Maximizing MLC NAND lifetime and reliability in the presence of write noise

Borja Peleato and Rajiv Agarwal
Stanford University
Email: {peleato, rajivag}@stanford.edu

Abstract—**The aggressive scaling of the NAND flash technology has led to write noise becoming the dominant source of disturbance in the currently shipping sub-30 nm MLC NAND memories. Write noise can be mitigated by reducing the magnitude of the voltage levels programmed into the cells, which additionally translates to longer flash memory lifetime. However, if all the target levels are small and close together, the probability of error could become excessively high. It is therefore necessary to optimize the target level placement in order to achieve a trade-off between flash lifetime and error probability. This paper proposes a method to maximize flash lifetime subject to reliability constraints, and vice versa. Simulation results show that the proposed method doubles flash lifetime in comparison to a naive scheme, for a 2% reliability constraint. It also comes very close to the optimal solution obtained by brute force search, while maintaining negligible computational complexity in comparison.**

I. INTRODUCTION

Recently, flash memories have emerged as a faster and more efficient alternative to hard drives. However, their higher cost is still an obstacle for their widespread use. Manufacturers have significantly reduced the cost by aggressively scaling the technology and using multi-level per cell (MLC) techniques, *i.e.*, storing several bits in each transistor, but this reduction has come at a price in lifetime and reliability [1][2][3].

Flash memories are currently being used in many different products, each with its own requirements. In industrial applications, such as cars or robots, cost is not as important as in solid state drives, but they require very high reliability. Consumer memory sticks, on the other hand, are not as sensitive to lifetime or reliability, but seek higher density at lower cost. Hence, in consumer applications MLC techniques are used to increase the number of bits per cell. Memories with 4 bits per cell are already commercially available, and it is expected that there will soon be memories with 5 bits per cell.

As the cell size shrinks and more bits are programmed into each one of them, interaction between neighboring cells increases significantly leading to noisier writing and unreliable reading. It is thus a signal processing challenge to write and read back reliably. Some significant sources of noise are inter-cell interference (ICI), random telegraph noise (RTN), read noise, etc. [4],[5]; and all of these increase with NAND/flash (both terms will be used interchangeably) cell size reduction.

Besides inter-cell interference and random noise, there exists a third significant source of disturbance introduced during the writing phase. Typical NAND cells are written by sending a sequence of voltage pulses that aim to increase the voltage of the target cell to a desired level [6]. However, in this process, the voltage level of other neighboring cells is also affected in a random manner. This noise source is referred to as write noise. In the currently shipping sub-30 nm and more so in the future sub-20 nm NAND memories, the write noise dominates over inter-cell interference, read noise, and RTN [7].

The write noise increases with the number and magnitude of programming pulses that are sent into the cell. On one hand, therefore, it is advantageous to have the target voltage levels low so that the number of programming pulses needed is small, leading to faster and less noisy writing. Additionally, if the voltage levels are large, the cell lifetime can be significantly reduced due to the wear and tear of program and erase. On the other hand, if the levels are too close together, there could be confusion between adjacent levels when they are read, leading to a larger error probability. It is therefore necessary to optimize the target levels to achieve a trade-off between cell lifetime (or write throughput, defined as bits written per second) and error probability. This paper tries to maximize the former subject to reliability constraints, and vice versa.

The rest of the paper will be organized as follows. Section II introduces the problem and presents the framework used throughout the rest of the paper. In Section III a model for the dependence of noise on the cell voltage levels is proposed, and validated through simulations. Section IV uses the proposed model to formulate the problem in a convex form, which is then solved in Section V. Finally, Section VI summarizes the paper and proposes directions for further research.

II. PROBLEM STATEMENT

Each cell in a NAND memory is a floating gate transistor which can be charged to a specific voltage value by injecting voltage pulses into the floating gate. In a b-bit MLC flash, each cell stores one of N distinct predefined voltage levels $v_1, v_2, \ldots, v_N$, where $N = 2^b$.

Cells in a NAND flash are organized in terms of pages, which is the smallest unit for writing. Writing to cells in a page is done through a program and verify approach, wherein small pulses are sent to the cell until its voltage level exceeds the desired one. Once this happens, the cell is inhibited from receiving any further pulses and the programming of other cells in the page continues. However, the inhibition mechanism is non-ideal and future pulses change the voltage level of inhibited cells, leading to significant write noise. The write

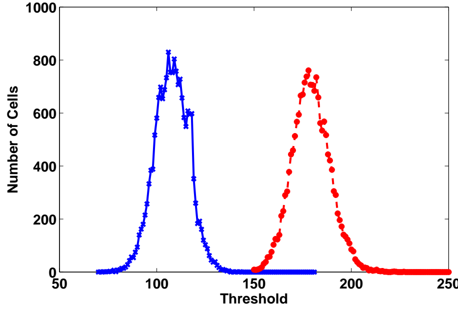

Fig. 1. Histogram of cell voltages in a 2-bit-per-cell flash memory storing random data. It can be observed that the noise is approximately Gaussian.

noise can be modeled to be Gaussian distributed, as shown by the lab collected data in Figure 1.

Another important source of disturbance is inter-cell interference (ICI). Pages in a block of NAND are written sequentially, so page $p+1$ is written after the programming of page p is complete. As a consequence, the voltage levels of cells on page p change during the programming of page $p+1$ and subsequent pages. The window over which the programming of one page affects another depends on the isolation mechanism between them. However, ICI is mostly deterministic, to the point that it is often useful for decoding adjacent pages [8]. For the work in this paper it is assumed that ICI is either avoided at write time with very good isolation mechanisms or subtracted at read time using joint decoding of adjacent pages; hence it is not studied in this paper.

Yet another source of noise is due to retention. Cells suffer voltage leakage from passage of time between the time they are written and the time they are read back, which shifts their voltage values towards zero. The process of reading back the voltage written on a cell is also subject to some noise, but it is negligible in comparison with the one introduced by the previously mentioned sources. The read noise is therefore neglected for the remainder of the discussion in this paper.

Let y_i denote the random variable representing the read back voltage from a cell that was written to a level i. Then

$$y_i = v_i + z_{w,i} + z_{\mathrm{ret},i} + z_{\mathrm{read},i}, \tag{1}$$

where v_i is the voltage level intended to be programmed into the cell and $z_{w,i}$, $z_{\mathrm{ret},i}$, and $z_{\mathrm{read},i}$ respectively represent the write, retention, and read noise for the ith level. The retention noise is modeled as a downward voltage shift with mean $-d_i$.

In current sub-30 nm and in future sub-20 nm NAND technology, write noise has become the dominant source of disturbance. As NAND manufacturers reduce the size of the transistors to increase power and cost efficiency, cells get closer together and the interaction between them is enhanced. This leads to write noise by far dominating all the others, as will be the consideration for the rest of the paper. Therefore, Eq (1) can be changed to

$$y_i = v_i + z_{w,i} - d_i.$$

Using x_i to denote $v_i - d_i$, the above equation can be further simplified to

$$y_i = x_i + z_{w,i}.$$

As discussed previously, write noise is caused by the incremental pulse programming mechanism used for writing each NAND page. The number of pulses and their voltage values increases with the desired voltage level to be programmed. It is therefore desirable to choose small target voltage levels in order to reduce the write noise . Additionally, smaller voltage levels lend themselves to less damage during re-writes, thereby prolonging lifetime. The flipside is that in MLC flash the number of levels N is large and sufficient separation is required between them to read back reliably in the presence of noise, thus requiring large values of target voltage levels.

The goal of this paper is to characterize this trade-off between reliability and flash lifetime in terms of allowed number of program erase cycles and propose an algorithm to find the optimal set of target levels under reliability constraints. In a similar manner, the target levels can be optimized to maximize reliability for a given lifetime requirement. The analytical framework for the work is described next.

III. System model

Let $x \in \mathbf{R}^N$ be the vector of target voltage levels (conveniently adjusted to compensate for retention). The read back voltage value for the ith voltage level is a random variable given by $y_i = x_i + z_{w,i}$, where x_i is the ith component of x and $z_{w,i}$ represents the write noise at the ith level. The write noise for a specific target level can be modeled to be Gaussian, with a mean and variance that depend on its voltage as well as on the voltages of the other target levels. This paper captures such dependence using a simple linear model with different parameters for each level. Specifically, $z_{w,i}$ is modeled as a Gaussian noise process with mean $a_i^T x + c_i - x_i$ and standard deviation $p_i^T x + q_i$, where $a_i, p_i \in \mathbf{R}^N$ and $c_i, q_i \in \mathbf{R}$ are the model parameters. Despite the simplicity of this model, it is possible to achieve highly accurate solutions by using a successive approximations approach, as will be explained in Section IV.

During read back, a set of fixed thresholds is used to distinguish between the different levels. Specifically, there are $N-1$ fixed thresholds chosen to minimize error probability. Let t_i represent the threshold between levels i and $i+1$ so that cells with a voltage between t_{i-1} and t_i will be considered to store level i and cells with a voltage between t_i and t_{i+1} will be considered to store level $i+1$.

Assuming that a cell is equally likely to store any of the levels (which is a reasonable assumption when random data is written onto the flash), the probability of mistaking level i for $i+1$ or level $i+1$ for i is given by

$$P_e^{(i)} = \frac{1}{N} Q\left(\frac{t_i - \mu_i}{\sigma_i}\right) + \frac{1}{N} Q\left(\frac{\mu_{i+1} - t_i}{\sigma_{i+1}}\right), \tag{2}$$

where $\mu_i = a_i^T x + c_i$ and $\sigma_i = p_i^T x + q_i$ are deterministic functions of x representing the mean and standard deviation of the voltage at the ith level and

$$Q(u) = \int_u^\infty \frac{1}{\sqrt{2\pi}} e^{-(1/2)s^2} ds.$$

For a given set of target levels x, the thresholds $t_i^\star$, $i = 1, \ldots, N-1$, which minimize the average probability of error can be obtained by solving the equations

$$\frac{1}{\sigma_i} \exp\left(-\frac{\|t_i^\star - \mu_i\|^2}{2\sigma_i^2}\right) = \frac{1}{\sigma_{i+1}} \exp\left(-\frac{\|\mu_{i+1} - t_i^\star\|^2}{2\sigma_{i+1}^2}\right) \quad (3)$$

for $i = 1, \ldots, N-1$, respectively. Eq (3) can be transformed into a quadratic in $t_i^\star$ and has two solutions: one minimizing Eq (2) and the other maximizing it. The former should be picked. Thus, $t_i^\star$ can be obtained as a function of x.

The claim that these $t_i^\star$ are indeed optimal can be proved by first observing that $P_e^{(i)}$ only depends on the ith threshold t_i and not on any of the others. Therefore, each threshold can be optimized independently as $t_i^\star = \operatorname{argmin}_{t_i} P_e^{(i)}$. Since $P_e^{(i)}$ is a smooth function of t_i with value $1/N$ for $t_i = \pm\infty$ and strictly smaller than $1/N$ for $t_i = \frac{1}{2}(\mu_i + \mu_{i+1})$, it must have a global minimizer fulfilling $\nabla_{t_i} P_e^{(i)} = 0$. Making the gradient with respect to t_i of Eq (2) equal to zero gives Eq (3).

The next section introduces reliability constraints based on minimizing the probability of error with respect to the target values x. $P_e^{(i)}$ depends on both x and $t_i^\star$, but $t_i^\star$ can be put as a function of x by solving Eq (3). However, this function is rather complicated. For this reason, the following approximation on the value of $t_i^\star$ is used for the rest of the discussion in this paper:

$$t_i^\star \simeq \frac{\sigma_i \mu_{i+1} + \sigma_{i+1}\mu_i}{\sigma_{i+1} + \sigma_i}.$$

This approximation has a merit because if all σ_i are equal, it is exact. If the σ_i are unequal, the approximate value of $t_i^\star$ will yield a probability of error higher than the minimum value of $P_e^{(i)}$. Thus, $\min_{t_i} P_e^{(i)}$ will be upper bounded by

$$\tilde{P}_e^{(i)} = \frac{2}{N} Q\left(\frac{\mu_{i+1} - \mu_i}{\sigma_i + \sigma_{i+1}}\right). \quad (4)$$

This approximation has the additional advantage of expressing the probability of mistaking levels i and $i+1$ in terms of a single Q function, the argument of which does not depend on the thresholds. This will help to formulate the problems of interest as convex optimization problems, as described in the next section.

IV. Problem formulations

This section presents a method for finding the target levels x which maximize NAND lifetime under reliability constraints, using the framework described in the previous sections. NAND lifetime generally depends on the target voltage levels chosen for programming the cells. The higher the target voltage levels, the more is the damage caused to the cells during programming, erasing, and reprogramming. Thus, NAND lifetime is modeled to be proportional to $x^T x$. However, this method can be directly applied to any other lifetime model which is a convex function of the levels x.

The probability of error can be defined in different ways depending on the application, coding scheme, etc. This paper will focus on the two most typical scenarios: level error rate and bit error rate, described below.

A. *Level error rate*

The level error rate is defined as the probability that the read level is different from the one that was written to the cell. The trade-off between level error rate and NAND lifetime can be characterized by solving

$$\begin{array}{lll} \text{Minimize} & (1/2)x^T x & (5) \\ \text{subject to} & \sum_{i=1}^{N-1} P_e^{(i)} \leq \gamma & (6) \\ & x_j \geq 0, \quad j = 1, \ldots, N, & (7) \end{array}$$

with variables $x \in \mathbf{R}^N$. Unfortunately, this problem is non-convex and cannot be solved easily.

To make the problem convex two approximations are made. The first was already described in the previous section, wherein $P_e^{(i)}$ was replaced by $\tilde{P}_e^{(i)}$. The second approximation is to replace constraint (6) on the average probability of error by a constraint on each individual probability of error term. The second approximation is justified because in practice, it is desired that all levels have similar probability of error . Specifically, (6) is modified as $\tilde{P}_e^{(i)} \leq \frac{\gamma}{N-1}$, $i = 1, \ldots, N-1$. The problem can then be transformed into

$$\begin{array}{lll} \text{Minimize} & (1/2)x^T x & (8) \\ \text{subject to} & \mu_{i+1} - \mu_i \geq (\sigma_i + \sigma_{i+1}) Q^{-1}\left(\frac{N\gamma}{2(N-1)}\right) & \\ & \qquad i = 1, \ldots, N-1 & (9) \\ & x_j \geq 0, \quad j = 1, \ldots, N, & (10) \end{array}$$

where $\mu_i = a_i^T x + c_i$, $\sigma_i = p_i^T x + q_i$ and γ is a constant. This is a linearly constrained quadratic program in the voltage levels x, which can be solved using, for example, an interior point method. Since (9) is more restrictive than (6), the resulting optimal target levels x obtained as a solution of this problem will satisfy the original constraint (6) as well.

B. *Bit error rate*

In this subsection, the goal is to characterize the trade-off between bit error rate and NAND lifetime. The mapping of bits to voltage levels in the cells is done so that neighboring levels only differ in one bit (Gray coding). Each threshold can therefore be associated with one bit changing. As an example consider the case for $N = 8$ illustrated in Figure 2. If the probability of error between adjacent levels is uniform, as in Section IV-A, some bits suffer significantly higher probability of error than others. Specifically, the LSBs will have a significantly higher error than the MSBs. Finding a set of target levels that offers identical probability of error to all the bits therefore requires solving a problem of the form

$$\begin{array}{ll} \text{Minimize} & (1/2)x^T x \\ \text{subject to} & P_e^{(1)} + P_e^{(3)} + P_e^{(5)} + P_e^{(7)} \leq \gamma \\ & P_e^{(2)} + P_e^{(6)} \leq \gamma \\ & P_e^{(4)} \leq \gamma \\ & x_j \geq 0, \quad j = 1, \ldots, 8, \end{array} \quad (11)$$

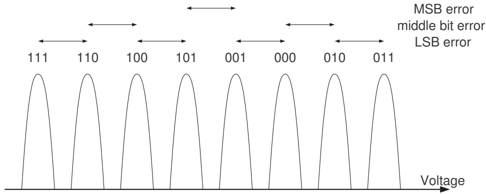

Fig. 2. Gray mapping of bits to levels. Observe that, assuming that errors only happen between adjacent levels, there is only one possible way of bit 1 being wrong: if a mistake is made between levels 4 and 5. However, bit 3 can be wrong regardless of which level was written.

where it has been assumed that there are $N = 8$ levels with Gray mapping for the bits, as in Figure 2. (A similar problem could be written for any other number of levels.) The same two approximations as in Section IV-A are made to transform the above problem into a convex one, now given as

$$\begin{array}{ll} \text{Minimize} & (1/2)x^T x \\ \text{subject to} & \mu_{i+1} - \mu_i \geq (\sigma_i + \sigma_{i+1})Q^{-1}(\gamma) \quad i = 1,3,5,7 \\ & \mu_{i+1} - \mu_i \geq (\sigma_i + \sigma_{i+1})Q^{-1}(2\gamma) \quad i = 2,6 \\ & \mu_5 - \mu_4 \geq (\sigma_4 + \sigma_5)Q^{-1}(4\gamma) \\ & x_j \geq 0, \qquad j = 1,\ldots,8, \end{array} \tag{12}$$

where $\mu_i = a_i^T x + c_i$ and $\sigma_i = p_i^T x + q_i$.

In sections IV-A and IV-B the objective is to maximize lifetime for a given reliability constraint. It should be mentioned that both the optimization problems in sections IV-A and IV-B can be modified to solve the reverse problem of minimizing the probability of error for a given lifetime constraint. The objective would then be to minimize γ instead of $(1/2)x^T x$ and the constraint would be of the form $(1/2)x^T x \leq K$ for some given K.

The method in this section assumes a linear dependence of the noise mean and variance on the target levels, but this may not be the case in general. If this dependence is not linear, then a sequential convex programming method can be used to find the optimal target levels x. Sequential convex programming methods are iterative algorithms that sequentially approximate the non convex parts of the problem with convex functions during their execution. (In this case, the noise model would be approximated with a linear function.) There exists a large body of literature under the name of Sequential Convex Programming or Trust Region methods [9][10][11] devoted to deciding how and how often the approximations should be updated, but these methods are beyond the scope of this paper.

V. SIMULATION RESULTS

This section gives simulation results to evaluate the performance of the optimization problems in Section IV. These problems need an accurate model of the noise dependency on target levels. Such a model can be obtained by taking enough measurements on the NAND to find values of p_i, q_i, a_i, and c_i for $i = 1,\ldots,N$. It is expected that the noise dependence on target levels will vary from one vendor of NAND to another. Hence, the simulations in this section are performed for representative values of p_i, q_i, a_i, and c_i, instead of using values applicable for a particular vendor's NAND

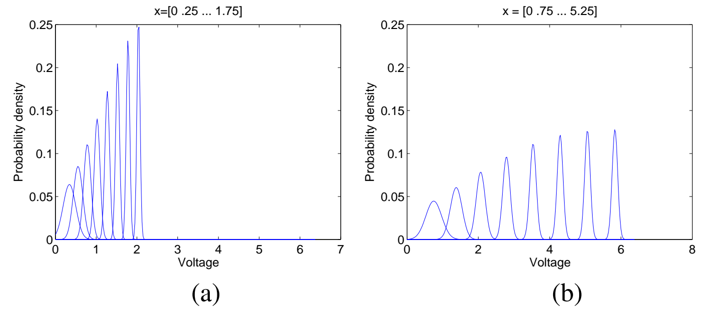

Fig. 3. Voltage probability density functions for uniformly spaced levels, with a spacing of 0.25 volts in (a) and a spacing of 0.75 volts in (b).

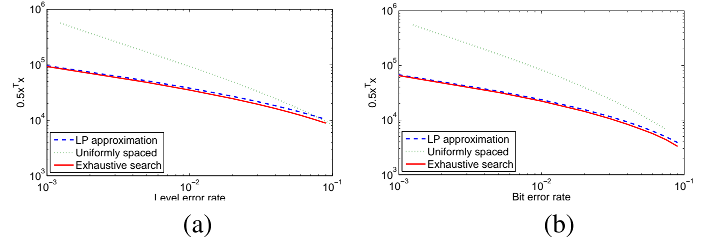

Fig. 4. Trade-off curves between memory lifetime (assumed proportional to $(1/2)x^T x$) and level error rate (a) or bit error rate (b) for a 3-bits-per-cell flash with levels chosen by the proposed scheme, uniform spacing, and brute force search.

only. Figure 3a shows the resulting noise variances and mean shifts for the choice of parameters made in this paper when $x_1 = 0$ and the levels are uniformly spaced with 0.25 volts difference between them. Figure 3b shows the change in noise mean and variance when the level separation is increased to 0.75 volts. As seen in Figure 3, both the noise mean and variance increase when the target level separation increases, which is intuitive.

Figure 4 shows the trade-off curve between probability of error and NAND lifetime for a 3-bits-per-cell MLC flash, obtained by solving problems (8) and (12) for several values of γ. NAND lifetime is characterized in terms of $(1/2)x^T x$. Figures 4a and 4b plot level error rate and bit error rate respectively as discussed in sections IV-A and IV-B. Also included in the figure is the performance for a naive scheme, which chooses uniformly spaced target levels. As seen in the figure, the proposed algorithm outperforms the naive scheme. For a target level error rate constraint of 10^{-2}, NAND lifetime is seen to increase by 100% when using the proposed method.

Brute force search for the best target levels is also performed using grid search[1] and the resulting curves are also shown in Figure 4. However, running the grid search required 4 hours using a 2.4 GHz computer running Matlab R2011a with CVX 1.21 [12]. This is significantly more complex and time taking than the proposed method, which took 2 seconds to run on the same system. As seen in Figure 4 the proposed method performs very close to exhaustive search with negligible computational complexity in comparison.

If the number of bits per cell were greater than 3, as it is

[1]The exhaustive search is performed over a grid size of 75 mV and the best solution x_{exh} is then refined through a secondary grid search where level x_i varies between $x_i - 0.1$ volt and $x_i + 0.1$ volt in 25 mV intervals.

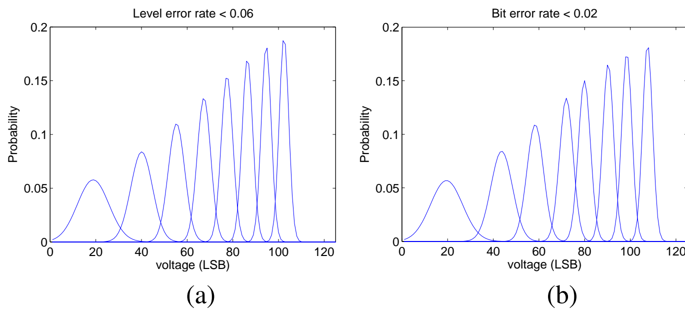

Fig. 5. Voltage distributions maximizing lifetime for a 3-bits-per-cell MLC with level error rate below 0.06 (a) or bit error rate below 0.02 (b).

the case in some commercial products, an exhaustive search would be practically impossible. The complexity increases exponentially with the number of target levels so it may require several weeks to finish. Furthermore, in practice the noise model needs to be updated in an iterative manner when the choice of target levels changes, hence exhaustive search may be ruled out for even lower density MLC flash.

The significant improvement provided by the proposed method over the naive scheme becomes more apparent by looking at Figure 5, which plots the resulting voltage distribution for the 8 levels. Figure 5a shows the results when the level error rate is required to be below 0.06 and 5b shows the results when the bit error rate is required to be below 0.02. It can be observed that the fourth and fifth levels are closer in Figure 5b than in Figure 5a, and the opposite happens with the first and second levels. This is due to the fact that mistaking the fourth and fifth levels would cause an error on the MSB (according to the mapping in Figure 2), which has negligible probability of error when any of the other levels is written. Mistaking the first and second levels, on the other hand, would cause an error on the LSB, which suffers significant probability of error regardless of which level was written and therefore requires higher reliability in each of them.

Finally, to illustrate the wide applicability of the proposed method, it is used for a different noise model, where the noise at the lowest level is larger than that for all other levels, which have the same noise variance. This model for write noise is similar to that used in Figure 6 of [4]. The resulting trade-off curves and voltage distributions are shown in Figure 6. It is important to note how, when high level error rate is tolerated, uniform spacing of the levels provides better results than the proposed scheme. This is due to the fact that the approximations made in Section IV-A assume that all levels suffer approximately the same probability of error. In this case, however, the optimal arrangement of the levels at high error rates yields significantly higher probability of error for the first level than for the others.

VI. Conclusions

This paper proposes a method to choose the target voltage levels of a MLC flash memory that maximize lifetime for a given reliability constraint (or vice versa). Using convex approximations to the probability of error, the method finds an approximate solution to the problem. Simulations have shown that this solution is very close to the one found through an exhaustive search.

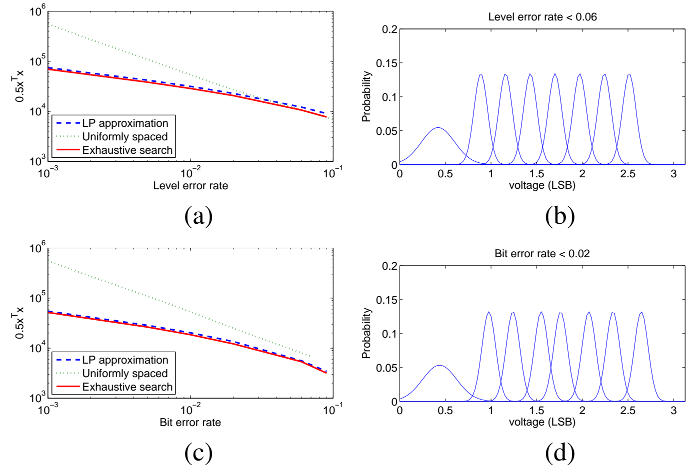

Fig. 6. Trade-off curves between lifetime and level error rate (a) or bit error rate (c) for an alternative noise model. Plots (b) and (d) show the voltage level distributions when the level error rate is required to be below 0.06 and when the bit error rate is required to be below 0.02, respectively.

Write noise is the dominant source of disturbance in most current sub-30 nm flash technologies and it is expected to become even more so in future sub-20 nm technologies. The method in this paper can help significantly alleviate this noise by providing an appropriate choice of target voltage levels in MLC NAND cells.

References

[1] L. Grupp, A. Caulfield, J. Coburn, S. Swanson, E. Yaakobi, P. Siegel, and J. Wolf, "Characterizing flash memory: anomalies, observations, and applications," in Proceedings of the 42nd Annual IEEE/ACM International Symposium on Microarchitecture. ACM, 2009, pp. 24–33.

[2] R. Agarwal and M. Marrow, "A closed-form expression for write amplification in nand flash," in GLOBECOM Workshops (GC Wkshps), 2010 IEEE. IEEE, 2010, pp. 1846–1850.

[3] K. Kim and J. Choi, "Future outlook of nand flash technology for 40nm node and beyond," in Non-Volatile Semiconductor Memory Workshop, 2006. IEEE NVSMW 2006. 21st. IEEE, 2006, pp. 9–11.

[4] N. X. G. Dong and T. Zhang, "Techniques for embracing intra-cell unbalanced bit error characteristics in mlc nand flash memory," in Workshop on Application of Communication Theory to Emerging Memory Technologies, 2010.

[5] A. Ghetti, C. Compagnoni, A. Spinelli, and A. Visconti, "Comprehensive analysis of random telegraph noise instability and its scaling in deca–nanometer flash memories," IEEE Transactions on Electron Devices, vol. 56, no. 8, pp. 1746–1752, 2009.

[6] M. Grossi, M. Lanzoni, and B. Ricco, "Program schemes for multilevel flash memories," Proceedings IEEE, vol. 91, no. 4, pp. 594–601, 2003.

[7] A. Torsi, Y. Zhao, H. Liu, T. Tanzawa, A. Goda, P. Kalavade, and K. Parat, "A program disturb model and channel leakage current study for sub-20 nm nand flash cells," IEEE Transactions on Electron Devices,, vol. 58, no. 1, pp. 11–16, 2011.

[8] G. Dong, S. Li, and T. Zhang, "Using data postcompensation and pre-distortion to tolerate cell-to-cell interference in mlc nand flash memory," IEEE Transactions on Circuits and Systems I: Regular Papers, vol. 57, no. 10, pp. 2718–2728, 2010.

[9] S. Boyd, "Sequential convex programming," http://www.stanford.edu/class/ee364b/lectures/seq_slides.pdf, 2008.

[10] D. Bertsekas, Nonlinear programming. Athena Scientific, 1999.

[11] J. Nocedal and S. Wright, Numerical Optimization, Series in Operations Research and Financial Engineering. Springer, New York, 2006.

[12] M. Grant and S. Boyd, "CVX: Matlab software for disciplined convex programming, version 1.21," http://cvxr.com/cvx, Apr. 2011.