Vulnerability Analysis of On-Chip Access-Control Memory

Chintan Chavda[†], Ethan C. Ahn[†], Yu-Sheng Chen[⋆], Youngjae Kim[‡*]
Kalidas Ganesh[†], Junghee Lee[†]
[†]*University of Texas at San Antonio, USA,* [⋆]*Industrial Technology Research Institute, Taiwan*
[‡]*Sogang University, Seoul, Republic of Korea*

Abstract

Encryption is often employed to protect sensitive information stored in memory and storage. It is the most powerful countermeasure against data breach, but it has performance overhead. As a low-cost alternative to encryption, an access-control memory (ACM) has been introduced, which integrates an access-control mechanism with memory. While ACM minimizes the performance overhead of encryption, it provides similar levels of security as to encryption method. ACM reveals information only when the access codes are correct. However, if an adversary attempts to access data directly from memory cells through a physical attack without going through a standard interface, the vulnerability could occur. This paper discusses feasibility and countermeasures for physical attacks, including fault injection attack, power analysis attack, chip modification, microprobing, and imaging for ACM. Moreover, as a concrete example of ACM, we compare the security aspects of SSDs when the write buffers in the SSDs employ ACM with emerging non-volatile memories such as STT-RAM, PRAM, and RRAM.

1 Introduction

Encryption is the most popular and strong method to protect sensitive information in storage [11]. Without a matching description key, the information cannot be revealed even if it is stolen from storage. However, encryption and decryption are not free. Performance overhead is caused by performing encryption or decryption whenever data is accessed. Previous studies have investigated hardware accelerators [17] and light-weight algorithms [4, 16, 2, 15] to mitigate the encryption overhead. These techniques can help reduce the performance overhead of encryption to a certain extent. However, since encryption overhead occurs during every data access, even a small overhead can not be ignored. Particularly, in the data intensive environments, where I/O requests arrive in a short time period, the accumulation of the encryption/decryption overhead will seriously degrade performance [10]. Moreover, the encryption overhead is relatively large, especially as data access time is reduced in memory-based storage.

Access-control memory (ACM) has recently been introduced as a low-cost alternative to encryption [10]. The key idea is to integrate the access-control mechanism with memory. Any type of memory can be used for access-control memory. It allows memory access only if the access code of data is matched. ACM offers the same security effect as encryption. As a use case, ACM can be employed for the write buffer in an SSD. By employing the ACM, it will not require encryption but data will be encrypted in background when it is transferred from the write buffer to flash memory. The use of non-volatile memory as write buffers for ACM in SSD, will act against data loss from power-failures.

Spin-Transfer-Torque Random Access Memory (STT-RAM) technology can be chosen as an access-control memory. STT-RAM is one of the most promising non-volatile memory candidates due to its excellent scalability and superior performance [3, 7]. In STT-RAM, a spin-polarized electric current is used in order to exert a torque to change the magnetization direction of the magnetic free layer in a magnetic tunnel junction (MTJ) element. The resultant resistance difference of the MTJ is used for information readout. STT-RAM offers fast read and write access latencies in about 10 nanoseconds [6] and practically unlimited lifetime with programming endurance of up to 10^{15} cycles [8].

In this paper, we discuss the safety of the access-control memory for various physical attacks such as fault injection attack, power analysis attack, chip modification, microprobing, and imaging. In addition, we confirm the difficulty of distinguishing ACM data according to the imaging method through actual experi-

[*] Corresponding author

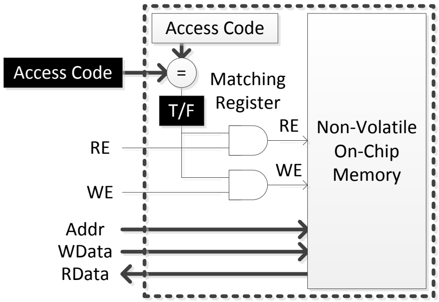

Figure 1: A block diagram of an access-control memory. Reprinted from [10], Copyright 2017 by IEEE.

ments. This paper also compares three types of emerging non-volatile memory technologies as candidates for the access-control memory in terms of security.

2 Access-Control Memory

Figure 1 depicts a block diagram of the access-control memory. The key idea of the ACM is to integrate an access-control mechanism with memory [10]. If the memory is ACM, the access-control mechanism can neither bypassed nor disabled. Any type of memory can be used for ACM. ACM grants memory access only if the access code is matched. In other words, it will return valid data or allow to write data, only when the access code is approved. In fact, this grant mechanism is exactly what encryption techniques intend to do. Encryption techniques allow access to read data only if a correct decryption key is accompanied. ACM performs the same function as encryption techniques, while it does not employ data encryption.

The matching register is employed to reduce the critical path delay. If it is not employed and the access code is compared for every access, the comparison logic affects critical path delay. Since the access code should be at least 128-bit long to prevent the exhaustive search, the comparison logic may incur non-negligible delay to critical path delay. To minimize the impact on the critical path delay, the matching register is added.

3 Using ACM as the Write Buffer for Encrypted SSDs

Unencrypted write buffers can be used to minimize performance overhead due to the encryption of fully encrypted SSDs. However, if the write buffer is not encrypted, the adversary can read the sensitive data from the write buffer. ACM can be used as a write buffer in an encrypted SSD to minimize encryption overhead and provide the same level of security as a fully encrypted SSD [10]. The ACM-based write buffer achieves the same security effect as encryption without full data encryption because the data can be accessed only when the access code is matched in read and write.

3.1 Threat Model

Attackers can compromise the SSD firmware. Once the firmware is compromised, an attacker can change the SSD hardware configuration and read and write all data in memory and registers. Under this threat model, any unencrypted on-chip memory is not secure because they can be accessed by the compromised firmware. However, if ACM is employed, even if the firmware is compromised, data can not be read from the ACM unless there is a valid access code [10].

The access code given by the host is stored in the *volatile* register. If the SSD power is powered on, the attacker may attempt to read the access code in a different way than the normal way through the compromised firmware, because the access code in the volatile register has not been erased. This is possible only if the power cable of an SSD is kept connected to a host while its bus interface (e.g. SATA or SCSI) is reconnected to the adversary's machine. This is a common issue with self-encrypting drives. It is a type of hot plug attacks, however, it can be prevented by introducing connection sensitivity to the bus interface [11]. When the bus interface is disconnected, an SSD will be locked, thus it cannot be reconnected to other machines [11].

SSD data encryption is a technique typically used to prevent data breach when the SSD is stolen. That is, before the attacker attacks, the SSD is disconnected from the host and power supply is disconnected at least once. If power is disconnected, all data in volatile memories and registers will be erased before an attack. All data in volatile memory is assumed to be erased before the attack if the SSD is stolen.

3.2 Write Buffer Design with ACM

The purpose of employing the ACM as a write buffer is to protect sensitive data stored in write buffer from *compromised firmware*. To be used as a write buffer, a non-volatile memory should be used to avoid data loss by sudden power fail. If the ACM is used as a write buffer, the access-control mechanism can neither bypassed nor disabled even if the firmware is compromised. It is because the access-control mechanism is not under control from the firmware [10].

While booting, the firmware receives an access code from the host and stores it in a *volatile* register (in the black box of Figure 1). Since it is volatile, it will be erased when power is switched off. When a hardware-based full disk encryption is employed, the key is usu-

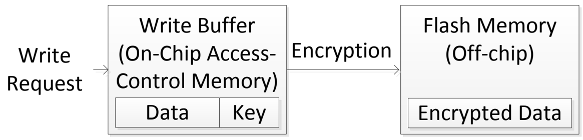

Figure 2: A data path when an access-control memory is employed as a write buffer. Reprinted from [10], Copyright 2017 by IEEE.

ally stored in the device (SSD) [11]. However, under our threat model where the firmware can be compromised, the access code cannot be permanently stored in an SSD because it could be revealed by compromised firmware.

The secure SSD still employs encryption to protect data stored in flash memory. By employing the ACM as a write buffer, it can hide encryption delay which results in performance improvement. A typical data path for an SSD is shown in Figure 2 when the ACM is employed as a write buffer. When an incoming write request arrives, it is stored in the ACM and immediately committed. While the request is stored, the key to encrypt data should also be stored. Note that even though the key is stored in the ACM, it is safe under our threat model. In the background, the firmware encrypts the stored data and transfers the encrypted data to a flash memory. If the ACM is not employed and the write buffer is a plain on-chip memory, the data stored in the write buffer should be encrypted. Therefore, the write request can be committed *after* encryption completes. In contrast, the ACM allows performing encryption in background while guaranteeing the same level of security with encryption.

3.3 Performance and Security Concern

In order to evaluate the effectiveness of STT-RAM based access-control memory, we used the DiskSim augmented SSD simulator developed by Microsoft Research [1]. The details of the simulation parameters are given in Table 1.

Total capacity	8 GB	Pages per block	64
Reserved free blocks	15 %	Page size	4 KB
Minimum free blocks	5 %	Page read latency	0.025 ms
Cleaning policy	Greedy	Page write latency	0.200 ms
Flash chip packages	4	Block erase latency	1.5 ms
Planes per package	4	Encryption delay	1,720 ns / sector
Blocks per plane	512	Decryption delay	1,720 ns / sector

Table 1: SSD model parameters.

Figure 3 shows the results to compare the performance of an encrypted SSD employing the ACM as a write buffer with that of an encrypted SSD with an encrypted write buffer. The results show that the ACM approach increase I/O throughputs by 16%, 57%, 16% and 11%, for Exchange, Cell, Financial, and TPCH, respectively.

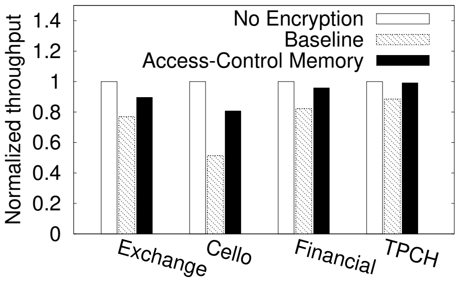

Figure 3: Throughput comparison of an SSD where an encryption technique is employed (Baseline) and where the access-control memory is employed as a write buffer (Access-Control Memory).

However, the access-control mechanism integrated with the non-volatile memory can prevent unauthorized access only if the request comes through the standard interface that the access-control mechanism monitors. If an adversary physically steals access-control memory and tries to access data directly by a physical attack, the data in the memory can be compromised. ACM has potential vulnerabilities that are not present in encryption techniques. Therefore, it is imperative to analyze the possibilities of physical attacks to ensure that the access-control memory is physically secure.

4 Security Analysis

This section is focused on discussing feasibility of possible physical attack mechanisms of fault injection, chip modification, power analysis, microprobing, and imaging attacks. In particular, we add our own experimental analysis for the imaging technique, because it is considered the simplest, non-destructive attack for information retrieval.

Fault Injection Attack: The matching register could be flipped by a fault injection attack. If the matching register is flipped, an adversary may access the access-control memory even if the access code is not matched. However, the matching register is recovered immediately at the following clock cycle as long as the access code remains unmatched. Since faults cannot be injected persistently, the fault in the matching register is unlikely to last for more than one clock cycle. If one wants to eliminate this potential vulnerability, one can employ redundant comparison logic and matching register [9]. Although repetition is considered the most expensive countermeasure for fault injection attacks [9], it does not incur much overhead for the access-control memory because the comparison logic and matching register are very small.

Chip Modification: An adversary may use a special equipment such as an ion beam station to alter the functionality of a chip. He may cut and reconnect interconnection wires. In case of the access-control memory, if he manages to cut the output of the matching register and tie it to true, the access-control memory will always grant access. To achieve this, the adversary needs to keep all other parts of the chip working correctly. In other words, he cannot destroy other parts. Therefore, if the interconnection wires of the access-control circuit of the access-control memory are placed in inter metal layers, it will be extremely difficult for an adversary to alter those wires without destroying other parts of the chip. One possible countermeasure to chip modification is to place interconnection wires of the access-control memory in inter metal layers.

Power Analysis Attack: By analyzing power consumed by the micro-controller chip, an adversary may acquire the access code. However, since the comparison logic of the access codes is very small compared to the entire chip, the power profile of the logic is within noise margin. The hardware logic to implement the access-control mechanism in the memory consumes 98.41 μW, which is estimated by Candence Encounter with Nangate 45 nm technology. Therefore, a power monitoring attack is infeasible for the access-control memory.

Microprobing: By using microprobes, an adversary may measure electronic characteristics to determine the value stored in a memory cell. Because microprobing requires to measure memory cells one by one manually, microprobing is almost impossible and getting harder for modern memories for two reasons; (1) there are too many cells in a memory nowadays and (2) as the feature size shrinks, it is getting more difficult to attach probes to memory cells. Considering the capacity of memory is well beyond thousands or even millions of bits, it will take a great amount of time to measure all of them manually.

Imaging: The fabricated memory cells can be examined in scanning electron microscope (SEM) to reveal the bi-stable bit information (1 or 0) of the memory cell. It is noted that state-of-the-art emerging non-volatile memory cells are typically in the order of a few tens to hundreds of nanometers, requiring the magnification of at least 5,000x to obtain the visible image; it is well within the capability of commercial SEM equipment. Obviously, another imaging technique of optical microscope that is much simpler and cheaper, is not a suitable choice for imaging the modern, nanoscale memory cells.

In order to investigate the vulnerability of the STT-RAM cell for the non-invasive imaging attack, we have examined the fabricated STT-RAM cells in a Hitachi 30kV Variable Pressure-SEM (SU1510). First, two different memory cells were prepared; one has been programmed to a high-resistance state (HRS, or '0' state) and the other to a low-resistance state (LRS, or '1' state). These bi-stable magnetic bits typically generate a relatively large contrast in electrical resistance (or current) of around or higher than 100% tunneling magnetoresistance (TMR) ratio, so if adversaries manage to electrically probe (e.g., microprobing) these memory cells, the individual bit information (whether '1' or '0') can be revealed. However, as discussed before, microprobing is very difficult for modern memory technologies. The purpose of this experiment is to see if adversaries can extract the stored bit information for the case of imaging attack.

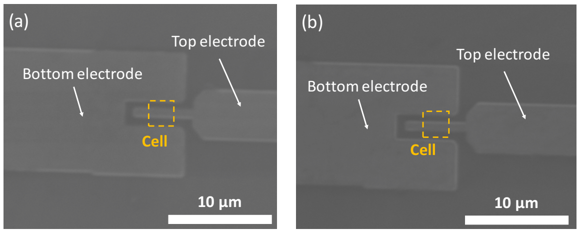

(a) HRS ('0') (b) LRS ('1')

Figure 4: Comparison of two different STT-RAM cells in the different bit-status ((a) for HRS or '0' state and (b) for LRS or '1' state) in the 30kV SEM imaging.

After the samples were prepared, the SEM technique was applied to each STT-RAM cell to identify any notable differences in the SEM image between HRS and LRS memory cells. Figure 4 compares two different STT-RAM cells in the different bit-status in the 30kV SEM imaging. Magnification has been set to be around 5,000x to be well within the range of most of the commercially available (cost-effective) SEM equipment. Figure 4(a) and (b) show STT-RAM in the HRS ('0') and STT-RAM in the LRS ('1'), respectively. The active memory storage layer is between the marked top and bottom electrodes (its location is marked by the yellow dashed square), and it is not visible in the SEM imaging.

As show in Figure 4, for a wide range of magnification (up to 15,000x), no significant difference in the SEM image was observed. This is attributed to the fact that the secondary electrons, the primary source of the SEM imaging, are largely generated from the top electrode metal layer, which is far from the active data storage (magnetic free) layer. The STT-RAM cell typically consists of a large number of magnetic and non-magnetic layers (i.e., the total number of layers in the modern MTJ easily exceeds 10), making the magnetic free layer positioned nearly in the middle of the entire device stack and thus less exposed to the direct imaging attack.

A cross-sectional SEM imaging technique can be employed to directly observe the active layer from the side edge, but it requires additional efforts in preparing samples (e.g., cleaving) and advanced focusing and stigmation techniques. Furthermore, it is impossible to examine all memory cells without destructing nearby cells.

We have also tried a backscattered electron (BSE) image in SEM, because the BSEs are higher energy particles than the secondary electrons and may generate a different image contrast. Yet, no clear difference between the HRS and LRS cells has been observed. Therefore, it is concluded from our SEM imaging studies that the STT-RAM technology still preserves a high level of security for the non-invasive attack with minimum efforts.

Other imaging techniques that have been widely used for observing magnetic microstructures, such as the MOKE (magneto-optical Kerr effect) magnetometry, may also be adopted as a non-destructive method for STT-RAM information readout. The MOKE effect [12] measures the change in the polarization of the reflected light from a magnetic material surface, thus giving access to the magnetization state of the sample. However, in general, since this Kerr effect is based on the surface magnetism within a relatively small skin depth (less than 20 nm in most metals for the visible wavelength range used for conventional MOKE), it is still very hard to directly probe using the MOKE technique the magnetic free (storage) layer that is located deep inside the whole STT-RAM device stack. Another limiting factor that may also impact the vulnerability of STT-RAM cells for the MOKE-based imaging attack is the spatial resolution. Since MOKE is basically an optical technique, its spatial resolution is strongly limited by the diffraction limit of the light beam used. For example, the MOKE system that uses a typical visible light (wavelength range of 400 nm to 700 nm) has the diffraction limit of nearly 1 micron, which is too large to decode individual memory cells in a modern memory array architecture of nanoscale cell dimensions and cell-to-cell pitches. More recently, the x-ray magnetic circular dichroism (XMCD) spectroscopy has been utilized as an alternative magnetic imaging technique with nanometer resolution and possibly higher penetration depth [5], but the development of reliable, efficient x-ray sources still remains a great challenge.

5 ACM with Other Non-Volatile Memory

For this work, the STT-RAM has been chosen for the access-control memory of write buffers in SSD, considering not only performance and endurance but also security. However, it is possible to employ other types of non-volatile memory. In this section, we discuss two other representative memory types, as an alternative memory for write buffers of the SSD to STT-RAM.

Phase-change RAM (PRAM) is an important class of data storage, with its history dating back to 1960's. The large electrical contrast between amorphous and crystalline phases of chalcogenide alloys is utilized for bistable memory operation. Resistive RAM (RRAM) has a simple capacitor-like metalinsulatormetal structure, typically composed of transition metal oxides sandwiched between two metal electrodes. The RRAM's structural simplicity and excellent scalability ($<$ 10 nm) made it one of the most promising candidates for next-generation mass-storage applications.

Adversaries may consider removing top electrode layers to reach the active memory storage layer (magnetic free layer, chalcogenide phase-change, and metal-oxide layer in STT-RAM, PRAM, and RRAM, respectively) and then apply the discussed imaging techniques. In this case, the STT-RAM provides the highest level of security due to the largest number of layers on top of the storage layer (about 3-5 for STT-RAM vs. one or two for PRAM and RRAM). However, we also note that removing the layers is not a practically viable approach as optimization of the etching conditions with a very good selectivity (not to damage the storage layer) will be a very challenging task. In addition, the PRAM cells can be more vulnerable to the physical attack that comes in the form of cross-sectional imaging; the cell's crystalline structure (amorphous for '0' vs. crystalline for '1') might be easily revealed in the SEM imaging.

If we only consider performance, RRAM is the best choice. At similar technology nodes, STT-RAM [6] offers 8 ns read and 12 ns write access times, PRAM [13] offers 68 ns read and 180 ns write access times, and RRAM [14] offers 7.2 ns read and 7.2 ns write access times. Though RRAM offers the best performance, STT-RAM has been chosen considering potential risks of imaging, removal of top layers, and corss-sectional imaging attacks.

6 Conclusion

In this paper, we discuss various physical attacks to the access-control memory. We believe the power analysis attack and the microprobing are not feasible and we have countermeasures to the fault injection attack and the chip modification. We also demonstrated that imaging with SEM does not reveal the bit information of memory cells. However, if adversaries remove layers until they expose the layer where the actual data is stored, they may manage to find the bit information by imaging. Though we consider this is practically very difficult, we will examine this by experiments in our future work. Among STT-RAM, PRAM, and RRAM, the STT-RAM is chosen for the access-control memory. Though STT-RAM does not offer the best performance among them, it provides the highest level of security from physical attacks.

Acknowledgment

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea Government (MISP) (No. 2015R1C1A1A0152105).

References

[1] Agrawal, N., Prabhakaran, V., Wobber, T., Davis, J. D., Manasse, M., and Panigrahy, R. Design Tradeoffs for SSD Performance. In *USENIX 2008 Annual Technical Conference* (2008), pp. 57–70.

[2] Awad, A., Manadhata, P., Haber, S., Solihin, Y., and Horne, W. Silent shredder: Zero-cost shredding for secure non-volatile main memory controllers. In *Proceedings of the Twenty-First International Conference on Architectural Support for Programming Languages and Operating Systems* (2016), ASPLOS '16, ACM, pp. 263–276.

[3] Chen, E., Apalkov, D., Diao, Z., Driskill-Smith, A., Druist, D., Lottis, D., Nikitin, V., Tang, X., Watts, S., Wang, S., Wolf, S. A., Ghosh, A. W., Lu, J. W., Poon, S. J., Stan, M., Butler, W. H., Gupta, S., Mewes, C. K. A., Mewes, T., and Visscher, P. B. Advances and future prospects of spin-transfer torque random access memory. *IEEE Transactions on Magnetics 46*, 6 (June 2010), 1873–1878.

[4] Chhabra, S., and Solihin, Y. i-NVMM: a secure non-volatile main memory system with incremental encryption. In *38th Annual International Symposium on Computer Architecture* (2011), pp. 177–188.

[5] Eisebitt, S., Lüning, J., Schlotter, W. F., Lörgen, M., Hellwig, O., Eberhardt, W., and Stöhr, J. Lensless imaging of magnetic nanostructures by X-ray spectro-holography. *Nature 432* (Dec. 2004), 885–888.

[6] Halupka, D., Huda, S., Song, W., Sheikholeslami, A., Tsunoda, K., Yoshida, C., and Aoki, M. Negative-resistance read and write schemes for stt-mram in 0.13 um cmos. In *2010 IEEE International Solid-State Circuits Conference - (ISSCC)* (Feb 2010), pp. 256–257.

[7] Jin, Y., Shihab, M., and Jung, M. Area, power, and latency considerations of STT-MRAM to substitute for main memory. In *Proceedings of Memory Forum* (2014).

[8] Kan, J. J., Park, C., Ching, C., Ahn, J., Xue, L., Wang, R., Kontos, A., Liang, S., Bangar, M., Chen, H., Hassan, S., Kim, S., Pakala, M., and Kang, S. H. Systematic validation of 2x nm diameter perpendicular mtj arrays and mgo barrier for sub-10 nm embedded stt-mram with practically unlimited endurance. In *2016 IEEE International Electron Devices Meeting (IEDM)* (Dec 2016), pp. 27.4.1–27.4.4.

[9] Karaklaji, D., Schmidt, J. M., and Verbauwhede, I. Hardware designer's guide to fault attacks. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems 21*, 12 (Dec 2013), 2295–2306.

[10] Lee, J., Ganesh, K., Lee, H. J., and Kim, Y. FESSD: A fast encrypted ssd employing on-chip access-control memory. *IEEE Computer Architecture Letters PP*, 99 (2017), 1–1.

[11] Mller, T., and Freiling, F. C. A systematic assessment of the security of full disk encryption. *IEEE Transactions on Dependable and Secure Computing 12*, 5 (Sept 2015), 491–503.

[12] Qiu, Z. Q., and Bader, S. D. Surface magneto-optic kerr effect. *Review of Scientific Instruments 71*, 3 (2000), 1243–1255.

[13] rok On, H., hyung Cho, B., Cho, W. Y., Kang, S., gil Choi, B., jin Kim, H., sung Kim, K., eung Kim, D., keun Kwak, C., geun Byun, H., tae Jeong, G., silk Jeong, H., and Kim, K. Enhanced write performance of a 64 mb phase-change random access memory. In *ISSCC. 2005 IEEE International Digest of Technical Papers. Solid-State Circuits Conference, 2005.* (Feb 2005), pp. 48–584 Vol. 1.

[14] Sheu, S. S., Chang, M. F., Lin, K. F., Wu, C. W., Chen, Y. S., Chiu, P. F., Kuo, C. C., Yang, Y. S., Chiang, P. C., Lin, W. P., Lin, C. H., Lee, H. Y., Gu, P. Y., Wang, S. M., Chen, F. T., Su, K. L., Lien, C. H., Cheng, K. H., Wu, H. T., Ku, T. K., Kao, M. J., and Tsai, M. J. A 4mb embedded slc resistive-ram macro with 7.2ns read-write random-access time and 160ns mlc-access capability. In *2011 IEEE International Solid-State Circuits Conference* (Feb 2011), pp. 200–202.

[15] Yan, C., Englender, D., Prvulovic, M., Rogers, B., and Solihin, Y. Improving cost, performance, and security of memory encryption and authentication. In *Proceedings of the 33rd Annual International Symposium on Computer Architecture* (2006), ISCA '06, IEEE Computer Society, pp. 179–190.

[16] Young, V., Nair, P. J., and Qureshi, M. K. Deuce: Write-efficient encryption for non-volatile memories. In *Proceedings of the Twentieth International Conference on Architectural Support for Programming Languages and Operating Systems* (2015), pp. 33–44.

[17] Zhang, X., Li, H., Yang, S., and Han, S. On a high-performance and balanced method of hardware implementation for aes. In *Software Security and Reliability-Companion (SERE-C), 2013 IEEE 7th International Conference on* (June 2013), pp. 16–20.