

Demonstration of Reliable Triple-Level-Cell (TLC) Phase-Change Memory

M. Stanisavljevic, H. Pozidis, A. Athmanathan, N. Papandreou, T. Mittelholzer, and E. Eleftheriou
IBM Research – Zurich, Säumerstrasse 4, CH-8803, Rüschlikon, Switzerland

Abstract—Although phase-change memory is admittedly the most mature of the emerging nonvolatile memory technologies, its eventual mass production and market adoption may depend on its cost, in particular in comparison to DRAM and to NAND Flash. In addition to process complexity, another major factor that affects the cost of a memory technology is the capability to store multiple bits per memory cell. As a notable example, Triple-Level-Cell (TLC) NAND Flash is currently leading the Flash capacity shipments. With this as motivation, we present a combination of electrical sensing techniques and signal processing technologies to demonstrate, for the first time, the viability of reliable, nonvolatile, TLC storage in phase-change memory cells after extended endurance cycling and temperature stress.

I. INTRODUCTION

Phase-Change Memory (PCM) has developed into a mature technology and is considered the top contender for realizing storage-class memory, i.e., a technology that bridges the gap between DRAM and Flash in the memory hierarchy. One notable feature of PCM is the capability to store more than 1 bit of information per cell, also known as Multi-Level-Cell (MLC) storage. Reliable 2 bits/cell storage under temperature and cycling stress has recently been demonstrated in PCM cell arrays [1]. However, there is an inexorable demand for higher memory capacity, driven mainly by applications such as storage and processing of entire databases in memory. Accordingly, we establish in this article that reliable TLC (3 bits/cell) storage and moderate data retention are viable in PCM cell arrays, at elevated ambient temperatures and even after heavy endurance cycling. This is the first time that TLC storage is demonstrated experimentally in actual PCM arrays.

Reliable data storage and retention represent a formidable challenge for any memory technology. This is the case even for the ubiquitous NAND Flash memory. In particular, although NAND Flash manufacturers typically guarantee endurance of 3,000 program/erase cycles in MLC parts, this number drops drastically to a mere 300-500 cycles for TLC parts. As a consequence, at least until very recently, TLC Flash has been used predominantly in commercial applications that do not require high durability.

In PCM in particular, TLC storage and retention are hampered by the limited signal margin, the inherent PCM noise, and the shifting and broadening of the level distributions over time due to resistance drift. The capping of the signal margin is dictated by requirements such as minimization of programming power and avoidance of thermal interference between adjacent memory cells. The low-frequency noise that phase-change materials exhibit [2] distorts the distributions of

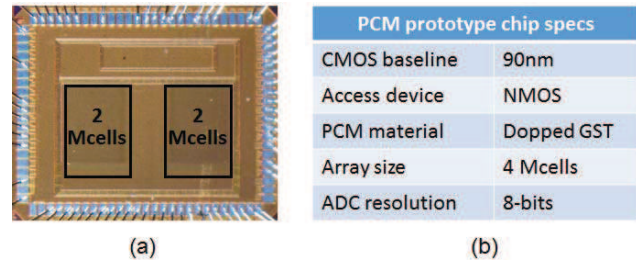


Fig. 1. (a) Micrograph of the prototype PCM chip used in this study and (b) specifications of the prototype chip.

adjacent signal levels and therefore leads to reliability errors. Resistance drift, in turn, is the predominant source of errors in multi-bit PCM as it causes shifting and broadening of the level distributions over time shortly after programming [3].

All of the above reliability limitations are exacerbated in worn-out cells, i.e., cells that have been programmed many times and are close to their end of life. To make things worse, as the electrical conductivity of phase-change materials is temperature dependent, ambient temperature variations cause large deviations in the cell resistance and thus may lead to level crossing and bit-detection errors.

In this paper, we address all of the above reliability issues and demonstrate experimentally the feasibility of TLC storage and moderate data retention in PCM cell arrays. The key components of this demonstration are (a) an advanced, non-resistance cell-state metric that exhibits robustness to drift and PCM noise, and (b) an adaptive level-detection and modulation-coding framework that enables further resilience to drift, noise and temperature variation effects.

II. EXPERIMENTAL PLATFORM

All experiments are conducted on a prototype PCM chip of 4 Mcells with a 4-bank interleaving architecture, shown in Fig. 1. The PCM cells are based on doped $\text{Ge}_2\text{Sb}_2\text{Te}_5$ and are integrated with the peripheral Read/Write circuitry at 90nm CMOS baseline technology [4]. Our experimental platform comprises an FPGA with an embedded processor for system control and data acquisition, as well as a heater to elevate the temperature of the chip. A PID controller monitors a temperature sensor and regulates the power supply to the heater attached under the chip.

III. READ METRIC FOR TLC STORAGE

Conventionally, the stored information in PCM is read by biasing the cell at low voltage and measuring the current

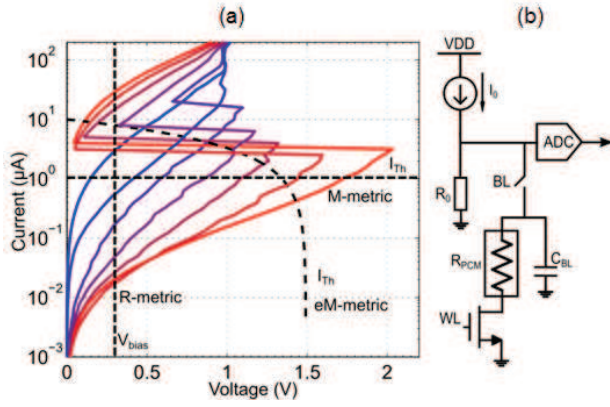


Fig. 2. (a) Measured I - V characteristics of a typical PCM cell programmed at various states. The dashed lines depict the principles of the detection schemes used to extract different read metrics. (b) Circuit diagram of the eM-metric.




Fig. 3. Adaptive level-estimation and decoding scheme adopted for 3 bits/cell data detection. Input y denotes a vector of read values from a collection of PCM cells, while \tilde{y} is the sorted version of y .

flowing through it. This is known as the low-field resistance metric (R-metric). However, the R-metric is hampered by resistance drift, whereby the cell resistance fluctuates around an average value that increases over time. Drift causes time-dependent broadening of the stored level distributions, which limits TLC data retention. In addition, the R-metric exhibits compression and low SNR at highly resistive states, and hence it is not well-suited for reliable TLC storage [5].

Given the limitations of the R-metric, alternate cell-state metrics that are more suitable for multi-bit storage have been explored. Recently, it has been shown that a non-resistance-based cell-state metric (M-metric) offers significant drift immunity compared with the R-metric [5], [6]. However, the drift robustness of the M-metric comes at the cost of latency penalty compared with the R-metric [7]. Also, because of the constant detection current used in the M-metric, the intermediate levels close to SET states are not efficiently detected. Therefore, a new readout metric was proposed in [1], that combines the advantages of the R-metric for intermediate states closer to SET and those of the M-metric for states closer to RESET, along with low latency. The idea behind the so-called enhanced M-metric (eM-metric) is to apply a variable detection current rather than a constant voltage bias (R-metric) or a constant detection current (M-metric). The principle of the eM-metric is illustrated in Fig. 2 and compared with the R- and M-metric principles in I - V curves from a PCM cell that has been programmed at various cell states. The eM-metric has been integrated in circuitry in 25nm technology as part of a multilevel PCM chip [8]. A read latency of 450 ns (for 2 bits/cell readout) has been experimentally verified.




Fig. 4. Programming eM- I curves averaged across 64k cells as a function of SET/RESET endurance cycles.

IV. DRIFT-TOLERANT CODING AND SIGNAL PROCESSING

Despite the largely improved drift resilience offered by the eM-metric, the stored levels may still exhibit some residual time-dependence. Moreover, temperature variations cause further shifts and broadening of the PCM level distributions. To reliably detect the stored TLC levels, appropriate level detection thresholds have to be placed between the distributions of adjacent levels, and those thresholds need to be adjusted over time according to the shift of levels due to drift or other effects. One way to adjust the detection thresholds is via “reference” cells, i.e., cells with known stored data, which are used to estimate the changing levels over time. However, the use of reference cells entails a loss of storage capacity and is not effective. Instead, coding of the levels stored in TLC PCM is a very efficient means for increasing drift immunity. Unions of permutation modulation codes are used for this purpose as they have been shown to offer far superior performance compared with reference cells [9]. These codes offer further immunity to level variations by storing information in the relative order of levels within a set of cells rather than the absolute values of the levels themselves. In this work, a permutation code of length 32 has been used. A block diagram of the adaptive level-detection process is shown in Fig. 3.

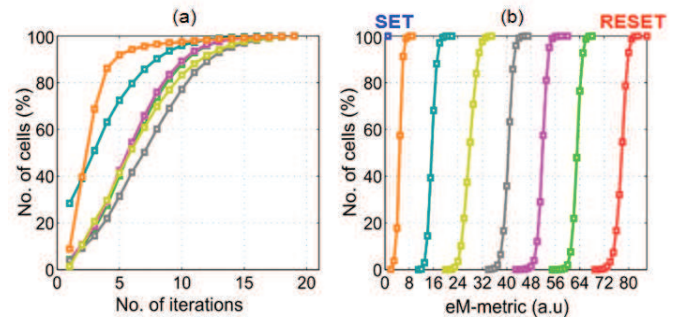


Fig. 5. Iterative programming in TLC PCM: Cumulative distributions of the (a) number of program-verify iterations and (b) programmed levels.

V. RELIABLE TLC DATA STORAGE AND RETENTION

To demonstrate the feasibility of TLC storage in PCM devices, a sub-array of 64k cells is first cycled one million

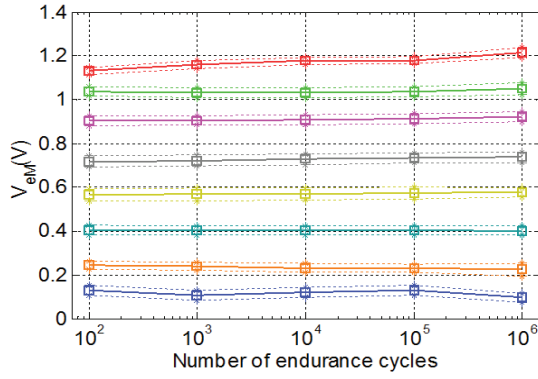


Fig. 6. Evolution of the eM-metric values of the 8 TLC levels as a function of endurance cycles, measured 100 s after cell programming. Solid and dotted lines correspond to mean and $\pm 1\sigma$ values, respectively.

times. The SET/RESET cycling is performed using single voltage pulses having a long trailing edge and high amplitude, respectively. The average programming (eM- I) curves of the cycled cells are shown in Fig. 4. It can be seen that the cell characteristics change moderately as a function of cell wear. After cycling has been completed, TLC data storage follows. Eight levels are optimally placed within the available sensing range between the boundary SET and RESET levels, so as to maximize data retention. Intermediate levels are written with an iterative program-verify (PV) algorithm, that, at every step, adapts the programming current flowing to the cell based on the error between the current eM-metric value of the cell and a pre-defined target value. The effectiveness of this programming scheme is illustrated in Fig. 5, where the number of PV iterations required to program the six intermediate levels, and the TLC level distributions immediately after programming are shown. It can be observed that 99% of the cells converge in up to 15 iterations and that very tight level distributions can be achieved.




Fig. 7. Distributions in eM-metric of the 8 programmed levels across 64k cells 69.8 s and 9.8 days after cell programming.

Fig. 6 shows the TLC level statistics (mean ± 1 standard deviation) as a function of endurance cycles. Despite the changing cell behavior (Fig. 4), the programming scheme is always able to achieve tight level distributions.

In Fig. 7, the effects of drift on the eM-metric are illustrated.

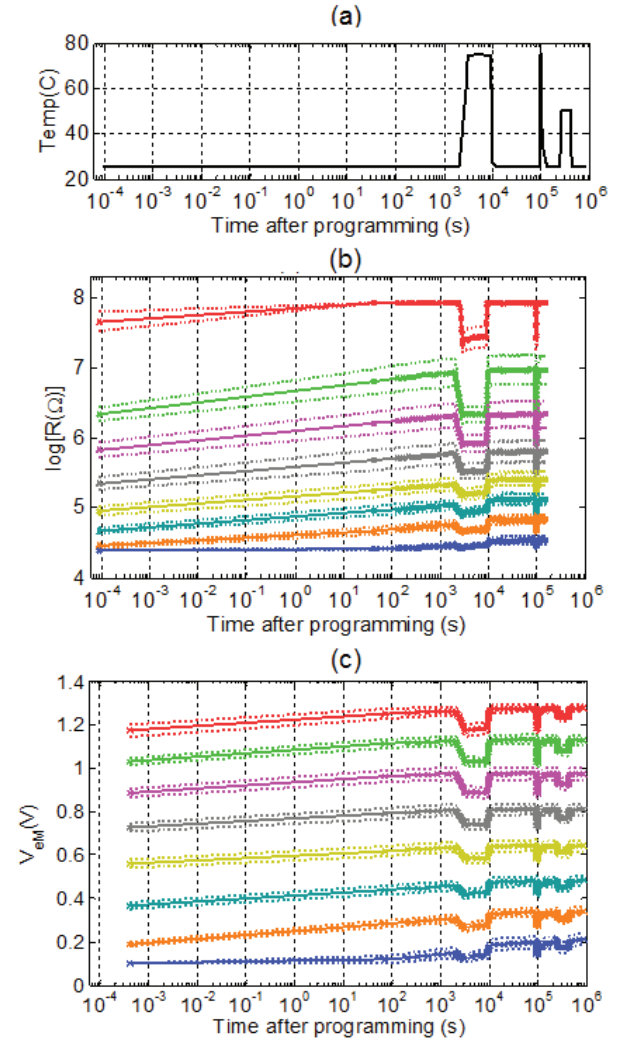


Fig. 8. (a) Time-temperature profile and corresponding trajectories of 8 stored levels for (b) R-metric and (c) eM-metric. Solid and dotted lines depict mean and $\pm 1\sigma$ values, respectively.

As expected, the TLC level distributions shift and broaden over time, resulting in a decrease of the signal margin and of the signal-to-noise ratio. Despite that, the eM-metric is able to maintain a clear separation between adjacent levels even 10 days after cell programming. As will be shown later, this is not the case when the R-metric is used, making it impossible to retain the programmed data even for short times.

Once programmed, the cells are periodically read over time to assess the effect of drift on the read metric (R or eM). At the same time, the chip is exposed to temperatures as high as 75°C to study TLC data retention in the presence of ambient temperature variations. Fig. 8(a) illustrates the time-temperature profile the chip is subjected to. Figs. 8(b) and (c) show the time evolution of the 8 levels in R-metric and eM-metric, respectively. Each figure corresponds to a different experiment, where the 8 levels are optimally placed according to the respective cell-state metric used. Clearly, the eM-metric offers much larger contrast between the 8 levels




Fig. 9. I - V curves for 8k cells and 8 written levels measured 1,000 s after programming at two different ambient temperatures: (a) 25°C and (b) 75°C. Dotted lines indicate $1.2 \mu\text{W}$ read power, a safe limit to avoid read disturb.

even at high temperature excursions. This is also evident from the I - V curves shown in Fig. 9, which exhibit a clear separation between the eM levels. In contrast, the R-metric suffers from high variance and sensitivity of the RESET-like states to drift and temperature variations (Fig. 8(b)).

The suitability of the eM-metric in combination with the adaptive detection scheme for TLC PCM is quantified in Fig. 10, where the bit-error rate (BER) is shown as a function of time. Note that the x -axis corresponds to the time-temperature profile of Fig. 8(a) and that the 64k cells have been pre-cycled one million times. The BER of the eM-metric remains below 3×10^{-4} throughout the 10 days of the experiment, whereas the R-metric exhibits almost 2 orders of magnitude degradation. At this raw BER level, standard low-redundancy ECC schemes can offer highly reliable data retrieval. It should be noted here that, while the eM-metric offers drift and noise immunity, the adaptive detection scheme is instrumental in tracking not only the uni-directional drift effects, but also the more abrupt bi-directional shifts of the levels caused by temperature variations.

VI. CONCLUSIONS

Storage of multiple bits per cell is one of the most desirable features of nonvolatile memory technologies because of the

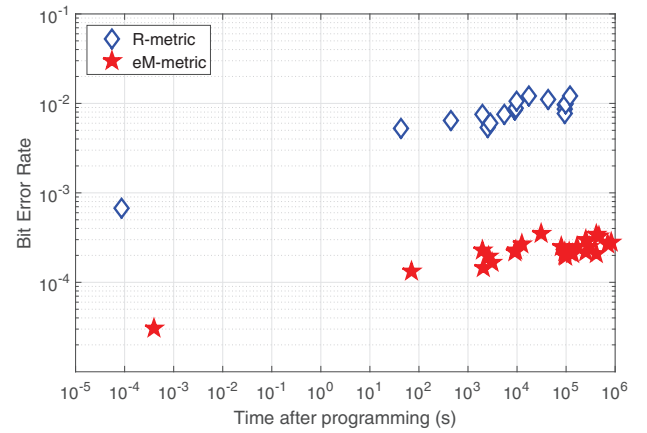


Fig. 10. Bit error rate for R-metric and eM-metric for 3 bits/cell storage. The cells have been subjected to 1 million write cycles and to the time-temperature profile of Fig. 8(a). The level detection process of Fig. 3 is used.

associated cost/bit advantages. This work has demonstrated the feasibility of Triple-Level Cell (TLC) phase-change memory. In particular, we have shown that an array of 64k PCM cells that have been pre-cycled one million times can store 8 levels per cell (3 bits/cell) and retain the stored data over 10 days amidst ambient temperature variations between 25°C and 75°C. The bit error rate remains below 3×10^{-4} throughout the 10-day retention period. Key enabling technologies for this demonstration are a drift-immune readout metric and an efficient modulation coding and adaptive detection framework.

ACKNOWLEDGMENTS

We gratefully acknowledge the support of the PCM teams at IBM Research – Zurich and at the IBM T.J. Watson Research Center, and the IBM/Macronix PCRAM Joint Project team.

REFERENCES

- [1] M. Stanisavljevic, A. Athmanathan, N. Papandreou, H. Pozidis, and E. Eleftheriou, “Phase-change memory: Feasibility of reliable multilevel-cell storage and retention at elevated temperatures,” in *Proc. IRPS*, 2015, pp. 5B.6.1–5B.6.6.
- [2] D. Fugazza, D. Ielmini, S. Lavizzari, and A. L. Lacaita, “Experimental investigation of transport properties in chalcogenide materials through $1/f$ noise measurements,” *Appl. Phys. Lett.*, vol. 88, p. 263506, 2006.
- [3] D. Ielmini, D. Sharma, S. Lavizzari, and A. L. Lacaita, “Reliability impact of chalcogenide-structure relaxation in phase-change memory (PCM) cells – Part I: Experimental study,” *IEEE Trans. Electron Devices*, vol. 56, no. 5, pp. 1070–1077, May 2009.
- [4] J.Y. Wu, M. Breitwisch, S. Kim, et al., “A low power phase change memory using thermally confined tan/tin bottom electrode,” in *IEDM Tech. Dig.*, 2011, pp. 43–46.
- [5] A. Sebastian, N. Papandreou, A. Pantazi, H. Pozidis, and E. Eleftheriou, “Non-resistance-based cell-state metric for phase-change memory,” *J. Appl. Phys.*, vol. 110, p. 084505, 2011.
- [6] N. Papandreou, A. Sebastian, A. Pantazi, et al., “Drift-resilient cell-state metric for multilevel phase-change memory,” in *IEDM Tech. Dig.*, 2011, pp. 55–58.
- [7] A. Athmanathan, M. Stanisavljevic, J. Cheon, et al., “A 6-bit drift-resilient readout scheme for multi-level phase-change memory,” in *Proc. Asian Solid-State Circ. Conf. (A-SSCC)*, 2014, pp. 137–140.
- [8] J. Cheon, I. Lee, C. Ahn, et al., “Non-Resistance Metric based Read Scheme for Multi-level PCRAM in 25nm Technology,” in *Proc. Custom Integrated Circuits Conf.*, 2015, pp. 1–4.
- [9] H. Pozidis, T. Mittelholzer, N. Papandreou, T. Parnell, and M. Stanisavljevic, “Phase change memory reliability: A signal processing and coding perspective,” *IEEE Trans. Magnetics*, vol. 51, no. 4, pp. 1–7, April 2015.