

Quality of Service implications of Enhanced Program Algorithms for Charge Trapping NAND in future Solid State Drives

Alessandro Grossi, Lorenzo Zuolo, Francesco Restuccia, Cristian Zambelli, *Member, IEEE*, and Piero Olivo

Abstract— 3D-NAND memories based on Charge Trapping (CT) technology represent the most promising solution for hyper-scaled Solid State Drives (SSD). However, the intrinsic low reliability offered by that storage medium leads to a high number of errors requiring an extensive use of complex Error Correction Codes (ECC) and advanced read algorithms such as Read Retry. This materializes in an overall SSD's Quality of Service (QoS) reduction. In order to limit the errors number, enhanced program algorithms able to improve the reliability figures of CT memories have been introduced. In this work, the impact of such program algorithms combined with Read Retry and ECC is experimentally characterized on CT-NAND arrays. The results are then exploited for co-simulations at system level, assessing reliability, performance and QoS of future SSD integrating CT-based memories.

Index Terms—CT-NAND; bit error rate; read retry; program algorithms; reliability; performance; quality of service; solid state drives

I. INTRODUCTION

SSDs are now the most effective solution for fast mass storage systems in cloud services and high performance computing [1]. One main SSDs' limitation is their reliability, which is dependent upon the non-volatile NAND memories used as the storage medium. This reliability progressively decreases because of their intrinsic wear-out [2]. A direct indication of this phenomenon is an increase in the Bit Error Rate (BER) in a NAND memory. The BER is the percentage of bits in error after a single read operation [2]. Such an increase translates into the inability to correct data after a number of Program/Erase operations (i.e., P/E cycles) or after long retention times.

To deal with increasing BERs, NAND vendors have introduced new read techniques, such as the Read Retry (RR) algorithms [3], bridging the gap between the internal memory algorithms and ECC: if, during a read operation, the number of errors in a page exceeds the ECC's correction capability ECC_{th} , the operation is repeated modifying read parameters with the aim of reducing the number of read errors. However, the joint use of sophisticated RR algorithms and advanced ECC correction engines such as BCH or Low Density Parity Check (LDPC) deeply impact on the overall SSD performance since they reduce read bandwidth while increasing power consumption [4]. The overall metric considering latency,

bandwidth, power consumption, endurance and retention is represented by the Quality of Service (QoS) [5].

The wide adoption of SSDs also in enterprise scenarios requires newly scaled NAND memories with higher storage densities, while complying to the same QoS requirements for the former storage medium generations. Moreover, since SSD's form factors are limiting the space allocation for memory chips, there are new constraints on the development of the Printed Circuit Boards hosting the SSDs.

3D-NAND solutions are to be rapidly adopted in order to take into account all the previous considerations [6]. The 3D-NAND concept is based on the CT technology [7], [8], that is preferred with respect to the standard floating gate NAND Flash since the stored information depends on the amount of charge trapped into discrete traps present in a storage layer. In such a way the cells scalability is increased while reducing the coupling effects and the typical disturbances of a floating gate-based technology.

CT memories however, suffers for a relatively low reliability since part of the trapped charge constituting the storage information is located in shallow traps that can be rapidly emptied. Therefore, the charge measured in read operations, and hence the threshold voltage determining the logical stored information, may be different to the one determined during program [9]–[11]. This problem is further aggravated by the oxide degradation related to the physical mechanisms exploited to introduce or remove charge into/from the storage layer, thus resulting in a low number of P/E cycles (i.e., the endurance), a reduced retention time, and a high BER [12]–[14].

To overcome these reliability issues in future CT-based SSDs, a massive exploitation of RR algorithms and data correction by ECC engines would be mandatory, thus resulting in a degraded QoS. As a consequence, in order to make CT technology appealing in hyper-scaled SSDs, it is necessary to reduce memories' BER and, consequently, the requests for ECC and RR intervention. Since the high BER in CT-NAND is mainly due to charge loss in shallow traps, the straightforward solution is to design program algorithms able to stabilize the trapped charge. The enhanced program algorithms presented in literature [11], [15] will result in a programming time increase compared to the standard program algorithms [16], but their advantages in terms of reliability and overall QoS are terrific.

In this work, by characterizing Multi-Level-Cells (MLC) CT-NAND arrays it was possible to understand the impact of different program algorithms in terms of performance and

A. Grossi, L. Zuolo, F. Restuccia, C. Zambelli and P. Olivo are with Dipartimento di Ingegneria, Università degli Studi di Ferrara, Via Saragat 1, Ferrara (Italy), 44122.




Fig. 1. MLC target distributions and LP, UP discrimination levels (V_1 , V_2 , V_3)

reliability at a single memory level. The results are then used in simulations using a dedicated SSD co-simulation environment [17] to assess the QoS implications in future SSD architectures integrating multiple CT-based memories.

II. READ OPERATIONS BACKGROUND

A. MLC Read operation

In this paper we consider MLC architectures, where each Word-Line (WL) contains an upper page (storing the MSB) and a lower page (storing the LSB). By considering the standard MLC NAND Flash coding (see table I), lower page (LP) is read by applying a read voltage V_2 , whereas upper page (UP) is read by applying a read voltage pair (V_1, V_3) as depicted in Fig. 1. Therefore, when reading a LP, an error occurs when a cell in L1 moves to L2 and vice versa. On the contrary, when reading an UP, two different errors are possible: a bit flip either between E and L1 or L2 and L3.

TABLE I
MLC STANDARD NAND FLASH CODING

	E	L1	L2	L3
UP	1	0	0	1
LP	1	1	0	0

B. Read Retry procedure

Among several optimized reading techniques, RR allows a dynamic adaptation of the read reference voltages: when $BER > ECC_{th}$, the algorithm shifts upwards or downwards the read reference voltage and repeats the read operation until $BER \leq ECC_{th}$ as sketched in Fig. 2. If after a maximum number of attempts BER is still higher than ECC_{th} , the page is considered as failed, and therefore unrecoverable. Since any page have a different BER and since it is not known *a priori* whether an up-shift or a down-shift must be applied (the former required to deal with endurance effects [18], the latter to take into account retention problems [18], [19]), it is not predictable which solution provides the best results, eventually burdening on the read time predictability and on the reliability.

III. EXPERIMENTAL BER MEASUREMENT

A. Experimental setup

The different programming techniques for CT-based memories have been experimentally tested on 4Mbits 2D CT-NAND test vehicles manufactured in a sub-4X technology node. The memory cells feature a p-Si/SiO₂/Si₃N₄/SiO₂/Si₃N₄/Al₂O₃ stack overwhelmed by a high work function TaN/Ti/TaN metal




Fig. 2. Read Retry technique schematic: a) Threshold voltage shift induced by increased number of writing operations. Several cells may result in a threshold voltage higher than the reference read voltage V_{READ} , thus producing a read error. b) When the number of erroneous bits is too high to be corrected by ECC the read reference voltage is shifted.




Fig. 3. Schematic representation of the CT cell tested in this work.

gate (Fig. 3). Such a stack will likely be present also in 3D-NAND architectures [7], [8], and therefore the issues retrieved on a traditional 2D technology are inherited by 3D architectures. The array architecture consists of a standard NAND array, whose pages organization is indicated in Fig. 4. The program and the read operations are performed page-wide. The erase operation is performed block-wide with a single voltage pulse featuring 19 V amplitude and 100 μ s duration.

B. Program Algorithms

The standard Incremental Step Pulse Program (ISPP) [16] algorithm and the enhanced one to reduce the charge loss suffered by CT-NAND, hereafter denoted as Recovery (REC) [15], [18], are depicted in Fig. 5. MLC paradigm has been performed by defining three target program distributions (L1, L2 and L3). Program operation has been performed by applying the Full-sequence paradigm [20]: after every program pulse a read-verify operation has been performed in order to check the cells state and to stop the algorithm execution when the target distribution is reached. ISPP algorithm has been

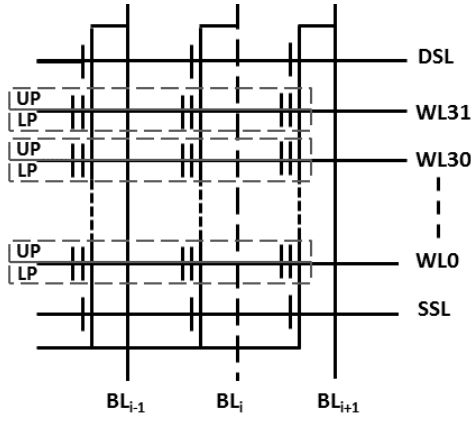


Fig. 4. Single block of the 4 Mbits CT-NAND array considered in this work.

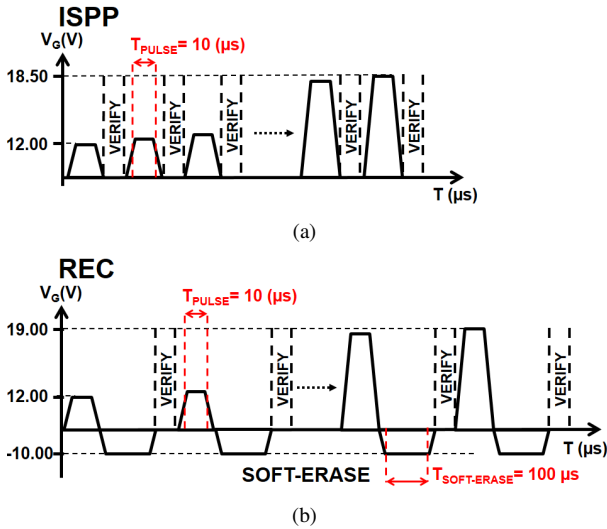


Fig. 5. Schematic of ISPP (a) and REC (b) algorithms.

performed by increasing the pulse voltage from 12 V up to 18.5 V (depending on the target level) with 0.25 V steps and 10 μ s duration (Fig. 5a), with a maximum programming time of 1.23 ms for L3 distribution.

REC has been performed by increasing the pulse voltage from 12 V up to 19 V with 0.25 V steps and 10 μ s duration and by applying a soft erase pulse after every program pulse with a constant voltage of -10 V and 100 μ s duration (Fig. 5b), with a maximum programming time of 6.97 ms for L3 distribution. These values have been chosen to minimize the disturbs and other unwanted phenomenon due to the soft-erase operation applied on all the cells within a common block [21].

In program operation, electrons cross the equivalent oxide barrier and are randomly captured by deep and shallow traps of the nitride storage layer [22]. In cells programmed by ISPP, electrons in shallow traps are easily de-trapped during storage period, thus charge loss is observed (Fig. 6a) [11]. On the contrary, REC is able to stabilize the stored charge during programming, by reducing the presence of electrons trapped in shallow traps with a high de-trapping probability even at low electric fields (Fig. 6b). The soft programming pulses applied during REC are able to remove charge from shallow traps

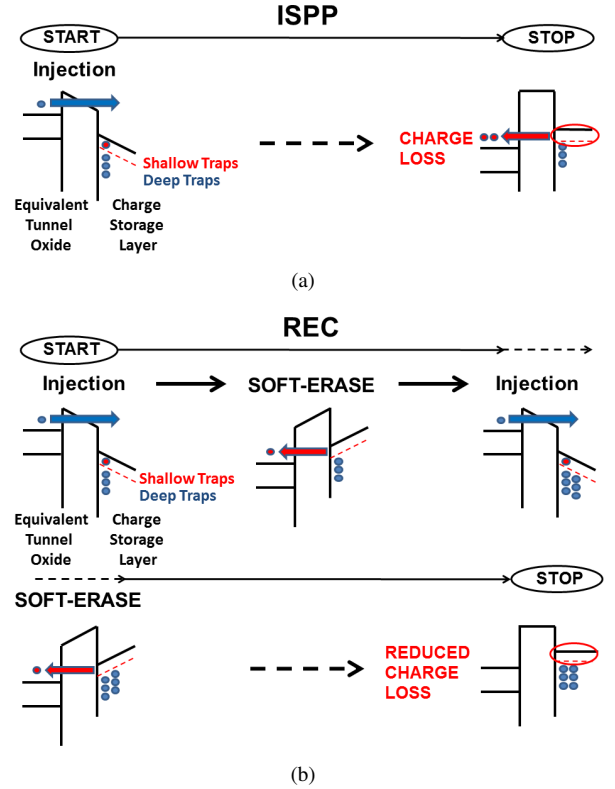


Fig. 6. Illustration of trapped charge distribution during ISPP (a) and REC (b) program algorithms and charge loss mechanisms after writing.

before verify operations occur. Therefore, the target voltage threshold mainly depends on charge stabilized in deep traps [15].

C. BER results

Fig. 7 shows V_T distributions shifts and broadening at different endurance cycles when REC algorithm is performed. Results are even worse for the ISPP algorithm. As it can be seen, the most dangerous effect is due to L1 distribution crossing the V2 read reference voltage. For this reason, in the rest of the paper we will consider LP errors as source of reliability decrease, which represent the worst case without lack of generality.

Fig. 8 shows the BER calculated on LP during endurance cycles obtained with ISPP before (a) and after (b) RR application. Thanks to RR, a BER reduction can be observed. Nevertheless a rapid BER increase occurs after 6k P/E cycles. Moreover, due to the edge wordline effects [23], all the cells on WL0 of a memory string show significantly higher BER compared to the average value, further increasing the average BER. Fig. 9 shows the BER calculated on LP during endurance cycles obtained with REC before (a) and after (b) RR application: in both cases the BER is reduced compared to ISPP and it is shown to be lower than 10^{-3} up to 20k cycles thanks to RR procedure. Fig. 10 shows the percentages of uncorrectable pages calculated during cycling for both ISPP and REC algorithms: a page has been considered uncorrectable by the ECC when more than 100 errors are detected on




Fig. 7. REC V_T distributions after a Recovery algorithm at P/E=1, 2k, 10k.




Fig. 8. ISPP programmed cells LP-BER vs. P/E cycle calculated without (a) and with (b) RR procedure, respectively. ECC_{TH} limit corresponding to 100 errors per read page is shown.

the page ($BER > ECC_{TH}$) and all RR attempts have been applied. Although RR procedure allows improving ISPP performances, all pages become uncorrectable after 5k cycles. The usage of REC program algorithm combined with RR procedure, on the contrary, allows keeping null the percentage of uncorrectable pages up to 20k cycles. As a comparison, the same uncorrectable pages percentages have been calculated not considering the contributions coming from cells on WL0 pages. However, no significant advantages are obtained in this case. In fact, when ISPP with RR is considered, all the pages are shown to suddenly cross the ECC_{TH} limit after 6k P/E cycles, hence the failures are almost simultaneous. On the contrary, when REC with RR is considered, all pages' BER is below ECC_{TH} limit up to 20k P/E cycles, hence there is no perceived impact.

IV. IMPACT ON SYSTEM LEVEL

In this section, the different programming algorithms performance/reliability figures obtained through the characterization of the CT-NAND arrays are evaluated from a SSD-QoS perspective. This task has been performed exploiting a co-simulation framework able to extract performance and latency of a target disk architecture as a function of memory wearout, while allowing detailed reliability analysis [17], [24], [25]. Fig. 11 shows the baseline architecture modeled by the

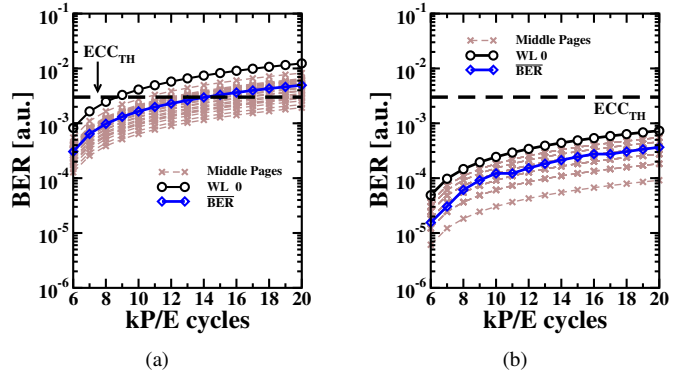


Fig. 9. REC programmed cells LP-BER vs. P/E cycle calculated without (a) and with (b) RR procedure, respectively. ECC_{TH} limit corresponding to 100 errors per read page is shown.

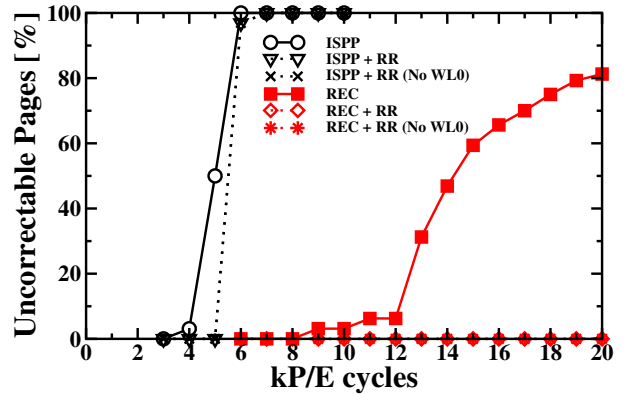


Fig. 10. Uncorrectable pages percentage calculated at different endurance cycles.

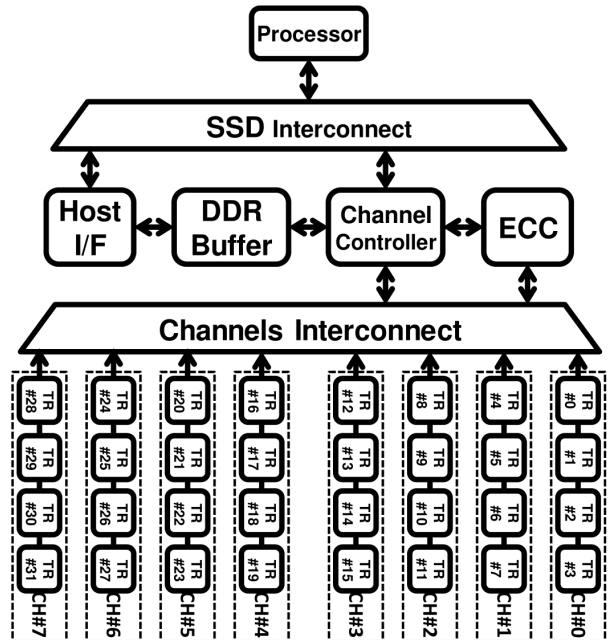


Fig. 11. SSD baseline architecture modeled by the simulator.

simulator. As it can be seen it is composed by a processor, a host interface, a DRAM buffer, a channel controller, a multi-

TABLE II
ARCHITECTURE CONFIGURATIONS

	A	B	C	D	E
Channels	8	8	4	2	1
Targets	4	2	4	8	8

threaded BCH ECC engine [26] and a regular matrix of non-volatile memory targets. Different CT-NAND chips sharing the same bus are defined as channels. Each single channel is connected in parallel on the same channel controller.

In the following system level analysis both the ISPP and the REC algorithms embodying RR algorithms will be considered since this represents a realistic study case used in all SSD platforms [27], [28]. A QoS threshold value has been set for an enterprise scenario, in terms of host interface bandwidth. For a standard interface like the PCIexpress Gen2 x8 [29] the limits corresponds to 4 GB/s which is further reduced to 1150 MB/s due to the host system I/O drivers overhead [30], [31].

In the simulated architectures a SSD is considered to miss the QoS target when the achieved bandwidth is below such threshold. It is worth to point out that if the bandwidth achieved by the SSD is greater than the target QoS, any performance fluctuation introduced by the programming algorithm, the ECC, and the RR are not exposed to the end user, since the perceived overall SSD bandwidth is the one imposed by the QoS limit.

Program algorithm's QoS implications have been investigated on multiple SSD configurations exploiting different number of channels and targets (see table II). The results of the simulations on different SSD architectures, performed by using a 100 % sequential read workload after a program operation at different endurance cycles using either the ISPP or the REC, are shown in Fig. 12. As it can be seen, only the configuration A is able to satisfy enterprise QoS requirements for a defined number of P/E cycles that varies by using one of the two programming algorithms. When considering SSD architectures using the ISPP algorithm, in all cases a sudden performance drop is experienced after 3k P/E cycles because of the too high bit error density. Configurations A, B, and C show a faster QoS degradation around 3k P/E cycles with respect to D and E. This effect is due to the high number of flash targets saturating the channel bandwidth, hence masking the initial performance drop. A similar behavior is observed for SSDs using the REC algorithm, although a larger number of P/E cycles can be experienced before performance degradation.

Since the scope of this work is to show CT-NAND-based SSDs' exploitation in hyper-scaled systems, the following analysis will be focused on enterprise-class SSD's, hence only configuration A will be considered further on. This configuration meets the high density requirements provided by 3D memories exploiting CT technology that is leading to use these solutions for cold storage applications, in which data are written once and read many times [32]. As a consequence, read-intensive workload is the most expected use case for such memories, whereas write-intensive workload represent a corner case.

The 100 % sequential read throughput performances ob-

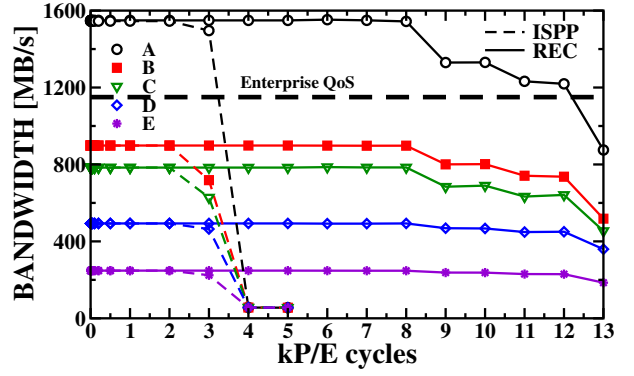


Fig. 12. 100 % sequential read throughput calculated at different endurance cycles with different architecture configurations for ISPP and REC, respectively.




Fig. 13. Sequential read throughput calculated at different endurance cycles.

tained with ISPP and REC algorithms are reported in Fig. 13: while with ISPP no more than 3k P/E cycles are possible before QoS degradation, the endurance obtained with REC is improved by a factor of 4, reaching almost 12k P/E cycles. The experienced bandwidth degradation is due to the joint effect of the RR and ECC engines whose execution time increases with the number of errors to be corrected. Moreover, by considering that WL0 pages usually show higher BER than others, if it is used a Flash Translation Layer in the SSD that excludes those pages from programming, REC throughput can get a further endurance boost of ≈ 500 P/E cycles. On the contrary, when ISPP is considered, no relevant advantages are obtained in terms of bandwidth by using this approach, because of the too high bit error density. Finally, in the case of a 100 % random read workload, results obtained perfectly matched those achieved with the 100 % sequential read workload. This phenomenon is due to the high parallelism offered by the simulated SSD architectures, which is able to sustain the output bandwidth even when random operations are issued. To this extent, random results are not shown in the paper.

The usage of enhanced program algorithms like the REC in SSD architectures carries the drawback of an increased program time. To this extent, in order to understand if a longer program time could impact the perceived SSD's bandwidth, two mixed traffic scenarios have been considered: 75% write and 25% read, 75% read and 25% write (both using 4 kB

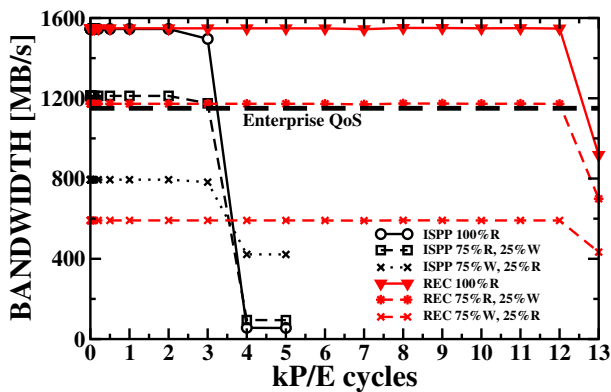


Fig. 14. Random mixed traffic throughput calculated at different endurance cycles.

interleaved random read and write operations) [33]. The latter is the one closer to the scenario that are targeted in this paper. SSD throughput results are shown in Fig. 14. Even if in write-intensive conditions ISPP shows a slightly higher throughput compared to REC, both algorithms do not satisfy the QoS enterprise requirements. REC allows satisfying enterprise QoS requirements up to 12k P/E cycles if read-intensive or 100 % read conditions are considered. Bandwidths obtained with mixed scenarios involving more than 25% write operations cross the QoS limit for both ISPP and REC algorithms. However, it is worth pointing out that such workloads represent a worst-case corner for hyper-scaled SSDs which are foreseen in cold storage scenarios.

V. CONCLUSIONS

In this work the results obtained through CT-NAND test vehicles characterization have been used to investigate the program algorithm implications in terms of SSD's QoS. Thanks to a paved charge loss reduction, an enhanced program algorithm is shown to increase the overall memory reliability of the memory compared to the standard programming algorithm. The endurance gain in different SSD architectures for enterprise environments is quantified by a factor four. As a consequence, the advantages in terms of SSD's QoS are demonstrated to be outstanding for cold storage scenarios.

REFERENCES

- [1] R. Micheloni, A. Marelli, and K. Eshghi, Ed., *Inside Solid State Drives (SSDs)*. Springer London, Limited, 2012.
- [2] N. Mielke, T. Marquart, W. Ning, J. Kessenich, H. Belgal, E. Schares, F. Trivedi, E. Goodness, and L. Nevill, "Bit error rate in NAND flash memories," in *IEEE International Reliability Physics Symposium (IRPS)*, Apr 2008, pp. 9–19.
- [3] Y. Cai, E. F. Haratsch, O. Mutlu, and K. Mai, "Threshold voltage distribution in MLC NAND flash memory: Characterization, analysis, and modeling," in *Design, Automation Test in Europe Conference Exhibition (DATE)*, Mar 2013, pp. 1285–1290.
- [4] K. Zhao, W. Zhao, H. Sun, T. Zhang, X. Zhang, and N. Zheng, "Ldpc-in-ssd: Making advanced error correction codes work effectively in solid state drives," in *USENIX Conference on File and Storage Technologies (FAST'13)*, 2013, pp. 243–256.
- [5] Intel, "Intel SSD Data Center S3700 Series Quality of Service," Jun 2013. [Online]. Available: <http://www.intel.com/content/www/us/en/solid-state-drives/ssd-dc-s3700-quality-service-tech-brief.html>
- [6] B. Prince, Ed., *Vertical 3D Memory Technologies*. Wiley, 2014.
- [7] J.-W. Im, W.-P. Jeong, D.-H. Kim, S.-W. Nam, D.-K. Shim, M.-H. Choi, H.-J. Yoon, D.-H. Kim, Y.-S. Kim, H.-W. Park, D.-H. Kwak, S.-W. Park, S.-M. Yoon, W.-G. Hahn, J.-H. Ryu, S.-W. Shim, K.-T. Kang, S.-H. Choi, J.-D. Ihm, Y.-S. Min, I.-M. Kim, D.-S. Lee, J.-H. Cho, O.-S. Kwon, J.-S. Lee, M.-S. Kim, S.-H. Joo, J.-H. Jang, S.-W. Hwang, D.-S. Byeon, H.-J. Yang, K.-T. Park, K. Hyun Kyung, and J.-H. Choi, "A 128Gb 3b/cell V-NAND flash memory with 1Gb/s I/O rate," in *IEEE International Solid-State Circuits Conference (ISSCC)*, Feb 2015, pp. 1–3.
- [8] A. Nitayama and H. Aochi, "Bit Cost Scalable (BiCS) flash technology for future ultra high density storage devices," in *International Symposium on VLSI Technology Systems and Applications (VLSI-TSA)*, Apr 2010, pp. 130–131.
- [9] H. Park, G. Bersuker, D. Gilmer, K. Lim, M. Jo, H. Hwang, A. Padovani, L. Larcher, P. Pavan, W. Taylor, and P. Kirsch, "Charge loss in tanos devices caused by vt sensing measurements during retention," in *IEEE International Memory Workshop (IMW)*, May 2010, pp. 1–2.
- [10] W. Tsai, N. Zous, C. Liu, C. Chen, T. Wang, S. Pan, C.-Y. Lu, and S. Gu, "Data retention behavior of a sonos type two-bit storage flash memory cell," in *IEEE International Electron Devices Meeting (IEDM)*, Dec 2001, pp. 32.6.1–32.6.4.
- [11] C.-P. Chen, H.-T. Lue, C.-C. Hsieh, K.-P. Chang, C.-C. Hsieh, and C.-Y. Lu, "Study of fast initial charge loss and its impact on the programmed states V_T distribution of charge-trapping NAND Flash," in *IEEE International Electron Devices Meeting (IEDM)*, Dec 2010, pp. 5.6.1–5.6.4.
- [12] G. Ghidini, C. Scozzari, N. Galbiati, A. Modelli, E. Camerlenghi, M. Alessandri, A. D. Vitto, G. Albin, A. Grossi, T. Ghilardi, and P. Tessorio, "Cycling degradation in TANOS stack," *Microelectronics Engineering*, vol. 86, pp. 1822–1825, 2009.
- [13] H. Park, G. Bersuker, M. Jo, D. Veksler, K. Lim, D. Gilmer, N. Goel, C. Kang, C. Young, M. Chang, H. Hwang, H. Tseng, P. Kirsch, and R. Jammy, "Tunnel oxide degradation in TANOS devices and its origin," in *International Symposium on VLSI Technology Systems and Applications (VLSI-TSA)*, Apr 2010, pp. 50–51.
- [14] H. Park, G. Bersuker, D. Gilmer, K. Lim, M. Jo, H. Hwang, A. Padovani, L. Larcher, P. Pavan, W. Taylor, and P. Kirsch, "Charge loss in TANOS devices caused by V_T sensing measurements during retention," in *IEEE International Memory Workshop (IMW)*, May 2010, pp. 1–2.
- [15] C. C. Yeh, W. J. Tsai, T. C. Lu, H. Y. Chen, H. C. Lai, N. Zous, Y. Y. Liao, G. D. You, S. Cho, C. Liu, F. S. Hsu, L. T. Huang, W. S. Chiang, C. J. Liu, C. F. Cheng, M. H. Chou, C. H. Chen, T. Wang, W. Ting, S. Pan, J. Ku, and C.-Y. Lu, "Novel operation schemes to improve device reliability in a localized trapping storage SONOS-type flash memory," in *IEEE International Electron Devices Meeting (IEDM)*, Dec 2003, pp. 7.5.1–7.5.4.
- [16] K.-D. Suh, B.-H. Suh, Y.-H. Um, J.-K. Kim, Y.-J. Choi, Y.-N. Koh, S.-S. Lee, S.-C. Kwon, B.-S. Choi, J.-S. Yum, J.-H. Choi, J.-R. Kim, and H.-K. Lim, "A 3.3 V 32 Mb NAND flash memory with incremental step pulse programming scheme," in *IEEE International Solid-State Circuits Conference (ISSCC)*, Feb 1995, pp. 128–129.
- [17] L. Zuolo, C. Zambelli, R. Micheloni, S. Galfano, M. Indaco, S. Di Carlo, P. Prinetto, P. Olivo, and D. Bertozzi, "SSDExplorer: A virtual platform for fine-grained design space exploration of solid state drives," in *Design, Automation Test in Europe Conference Exhibition (DATE)*, Mar 2014, pp. 1–6.
- [18] A. Grossi, C. Zambelli, and P. Olivo, "Bit error rate analysis in charge trapping memories for SSD applications," in *IEEE International Reliability Physics Symposium (IRPS)*, Jun 2014, pp. MY.7.1–MY.7.5.
- [19] S. Amoroso, A. Mauri, N. Galbiati, C. Scozzari, E. Mascellino, E. Camozzi, A. Rangoni, T. Ghilardi, A. Grossi, P. Tessorio, C. Monzio Compagnoni, A. Maconi, A. Lacaita, A. Spinelli, and G. Ghidini, "Reliability constraints for tanos memories due to alumina trapping and leakage," in *IEEE International Reliability Physics Symposium (IRPS)*, May 2010, pp. 966–969.
- [20] L. Crippa and R. Micheloni, "MLC Storage," in *Inside NAND Flash memories*. Springer-Verlag, 2010, pp. 261–298.
- [21] C. Zambelli and P. Olivo, "Statistical investigation of anomalous fast erase dynamics in charge trapping nand flash," *IEEE Electron Device Letters*, vol. 34, no. 4, pp. 514–516, 2013.
- [22] A. Padovani, L. Larcher, D. Heh, and G. Bersuker, "Modeling TANOS Memory Program Transients to Investigate Charge-Trapping Dynamics," *IEEE Electron Device Letters*, vol. 30, no. 8, pp. 882–884, 2009.
- [23] C. Zambelli, A. Chimenton, and P. Olivo, "Analysis of Edge Wordline Disturb in multimegabit charge trapping flash NAND arrays," in *IEEE International Reliability Physics Symposium (IRPS)*, Apr 2011, pp. MY.4.1–MY.4.5.

- [24] L. Zuolo, C. Zambelli, R. Micheloni, D. Bertozzi, and P. Olivo, "Analysis of reliability/performance trade-off in solid state drives," in *IEEE International Reliability Physics Symposium (IRPS)*, Jun 2014, pp. 4B.3.1–4B.3.5.
- [25] L. Zuolo, C. Zambelli, R. Micheloni, S. Galfano, M. Indaco, S. Di Carlo, P. Prinetto, P. Olivo, and D. Bertozzi, "SSDEXplorer: a Virtual Platform for Performance/Reliability-oriented Fine-Grained Design Space Exploration of Solid State Drives," to appear on *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)*, 2015.
- [26] Y. Lee, H. Yoo, I. Yoo, and I.-C. Park, "6.4Gb/s multi-threaded BCH encoder and decoder for multi-channel SSD controllers," in *IEEE International Solid-State Circuits Conference (ISSCC)*, Feb 2012, pp. 426–428.
- [27] R. Frickey, Intel-Corporation, "Data Integrity on 20nm SSDs," 2012.
- [28] J. Yang, Silicon Motion, "High-Efficiency SSD for Reliable Data Storage Systems," 2012.
- [29] PCI-SIG, "PCIe Base 2.1 Specification," 2005. [Online]. Available: <https://www.pcisig.com/specifications/pciexpress/base2>
- [30] D. Cobb and A. Huffman, "NVM Express and the PCI Express SSD Revolution," 2012. [Online]. Available: <http://www.nvmexpress.org/wp-content/uploads/2013/04/IDF-2012-NVM-Express-and-the-PCI-Express-SSD-Revolution.pdf>
- [31] "Flexible I/O tester." [Online]. Available: <http://freecode.com/projects/fio>
- [32] C. Sun, A. Soga, T. Onagi, K. Johguchi, and K. Takeuchi, "A workload-aware-design of 3D-NAND flash memory for enterprise SSDs," in *International Symposium on Quality Electronic Design (ISQED)*, Mar 2014, pp. 554–561.
- [33] "Standard JESD219, Solid-State Drive (SSD) Endurance Workloads," 2010.




Cristian Zambelli received the M.Sc., and the Ph.D. degrees in Electronic Engineering (with honors) from Università degli Studi di Ferrara respectively in 2008, and 2012. Since 2015 he holds an Assistant Professor position with the Dipartimento di Ingegneria of the same institution. His main research interests are focused on the characterization, physics and modeling of non-volatile memories reliability, and algorithmic solutions for reliability/performance trade-off exploitation in Solid State Drives.



Alessandro Grossi received the M.Sc. degree in electronic and telecommunications engineering from the University of Ferrara, Ferrara, Italy, in 2013. He is currently pursuing the Ph.D. degree in Engineering Science within the Department of Engineering, University of Ferrara. His main research interests are focused on the characterization, physics and modeling of emerging non-volatile memories.



Piero Olivo graduated in Electronic Engineering in 1980 at the University of Bologna, where he received the PhD degree in 1987. Since 1994 he is Full Professor of electronics at the University of Ferrara (Italy). His scientific activity concerns theoretical and experimental aspects of microelectronics, with emphasis on physics, reliability and characterization of electron devices and non volatile memories.



Lorenzo Zuolo received the Laurea Magistrale degree (M.Sc.) in Technology for Telecommunications and Electronic Engineering from Università degli Studi di Ferrara respectively in 2012. Currently, he is a Ph.D student in the Dipartimento di Ingegneria of the same institution. His main research interests are focused on architectural/physical simulation of Solid State Disks (SSD) and emerging non-volatile memories.



Francesco Restuccia received the B.Sc. in Information Engineering in 2014 from Università degli Studi di Ferrara. He is currently pursuing the M.Sc. in Electronic and Telecommunications Engineering at the same university.