Data-Driven Low-Cost On-Chip Memory with Adaptive Power-Quality Trade-off for Mobile Video Streaming

Dongliang Chen, Jonathon Edstrom, Xiaowei Chen, *Wei Jin, Jinhui Wang and Na Gong

Department of Electrical and Computer Engineering, *Department of Computer Science

North Dakota State University, Fargo, North Dakota, USA

Emails: {dongliang.chen, Jonathon.edstrom, xiaowei.chen, wei.jin, jinhui.wang.1, na.gong}@ndsu.edu

ABSTRACT

Nowadays, people enjoy watching mobile videos more than ever and mobile video streaming contributes to the majority of the total mobile data traffic. However, due to the high power consumption of mobile video decoders, especially the on-chip memories, short battery life represents one of the biggest contributors to user dissatisfaction. Various mobile embedded memory techniques have been investigated to reduce power consumption and prolong battery life. Unfortunately, the existing hardware-level research suffers from high implementation complexity and large overhead. In this paper, by introducing advanced data-mining techniques, we investigate meaningful data patterns hidden in mobile video data and apply the identified patterns to implement a low-power flexible hardware design with dynamic power-quality trade-off. A 45nm 32kb SRAM is presented that enables three levels of power-quality trade-off (up to 43.7% power savings) with negligible area overhead (0.06%).

CCS Concepts

• Hardware → Semiconductor memory and Circuit power issues

Keywords

Mobile video streaming, on-chip memory, data mining, power efficient, dynamic trade-off.

1. INTRODUCTION

In recent years, the growing popularity of smart phones and the powerful accessibility to high speed networks have resulted in exponentially increasing demand for video services on mobile devices. According to market research, by 2020, the amount of data that is created, replicated, and consumed, will be as large as 40ZB (Zettabyte, or 10^{21} B) [14]; and more than half of the data traffic will be video data [15]. However, as people enjoy watching videos anytime and anywhere, video processing has become the most important energy-intensive application used in mobile devices [16]. In particular, the intensive computations of processing these video streams need to frequently access on-chip memory, which contributes to over 30% of system power consumption and occupies more than 65% of video decoder area [17, 18]. This situation is only expected to grow with the emerging high quality mobile video formats such as 8K Ultra HD applications due to the increased complexity per pixel and a huge requirement for on-chip memory as pipeline buffers [19].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

ISLPED '16, August 08-10, 2016, San Francisco Airport, CA, USA

© 2016 ACM. ISBN 978-1-4503-4185-1/16/08...\$15.00 DOI:

<http://dx.doi.org/10.1145/2934583.2934619>

Various low power SRAM designs have been developed for mobile video applications. In [10] and [24], two hybrid SRAM structures with 6T/8T and 8T/10T are developed to optimize the power efficiency for mobile video streaming. In [16], a heterogeneous sizing scheme was presented to reduce the failure probability of conventional 6T bitcells. In [11], a two-port SRAM with majority logic and data reordering is presented to minimize power consumption. However, all of those existing techniques suffer from large silicon area overhead. Also, the power-quality tradeoff is set during design time, which can no longer automatically guarantee maximum power efficiency for different video applications. Recently, Frustaci *et al.* [12] presents a voltage-scaled SRAM that can dynamically manage the trade-off between power and video quality, but the utilized write assist technique and Error Correcting Code (ECC) encoder and decoder circuits result in large penalties in computation complexity and silicon area.

In this paper, we propose a low-cost Data-Driven power efficient Adaptable SRAM Hardware (D-DASH) design with dynamic power-quality tradeoff for mobile video applications. By introducing advanced data mining techniques, we investigate meaningful data patterns hidden in video data and incorporate these key findings into our hardware design. D-DASH enables three levels of power-quality management (up to 43.7% power savings) with negligible area overhead (0.06%).

This paper is organized as follows. In Section 2, data-mining enabled mobile video data patterns are analyzed. In Section 3, we present D-DASH. The evaluation results are provided in Section 4. Finally, the conclusion is drawn in Section 5.

2. Mobile Video Data Pattern Analysis

2.1 Mobile Video Data Characteristics

Mobile video application characteristics imply that it is possible to incorporate application-level video data behaviors to the hardware-level design process. Although mobile videos are delivered over different networks and are visualized in various mobile terminals, there are three common characteristics that may potentially contribute to hardware-level memory design [20]: (1) *inputs*: the video data is noisy and redundant; (2) *outputs*: the videos on mobile devices are generated for humans and minor variations cannot be discerned by humans' eyes; and (3) *computation patterns*: statistical computations during the video decoding process potentially result in specific data patterns, which can contribute to low-power hardware design. However, it is difficult for hardware designers to observe the inherent video data behaviors directly from the large volume of video data. To achieve this, we introduce data-mining techniques to comprehensively explore mobile video storage data characteristics. In particular, we use an association rule mining technique to explore the relationship between different video data bits and obtain data patterns for efficient hardware design.

Figure 1. YUV Format, Dataset Implementation, and Discovered Rules.

2.2 Data Mining Assisted Video Analysis

Today's mobile video frames are typically stored and processed in YUV format. It includes one *luma* (Y) component, which contains the brightness information of the image and two *chroma* components which contain the blue-difference (Cb) and red-difference (Cr) color information. Fig. 1 shows a typical frame of video data stored in embedded memory using a 352×288 resolution YUV 4:2:0 video as an example. As shown, each pixel has 8-bit luma data and 8-bit subsampled chroma data. Since video data is stored in on-chip memory as binary bits, we utilize an association data mining technique to identify the bit-level data patterns.

Association rule mining was introduced in 1993 to discover *relationships* between different variables, called items, in a dataset or database [9]. A complete dataset is made up of many transactions where each transaction contains a set of items. Each item can be associated with a binary attribute, 0 or 1, that is used to distinguish that item is present or not in its corresponding transaction. This type of data organization is illustrated in Fig. 1. Each resulting rule, generated from the association rule mining process, is an implication of the form $X \rightarrow Y$, where X and Y are disjoint sets of, or individual, items. Each rule is also accompanied by collected statistics from the dataset called support and confidence values. The support value for a set of items is the proportion of transactions in the dataset that contains such set of items. The confidence value for an association rule, $X \rightarrow Y$, is the proportion of transactions that contain X which also contain Y , or the conditional probability $P(Y/X)$.

To enable association data mining, we use different video benchmarks with various characteristics (e.g. motion, scene complexity) to build a dataset, including 12 videos from [3] and 4 videos from [4]. In total, the video data size is 415,600,000 bytes. 12 benchmark videos from [3] were combined into a single YUV file. This combined video file contains a total of 3470 frames with 352×288 resolution. Each chroma bit is defined as an individual item and a sample of the *.arff* format with descriptions of its different parts can be seen in Fig.1. We used *Weka* [2] to perform the well-known association rule mining algorithm - *Apriori* on our large video dataset. To evaluate the impact of compression on the identified data patterns, the two formats including the non-sampled

YUV 4:4:4 format and the subsampled YUV 4:2:0 format are investigated using data mining techniques. They both contain eight bits of luma (Y) data and eight bits of each chroma component (Cb and Cr) that are shared among 4 pixels.

One interesting data pattern we obtained from the results is the strong association of chroma bits to Cr's most significant bit (MSB), Cr1. As shown in Fig. 2, if the value of Cr1 is equal to 1 (0) the remaining Cr bits have a larger probability to be 0 (1). Other than Cb1, the majority of Cb bits have a larger probability to be 0 (1) when Cr1 equals 0 (1). Based on such identified data patterns, a flexible D-DASH with power-quality adaption is implemented which will be discussed in the following section.

Values	Bit 2	Bit 3	Bit 4	Bit 5	Bit 6	Bit 7	Bit 8
Cr1=0, CrX=1	99.9%	92.2%	89.7%	71.8%	71.4%	64.3%	56.3%
Cr1=1, CrX=0	99.2%	97.8%	83.3%	67.5%	59.8%	57.3%	53.2%
Cr1=0, CbX=0	60.5%	60.4%	64.0%	56.6%	52.8%	48.6%	50.5%
Cr1=1, CbX=1	86.6%	82.9%	55.3%	51.4%	54.8%	49.5%	50.4%

Figure 2. Rules Diagram and Proposed Method. The percentage represents the probability of next bit to be 0 or 1.

2.3 Video Quality Metrics

The well-known peak signal-to-noise ratio (PSNR) metric is applied widely to evaluate video quality [10, 12, 23], which is defined as [21]

$$PSNR = 10 \log_{10} \left(\frac{255^2}{MSE} \right) \quad (1)$$

where MSE is the mean square error between the original videos (Org) and the degraded videos (Deg), expressed as

$$MSE = \frac{1}{mn} \sum_{i=0}^{m-1} \sum_{j=0}^{n-1} [Org(i, j) - Deg(i, j)]^2 \quad (2)$$

However, recent research shows that the PSNR cannot describe the true human perception of videos since it only takes the amount of errors into account, not necessarily the effect that errors have on the user's perception of the image being displayed [1]. Accordingly, the structural similarity (SSIM) metric is developed to predict the perceived image quality and it combines separate calculations for luminance, contrast, and structure changes all together, as expressed in (3) [1]

$$SSIM(x, y) = [l(x, y)]^\alpha \cdot [c(x, y)]^\beta \cdot [s(x, y)]^\gamma$$

$$= \frac{(2\mu_x\mu_y + C_1)(2\sigma_{xy} + C_2)}{(\mu_x^2 + \mu_y^2 + C_1)(\sigma_x^2 + \sigma_y^2 + C_2)} \quad (3)$$

where the luminance comparison function $l(x, y)$, is a function of the mean intensities, μ_x and μ_y , the contrast comparison function $c(x, y)$, is a function of the standard deviations, σ_x and σ_y , and the structural comparison function $s(x, y)$, is a function of the correlation between x and y , σ_{xy} . Setting $\alpha = \beta = \gamma = 1$ in the original equation results in the second equation. C_1 (C_2) is a constant that is included to avoid instabilities when the sum of the means (standard deviations) squared is equal to values near zero.

In our analysis, we use both PSNR and SSIM to evaluate the video output quality. The quality reduction ($PSNR(SSIM)\%$ Reduction) can be calculated using (4):

$$PSNR(SSIM) \% Reduction = \left(1 - \left(\frac{Output\ PSNR(SSIM)}{Original\ PSNR(SSIM)} \right) \right) \times 100\% \quad (4)$$

3. PROPOSED D-DASH

This section presents the identified data patterns enabled D-DASH. D-DASH is highly flexible with three design schemes (D-DASH-I to D-DASH-III), providing a run-time dynamic power-quality trade-off for mobile video streaming.

3.1 D-DASH-I

D-DASH-I enables zero-cost power-efficient storage, using *data-aware* low-power readout buffer connections based on the obtained data patterns in mobile video data, as shown in Fig. 3. In conventional SRAM, a readout buffer consisting of two NMOS transistors is used to access the stored value by connecting it to reversed storage node (QB) [22], as shown in Fig. 3 (a). During the reading process, the read bit-line (RBL) is precharged to supply voltage (Vdd) before RWL is asserted. In the case a 1 is stored in the SRAM cell and RWL is asserted, RBL will stay at Vdd as the bottom NMOS is turned off and there is no switching activity in RBL, enabling low-power reading process. In the case that a 0 is stored in the SRAM cell, the RBL will be discharged to ground (GND), resulting in large power consumption. The discharging activity during readout process contributes significant power consumption in mobile video memory [10-12, 24]. In this paper, data-aware low-power readout buffer connections are proposed. The traditional connection as shown in Fig. 3 (a) are referred to as type-1 bitcell. Alternatively, a type-0 bitcell is presented to achieve low-power reading 0 process, by connecting readout buffers to Q, as shown in Fig. 3 (b). Note that, as compared to the conventional

Figure 3. Data-aware D-DASH bitcells (a) type-1 bitcell without discharging when reading 1; (b) type-0 bitcell without discharging when reading 0.

SRAM bitcell (type-1), the type-0 bitcell does not cause any silicon area overhead.

The obtained data patterns in Section 2.2 show that, in the chroma data for each pixel, there is over 70% probability that the Cr and Cb data will be 100000000111111 in binary. Based on this, type-1 and type-0 bitcells are applied to store/load 1 and 0, respectively. Accordingly, for the majority of chroma data, the switching activity during the readout process is significantly reduced, achieving power savings without area overhead. Fig. 3 (c) shows the structure of one wordline in D-DASH-I. Based on the identified data pattern, it uses an optional combination of type-0 and type-1 bitcells to enable zero-overhead power efficient mobile video on-chip storage.

3.2 D-DASH-II

Based on D-DASH-I, we implement D-DASH-II for additional power savings by considering the rest of the data with the pattern of 100000000111111 . Since the most significant bit (MSB) of the Cr data determines the value of lower order bits (LSBs) as shown in the identified patterns, we implement a write circuit and a read circuit to the SRAM to maximize the read bitline power saving, as shown in Fig. 4. In the write circuit, we first detect the MSB of the Cr data and then determine whether to invert the input data or not, and use a similar flag-bit scheme as [22] to indicate if the data is flipped or not: 1 means the data is not inversed and 0 means data is inversed. The scheme implemented in [22] uses an additional bitcell to store the flag bit, causing 7% area overhead. To minimize the overhead, we utilize the least significant bit (LSB) of the chroma data to store the flag bit. Our results in Table 1 show that using the Cr LSB induces negligible video quality degradation (0.044% PSNR reduction and 0.058% SSIM reduction) and therefore we use the Cr LSB to store the flag bit, as shown in Fig. 4. Accordingly, D-DASH-II enables more power savings as compared to the implementation cost and video quality degradation of D-DASH-I.

3.3 D-DASH-III

To meet the high power efficiency requirement of some video applications, we further design D-DASH-III to maximize the power savings based on bit-truncation (dropping) technique. Bit-

Figure 4. D-DASH with real-time adjustment between three schemes

truncation technique has been widely used in mobile video memory area [13]. To determine the number of bits for truncation, we evaluate the video quality with the bit-truncation technique and the results are listed in Table 1. As the number of truncated bits is larger than 7, PSNR and SSIM reduction are significant (PSNR reduction $\geq 24.395\%$ and SSIM reduction $\geq 2.512\%$), indicating large video quality degradation. Accordingly, the number of bits for truncation can be 5 or 7. We further evaluate the video output quality using the *sign_irene* video benchmark as shown in Fig. 5. The video *sign_irene* contains blue and red colors that would be directly affected by the corruption of chroma data [5]. It shows that, *truncating 5 LSB bits (2 LSBs of Cr and 3 LSBs of Cb) is an optimized trade-off between video quality and power saving*. To enable the bit-truncation technique, we implement the control circuit as shown in Fig. 4. The control bit SW is connected to Write Enable (WE) and Read Enable (RE) of a wordline. When SW is 0, all of the 32 bitcells connected to the wordline work as traditional bitcells; when SW is 1, D-DASH-III is enabled and the outputs (WE_out and RE_out) will disable the WE and RE of the truncated bitcells (yellow bitcells shown in Fig. 4) and therefore, bit truncation is enabled to achieve additional power savings.

Table 1. PSNR and SSIM Calculations.

# of drop bits	PSNR	SSIM	PSNR % Reduction	SSIM % Reduction
0 (Original)	36.868	0.93436	0.000%	0.000%
0 (Flag Bit)	36.852	0.93382	0.044%	0.058%
1	36.848	0.93321	0.054%	0.124%
3	36.641	0.92989	0.616%	0.479%
5	35.544	0.92256	3.592%	1.264%
7	32.593	0.91089	11.595%	2.512%
9	27.874	0.89529	24.395%	4.182%
11	21.506	0.86243	41.667%	7.699%
13	14.851	0.75625	59.718%	19.062%
15	10.782	0.62366	70.756%	33.253%

4. SIMULATION RESULTS

To evaluate the effectiveness of the proposed technique, a 32kb SRAM is implemented using a high-performance 45-nm FreePDK CMOS process to meet the multi-megahertz performance requirement of today's mobile video decoders.

4.1 Performance

We first evaluate the performance of the proposed D-DASH. Table 2 lists the detailed performance parameters. Both write delay and read delay are approximately 0.15 ns, which successfully delivers high-quality video format such as 8K Ultra HD applications [19].

4.2 Layout

We also evaluate the area overhead of the proposed design. Fig. 6 shows the layout of D-DASH. For scheme I, there is no area overhead; for scheme II, the area overhead is 0.64%; for scheme III, after careful layout design, we integrate the control circuit with read decoder, without additional area overhead.

4.3 Output quality

We use different videos to verify the output quality based on the proposed SRAM scheme. Fig. 7 shows three video outputs as examples. D-DASH-I and D-DASH-II can deliver good video output quality and D-DASH-III results in negligible video degradation to achieve optimal power efficiency.

Figure 5. *sign_irene* output with bit-truncation technique

D-DASH	Akiyo	Coastguard	Foreman
Original/D-DASH-I	PSNR: 41.250858, SSIM: 0.961865	PSNR: 35.638277, SSIM: 0.919692	PSNR: 37.355170, SSIM: 0.923900
D-DASH-II	PSNR: 41.194586, SSIM: 0.961432	PSNR: 35.638277, SSIM: 0.919196	PSNR: 37.335870, SSIM: 0.923268
D-DASH-III	PSNR: 37.971173, SSIM: 0.950497	PSNR: 34.781284, SSIM: 0.914910	PSNR: 35.829443, SSIM: 0.913802

Figure 7. Video Output

Figure 6. Layout of D-DASH

4.4 Power savings

To evaluate the power efficiency of D-DASH, we model the read bitline (RBL) power consumption of mobile video memory as:

$$P_r = \sum_{k=0}^{16} \sum_{i=0,1} [F_k(i) \cdot P_{rkt}(i) \cdot Z(k)] \quad (5)$$

where P_r is the power consumption on read operation; k is the bit number; t is the SRAM type; i is the value stored in SRAM, $F(i)$ indicates the probabilities of a bit to be 0 and 1, which is shown in Table 4; $Z(i)$ indicates if the bit will be truncated (if truncated, $Z(i)$ will be 0, if not truncated, $Z(i)$ will be 1).

Table 3 lists the read power consumption for two types of D-DASH bitcells. We also extract the probability of each bit being 0 or 1 from the 12 benchmarks and the results are listed in Table 4. The bits marked in grey are truncated bits in D-DASH-III.

Table 5 concludes the power saving of our proposed technique over a standard SRAM design. D-DASH enables power savings from 7.82% (D-DASH-I) to 43.07% (D-DASH-III).

Table 3. RBL power consumption for different type of SRAM

SRAM type	Read 1	Read 0
Power of Type 1	3.22e-8	5.80e-8
Power of Type 0	5.88e-8	3.13e-8

4.5 Comparison with prior work

Table 6 compares the D-DASH's performance with the state-of-the-art. D-DASH exhibits the lowest implementation cost (0.06%) with dynamic power-quality tradeoff. D-DASH-I and D-DASH-II exhibit the best video quality, except for reference [11], which, however, is realized with large area overhead (~14%). D-DASH-III demonstrates highest power efficiency, except for [24], but the hybrid 8T/10T structure in [24] requires bitcell array modification, resulting in as high as 52% silicon area overhead.

Table 4. Probability of Each bit being 0 or 1 probability ($F(i)$)

bitcell type	Bit	D-DASH-I		D-DASH-II	
		0	1	0	1
1	Cr1	0.2352	0.7648	0	1
0	Cr2	0.7588	0.2412	0.9940	0.0060
0	Cr3	0.7663	0.2337	0.9649	0.0351
0	Cr4	0.6615	0.3385	0.8482	0.1518
0	Cr5	0.5826	0.4174	0.6853	0.3147
0	Cr6	0.5247	0.4753	0.6254	0.3746
0	Cr7	0.5221	0.4779	0.5892	0.4108
0	Cr8	0.5097	0.4903	0.5097	0.4903
0	Cb1	0.7493	0.2507	0.8331	0.1669
1	Cb2	0.2456	0.7544	0.1958	0.8042
1	Cb3	0.2732	0.7268	0.2240	0.7760
1	Cb4	0.4923	0.5077	0.4265	0.5735
1	Cb5	0.5046	0.4954	0.4735	0.5265
1	Cb6	0.4702	0.5298	0.4568	0.5432
1	Cb7	0.5011	0.4989	0.5069	0.4931
1	Cb8	0.4986	0.5014	0.4956	0.5044

Table 5. Power savings of different schemes

SRAM Designs	Power (W)	Power saving
Conventional SRAM	7.29E-07	-
D-DASH-I	6.72E-07	7.82%
D-DASH-II	6.33E-07	13.17%
D-DASH-III	4.15E-07	43.07%

Table 6. Comparison with prior art on low-power mobile video SRAM.

	<i>TVLSI'08</i> [11]	<i>TCASVT'11</i> [10]	<i>TCASII'12</i> [24]	<i>JSCC'15</i> [12]	This work		
					D-DASH-I	D-DASH-II	D-DASH-III
video specific characteristics	correlation of MSB	contribution of MSB and LSB	different contribution of MSB and LSB	contribution of MSB and LSB	data-mining assisted video data patterns exploration		
dynamic adaption	No	No	No	Yes	Yes		
low-power technique	data flipping	6T+8T bitcells	8T+10T bitcells	ECC	data-aware bitline connection	data-aware bitline connection and data flipping	data-aware bitline connection, data flipping and bit truncation
bitcell array modification	Yes	Yes	additional word line	No	No	No	No
additional hardware needed	majority logic and data flipping block	single-ended 6T, peripheral circuitries	No	ECC encoder and decoder, write assisted circuit	No	control circuit consists simple gates	control circuit consists simple gates
power penalty for extra bits	Yes	No	No	No	No	No	No
readout power ¹	-14%	-32%	-95%	-28%	-7.82%	-13.17%	-43.07%
video quality	good	acceptable	acceptable	acceptable	good	good	acceptable
area	+14%	+11.64%	+52%	+1.5%	0%	+0.06%	+0.06%
technology	90nm	90nm	45nm	28nm	45nm	45nm	45nm

5. CONCLUSION

This paper presented a data-driven flexible mobile video SRAM with negligible silicon area overhead (0.06%) that can be adjusted between three different levels of power-quality tradeoff. Based on the data patterns obtained by data-mining techniques, a low-power data-aware readout buffer connection technique is applied to enhance the power efficiency without implementation overhead; based on this, a bit-flipping technique is used to further improve the power efficiency by introducing 0.06% area overhead; finally, a bit-truncation technique is utilized to maximize the power efficiency with negligible video quality degradation. The developed D-DASH provides three flexible low-cost schemes, each with different power-quality tradeoff, to meet the requirements of various applications.

6. ACKNOWLEDGMENTS

This work was supported in part by NSF under grant CCF-1514780, and the Beijing Municipal Natural Science Foundation under Grant 4152004.

REFERENCES

- [1] Wang, Z. *et al.* 2004. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13, 4, 600-612.
- [2] Weka 3: <http://www.cs.waikato.ac.nz/ml/weka/>
- [3] YUV Video Sequences. <http://trace.eas.asu.edu/yuv>
- [4] Xiph.org Video Test Media. <https://media.xiph.org/video/derf/>
- [5] Yueh, W., Cho, M. and Mukhopadhyay, S. 2013. Perceptual Quality Preserving SRAM Architecture for Color Motion Pictures. In *Proceedings of DATE, 2013*. Grenoble, France, 103-108.
- [6] Vijayakumar, V. and Nedunchezian, R. 2011. Recent Trends and Research Issues in Video Association Mining. *The International Journal of Multimedia & Its Applications (IJMA)*. 3, 4, 49-61.
- [7] Frustaci, F *et al.* 2014. 13.8 A 32kb SRAM for error-free and error-tolerant applications with dynamic energy-quality management in 28nm CMOS. In *Proceedings of ISSCC*. San Francisco, CA, USA, 244-246.
- [8] Agrawal, R. and Srikant, R. Fast Algorithms for Mining Association Rules. *IBM Almaden Research Center*, 650 Harry Road, San Jose, CA 95120.
- [9] Agrawal, R., Imielinski, T. and Swami, A. Mining Association Rules between Sets of Items in Large Databases. *IBM Almaden Research Center*, 650 Harry Road, San Jose, CA 95120.
- [10] Chang, I. J., Mohapatra, D. and Roy, K. 2011. A priority-based T/8 T hybrid SRAM architecture for aggressive voltage scaling in video applications. *IEEE Trans. Circuits Syst. Video Technol.*, 21, 2, 101-112.
- [11] Fujiwara, H. *et al.* 2006. A Two-Port SRAM for Real-Time Video Processor Saving 53% of Bitline Power with Majority Logic and Data-Bit Reordering. In *Proceedings of ISLPED*, 61-66.
- [12] Frustaci, F. *et al.* 2015. SRAM for error-tolerant applications with dynamic energy-quality management in 28 nm CMOS. *IEEE J. Solid-State Circuits*. 50, 5, 1310-1323.
- [13] Chen, D., Wang, X., Wang, J., and Gong, N. 2015. VCAS: Viewing context aware power-efficient mobile video embedded memory. In *Proceedings of SOCC*. 333-338.
- [14] IDC. 2012. The Digital Universe in 2020: Big Data, bigger Digital Shadows, and Biggest Growth in the Far East. <https://www.emc.com/collateral/analyst-reports/idc-digital-universe-united-states.pdf>
- [15] Cisco Systems, Inc. 2013. Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2012-2017. http://broadcastengineering.com/automation/http_abr_streaming
- [16] Kwon, J., Lee I., and Park J. 2012. Heterogeneous SRAM Cell Sizing for Low Power H.264 Applications. *IEEE Transactions on Circuits and Systems I*, 99, 2, 1-10.
- [17] Hoque, M. *et al.* 2014. Energy Efficient Multimedia Streaming to Mobile Devices - A Survey. *IEEE Communications Surveys & Tutorials*. 16, 1, 579-597.
- [18] Benmoussa, Y. *et al.* 2013. Energy Consumption Modeling of H.264/AVC Video Decoding for GPP and DSP. In *Proceedings of 16th Euromicro Conference on Digital System Design*. 890-896.
- [19] Zhou, D. *et al.* 2016. A 4Gpixel/s 8/10b H.265/HEVC Video Decoder Chip for 8K Ultra HD Applications. In *Proceedings of ISSCC*, 266-267.
- [20] Venkataramani, S. *et al.* 2015. Approximate Computing and the Quest for Computing Efficiency. In *Proceedings of DAC*. 120-126.
- [21] Hore, A. and Ziou, D. 2010. Image Quality Metrics: PSNR vs. SSIM. In *Proceedings of ICPR*, 2366-2369.
- [22] Fujiwara, H. *et al.* 2008. Novel Video Memory Reduces 45% of Bitline Power Using Majority Logic and Data-Bit Reordering. *IEEE Trans. on VLSI Systems*. 16, 6, 620-627.
- [23] Xue, J. and Chen, C. W. 2014. Mobile Video Perception: New Insights and Adaptation Strategies. In *IEEE Journal of Selected Topics in Signal Processing*. 8, 3, 390-401.
- [24] Gong, N. *et al.* 2012. Ultra-Low Voltage Split-Data-Aware Embedded SRAM for Mobile Video Applications. *IEEE Transactions on Circuits and Systems II*. 59, 12, 883-887.