bhanja received B.Sc. in EE from Jadavpur University, Calcutta in 1991 and M.Sc. from Indian Institute of Science, Bangalore in 1994. She finished her PhD degree in CSE in 2002 from the University of South Florida, Tampa. She is currently a Professor of EE and Executive Associate Dean, College of Engineering, at the University of South Florida. She has published more than 100 publications in VLSI and nano-electronics. She has been an AE of IEEE TVLSI and ACM JETC and is a Steering committee member of ACM GLSVLSI and IEEE ISVLSI. She is the recipient of NSF CAREER award 2007-2014; USF Tau Beta Pi "Outstanding Engineering Faculty Researcher" award, 2007; USF "Outstanding Faculty Research Achievement Award" in 2008; USF Outstanding Undergraduate Teaching award 2010; Florida Education Foundation (F.E.F) William Jones Outstanding Mentor award 2010; Honorable mention award Outstanding Graduate faculty mentor in 2013.



Alex K. Jones received the BS degree in 1998 in physics from the College of William and Mary in Williamsburg, VA, USA, and the MS and PhD degrees in 2000 and 2002, respectively, in ECE from Northwestern University, Evanston, IL, USA. He is a Professor of ECE and CS at the University of Pittsburgh, Pittsburgh, PA, USA. He is currently serving as a Program Director at the US NSF in the CNS Division of the CISE Directorate. Dr. Jones' research interests include compilation for configurable systems and architectures, scaled and emerging memory, reliability, fault tolerance, and sustainable computing. He is the author of more than 200 publications in these areas. His research is funded by the NSF, DARPA, NSA, and industry. Dr. Jones received a top 25 paper award from the first 20 years of FCCM. He is a senior member of the IEEE and the ACM.