Estimating Information-Theoretical NAND Flash Memory Storage Capacity and its Implication to Memory System Design Space Exploration

Guiqiang Dong, *Student Member, IEEE*, Yangyang Pan, *Student Member, IEEE*, Ningde Xie, Chandra Varanasi, and Tong Zhang, *Senior Member, IEEE*

***Abstract*—Today and future NAND flash memory will heavily rely on system-level fault-tolerance techniques such as error correction code (ECC) to ensure the overall system storage integrity. Since ECC demands the storage of coding redundancy and hence degrades effective cell storage efficiency, it is highly desirable to use more powerful coding solutions that can maintain the system storage reliability at less coding redundancy. This has motivated a growing interest in the industry to search for alternatives to BCH code being used in today. Regardless to specific ECCs, it is of great practical importance to know the theoretical limit on the achievable cell storage efficiency, which motivates this work. We first develop an approximate NAND flash memory channel model that explicitly incorporates program/erase (P/E) cycling effects and cell-to-cell interference, based on which we then develop strategies for estimating the information-theoretical bounds on cell storage efficiency. We show that it can readily reveal the tradeoffs among cell storage efficiency, P/E cycling endurance, and retention limit, which can provide important insights for system designers. Finally, motivated by the dynamics of P/E cycling effect revealed by the information-theoretical study, we propose two memory system design techniques that can improve the average NAND flash memory programming speed and increase the total amount of user data that can be stored in NAND flash cell over its entire lifetime.**

***Index Terms*—Endurance, information theory, interference, model, NAND flash, retention, storage capacity, tradeoff.**

I. Introduction

As one of the fastest growing segments in the global semiconductor industry, NAND flash memory has been entering increasingly diverse real-life applications from consumer electronics to personal and enterprise computing, due to its steady bit cost reduction. The continuous bit cost reduction of NAND flash memory mainly relies on aggressive technology scaling and use of multi-level per cell (MLC) technique. Most MLC NAND flash memories store 2 bits per cell, while 3 and even 4 bits per cell NAND flash memories have been recently reported [1]–[5]. As technology continues to scale down, it becomes increasingly challenging to ensure NAND flash memory storage reliability and maintain historical values of important performance metrics such as endurance and retention limit. This is because, at smaller device feature size, NAND flash memory cells are more severely subject to various device and circuit level noises such as program/erase (P/E) cycling effects and cell-to-cell interference, leading to increasingly worse memory cell storage reliability. This forces designers to use more and more sophisticated system-level fault-tolerance techniques such as error correction code (ECC) to embrace the worse reliability of the underlying memory cells. The use of ECC comes with the overhead of storing its coding redundancy, leading to reduced memory storage efficiency. In this work, we use a metric called *cell storage efficiency*, defined as the average number of real user bits per cell, to represent the memory storage efficiency. For example, if we use a BCH code that requires 28-byte coding redundancy to protect each 512-byte user data in a 2 bits/cell NAND flash memory, then the cell storage efficiency is $512/512 + 28 \times 2 = 1.90$ bits/cell.

Regardless to specific ECC being used, it is of practical importance to know the theoretical limit on the achievable memory cell storage efficiency. This can be realized by mathematically modeling NAND flash memory as a communication channel that can capture the major data distortion noise sources, based on which we can apply Shannon's information theory [6] to estimate the theoretical cell storage efficiency limit. The major distortion sources in NAND flash memory have been intensively studied and modeled by the device research community over the past two decades, which provide solid foundation for us to develop such a mathematical flash memory channel model. In this work, we first develop an approximate NAND flash memory channel model that explicitly incorporates P/E cycling effects and cell-to-cell interference, the two most important cell storage distortion noise sources. For P/E cycling effects, this model incorporates both random telegraph noise (RTN), which widens the threshold voltage distribution windows, and interface trap recovery and electron detrapping, which gradually destroy the stored data and set the data retention limit.

Based upon this mathematical memory channel model, we investigate how to estimate the information-theoretical limit of the memory cell storage efficiency. As elaborated later, since this channel is essentially a channel with memory, it is very difficult to directly calculate the theoretical capacity (i.e., the the-

Manuscript received October 12, 2010; revised February 25, 2011; accepted June 13, 2011. Date of publication August 04, 2011; date of current version July 05, 2012. This work was supported by the National Science Foundation under Grant 0937794.

G. Dong, Y. Pan, and T. Zhang are with the Electrical, Computer, and Systems Engineering Department, Rensselaer Polytechnic Institute (RPI), Troy, NY 12180 USA (e-mail: dongguiqiang@gmail.com; yyangpan@gmail.com; tzhang@ecse.rpi.edu).

N. Xie is with Intel Corporation, Hillsboro, OR 97124 USA (e-mail: ningdexie@gmail.com).

C. Varanasi is with the Micron Technology, San Jose, CA 95732 USA (e-mail: cvaranasi@micron.com).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TVLSI.2011.2160747

1063-8210/$26.00 © 2011 IEEE

oretical limit of the cell storage efficiency). Therefore, we instead develop strategies to estimate the upper and lower bounds of the theoretical limit. We also show that it can readily reveal the theoretical tradeoffs among cell storage efficiency, P/E cycling endurance, and retention limit. By no means we claim this memory channel model is absolutely accurate, but we believe that, by explicitly capturing several major memory cell storage noise sources, this approximate model can serve as a good vehicle to carry out information-theoretical investigation and reveal the tradeoffs among important memory system metrics. In addition, the developed strategies for estimating the upper and lower information-theoretical bounds of cell storage efficiency are still applicable even when the channel model is further improved by incorporating some other noise sources.

Moreover, the information-theoretical investigations reveal the significant impact of the inherent dynamics of P/E cycling. This motivates us to develop two new memory system design techniques that can exploit such dynamics to improve certain system performance metrics. The first technique aims to improve the average NAND flash memory programming speed. The key is to dynamically tune the instantaneous NAND flash memory programming speed adaptive to the present P/E cycling number. The second technique aims to improve the total amount of user data that can be programmed into NAND flash memory over its lifetime. The key is to jointly consider the cell storage efficiency and P/E cycling endurance. We apply the information-theoretical study to demonstrate the potential of these two design techniques, and further evaluate the effectiveness when BCH code is employed.

The reminder of this paper is organized as follows. Section II reviews the basics of NAND flash memory and presents one NAND flash channel model. Section III develops strategies for estimating the information-theoretical bounds of memory cell storage efficiency. Using hypothetical 2 bits/cell NAND flash memory as an example, we apply these strategies to derive the theoretical bounds and show the tradeoffs among cell storage efficiency, endurance, and retention limit. In Section IV, we present two design techniques that can improve certain NAND flash memory system performance metrics. Conclusions are drawn in Section V.

II. Background

A. NAND Flash Memory Basics

Each NAND flash memory cell is a floating gate transistor whose threshold voltage can be configured (or programmed) by injecting certain amount of charges into the floating gate. Before a flash memory cell is programmed, it must be erased (i.e., remove all the charges from the floating gate, which sets its threshold voltage to the lowest voltage window). It is well known that the threshold voltage of erased memory cells tends to have a wide Gaussian-like distribution [7]. Hence, we can approximately model the threshold voltage distribution of erased state as

$$p_e(x) = \frac{1}{\sigma_e\sqrt{2\pi}} e^{-(x-\mu_e)^2/2\sigma_e^2} \tag{1}$$

where μ_e and σ_e are the mean and standard deviation of the erased state.

Regarding memory programming, a tight threshold voltage control is typically realized by using incremental step pulse program (ISPP) [8], [9], i.e., memory cells on the same word-line are recursively programmed using a program-and-verify approach with a stair case program word-line voltage V_{pp}. Under such a program-and-verify strategy, each programmed state (except the erased state) associates with a verify voltage that is used in the verify operations and sets the target position of each programmed state threshold voltage window. Denote the verify voltage of the target programmed state as V_p, and program step voltage as ΔV_{pp}. The threshold voltage of the programmed state tends to have a uniform distribution over $[V_p,\ V_p + \Delta V_{pp}]$ with the width of ΔV_{pp} [10]. Denote V_p and $V_p + \Delta V_{pp}$ for the k-th programmed state as $V_l^{(k)}$ and $V_r^{(k)}$. We can model the ideal threshold voltage distribution of the kth programmed state as

$$p_p^{(k)}(x) = \begin{cases} \frac{1}{\Delta V_{pp}}, & \text{if } V_l^{(k)} \le x \le V_r^{(k)} \\ 0, & \text{else.} \end{cases} \tag{2}$$

Unfortunately, the above *ideal* memory cell threshold voltage distribution can be (significantly) distorted in practice, mainly due to P/E cycling effect and cell-to-cell interference, which will be discussed in the remainder of this section.

B. Effects of P/E Cycling

Flash memory P/E cycling causes damage to the tunnel oxide of floating gate transistors in the form of charge trapping in the oxide and interface states [11]–[14], which directly results in threshold voltage shift and fluctuation and hence gradually degrades memory device noise margin. Major distortion sources include the following.

1) Electrons capture and emission events at charge trap sites near the interface developed over P/E cycling directly result in memory cell threshold voltage fluctuation, which is referred to as random telegraph noise (RTN) [10], [15].
2) Interface trap recovery and electron detrapping [16], [17] gradually reduce memory cell threshold voltage, leading to the data retention limitation.

RTN causes random fluctuation of memory cell threshold voltage, where the fluctuation magnitude is subject to exponential decay. Hence, we can model the probability density function $p_r(x)$ of RTN-induced threshold voltage fluctuation as a symmetric exponential function [10]

$$p_r(x) = \frac{1}{2\lambda_r} e^{-|x|/\lambda_r}. \tag{3}$$

Let N denote the P/E cycling number, λ_r scales with N in an approximate power-law fashion, i.e., λ_r is approximately proportional to N^α.

Interface trap recovery and electron detrapping processes approximately follow Poisson statistics [14], hence threshold voltage reduction due to interface trap recovery and electron detrapping can be approximately modeled as a Gaussian distribution $\mathcal{N}(\mu_d, \sigma_d^2)$. Both μ_d and σ_d^2 scale with N in an approximate power-law fashion, and scale with the retention time t in a logarithmic fashion. Moreover, the significance of threshold voltage reduction induced by interface trap recovery and electron detrapping is also proportional to the

initial threshold voltage magnitude [18], i.e., the higher the initial threshold voltage is, the faster the interface trap recovery and electron detrapping occur and hence the larger threshold voltage reduction will be.

C. Cell-to-Cell Interference

In NAND flash memory, the threshold voltage shift of one floating gate transistor can influence the threshold voltage of its neighboring floating gate transistors through parasitic capacitance-coupling effect [19]. This is referred to as cell-to-cell interference, which has been well recognized as the one of major noise sources in NAND flash memory [20]–[22]. Threshold voltage shift of a victim cell caused by cell-to-cell interference can be estimated as [19]

$$F = \sum_{k} (\Delta V_t^{(k)} \cdot \gamma^{(k)}) \tag{4}$$

where $\Delta V_t^{(k)}$ represents the threshold voltage shift of one interfering cell which is programmed after the victim cell, and the coupling ratio $\gamma^{(k)}$ is defined as

$$\gamma^{(k)} = \frac{C^{(k)}}{C_{\text{total}}} \tag{5}$$

where $C^{(k)}$ is the parasitic capacitance between the interfering cell and the victim cell, and C_{total} is the total capacitance of the victim cell. Cell-to-cell interference significance is affected by NAND flash memory bit-line structure. In current design practice, there are two different bit-line structures, including conventional even/odd bit-line structure [23], [24] and emerging all-bit-line structure [25], [26]. In even/odd bit-line structure, memory cells on one word-line are alternatively connected to even and odd bit-lines and even cells are programmed ahead of odd cells in the same wordline. Therefore, an even cell is mainly interfered by five neighboring cells and an odd cell is interfered by only three neighboring cells [27], as shown in Fig. 1. Therefore, even cells and odd cells experience largely different amount of cell-to-cell interference. Cells in all-bit-line structure suffers less cell-to-cell inference than even cells in odd/even structure, and the all-bit-line structure can effectively support high-speed current sensing to improve the memory read and verify speed. Therefore, throughout the remainder of this paper, we mainly consider NAND flash memory with the all-bit-line structure. Finally, we note that the design methods presented in this work are also applicable when odd/even structure is being used.

D. Equivalent NAND Flash Memory Channel Model

Based on the above discussions, we can approximately model NAND flash memory device characteristics as shown in Fig. 2, using which we can simulate memory cell threshold voltage distribution and hence obtain memory cell raw storage reliability. Based upon (1) and (2), we can obtain the threshold voltage distribution function $p_p(x)$ right after ideal programming operation. Recall that $p_{pr}(x)$ denotes the RTN distribution function

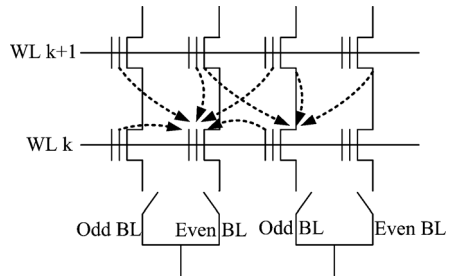

Fig. 1. Illustration of cell-to-cell interference in even/odd structure: even cells are interfered by two direct neighboring cells on the same wordline and three neighboring cells on the next wordline, while odd cells are interfered by three neighboring cells on the next wordline.

[see (3)], and let $p_{ar}(x)$ denote the threshold voltage distribution after incorporating RTN, which is obtained by convoluting $p_p(x)$ and $p_r(x)$, i.e.,

$$p_{ar}(x) = p_p(x) \bigotimes p_r(x). \tag{6}$$

The cell-to-cell interference is further incorporated based on (4). To capture inevitable process variability, we set both the vertical coupling ratio γ_y and diagonal coupling ratio γ_{xy} as random variables with bounded Gaussian distribution

$$p_c(x) = \begin{cases} \frac{c_c}{\sigma_c\sqrt{2\pi}} \cdot e^{-(x-\mu_c)^2/2\sigma_c^2}, & \text{if } |x-\mu_c| \le w_c \\ 0, & \text{else} \end{cases} \tag{7}$$

where μ_c and σ_c are the mean and standard deviation, and c_c is chosen to ensure the integration of this bounded Gaussian distribution equals to 1. In all the simulations in this paper, we set $w_c = 0.1\ \mu_c$ and $\sigma_c = 0.4\ \mu_c$.

Let p_{ac} denote the threshold voltage distribution after incorporating cell-to-cell interference. Denote the retention noise distribution as $p_t(x)$. The final threshold voltage distribution p_f is obtained as

$$p_f(x) = p_{ac}(x) \bigotimes p_t(x). \tag{8}$$

The above presented approximate mathematical channel model for simulating NAND flash memory cell threshold voltage is further demonstrated using the following example.

1) Example 2.1: Let us consider 2 bits/cell NAND flash memory. We set normalized σ_e and μ_e of the erased state as 0.35 and 1.4, respectively. For the three programmed states, we set the normalized program step voltage ΔV_{pp} as 0.2, and the normalized verify voltages V_p as 2.6, 3.2, and 3.93, respectively. For the RTN distribution function $p_r(x)$, we set the parameter $\lambda_r = K_\lambda \cdot N^{0.5}$, where K_λ equals to 0.00025. Regarding to cell-to-cell interference, according to [21], [28], we set the ratio between the means of γ_y and γ_{xy} as 0.08 and 0.0048, respectively. For the function $\mathcal{N}(\mu_d, \sigma_d^2)$ to capture trap recovery and electron detrapping during retention, according to [14] and [16], we set that μ_d scales with $N^{0.5}$ and σ_d^2 scales with $N^{0.6}$, and both scale with $\ln(1 + t/t_0)$, where t denotes the memory retention time and t_0 is an initial time and can be set as 1 hour. In addition, as pointed out earlier, both μ_d and σ_d^2 also depend on the initial threshold voltage. Hence we set that both approximately scale $K_s(x - x_0)$, where x is the initial

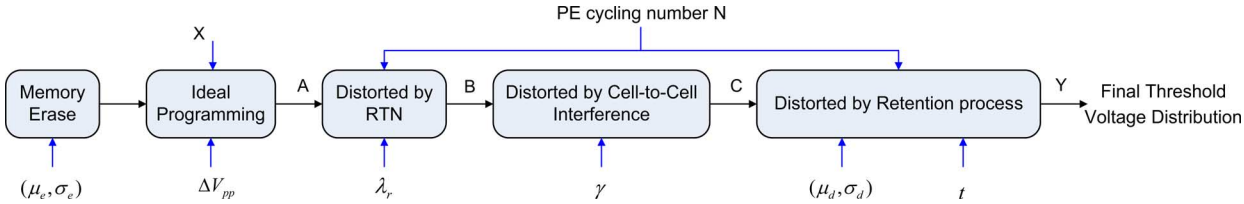

Fig. 2. Illustration of the approximate NAND flash memory device model to incorporate major threshold voltage distortion sources.

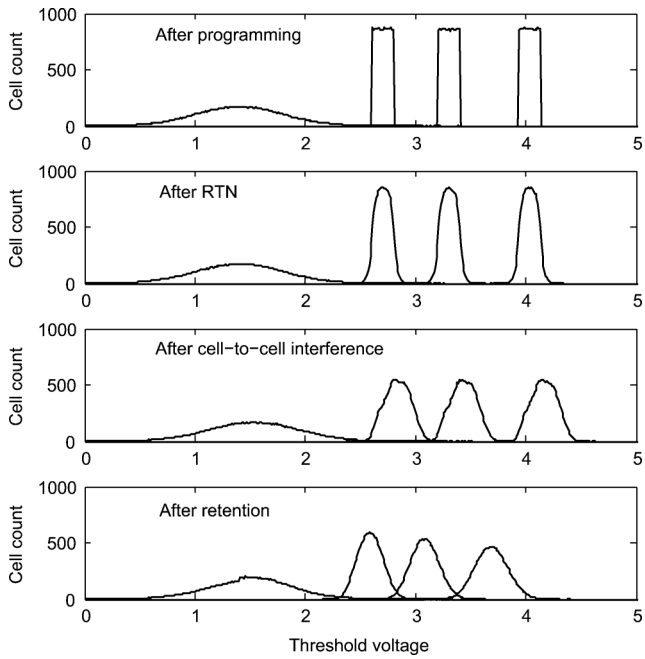

Fig. 3. Simulated results to show the effects of RTN, cell-to-cell interference, and retention on memory cell threshold voltage distribution after 10K P/E cycling and 10-year retention.

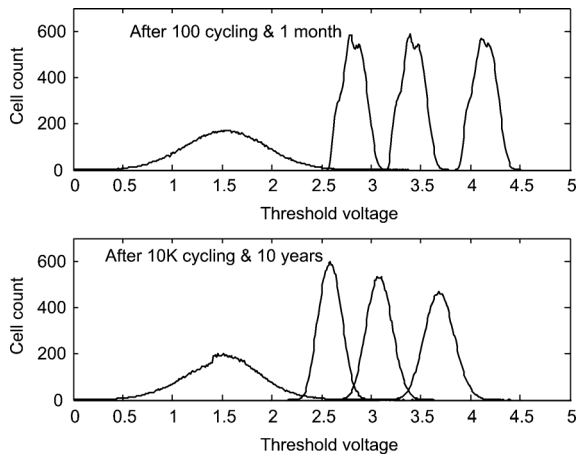

Fig. 4. Simulated threshold voltage distribution after 100 P/E cycling and 1-month retention and after 10K P/E cycling and 10-year retention, which clearly shows the dynamics inherent in NAND flash memory characteristics.

threshold voltage, and x_0 and K_s are constants. Therefore, we have

$$\begin{cases} \mu_d = K_s(x - x_0)K_d N^{0.5} \ln\left(1 + \frac{t}{t_0}\right) \\ \sigma_d^2 = K_s(x - x_0)K_m N^{0.6} \ln\left(1 + \frac{t}{t_0}\right) \end{cases} \tag{9}$$

where we set $K_s = 0.38$, $x_0 = 1.4$, $K_d = 4 \times 10^{-4}$, and $K_m = 4 \times 10^{-6}$ by fitting the measurement data presented in [14] and [16]. Accordingly, we carry out Monte Carlo simulations to obtain the cell threshold voltage distribution at different stages under 10K P/E cycling and with 10-year retention limit, as shown in Fig. 3. The final threshold voltage distributions after 100 P/E cycling and 1 month storage and after 10K P/E cycling and 10 years storage are both shown in Fig. 4. These results clearly show the dynamic characteristics of NAND flash memory.

III. Information-Theoretical Bounds on Memory Cell Storage Efficiency and Design Tradeoff

In this section, based upon the above approximate NAND flash memory channel model, we are interested in the theoretical limit of the NAND flash memory cell storage efficiency (i.e., the memory channel capacity), under which error-free storage can be realized in theory. Because NAND flash memory channel have different characteristics under different P/E cycling and retention limit as illustrated in Fig. 4, the information-theoretical memory cell storage efficiency limit varies with P/E cycling and retention limit. Therefore, we will further investigate the information-theoretical tradeoffs among cell storage capacity, P/E cycling endurance, retention limit.

Due to cell-to-cell interference, the NAND flash memory channel is essentially a communication channel with memory, for which the exact calculation of channel capacity is intractable. Therefore, instead of deriving the exact channel capacity, we aim to derive a pair of lower and upper bounds of this channel capacity. Throughout this paper, we denote random variables by upper case letters (X) and their particular realizations by lower case letters (x), and write the n-tuples $(X_1, X_2, \ldots, X_n)$ and $(x_1, x_2, \ldots, x_n)$ as X^n and x^n, respectively. Assume the input bits are statistically independent and have equal probability to be 0 and 1. We can formulate the channel capacity C as

$$C = \frac{1}{n} \lim_{n\to\infty} \{I(X^n; Y^n)\} \tag{10}$$

where $I(X^n; Y^n) = H(Y^n) - H(Y^n|X^n)$ is the mutual information between the input sequence X^n and output sequence Y^n, and $H(Y^n)$ is the entropy of output sequence Y^n, i.e.,

$$H(Y^n) = -\sum_{y^n} (p(y^n) \cdot \log_2 p(y^n)) \tag{11}$$

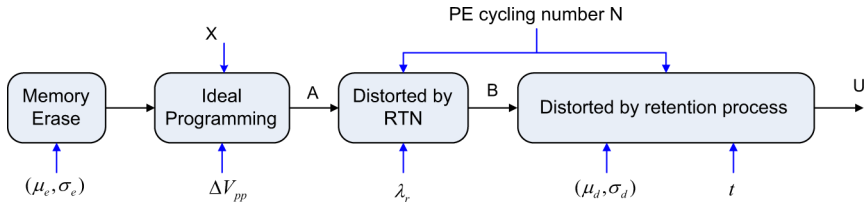

Fig. 5. Simplified channel used for calculating an upper bound of the NAND flash memory channel capacity C.

and $H(Y^n|X^n)$ is the conditional entropy, which can be expressed as

$$H(Y^n|X^n) = \sum_{x^n} p(x^n) \sum_{y^n} (p(y^n|x^n) \cdot \log_2 p(y^n|x^n)). \quad (12)$$

Due to the cell-to-cell interference, one output is correlated with many inputs (i.e., the channel is a channel with memory), which makes it very difficult to directly calculate the conditional entropy. In the following, we discuss the estimation of an upper and a lower bound of the channel capacity.

First, let us consider the upper bound of the channel capacity. In this context, we reduce the original memory channel into a memoryless channel by removing the cell-to-cell interference component, as shown in Fig. 5. Clearly, each output of this new reduced channel only depends on its corresponding input. Let C_U denote the capacity of this new channel, and according to [6], we have $C_U \geq C$, i.e., C_U is an upper bound of the theoretical limit C of memory cell storage efficiency. The input and output of this new reduced channel are denoted as X and U, respectively. Since this new channel is a memoryless channel, we have $C_U = I(X;U) = H(U) - H(U|X)$, which can be calculated by numerically estimating the distribution of output U through Monte Carlo simulations.

Next, let us consider the lower bound of the channel capacity. According to [29], for channel with memory we have

$$I(X^n;Y^n) \geq \sum_{i=1}^{n} I(X_i;Y_i) \quad (13)$$

where $I(X_i;Y_i)$ is the mutual information between each pair of input and output. Since all the X_i are independent and identically-distributed random variables, a lower bound of C can be estimated as $I(X;Y) = H(Y) - H(Y|X)$, which is represented as C_R. Nevertheless, this lower bound tends to be too loose, and we propose the following method to obtain a tighter lower bound. This is realized by creating an expanded channel as shown in Fig. 6, which concatenate the original NAND flash memory channel with a post-compensation module [27]. The post-compensation module aims to explicitly compensate cell-to-cell interference, i.e., if we know the threshold voltage shift of interfering cells, we can estimate the corresponding cell-to-cell interference strength according to (4) and subsequently subtract it from the threshold voltage of victim cells. According to the data processing theorem [6], the channel capacity C_z of this new expanded channel cannot be larger than the original channel capacity C. Hence, we can use C_z as a lower bound of C.

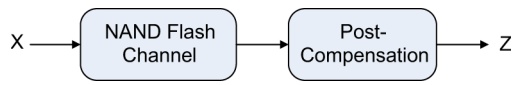

Fig. 6. Expanded channel used for calculating a tight lower bound of the memory channel capacity C.

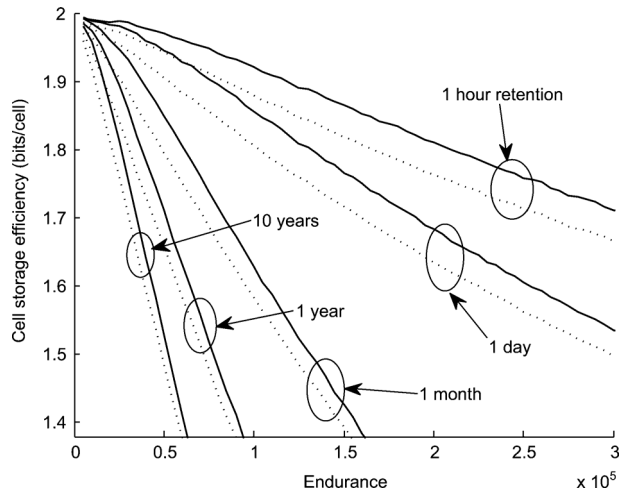

Fig. 7. Comparison of two lower bounds C_R (dashed line) and C_Z (solid line) of C.

1) Example 3.1: Let us again consider 2 bits/cell NAND flash memory and keep all the parameters the same as those in Example 2.1 in Section II-D. We carry out extensive simulations to estimate these bounds of flash memory channel capacity under various P/E cycling and retention limit. The two lower bounds C_R and C_Z are shown in Fig. 7, which clearly shows that the lower bound C_Z is tighter than C_R. Hence, we only consider the tighter lower bound C_Z in the remainder of this paper. Fig. 8 shows the upper and lower bounds of cell storage efficiency versus P/E cycling under different retention limit. It clearly shows that the cell storage capacity monotonically reduce as the P/E cycling and/or retention limit increases, which can intuitively be justified. Moreover, the results show that the gap between the upper and lower bounds monotonically increase as the retention time increases, which can be explained as follows. Under a relatively shorter retention time, the threshold voltage drop induced by trap recovery and electron detrapping is less, which makes the post-compensation process more accurately compensate the cell-to-cell interference in the expanded channel as shown in Fig. 6. As the retention time increases, the post-compensation process will become less effective, leading to a larger gap between the upper and lower bounds.

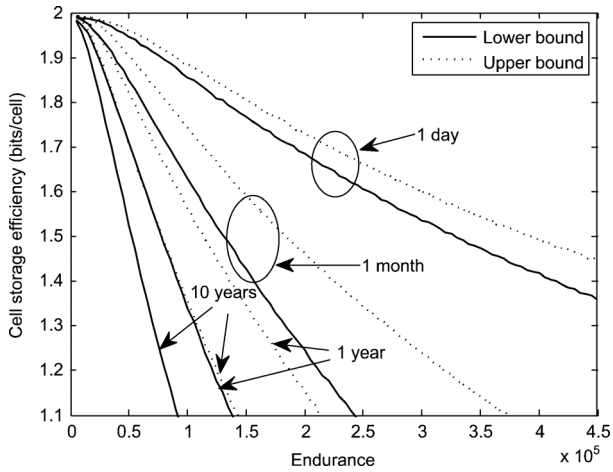

Fig. 8. Upper bounds and lower bounds of cell storage efficiency versus P/E cycling under different retention limit.

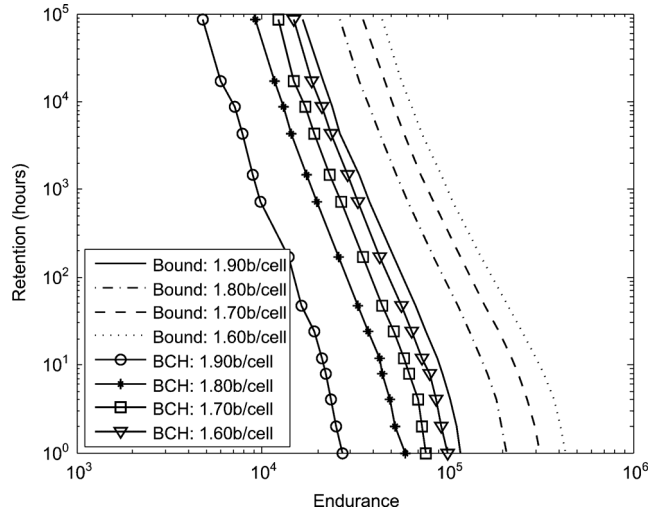

Fig. 9. Tradeoff between P/E cycling endurance and retention limit based on results obtained in the Example 3.1 and when using binary BCH code to protect each 4K-byte user data.

In conventional practice, most 2 bits/cell NAND flash memory chips are typically specified with a single pair of P/E cycling endurance and retention limit, e.g., endurance of 10K P/E cycling and 10-year retention. However, real-life applications may have diverse and even dynamically varying expectations on the endurance and retention limit. For example, let us consider its application in high-end computing. Modern large-scale high-performance computing systems use checkpoint/restart mechanism to realize system fault tolerance [30], where each server periodically take and store a snapshot of the current application state. If we use NAND flash memory to speed up the storage of checkpoints, we may demand much less retention limit (e.g., up to tens of hours or few days) since the periodic checkpointing interval tends to be short (e.g., 30 min or 1 hour) and we only need a limited number of previously stored checkpoints to restart the computing systems in case of failures. For this scenario, it is highly desirable to trade the retention limit for improving P/E cycling endurance and hence the NAND flash memory lifetime. Being able to quantitatively estimate the theoretical bounds on NAND flash memory cell storage efficiency under different P/E cycling and retention limit, this work makes it possible to investigate the tradeoffs among cell storage efficiency, P/E cycling endurance, and retention limit from an information-theoretical perspective. This can quantitatively reveal the potential effectiveness when we trade one metric for the other, and hence provide important insights for system designers.

Based upon the results obtained in the Example 3.1, Fig. 9 shows the tradeoff between P/E cycling endurance and retention limit, under different lower bounds of cell storage efficiency. Under the cell storage efficiency lower bound of 1.90 bits/cell, the endurance is only about 16K P/E cycling in case of 10-year retention, and if we can relax the retention limit to 1 year, the endurance can increase to about 24K P/E cycling, representing 44% improvement. Similarly, if we can further relax the retention limit to 1 month and 1 day, the endurance increases to about 38K and 77K P/E cycling, representing 127% and 362% improvement, respectively. If binary BCH code is being employed to realize NAND flash memory fault tolerance, where each BCH codeword protects 4K-byte user data and the target decoding block failure rate is below 10^{-15}, we estimate the tradeoffs among the achievable cell storage efficiency, P/E cycling endurance, and retention limit, as shown in Fig. 9. We note that the post-compensation technique is being used together with BCH code in order to better compensate the cell-to-cell interference. The results show that, although the BCH-based solution has a big gap from the theoretical bounds, they have the same trend on the tradeoffs. The results also suggest that more powerful signal processing and coding techniques are highly desirable to close this gap and approach the theoretical bounds.

Besides trading retention limit for P/E cycling endurance, we can also investigate the information-theoretical tradeoff between cell storage efficiency and P/E cycling endurance, given a fixed retention limit. For example, according to Fig. 8, the P/E cycling endurance corresponding to the cell storage efficiency of 1.90 bits/cell is about 16K PE cycles under 10-year retention limit. If we reduce the cell storage efficiency to 1.80 bits/cell, the P/E cycling endurance can accordingly improve to about 26K cycles. If we further reduce the cell storage efficiency to 1.70 and 1.6 bits/cell, the P/E cycling endurance can increase to about 35K and 44K, respectively. It is well known that, given a specific type of ECC such as BCH and RS codes, in order to improve the error correction capability, we have to increase the amount of coding redundancy and hence reduce memory cell storage efficiency. The above information-theoretical investigation on the cell storage efficiency versus P/E cycling endurance can readily reveal the potential of how much we may improve the NAND flash memory endurance when using stronger error correction capability.

Similarly, we can investigate the information-theoretical tradeoff between the retention limit and cell storage efficiency. Again, based on the results obtained in the Example 3.1, Fig. 10 shows how the achievable retention limit will increase as we reduce the cell storage efficiency under different P/E cycling endurance. It shows that, under the P/E cycling endurance of 40K cycles, the 1-day retention limit can enable the realization of up to 1.96 bits/cell, and as the retention limit increases to

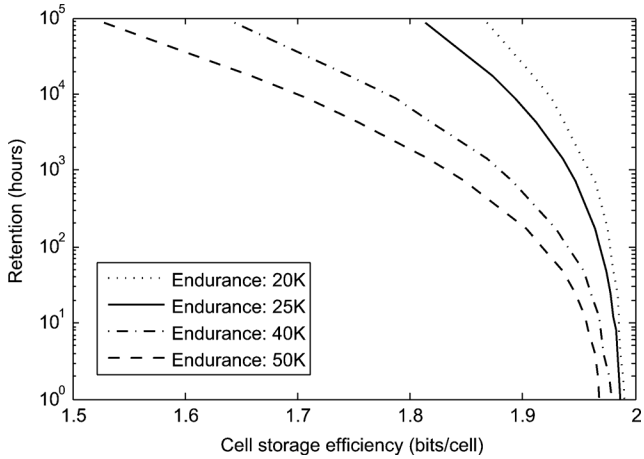

Fig. 10. Tradeoff between the cell storage efficiency and retention limit based on results obtained in the Example 3.1.

1 month and 10 years, the allowable cell storage efficiency drops to 1.89 and 1.64 bits/cell, respectively.

The above examples and discussions show that, under the information-theoretical framework, the developed channel model and channel capacity bound estimation methods can readily reveal the tradeoffs among the three important NAND flash memory system performance metrics: cell storage efficiency, P/E cycling endurance, and retention limit. Of course, for its practical use, we have to fine-tune the parameters in this channel model according to the measurement and characterization of underlying NAND flash memory technology. Finally, we emphasize that the main objective of this information-theoretical investigation is to reveal important insights for system designers on optimizing the use of NAND flash memory in various real-life applications, other than directly deriving the exact specifications of the end commercial products.

IV. Implications to Memory System Design Space Exploration

Besides revealing the fundamental tradeoffs among cell storage efficiency, P/E cycling endurance, and retention limit as discussed in the previous section, this information-theoretical study can inspire NAND flash memory system design innovations for improving certain system performance metrics. In this section, we present two such design techniques that can improve average NAND flash memory programming speed and maximize memory cell lifetime storage capacity.

A. Improving NAND *Flash Memory Programming Speed*

As discussed in Section II-A, NAND flash memory programming is carried out recursively by sweeping over the entire memory cell threshold voltage region with a program step voltage ΔV_{pp}. As a result, the memory programming latency is inversely proportional to ΔV_{pp}, and hence we can improve the memory programming speed by increasing ΔV_{pp}. However, a larger ΔV_{pp} directly results in a wider threshold voltage distribution of each programmed state, leading to less noise margin between adjacent programmed states and hence worse raw storage reliability. In current design practice, ΔV_{pp} is fixed and its value is sufficiently small so that the NAND flash memory can survive the specified P/E cycling endurance and retention limit values.

The value of ΔV_{pp} is an important parameter in the NAND flash memory channel model. Under the same P/E cycling and retention limit, the information-theoretical bounds on memory cell storage efficiency will be different when different values of ΔV_{pp} are used. Equivalently, if we fix the memory cell storage efficiency and retention limit, different values of ΔV_{pp} will enable different P/E cycling endurance. Intuitively, a larger ΔV_{pp} corresponds to a lower P/E cycling endurance. Very straightforwardly, such a tradeoff can directly enable the use of dynamic ΔV_{pp} tuning in the run time to improve the average NAND flash memory programming speed throughout the whole lifetime. Suppose the NAND flash memory can support M different values of ΔV_{pp}, denoted as $\Delta V_{pp}^{(i)}$ for $i = 1, \ldots, M$. We assume $\Delta V_{pp}^{(1)} > \Delta V_{pp}^{(2)} > \cdots > \Delta V_{pp}^{(M)}$. Given the target cell storage efficiency and retention limit, NAND flash memory with $\Delta V_{pp}^{(i)}$ can survive up to $N^{(i)}$ P/E cycling. Clearly, we have $N^{(0)} = 0 < N^{(1)} < \cdots < N^{(M)}$. Assume the NAND flash memory must survive $N^{(M)}$ P/E cycling, the conventional design practice tends to fix the program step voltage as $\Delta V_{pp}^{(M)}$ throughout its lifetime. Clearly, we can dynamically tune the program step voltage so that we use $\Delta V_{pp}^{(i)}$ when the present P/E cycling number falls into $(N^{(i-1)}, N^{(i)}]$. Therefore, the average NAND flash memory programming latency can approximately reduce by

$$l = 1 - \frac{\sum_{i=1}^{M} \frac{(N^{(i)} - N^{(i-1)})}{\Delta V_{pp}^{(i)}}}{\frac{N^{(M)}}{\Delta V_{pp}^{(M)}}}. \tag{14}$$

The potential of above presented design approach is further quantitatively demonstrated by the following example.

1) Example 4.1: Again, let us consider 2 bits/cell NAND flash memory and keep all the parameters the same as those in Example 2.1. We estimate the lower bounds of information-theoretical memory cell storage efficiency versus P/E cycling endurance under different retention limit and three different values of normalized program step voltage ΔV_{pp}, including 0.4, 0.3, and 0.2. The results are shown in Fig. 11. Assume we set the cell storage efficiency as 1.80 bits/cell, and according to the results shown in Fig. 11, we can obtain three $N^{(i)}s$ for each retention limit. Suppose the target retention limit is 1 year, we have that the three $N^{(i)}s$ are about 30K, 35K, and 38K, respectively. Therefore, according to (14), we can estimate that the proposed dynamic program step voltage tuning scheme can reduce average memory programming latency by about 44% over the entire lifetime.

Encouraged by this information-theoretical investigation, we further evaluate the potential gain when using binary BCH code in the 2 bits/cell NAND flash memory. The target cell storage efficiency of 1.80 bits/cell enables us to use a BCH code with the code rate of 0.9. Assume each BCH codeword protects 4K-byte

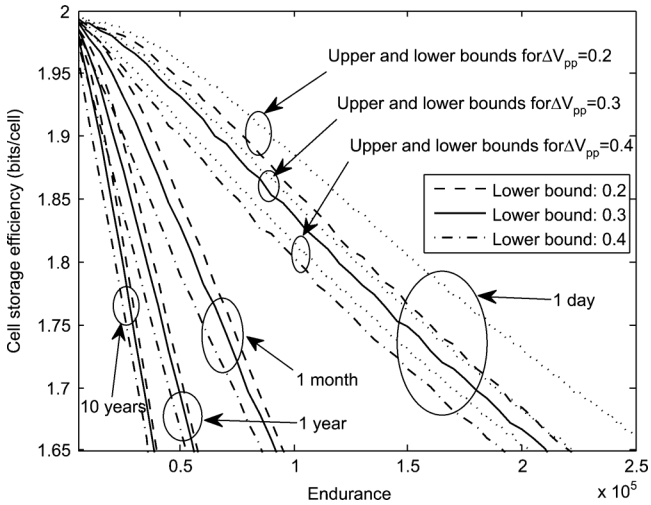

Fig. 11. Lower bounds of cell storage efficiency in case of three different program step voltage ΔV_{pp}.

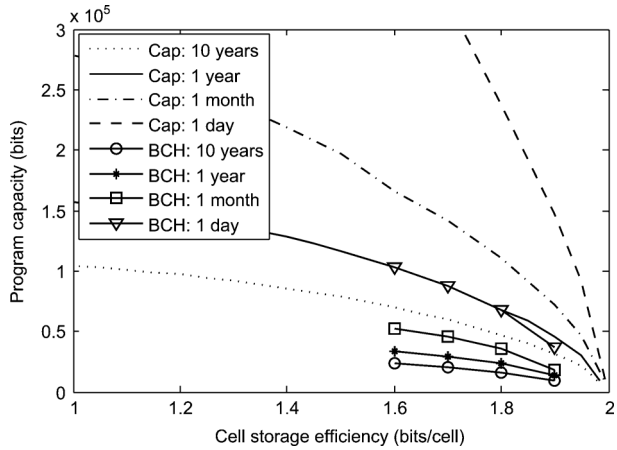

Fig. 12. Estimated memory cell lifetime program capacity versus cell storage efficiency under different retention limit.

user data and the target decoding block failure rate is below 10^{-15}, we estimate that, when the normalized program step voltage ΔV_{pp} is 0.4, 0.3, and 0.2, the corresponding P/E cycling endurance is about 5K, 10K, and 13K, respectively. This leads to about 32.7% of average memory programming latency reduction.

B. Improving Lifetime Program Capacity

In this section, we are interested in improving the total user data volume that can be stored (or programmed) in each NAND flash memory cell over its whole lifetime, which is referred to as lifetime program capacity. Clearly, the memory cell lifetime program capacity depends on both cell storage efficiency and P/E cycling endurance. Recall that the cell storage efficiency is the number of user bits per cell considering the storage of coding redundancy, and a larger cell storage efficiency does not necessarily translate to a larger memory cell lifetime program capacity, because a larger cell storage efficiency tends to results in a lower P/E cycling endurance. Therefore, we have to jointly consider the cell storage efficiency and P/E cycling endurance in order to truly maximize the memory cell lifetime program capacity.

In conventional design practice, NAND flash memory uses a fixed ECC and hence has a fixed cell storage efficiency and P/E cycling endurance, leading to a fixed lifetime memory program capacity. Hence, given a specified retention limit, we can apply the above presented method to derive the tradeoff between cell storage efficiency and P/E cycling endurance, based on which we can search for the pair of cell storage efficiency and corresponding P/E cycling endurance within a reasonable range that can maximize the memory cell lifetime program capacity.

1) Example 4.2: Let us again consider 2 bits/cell NAND flash memory and keep all parameters the same as those in Example 2.1. We can calculate the information-theoretical cell program capacity at different cell storage efficiency, as shown in Fig. 12. The results show that, in order to improve the memory cell lifetime program capacity, we should reduce the cell storage efficiency. This indicates that we should use stronger ECC with lower code rate in practice. Furthermore, we investigate the case when binary BCH code is being used, where we set each BCH codeword protects 4K-byte user data. The calculated cell lifetime program capacity is shown in Fig. 12, which shows the same trend as the information-theoretical results.

In the above discussion, we assume the cell storage efficiency (or ECC code rate) is fixed throughout the memory lifetime. In practice, if we are able to tune the ECC code rate (and hence cell storage efficiency), we can leverage the dynamic flash memory cell characteristics over P/E cycling to further improve the cell lifetime program capacity, i.e., we can dynamically tune the ECC code rate according to the present P/E cycling number in such a way that the overall lifetime cell program capacity can be maximized. In the early lifetime of NAND flash, the reliability of flash is relatively high and we can use ECC with higher code rate. As P/E cycling increases, we gradually reduce the code rate of ECC to overcome the gradually reduced reliability of NAND flash. Suppose we can choose the ECC code rate among S available code rates $r^{(1)} > r^{(2)} > \cdots > r^{(S)}$. Under each code rate $r^{(i)}$, we can estimate the corresponding P/E cycling endurance, denoted as $N_r^{(i)}$. Clearly, we have $N_r^{(1)} < N_r^{(2)} < \cdots < N_r^{(S)}$. In the run time, if the present P/E cycling number falls into $[N_r^{(i-1)}, N_r^{(i)})$, the NAND flash memory system uses the ECC code rate of $r^{(i)}$. Using such a dynamic code rate tuning scheme, we can improve the cell lifetime program capacity by

$$\frac{\sum_{i=1}^{S}(N_r^{(i)} - N_r^{(i-1)}) \cdot r^{(i)}}{N_r^{(S)} \cdot r^{(S)}} \tag{15}$$

where $N_r^{(0)} = 0$. We can use the results obtained in the above example to further evaluate the effectiveness of this dynamic code rate tuning scheme. Assume BCH code is being used and we can use four different code rates, including 0.95, 0.9, 0.85, and 0.8, corresponding to the cell storage capacity of 1.90, 1.80,

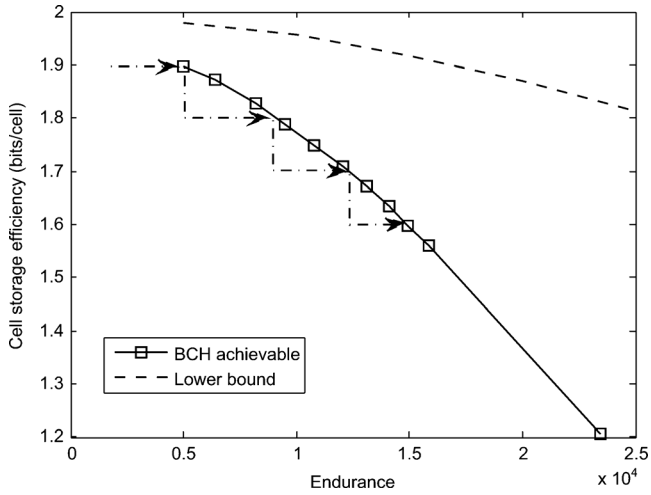

Fig. 13. Illustration of dynamically tuning ECC code rate to improve cell lifetime program capacity.

1.70, and 1.60 bits/cell, respectively. As illustrated in Fig. 13, the corresponding P/E cycling endurance is about 5K, 9K, 12K, and 15K, respectively. When P/E cycling is lower than 5K, BCH code with code rate of 0.95 is used; when P/E cycling count is larger than 5K but lower than 9K, BCH code with 0.90 code rate is used; when P/E cycling further increases to 12K and 15K, the code rate of BCH code is further decreased to 0.85 and 0.80, respectively, to accommodate the degraded reliability of NAND flash. According to (15), this leads to an 11% improvement of the memory lifetime program capacity.

V. Conclusion

By modeling NAND flash memory as a communication channel, the paper investigates how to apply information theory to estimate the theoretical bounds of the memory cell storage efficiency. This is motivated by the fact that NAND flash memory will heavily rely on system-level fault-tolerance techniques such as ECC to ensure overall system storage integrity, and it is highly desirable for these fault-tolerance techniques to maintain high cell storage efficiency. The developed NAND flash memory channel model incorporates those major noise sources including P/E cycling effects and cell-to-cell interference. Since this channel is a channel with memory, which makes it intractable to directly calculate the channel capacity, we instead develop methods to estimate tight upper and lower bounds through Monte Carlo simulation and numerical calculation. Using hypothetical 2 bits/cell NAND flash memory as an example, we carry out extensive simulations to demonstrate the estimation of the theoretical bounds of cell storage efficiency and reveal the inherent tradeoffs among cell storage efficiency, P/E cycling endurance, and retention limit. Moreover, motivated by the results of such an information-theoretical study, we develop two design techniques that can exploit the dynamics of P/E cycling to improve average NAND flash memory programming speed and increase the total amount of user data that can be stored in the memory. We expect that such an information-theoretical framework can be used to reveal important insights for designers to optimize future NAND flash memory systems for various real-life applications.

References

[1] G. Marotta, A. Macerola, A. D'Alessandro, A. Torsi, C. Cerafogli, C. Lattaro, C. Musilli, D. Rivers, E. Sirizotti, F. Paolini, G. Imondi, G. Naso, G. Santin, L. Botticchio, L. De Santis, L. Pilolli, M. L. Gallese, M. Incarnati, M. Tiburzi, P. Conenna, S. Perugini, V. Moschiano, W. Di Francesco, M. Goldman, C. Haid, D. Di Cicco, D. Orlandi, F. Rori, M. Rossini, T. Vali, R. Ghodsi, and F. Roohparvar, "A 3 bit/cell 32 Gb NAND flash memory at 34 nm with 6 MB/s program throughput and with dynamic 2 b/cell blocks configuration mode for a program throughput increase up to 13 mb/s," in *Proc. IEEE Int. Solid-State Circuits Conf.*, 2010, pp. 444–445.

[2] T. Futatsuyama, N. Fujita, N. Tokiwa, Y. Shindo, T. Edahiro, T. Kamei, H. Nasu, M. Iwai, K. Kato, Y. Fukuda, N. Kanagawa, N. Abiko, M. Matsumoto, T. Himeno, T. Hashimoto, Y.-C. Liu, H. Chibvongodze, T. Hori, M. Sakai, H. Ding, Y. Takeuchi, H. Shiga, N. Kajimura, Y. Kajitani, K. Sakurai, K. Yanagidaira, T. Suzuki, Y. Namiki, T. Fujimura, M. Mui, H. Nguyen, S. Lee, A. Mak, J. Lutze, T. Maruyama, T. Watanabe, T. Hara, and S. Ohshima, "A 113 mm^2 32 Gb 3 b/cell NAND flash memory," in *Proc. IEEE Int. Solid-State Circuits Conf.*, 2009, pp. 242–243.

[3] Y. Li, S. Lee, Y. Fong, F. Pan, T.-C. Kuo, J. Park, T. Samaddar, H. T. Nguyen, M. L. Mui, K. Htoo, T. Kamei, M. Higashitani, E. Yero, G. Kwon, P. Kliza, J. Wan, T. Kaneko, H. Maejima, H. Shiga, M. Hamada, N. Fujita, K. Kanebako, E. Tam, A. Koh, I. Lu, C. C.-H. Kuo, T. Pham, J. Huynh, Q. Nguyen, H. Chibvongodze, M. Watanabe, K. Oowada, G. Shah, B. Woo, R. Gao, J. Chan, J. Lan, P. Hong, L. Peng, D. Das, D. Ghosh, V. Kalluru, S. Kulkarni, R.-A. Cernea, S. Huynh, D. Pantelakis, C.-M. Wang, and K. Quade, "A 16 Gb 3-bit per cell (X3) NAND flash memory on 56 nm technology with 8 mb/s write rate," *IEEE J. Solid-State Circuits*, vol. 44, no. 1, pp. 195–207, Jan. 2009.

[4] N. Shibata, H. Maejima, K. Isobe, K. Iwasa, M. Nakagawa, M. Fujiu, T. Shimizu, M. Honma, S. Hoshi, T. Kawaai, K. Kanebako, S. Yoshikawa, H. Tabata, A. Inoue, T. Takahashi, T. Shano, Y. Komatsu, K. Nagaba, M. Kosakai, N. Motohashi, K. Kanazawa, K. Imamiya, H. Nakai, M. Lasser, M. Murin, A. Meir, A. Eyal, and M. Shlic, "A 70 nm 16 Gb 16-Level-Cell NAND flash memory," *IEEE J. Solid-State Circuits*, vol. 43, no. 4, pp. 929–937, Apr. 2008.

[5] C. Trinh, N. Shibata, T. Nakano, M. Ogawa, J. Sato, Y. Takeyama, K. Isobe, B. Le, F. Moogat, N. Mokhlesi, K. Kozakai, P. Hong, T. Kamei, K. Iwasa, J. Nakai, T. Shimizu, M. Honma, S. Sakai, T. Kawaai, S. Hoshi, J. Yuh, C. Hsu, T. Tseng, J. Li, J. Hu, M. Liu, S. Khalid, J. Chen, M. Watanabe, H. Lin, J. Yang, K. McKay, K. Nguyen, T. Pham, Y. Matsuda, K. Nakamura, K. Kanebako, S. Yoshikawa, W. Igarashi, A. Inoue, T. Takahashi, Y. Komatsu, C. Suzuki, K. Kanazawa, M. Higashitani, S. Lee, T. Murai, K. Nguyen, J. Lan, S. Huynh, M. Murin, M. Shlick, M. Lasser, R. Cernea, M. Mofidi, K. Schuegraf, and K. Quader, "A 5.6 MB/s 64 Gb 4 b/cell NAND flash memory in 43 nm CMOS," in *Proc. IEEE Int. Solid-State Circuits Conf.*, 2009, pp. 246–247.

[6] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York: Wiley, 1991.

[7] K. Takeuchi, T. Tanaka, and H. Nakamura, "A double-level-Vth select gate array architecture for multilevel NAND flash memories," *IEEE J. Solid-State Circuits*, vol. 31, no. 4, pp. 602–609, Apr. 1996.

[8] K.-D. Suh, B.-H. Suh, Y.-H. Lim, 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," *IEEE J. Solid-State Circuits*, vol. 30, no. 11, pp. 1149–1156, Nov. 1995.

[9] R. Bez, E. Camerlenghi, A. Modelli, and A. Visconti, "Introduction to flash memory," *Proc. IEEE*, vol. 91, no. 4, pp. 489–502, Apr. 2003.

[10] C. M. Compagnoni, M. Ghidotti, A. L. Lacaita, A. S. Spinelli, and A. Visconti, "Random telegraph noise effect on the programmed threshold-voltage distribution of flash memories," *IEEE Electron Device Lett.*, vol. 30, no. 9, pp. 984–986, Sep. 2009.

[11] P. Olivo, B. Ricco, and E. Sangiorgi, "High-field-induced voltage-dependent oxide charge," *Appl. Phys. Lett.*, vol. 48, p. 1135, 1986.

[12] P. Cappelletti, R. Bez, D. Cantarelli, and L. Fratin, "Failure mechanisms of flash cell in program/erase cycling," in *Proc. Int. Electron Devices Meet.*, 1994, pp. 291–294.

[13] H. Kurata, K. Otsuga, A. Kotabe, S. Kajiyama, T. Osabe, Y. Sasago, S. Narumi, K. Tokami, S. Kamohara, and O. Tsuchiya, "Random telegraph signal in flash memory: Its impact on scaling of multilevel flash memory beyond the 90-nm node," *IEEE J. Solid-State Circuits*, vol. 42, no. 6, pp. 1362–1369, Jun. 2007.

[14] N. Mielke, H. Belgal, I. Kalastirsky, P. Kalavade, A. Kurtz, Q. Meng, N. Righos, and J. Wu, "Flash EEPROM threshold instabilities due to charge trapping during program/erase cycling," *IEEE Trans. Device Mater. Reliab.*, vol. 4, no. 3, pp. 335–344, Sep. 2004.

[15] K. Fukuda, Y. Shimizu, K. Amemiya, M. Kamoshida, and C. Hu, "Random telegraph noise in flash memories—model and technology scaling," in *Proc. IEEE Int. Electron Devices Meet.*, 2007, pp. 169–172.

[16] N. Mielke, H. P. Belgal, A. Fazio, Q. Meng, and N. Righos, "Recovery effects in the distributed cycling of flash memories," in *Proc. IEEE Int. Reliab. Phys. Symp.*, 2006, pp. 29–35.

[17] J. D. Lee, J. H. Choi, D. Park, and K. Kim, "Degradation of tunnel oxide by fn current stress and its effects on data retention characteristics of 90 nm NAND flash memory cells," in *Proc. IEEE Int. Reliab. Phys. Symp.*, 2003, pp. 497–501.

[18] J. D. Lee, J. H. Choi, D. Park, K. Kim, R. D. Center, S. E. Co, and S. K. Gyunggi-Do, "Effects of interface trap generation and annihilation on the data retention characteristics of flash memory cells," *IEEE Trans. Device Mater. Reliab.*, vol. 4, no. 1, pp. 110–117, Mar. 2004.

[19] J.-D. Lee, S.-H. Hur, and J.-D. Choi, "Effects of floating-gate interference on NAND flash memory cell operation," *IEEE Electron. Device Lett.*, vol. 23, no. 5, pp. 264–266, May 2002.

[20] K. Kim, "Future memory technology: Challenges and opportunities," in *Proc. Int. Symp. VLSI Technol., Syst. Appl.*, 2008, pp. 5–9.

[21] K. Prall, "Scaling non-volatile memory below 30 nm," in *Proc. IEEE 2nd Non-Volatile Semicond. Memory Workshop*, 2007, pp. 5–10.

[22] H. Liu, S. Groothuis, C. Mouli, J. Li, K. Parat, and T. Krishnamohan, "3D simulation study of cell-cell interference in advanced NAND flash memory," in *Proc. IEEE Workshop Microelectron. Electron Devices*, 2009, pp. 1–3.

[23] K. Takeuchi, Y. Kameda, S. Fujimura, H. Otake, K. Hosono, H. Shiga, Y. Watanabe, T. Futatsuyama, Y. Shindo, M. Kojima, M. Iwai, M. Shirakawa, M. Ichige, K. Hatakeyama, S. Tanaka, T. Kamei, J.-Y. Fu, A. Cernea, Y. Li, M. Higashitani, G. Hemink, S. Sato, K. Oowada, S.-C. Lee, N. Hayashida, J. Wan, J. Lutze, S. Tsao, M. Mofidi, K. Sakurai, N. Tokiwa, H. Waki, Y. Nozawa, K. Kanazawa, and S. Ohshima, "A 56-nm CMOS 99-mm^2 8-Gb multi-level NAND flash memory with 10-MB/s program throughput," *IEEE J. Solid-State Circuits*, vol. 42, no. 1, pp. 219–232, Jan. 2007.

[24] K.-T. Park, M. Kang, D. Kim, S.-W. Hwang, B. Y. Choi, Y.-T. Lee, C. Kim, and K. Kim, "A zeroing cell-to-cell interference page architecture with temporary LSB storing and parallel MSB program scheme for MLC NAND flash memories," *IEEE J. Solid-State Circuits*, vol. 40, no. 4, pp. 919–928, Apr. 2008.

[25] Y. Li, S. Lee, Y. Fong, F. Pan, T.-C. Kuo, J. Park, T. Samaddar, H. Nguyen, M. Mui, K. Htoo, T. Kamei, M. Higashitani, E. Yero, G. Kwon, P. Kliza, J. Wan, T. Kaneko, H. Maejima, H. Shiga, M. Hamada, N. Fujita, K. Kanebako, E. Tam, A. Koh, I. Lu, C. Kuo, T. Pham, J. Huynh, Q. Nguyen, H. Chibvongodze, M. Watanabe, K. Oowada, G. Shah, B. Woo, R. Gao, J. Chan, J. Lan, P. Hong, L. Peng, D. Das, D. Ghosh, V. Kalluru, S. Kulkarni, R. Cernea, S. Huynh, D. Pantelakis, C.-M. Wang, and K. Quader, "A 16 Gb 3 b/cell NAND flash memory in 56 nm with 8 MB/s write rate," in *Proc. IEEE Int. Solid-State Circuits Conf. (ISSCC)*, 2008, pp. 506–632.

[26] R.-A. Cernea, L. Pham, F. Moogat, S. Chan, B. Le, Y. Li, S. Tsao, T.-Y. Tseng, K. Nguyen, J. Li, J. Hu, J. H. Yuh, C. Hsu, F. Zhang, T. Kamei, H. Nasu, P. Kliza, K. Htoo, J. Lutze, Y. Dong, M. Higashitani, J. Yang, H.-S. Lin, V. Sakhamuri, A. Li, F. Pan, S. Yadala, S. Taigor, K. Pradhan, J. Lan, J. Chan, T. Abe, Y. Fukuda, H. Mukai, K. Kawakami, C. Liang, T. Ip, S.-F. Chang, J. Lakshmipathi, S. Huynh, D. Pantelakis, M. Mofidi, and K. Quader, "A 34 Mb/s MLC write throughput 16 Gb NAND with all bit line architecture on 56 nm technology," *IEEE J. Solid-State Circuits*, vol. 44, no. 1, pp. 186–194, Jan. 2009.

[27] G. Dong, S. Li, and T. Zhang, "Using data post-compensation and pre-distortion to tolerate cell-to-cell interference in MLC NAND flash memory," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 57, no. 10, pp. 2718–2728, Oct. 2010.

[28] N. Shibata, , H. Maejima, K. Isobe, K. Iwasa, M. Nakagawa, M. Fujiu, T. Shimizu, M. Honma, S. Hoshi, T. Kawaai, K. Kanebako, S. Yoshikawa, H. Tabata, A. Inoue, T. Takahashi, T. Shano, Y. Komatsu, K. Nagaba, M. Kosakai, N. Motohashi, K. Kanazawa, K. Imamiya, H. Nakai, M. Lasser, M. Murin, and A. MeirA. Eyal and M. Shlick, "A 70 nm 16 Gb 16-level-cell NAND flash memory," in *Proc. IEEE Symp. VLSI Circuits*, 2007, pp. 190–191.

[29] R. J. McEliece, *The Theory of Information and Coding*. Cambridge, U.K.: Cambridge Univ., 2002.

[30] G. Gibson, B. Schroeder, and J. Digney, "Failure tolerance in petascale computers," *CTWatch Quarterly*, vol. 3, no. 4, pp. 4–10, Nov. 2007.

Guiqiang Dong (S'09) received the B.S. and M.S. degrees from the University of Science and Technology of China, Hefei, China, in 2004 and 2008, respectively. He is currently pursuing the Ph.D. degree from the Department of Electrical, Computer, and Systems Engineering, Rensselaer Polytechnic Institute, Troy, NY.

His research interests include coding theory, signal processing, and system and architecture design for data storage and memory systems.

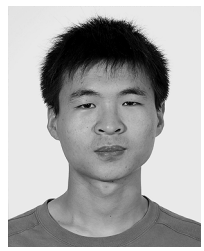

Yangyang Pan (S'12) received the B.S. degree in electrical engineering from the Zhejiang University, China, in 2007. He is currently pursuing the Ph.D. degree from the Electrical, Computer, and Systems Engineering Department, Rensselaer Polytechnic Institute, Troy, NY.

His current research areas are architecture for high performance storage system, signal processing and channel model for NAND flash SSD system.

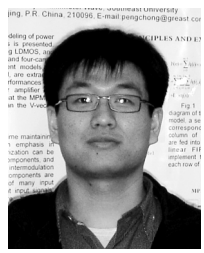

Ningde Xie received the B.S. and M.S. degrees in radio engineering from Southeast University, Nanjing, China, in 2004 and 2006, respectively, and the Ph.D. degree from the Electrical, Computer, and Systems Engineering Department, Rensselaer Polytechnic Institute, Troy, NY, in 2010.

After his Ph.D., he joined the Storage Technology Group, Intel Corporation, Portland, OR. His research interests include VLSI system and architecture design for storage and communication systems. Currently he is working on NAND and phase change memory (PCM)-based solid-state drives and non-volatile memory express host interface.

Chandra Varanasi received the Ph.D. degree in electrical engineering from the University of Minnesota, Minneapolis.

He is currently with Micron Technology, San Jose, CA. His current research interests include the area of coding and signal processing for flash-memory-based applications. Prior to that, he worked in read-channel development with Seagate Technology, Infineon Semiconductors, and Texas Instruments Storage Products Group.

Tong Zhang (M'02–SM'08) received the B.S. and M.S. degrees in electrical engineering from Xi'an Jiaotong University, Xi'an, China, in 1995 and 1998, respectively, and the Ph.D. degree in electrical engineering from the University of Minnesota, Minneapolis, in 2002.

He is currently an Associate Professor with the Department of Electrical, Computer, and Systems Engineering, Rensselaer Polytechnic Institute, Troy, NY. His research activities span over circuits and systems for various data storage and computing applications.

Dr. Zhang currently serves as an Associate Editor for the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS—PART II: EXPRESS BRIEFS and the IEEE TRANSACTIONS ON SIGNAL PROCESSING.