

Title:	SRAM bit-line swings optimization using generalized waterfilling
Archived version	Accepted manuscript: the content is identical to the published paper, but without the final typesetting by the publisher
Published version DOI :	
Conference homepage	https://www.isit2018.org/
Authors (contact)	Yongjune Kim (yongjune@illinois.edu) Mingu Kang (mkang17@illinois.edu) Lav R. Varshney (varshney@illinois.edu) Naresh R. Shanbhag (shanbhag@illinois.edu)
Affiliation	University of Illinois at Urbana Champaign

Article begins on next page

SRAM Bit-line Swings Optimization using Generalized Waterfilling

Yongjune Kim, Mingu Kang, Lav R. Varshney, and Naresh R. Shanbhag
University of Illinois at Urbana-Champaign
Email: {yongjune, mkang17, varshney, shanbhag}@illinois.edu

Abstract—We propose an information-theoretic approach to optimize non-uniform bit-line swings for static random access memories (SRAMs). We formulate convex optimization problems whose objectives are to minimize energy (for low-power SRAMs), maximize speed (for high-speed SRAMs), and minimize energy-delay product for a given constraint on mean squared error of retrieved words. We show that these optimization problems can be interpreted as *generalized water-filling* including classical water-filling, ground-flattening and water-filling, and sand-pouring and water-filling, respectively. Numerical results show that energy-optimal swing assignment reduces energy consumption by half at a peak signal-to-noise ratio of 30dB for an 8-bit accessed word.

I. INTRODUCTION

The historical scaling of transistor feature size per Moore’s Law [1] has driven the growth of the semiconductor industry since the 1960’s. Though feature sizes reduction leads to higher functional density, reduced energy consumption, and faster operation, reduction into the deep nanoscale regime today has also led, among other problems, to increased stochastic behavior of semiconductor devices and circuits, which in turn has slowed down Moore’s Law. The current scenario in the semiconductor sector presents a unique opportunity for the use of information-theoretic approaches to the design of reliable computing systems on unreliable nanoscale fabrics [2], [3].

We focus on the energy cost of memory accesses which tend to dominate the energy cost of data-centric workloads. An effective technique to reduce static random access memory (SRAM) read energy is to reduce the operating supply voltage and the internal signal swings. Doing so makes SRAM reads increasingly stochastic, e.g., bit error-rate (BER) increases. We address the problem of optimizing the SRAM read access energy while preserving a required level of fidelity via an information-theoretic approach (water-filling).

In many applications including signal processing and machine learning (ML), the impact of bit errors depends on bit position. For example, errors in the most significant bits (MSBs) of image pixels degrade overall image quality much more than errors in the least significant bits (LSBs). Likewise, an MSB error can cause a catastrophic loss in the inference accuracy of ML applications. Until now, the following tech-

niques have been proposed to address the different impacts of each bit position for energy efficiency:

- 1) Storing the MSBs in more robust bit cells and the LSBs in less robust cells [4],
- 2) Unequal error protection (UEP) [5],
- 3) LSB dropping (dropping the LSBs at the cost of reduced arithmetic precision) [6].

The first approach requires costly bit-cell redesign and manual array reorganization [4]. Fine-grained UEP [5] requires more complicated hardware implementations and provides usually two levels of granularity (e.g., strongly/weakly protected). LSB dropping [6] enables dynamic fidelity control by changing the number of dropped LSBs. However, LSB dropping allows two levels of granularity (dropped/undropped) for each bit position. In [6], selective error control codes were proposed by combining UEP and LSB dropping. Since parity bits are stored in dropped LSB-cells, the encoded data has the same length as the uncoded data. In [7], the authors proposed adaptive coding techniques for different computations inspired by selective ECC techniques.

In this paper, we propose an information-theoretic approach to determine the optimal BL swing assignments after building an abstract model of SRAMs. For a given constraint on mean squared error (MSE) of retrieved words, we formulate convex optimization problems whose objectives are

- C1. Minimize energy (low-power SRAMs),
- C2. Maximize speed (high-speed SRAMs),
- C3. Minimize energy-delay product (EDP).

Solutions to these convex problems yield optimal performance that is theoretically attainable.

Importantly, we provide *generalized water-filling* interpretations for our optimal solutions. This follows since accessing a B -bit word is equivalent to communicating information through B parallel channels. In classical water-filling, the ground represents the noise levels of parallel channels [8]. On the other hand, the importance of each bit position determines the ground level in our optimization problems. Each optimization problem has its own interpretation depending on its objective function: water-filling (C1), ground-flattening and water-filling (C2), and sand-pouring and water-filling (C3), respectively. We also observe interesting connections between our problems and variants on water-filling such as *constant water-filling* [9] and *mercury/water-filling* [10].

This work was supported in part by Systems on Nanoscale Information fabriCs (SONIC), one of the six SRC STARnet Centers, sponsored by MARCO and DARPA.

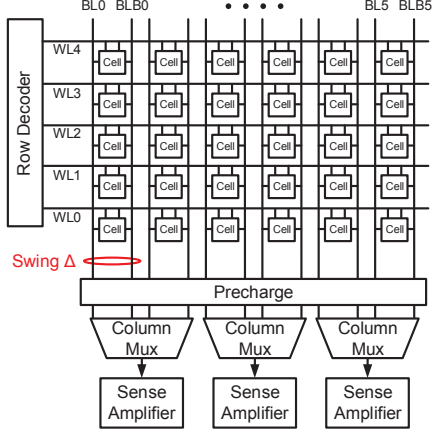


Fig. 1. A typical $N_{BL} \times N_{WL}$ SRAM block ($N_{BL} = 6$ and $N_{WL} = 5$).

The rest of this paper is organized as follows. Section II introduces key metrics of energy, delay, and fidelity. Section III formulates convex optimization problems to determine the optimum bit-level swings and provides generalized water-filling interpretations. Section IV gives numerical results and Section V concludes.

II. SRAM METRICS FOR RESOURCE AND FIDELITY

The read access energy E_{array} is the dominant contributor to energy consumption in high-density SRAMs during normal read operations [11], which is given by

$$E_{\text{array}} \propto N_{BL} N_{WL} C_{\text{bit}} V_{\text{dd}} \Delta \quad (1)$$

where N_{BL} and N_{WL} are the numbers of bit-lines (BLs) and word-lines (WLs) in a memory bank, respectively. C_{bit} is the BL capacitance per bit cell and V_{dd} is the supply voltage. Also, Δ denotes the BL voltage swing in read access. As shown in Fig. 1, the voltage swing Δ is the voltage difference between BL and BL-bar (BLB). This voltage difference occurs because either BL or BLB is discharged according to the stored bit. A sense amplifier detects which line (BL or BLB) has the higher voltage and decides whether the corresponding bit cell stores 1 or 0.

The swing Δ can be controlled by changing the WL pulse-width (i.e., WL activation time) T_{WL} since

$$\Delta \propto T_{WL} \quad (2)$$

From (1) and (2), we observe that E_{array} is directly proportional to T_{WL} . Also, T_{WL} affects the read access time [11].

Since larger voltage swing Δ improves noise margin, there are trade-off relations between reliability, energy, and delay. In the following subsections, we build an *abstract model* by taking into these SRAM properties.

A. Resource Metrics for Accessing B -bit Word

We define resource metrics for energy, delay, and EDP for accessing a B -bit word.

Definition 1: The read energy to access a B -bit word is

$$E(\Delta) = \sum_{b=0}^{B-1} \Delta_b = \mathbf{1}^T \Delta \quad (3)$$

where $\mathbf{1}$ denotes the all-one vector and the superscript T denotes transpose. Also, $\Delta = (\Delta_0, \dots, \Delta_{B-1})$ where Δ_b denotes the swing for the b th bit position in a B -bit word. Note that $E_{\text{array}} \propto E(\Delta)$.

Definition 2: The maximum swing corresponding to a B -bit word is

$$\rho = \max(\Delta) = \max\{\Delta_0, \dots, \Delta_{B-1}\}. \quad (4)$$

If we allot non-uniform swings for each bit position, the access time for a B -bit word depends on $T_{\text{max}} = \max\{T_{WL,0}, \dots, T_{WL,B-1}\}$ where $T_{WL,b}$ denotes the WL pulse-width for the b th bit position. Note that T_{max} is the pulse-width corresponding to the maximum swing ρ because of (2). Hence, the maximum swing ρ is a proper metric to be minimized to maximize read speed.

The EDP is a fundamental metric as it captures the trade-off between energy and delay [12]. We define the EDP for accessing a B -bit word based on Definitions 1 and 2.

Definition 3: The EDP to access a B -bit word is

$$\text{EDP}(\Delta) = E(\Delta) \cdot \rho = \mathbf{1}^T \Delta \cdot \rho. \quad (5)$$

B. Fidelity Metric for Accessing B -bit Word

We will define a fidelity metric for accessing a B -bit word. Suppose that a B -bit word $x = (x_0, \dots, x_{B-1})$ is stored in SRAM cells, where x_0 and x_{B-1} are the LSB and MSB, respectively. Note that x can be represented by $x = \sum_{b=0}^{B-1} 2^b x_b$ where $x_b \in \{0, 1\}$ and $x \in [0, 2^B - 1]$ (for integers i and j such that $i < j$, $[i, j] = \{i, \dots, j\}$). Also, $\hat{x} = (\hat{x}_0, \dots, \hat{x}_{B-1})$ denotes the retrieved B -bit word. A decision error flips the original bit x_b as follows: $\hat{x}_b = x_b \oplus \epsilon_b$ where \oplus denotes XOR operator and $\epsilon_b = 1$ denotes a bit error in b th bit position. The decimal representation of the retrieved word is $\hat{x} = \sum_{b=0}^{B-1} 2^b \hat{x}_b$. The decimal error e is given by $e = \hat{x} - x = \sum_{b=0}^{B-1} 2^b e_b$ where $e_b = \hat{x}_b - x_b \in \{-1, 0, 1\}$.

Since major noise sources of SRAMs are well modeled as Gaussian distributions [13], the error probability of the b th bit position is given by

$$p_b = \Pr(\epsilon_b = 1) = Q\left(\frac{\Delta_b}{\sigma}\right) \quad (6)$$

where Δ_b and σ^2 denote the swing of b th bit position and the noise variance in the corresponding BL, respectively. Note that $Q(x) = \int_x^\infty \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{t^2}{2}\right) dt$. By increasing Δ_b in (6), we can reduce p_b . However, larger Δ_b implies more energy consumption and slower speed (see Definitions 1 and 2).

To measure memory retrieval reliability, bit error probability (6) is not appropriate for many applications, since it does not distinguish the differential impact of MSB and LSB errors. Hence, we use the MSE as a fidelity metric.

Lemma 4: $\text{MSE}(x)$ is a convex function of Δ . Hence, $\text{MSE}(x)$ is given by

$$\text{MSE}(x) = \text{MSE}(\Delta) = \sum_{b=0}^{B-1} 4^b Q\left(\frac{\Delta_b}{\sigma}\right). \quad (7)$$

Proof: The proof is given in [14]. \blacksquare

TABLE I
RESOURCE AND FIDELITY METRICS FOR SINGLE-BIT AND B -BIT WORD ACCESS

	Single bit	B -bit word	Remarks
Energy	Δ	$E(\Delta) = \mathbf{1}^T \Delta$	Definition 1
Delay	Δ	$\rho = \max(\Delta)$	Definition 2
EDP	Δ^2	$EDP(\Delta) = E(\Delta) \cdot \rho$	Definition 3
Fidelity	$p = Q\left(\frac{\Delta}{\sigma}\right)$	$MSE(\Delta) = \sum 4^b Q\left(\frac{\Delta_b}{\sigma}\right)$	Lemma 4

Note that $MSE(x)$ is the nonnegative weighted sum of bit error probabilities. The weight 4^b represents the differential importance of each bit position. We show that $MSE(x)$ is convex. Table I summarizes the key resource and fidelity metrics for single-bit and B -bit word accesses.

III. OPTIMAL BIT-LEVEL SWINGS

We formulate convex optimization problems to determine the optimum swings. For a given constraint on MSE, we attempt to (1) minimize energy (low-power SRAMs), (2) maximize speed (high-speed SRAMs), and (3) minimize EDP. Also, we provide generalized water-filling interpretations of these optimization problems based on KKT conditions.

A. Energy Minimization

We formulate the following optimization problem to minimize the read access energy for a given constraint on MSE.

$$\begin{aligned} & \underset{\Delta}{\text{minimize}} && E(\Delta) = \mathbf{1}^T \Delta \\ & \text{subject to} && \sum_{b=0}^{B-1} 4^b Q\left(\frac{\Delta_b}{\sigma}\right) \leq \mathcal{V} \\ & && \Delta_b \geq 0, \quad b = 0, \dots, B-1 \end{aligned} \quad (8)$$

where \mathcal{V} is a constant corresponding to the given constraint of MSE.

Since the objective and constraints are convex, the optimization problem (8) is convex. The optimal solution can be derived by KKT conditions.

Theorem 5: The optimal swing Δ^* of (8) is given by

$$\Delta_b^* = \begin{cases} 0, & \text{if } \nu \leq \frac{\sqrt{2\pi}\sigma}{4^b}, \\ \sigma \sqrt{2 \log\left(\frac{4^b}{\sqrt{2\pi}\sigma} \cdot \nu\right)}, & \text{otherwise} \end{cases} \quad (9)$$

where ν is a dual variable.

Proof: We define the Lagrangian $L_1(\Delta, \nu, \lambda)$ associated with problem (8) as

$$L_1 = \mathbf{1}^T \Delta + \nu \left(\sum_{b=0}^{B-1} 4^b Q\left(\frac{\Delta_b}{\sigma}\right) - \mathcal{V} \right) - \sum_{b=0}^{B-1} \lambda_b \Delta_b \quad (10)$$

where ν and $\lambda = (\lambda_0, \dots, \lambda_{B-1})$ are the dual variables. The optimal solution (9) is derived from L_1 and the corresponding KKT conditions. The details of proof are given in [14]. ■

The optimal solution (9) can be interpreted as classical *water-filling* as shown in Fig. 2. Each bit position can be regarded as an individual channel among B parallel channels. In the water-filling interpretation, the ground levels depend on the importance of bit positions. We flood the bins to the water




Fig. 2. Water-filling interpretation of Theorem 5.

level of $\log \nu$. Since the MSB has the lowest ground level and the LSB has the highest ground level, larger swings are assigned to more significant bit positions. For a bit position b such that $\nu > \frac{\sqrt{2\pi}\sigma}{4^b}$, we can readily obtain the following equation (see [14]):

$$\log \nu = \log \frac{\sqrt{2\pi}\sigma}{4^b} + \frac{\Delta_b^2}{2\sigma^2} \quad (11)$$

where $\log \nu$, $\log \frac{\sqrt{2\pi}\sigma}{4^b}$, and $\frac{\Delta_b^2}{2\sigma^2}$ represent the water level, the ground level, and the water depth, respectively. The water level $\log \nu$ depends on \mathcal{V} in (8).

Remark 6 (LSB dropping and constant-power water-filling): Constant-power water-filling activates the subset of parallel channels but with a constant power allocation [9]. Constant-power water-filling in information theory is equivalent to LSB dropping in circuit theory [6] since LSB dropping allocates uniform swings for undropped bit positions.

B. Speed Maximization

Here, we maximize the speed of read access for a given constraint on MSE. The maximum speed can be achieved by minimizing ρ of (4) since ρ is proportional to the maximum pulse-width T_{\max} . By introducing an additional variable ξ , we can formulate the following convex optimization problem to maximize the speed.

$$\begin{aligned} & \underset{\Delta, \xi}{\text{minimize}} && \xi \\ & \text{subject to} && \sum_{b=0}^{B-1} 4^b Q\left(\frac{\Delta_b}{\sigma}\right) \leq \mathcal{V} \\ & && 0 \leq \Delta_b \leq \xi, \quad b = 0, \dots, B-1 \end{aligned} \quad (12)$$

From KKT conditions, we can find that $\xi = \rho$ [14].

Theorem 7: The optimal swing Δ^* of (12) is given by

$$\Delta_b^* = \rho = \xi = \sigma \sqrt{2 \log\left(\frac{4^B - 1}{3\sqrt{2\pi}\sigma} \cdot \nu\right)} \quad (13)$$

for all $b \in [0, B-1]$. Note that ν is a dual variable.

Proof: We define the Lagrangian $L_2(\Delta, \xi, \nu, \lambda, \eta)$ associated with problem (12) as

$$\begin{aligned} L_2 = & \xi + \nu \left(\sum_{b=0}^{B-1} 4^b Q\left(\frac{\Delta_b}{\sigma}\right) - \mathcal{V} \right) - \sum_{b=0}^{B-1} \lambda_b \Delta_b \\ & + \sum_{b=0}^{B-1} \eta_b (\Delta_b - \xi) \end{aligned} \quad (14)$$




Fig. 3. Ground-flattening and water-filling interpretation of Theorem 7.




Fig. 4. Sand-pouring and water-filling interpretations of Theorem 10.

where $\nu, \lambda = (\lambda_0, \dots, \lambda_{B-1})$ and $\eta = (\eta_0, \dots, \eta_{B-1})$ are dual variables. The optimal solution (13) can be derived from L_2 and corresponding KKT conditions. The details of the proof are given in [14]. ■

The optimal solution (13) can be interpreted as *ground-flattening* and *water-filling*. For any $b \in [0, B-1]$, we derive the following equation (see [14]):

$$\log \nu = \log \frac{\sqrt{2\pi\sigma}}{4^b} + \log \eta_b + \frac{\Delta_b^2}{2\sigma^2} \quad (15)$$

where η_b is a dual variable. Note that $\log \nu$, $\log \frac{\sqrt{2\pi\sigma}}{4^b}$, $\log \eta_b$, and $\frac{\Delta_b^2}{2\sigma^2}$ represent the water level, the ground level, the ground-flattening term, and the water depth, respectively. Compared with (11), we observe that (15) has an additional ground-flattening term $\log \eta_b$. From KKT conditions, we show that $\log \eta_b = \log \frac{3}{4^{B-1}} \cdot 4^b$ [14]. Hence, the *flattened ground level* (i.e., the sum of the ground level and the ground flattening term) is given by $\log \frac{\sqrt{2\pi\sigma}}{4^b} + \log \eta_b = \log \frac{3\sqrt{2\pi\sigma}}{4^{B-1}}$, which is constant for every bit position. Since the unequal ground levels are flattened by the flattening terms, the water depths of all bit positions are identical after water-filling (see Fig. 3).

Remark 8: Conventional uniform swing assignment maximizes the read access speed.

Remark 9: The overall BER is the sum of bit error probabilities of all bit positions, i.e., $\text{BER} = \sum_{b=0}^{B-1} Q\left(\frac{\Delta_b}{\sigma}\right)$. Since $Q(\cdot)$ is convex [14], the uniform swing assignment minimizes the overall BER.

C. EDP Minimization

We formulate the following convex optimization problem to minimize EDP for a given constraint on MSE.

$$\begin{aligned} & \underset{\Delta, \xi}{\text{minimize}} && \mathbf{1}^T \Delta \cdot \xi \\ & \text{subject to} && \sum_{b=0}^{B-1} 4^b Q\left(\frac{\Delta_b}{\sigma}\right) \leq \nu \\ & && 0 \leq \Delta_b \leq \xi, \quad b = 0, \dots, B-1 \end{aligned} \quad (16)$$

	Water-filling interpretations	Ground levels
Min energy	Water-filling	Unflattened
Max speed	Ground-flattening/water-filling	Perfectly flattened
Min EDP	Sand-pouring/water-filling	Partially flattened

which is derived by taking into account (5) and (12). We show that ξ is equal to ρ (see [14]).

Theorem 10: The optimal swing Δ^* of (16) is given by

$$\Delta_b^* = \begin{cases} 0, & \text{if } \log \frac{\nu}{\rho} \leq \log \frac{\sqrt{2\pi\sigma}}{4^b}, \\ \rho, & \text{if } \log \frac{\nu}{\rho} \geq \log \frac{\sqrt{2\pi\sigma}}{4^b} + \frac{\rho^2}{2\sigma^2}, \\ \sigma \sqrt{2 \log \left(\frac{4^b}{\sqrt{2\pi\sigma}} \cdot \frac{\nu}{\rho} \right)}, & \text{otherwise} \end{cases} \quad (17)$$

where ν is a dual variable.

Proof: We define the Lagrangian $L_3(\Delta, \xi, \nu, \lambda, \eta)$ associated with problem (16) as

$$\begin{aligned} L_3 = & \mathbf{1}^T \Delta \cdot \xi + \nu \left(\sum_{b=0}^{B-1} 4^b Q\left(\frac{\Delta_b}{\sigma}\right) - \nu \right) - \sum_{b=0}^{B-1} \lambda_b \Delta_b \\ & + \sum_{b=0}^{B-1} \eta_b (\Delta_b - \xi) \end{aligned} \quad (18)$$

where $\nu, \lambda = (\lambda_0, \dots, \lambda_{B-1})$, and $\eta = (\eta_0, \dots, \eta_{B-1})$ are dual variables. The optimal solution (17) can be derived from L_3 and corresponding KKT conditions. The details of the proof are given in [14]. ■

The following corollary shows the relation between the sand depth and other metrics.

Corollary 11: The sand depth s_b is given by

$$s_b = \begin{cases} 0, & \text{if } 0 \leq \Delta_b < \rho, \\ \log \left(1 + \frac{\eta_b}{\rho} \right), & \text{if } \Delta_b = \rho \end{cases} \quad (19)$$

where η_b is a dual variable.

Proof: The proof is given in [14]. ■

The optimal solution of (17) can be interpreted by *sand-pouring* and *water-filling* as shown in Fig. 4. For $\log \frac{\nu}{\rho} > \log \frac{\sqrt{2\pi\sigma}}{4^b}$, we derive the following equation (see [14]):

$$\log \frac{\nu}{\rho} = \log \frac{\sqrt{2\pi\sigma}}{4^b} + \log \left(1 + \frac{\eta_b}{\rho} \right) + \frac{\Delta_b^2}{2\sigma^2} \quad (20)$$

where $\log \frac{\nu}{\rho}$, $\log \frac{\sqrt{2\pi\sigma}}{4^b}$, $\log \left(1 + \frac{\eta_b}{\rho} \right)$, and $\frac{\Delta_b^2}{2\sigma^2}$ represent the water level, the ground level, the sand depth, and the water depth, respectively. *Pouring sand* suppresses the maximum water depth (i.e., the maximum swing) and *water-filling* allocates swings by taking into account energy efficiency.

Remark 12 (Sand-pouring and mercury-filling): *Sand-pouring* and *water-filling* has a connection to mercury/water-filling [10] because both are two-level filling. In the mercury/water-filling problem, the mercury is poured before water-filling to fill the gap between an ideal Gaussian signal and practical signal constellations, hence, each mercury depth depends only on the corresponding signal constellation. On

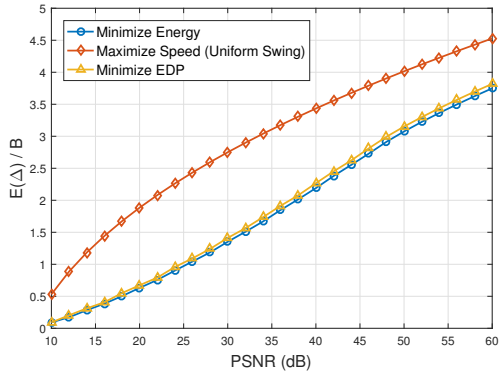


Fig. 5. Comparison of energy consumption ($B = 8, \sigma = 1$).

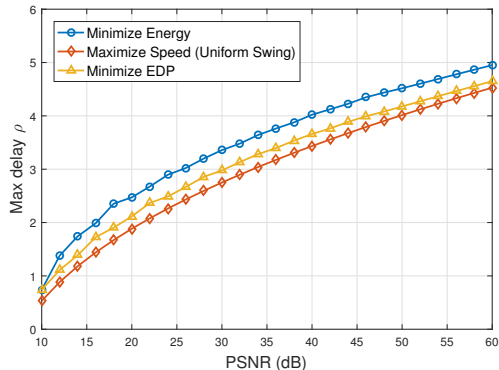


Fig. 6. Comparison of maximum delay ($B = 8, \sigma = 1$).

the other hand, sand-pouring depends on the ground levels and sand depths are correlated with each other since sand-pouring attempts to flatten the ground.

Table II summarizes generalized water-filling interpretations for our optimization problems.

IV. NUMERICAL RESULTS

We evaluate the solutions of the three optimization problems. Note that the solution of maximizing speed is equivalent to the conventional uniform swing. Fig. 5 compares the read energy consumption $E(\Delta)$ as in Definition 1 for a given constraint of peak signal-to-noise ratio (PSNR). The PSNR depends on the MSE, i.e., $\text{PSNR} = 10 \log_{10} \frac{(2^B - 1)^2}{\text{MSE}(\Delta)}$. At $\text{PSNR} = 30\text{dB}$, the optimal solution of (8) (i.e., minimizing energy) reduces the energy consumption by half for $B = 8$, compared to uniform swing (i.e., maximizing speed).

Fig. 6 compares the maximum delay ρ as in Definition 2 for a given PSNR. The conventional uniform swing minimizes the maximum delay; hence it is the speed-optimal solution. The swings minimizing energy achieve significant energy savings at the cost of speed (e.g., the maximum delay increase of 20% at $\text{PSNR} = 30\text{dB}$). The EDP-optimal swings increase only 8% of maximum delay at $\text{PSNR} = 30\text{dB}$.

Fig. 7 compares the EDP for a given PSNR. As formulated, the swings minimizing EDP show the best results. The EDP can be reduced by 45% for $B = 8$ at $\text{PSNR} = 30\text{dB}$. Note that slight loss of speed performance can result in significant energy and EDP savings.




Fig. 7. Comparison of EDP ($B = 8, \sigma = 1$).

V. CONCLUSION

We formulated convex optimization problems to determine the optimal swings for the objective functions of energy, maximum delay, and EDP. By treating each bit position as an individual channel, we interpreted bit-level swing optimization problems as generalizations of water-filling involving sand-pouring and ground-flattening.

REFERENCES

- [1] G. E. Moore, "Cramming more components onto integrated circuits," *Electronics*, vol. 38, no. 8, pp. 114–117, Apr. 1965.
- [2] N. R. Shanbhag, N. Verma, Y. Kim, A. D. Patil, and L. R. Varshney, "Shannon-inspired statistical computing for the nanoscale era," *Proc. IEEE*, 2018, under review.
- [3] R. A. Abdallah and N. R. Shanbhag, "An energy-efficient ECG processor in 45-nm CMOS using statistical error compensation," *IEEE J. Solid-State Circuits*, vol. 48, no. 11, pp. 2882–2893, Nov. 2013.
- [4] I. J. Chang, D. Mohapatra, and K. Roy, "A priority-based 6T/8T hybrid SRAM architecture for aggressive voltage scaling in video applications," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 21, no. 2, pp. 101–112, Feb. 2011.
- [5] X. Yang and K. Mohanram, "Unequal-error-protection codes in SRAMs for mobile multimedia applications," in *Proc. IEEE/ACM Int. Conf. Comput.-Aided Design (ICCAD)*, Nov. 2011, pp. 21–27.
- [6] F. Frustaci, D. Blaauw, D. Sylvester, and M. Alioti, "Approximate SRAMs with dynamic energy-quality management," *IEEE Trans. VLSI Syst.*, vol. 24, no. 6, pp. 2128–2141, Jun. 2016.
- [7] C. H. Huang, Y. Li, and L. Dolecek, "ACOCO: Adaptive coding for approximate computing on faulty memories," *IEEE Trans. Commun.*, vol. 63, no. 12, pp. 4615–4628, Dec. 2015.
- [8] C. E. Shannon, "Communication in the presence of noise," *Proc. IRE*, vol. 37, no. 1, pp. 10–21, Jan. 1949.
- [9] W. Yu and J. M. Cioffi, "On constant power water-filling," in *Proc. IEEE Int. Conf. Commun. (ICC)*, Jun. 2001, pp. 1665–1669.
- [10] A. Lozano, A. M. Tulino, and S. Verdú, "Optimum power allocation for parallel Gaussian channels with arbitrary input distributions," *IEEE Trans. Inf. Theory*, vol. 52, no. 7, pp. 3033–3051, Jul. 2006.
- [11] M. H. Abu-Rahma, M. Anis, and S. S. Yoon, "Reducing SRAM power using fine-grained wordline pulswidth control," *IEEE Trans. VLSI Syst.*, vol. 18, no. 3, pp. 356–364, Mar. 2010.
- [12] M. Horowitz, T. Indermaur, and R. Gonzalez, "Low-power digital design," in *Proc. IEEE Symp. Low Power Electr.*, Oct. 1994, pp. 8–11.
- [13] T. Mizuno, J. Okumtura, and A. Toriumi, "Experimental study of threshold voltage fluctuation due to statistical variation of channel dopant number in MOSFET's," *IEEE Trans. Electron Devices*, vol. 41, no. 11, pp. 2216–2221, Nov. 1994.
- [14] Y. Kim, M. Kang, L. R. Varshney, and N. R. Shanbhag, "Generalized water-filling for source-aware energy-efficient SRAMs," *arXiv preprint arXiv:1710.07153*, Oct. 2017.