Perceptual Quality Preserving SRAM Architecture for Color Motion Pictures

Wen Yueh, Minki Cho, Saibal Mukhopadhyay
School of Electrical and Computer Engineering
266 Ferst Drive, 2309 Klaus Advanced Computing Building
Atlanta, GA 30332
wyueh3@gatech.edu, mcho8@gatech.edu, saibal.mukhopadhyay@ece.gatech.edu

Abstract — This work proposes a low power methodology for video framebuffers to preserve the perceptual quality while reducing SRAM power. The bank-wise voltage scaling combined with error masking circuitry is proposed where voltage domains are separated according to the importance of luminous and color channels. The implementation may apply to standard embedded memory cores without redesigning specialized hardware within the SRAM bank. The simulation results showed that the proposed channel protection technique produced better energy-quality trade-off than the conventional higher-order-bit protection for the uncompressed as well as compressed motion image frames.

Keywords – parametric failure, color image protection, low power circuit, static random access memory (SRAM) .

I. INTRODUCTION

The demand for low power multi-media processing is growing along with the popularity of mobile/smart devices. The embedded video/audio processing units must balance between tight energy budget and quality expectation from the consumer. The memory system consumes a significant fraction of the total power for multimedia applications [1]. It is known that embedded video applications save memory power by reducing off-chip access [2]. Hosting the video frame-buffers on-chip using Static Random Access Memory (SRAM) for mobile application is an effective way to improve the overall energy efficiency [3]. However, the power dissipation of the on-chip SRAM buffer is nontrivial. It has been shown that even for mobile size video quality (QCIF@176x144) the buffer power may contribute as much as 70% of the chip power [4].

One effective approach to reduce the SRAM power is to scale down or reduce the operating voltage. Liu et al. have reported approximately 50% frame-buffer power reduction using voltage-scaling in a multimedia system [4]. However, one immediate concern to voltage scaling is that it reduces the SRAM margins, and causes the process variation induced parametric failures in SRAM cells become more prevalent. These failure may be due to access, read disturb and write failure [5]. All types of failures increase with the reduction of cell supply, the bitline, and the wordline voltage. In a frambuffer, increasing bit-error rate within the SRAM array translates to unpleasant artifacts during the video playback. Therefore, designing low-voltage SRAM with error tolerance and/or correction mechanisms to preserve quality is of crucial importance for low-power multimedia systems.

This paper presents a voltage scaling methodology for embedded video/image SRAMs; specifically we apply this technique to a video framebuffer as an application candidate. We guard the most predominant visual information channel, and allow error in the secondary channels. This is unlike the prior work which the higher-order-bits in the image pixels were protected (using higher voltage) to a confined error magnitude, see Fig. 1a. We propose to store image color channels in separate banks and maximize voltage scaling potential through exploiting the properties of the content distribution (Fig. 1b). The channel protection with low-cost filters, instead of bit-protection, achieves better energy-quality trade-off for *color images* without introducing specialized memory architecture to operate different bits at different voltages [5][6]. This paper demonstrates that the proposed channel protection technique enables effective error masking through low cost filtering, while the prior work may require a relative costly known-bad-cell table to achieve comparable masking quality. We study the trade-off associated with the voltage scaling of the memory, the generated bit-error rate, and the associated power overhead of the filter. We further evaluate the effect of the compression on both proposed and prior voltage scaling techniques to determine the quality to energy trade-off for a compressed video framebuffer.

II. RELATED WORK AND CONTRUBUTIONS

This paper relates to prior works in three primary areas as discussed below:

- Dynamic voltage scaling is an effective technique for low power multimedia application [5]. To limit the error associated with reduced margins and timing failures, Kang et al. and Cho et al. proposed to partially protect the most significant bits (MSBs) and exploit the arithmetic importance of higher order bits (HOBs) [5][7]. These

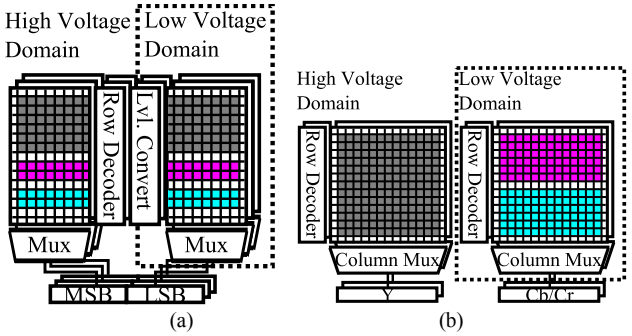

Figure 1. The figure shows (a) bit-wise protection. The technique requires level shifters on the row decoder to maximize power reduction (b) bank-wise protection essentially voltage scales the entier bank for simpler core design and reduced area overhead.

978-3-9815370-0-0/DATE13/©2013 EDAA

techniques save power by scaling voltage of the lower order bits and preserve acceptable video/image quality by protecting the HOB information. Krause et al. used delay detection latch and similar error metrics regulating the voltage level for a fix-error-rate operation [8].

- Data compression is an alternative way to minimize the transmission/storage energy and the memory footprint. Due to the disproportional energy cost communicating data with the off-chip storage, GPU texture caches and video decoder engines use compression to minimize bus utilization [9][2]. Form the de facto S3TC/DXTC to mobile specific iPACKMAN/ETC2, the truncation codes compress channels according to their relative importance for higher peak signal to noise ratio (PSNR) [10]. Chroma sampling and Huffman coding are also well known techniques in video/image standards (JPEG/MPEG) [11][13]. In the scope of the embedded framebuffers and caches, spatial subsampling and truncation coding are preferable over JPEG and similar algorithms because they do not require complex look up operations [12].
- Error masking/correction in video memory not only helps in power reduction during voltage scaling, it may also improve yield through error masking [14]. Kurdahi et al. proposed a defect-resilient filter targeting embedded video RAMs [6]. The technique masks faulty cells through spatially sampling neighboring cells, which relies on a known-bad-cell map operating at the selective error location. Shafique et al. designed a power-efficient error resilient H.264/AVC system [15]. The technique protects the variable length code table because it is crucial in multimedia decoding. The error resilient table may be power-gated based on partition for leakage saving.

The contributions of this work, in relation to the prior works described earlier, are discussed below. First, to the best of our knowledge, this is first dynamic voltage scaling technique that specifically considers the perceptual sensitivity between the color and luminance channels for low power optimization. Second, we apply the channel information redundancy, not for compression, but rather to enable quality-preserving voltage scaling in the framebuffer. Our goal is to minimize the on-chip pixel access power which deviates from the video/texture compression techniques that minimize the transmission energy on the bus. Nevertheless, we evaluated the effectiveness of our technique under compression and observed additional improvement on top of a compressed framebuffer. Third, we show how applying simple filtering may provide effective error masking when channels are stored separately, which may also be applied as a manufacturing defect-masking technique. Note that unlike [6], the proposed channel filtering does not require known-bad-cell pixel map as explained later in Section III, thereby simplifying the design. The trade-offs associated with the operating voltage in SRAM, the associated bit-error rates, and the corresponding power overhead of the filter were analyzed.

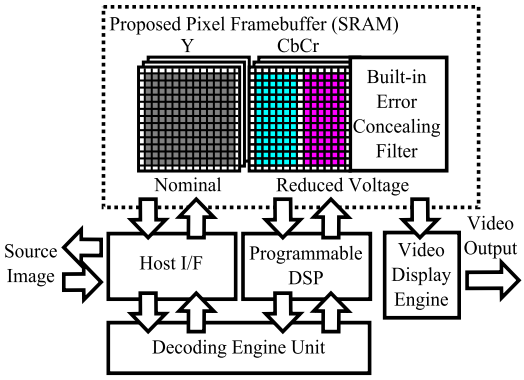

Figure 2. High level pixel framebuffer integrated flow. Dotted region represents our proposed modification. The error concealing filter may be implemented with a simple 1-D sliding window median filter.

III. SRAM FRAME-BUFFER WITH CHANNEL PROTECTION

The quality preservation methodology in low-voltage SRAM proposed in this work stems from the understanding of the perceptual sensitivity to image property. Conventional image storage treats color image for primary RGB equally, maximizing the richness of the image color. It was chosen to match human vision trichromacy achieving maximum spectra transfer to retina. In additive-color display (e.g. CRT, LCD, plasma display), the trichromatic color system optimized for pixel density and energy efficiency. Leading to the fact that image data storage has been in the RGB format for ease of conversion and processing. However, it is also well known in video encoding that human sensitivity to chrominance (or *chroma*) is less than the sensitivity to luminance (or *luma*) [12]. Chroma subsampling hence plays an important role in many well-known video compression algorithms. The technique separates color into YCbCr channels and has been used extensively in the area of digital video processing to ensure high quality lossy compression (MPEG, DVCAM, MJPEG, etc.). An alternative application to the information redundancy in the chroma channels (CbCr) is to use it to protect against bit-errors. Yet, this property has been largely overlooked in the content-aware memory design for low power applications. It is, hence, extremely worthwhile to investigate its application in energy reduction at the hardware level.

We propose to protect the embedded SRAM with understanding of perceptual relevance of the color and brightness channels. The SRAM bank will be arranged differently from the traditional YCbCr sequence. Color information (CbCr) will be stored separately from the main (luma) bank. The Y component is placed in the bank with nominal voltage while the CbCr components are stored in aggressively voltage scaled banks. Fig. 2 shows an integrated framework with the proposed frame-buffer architecture. Each bank can potentially operate at high or low voltage levels. Since the voltage scaling is applied at the bank level, the proposed architecture is simpler compared to the prior technique that utilizes bit-level protection [5][6]. For example, the low power bank presented by Cho et al. requires separate power supply lines for different columns and additional inverters on the word-line, see Fig. 1a [5]. The proposed architecture does not require such modification and may incorporate SRAM banks generated through a conventional SRAM compiler.

Protecting Y channel associated with the brightness is of crucial importance, and HOB flips in the secondary channels affects overall quality as well. The pixel error must be masked to conceal the effects of HOB errors in the chroma channels. We propose to add additional noise reduction filters at the output of the chroma channels to address this issue. Different from the bit-protection technique that limits the error magnitude, the masking filter may be ineffective removing noise due to the channel property. Fortunately, the color channels exhibit smooth transitions between opposite colors. Quantitatively, the delta differences between neighboring pixels for test sequence "bus" suggest the same in Fig. 3. The histograms for CbCr in decibel (for better visibility in high order magnitude difference) between neighboring pixels are generally within four least significant bits. This property may be used to mask significant bit errors relying on neighboring redundant information. However, this does not apply to the luma channel which the trailing tails sum up to be significant at each end. In our experiment, we evaluated two common post noise reduction filters for the chroma channels: low-pass filter and median filter. These two filters were chosen to represent the most primitive form of a linear filter and a non-linear filter. The filter computes the MSB results and selects the read LSB after the computation, effectively reduces the computation effort and improves power efficiency. The filters were designed to be 1-D moving windows of order eight under the assumption that the video display engine accesses pixel in form of horizontal scan. In a random access scenario (such as decoding engine access), the access should by-pass the filtering to avoid unnecessary energy overhead, because a generic post noise reduction filtering consumes more time than the actual decoding stage [3]. An applicable power optimization technique applies to compute/compare the *n* MSB for filtering. The filter power may be accurately captured through the Matlab driven switching activity for a fair power evaluation. Understanding the filter power is analogous to texture compression efficiency, when the power evaluation moves on-chip, the decoder energy may be within one to two order-of-magnitude per pixel access, a high packing ratio bits per pixel (BPP) for off-chip transmission may not be optimal on-chip.

In summary, we allow data corruption to occur in the chroma channels and then filter out the error. As for luminance channel, the image relies on its brightness transition to define forms. Applying 2D filtering on the luminance channel blurs the image; hence it requires a known-bad-pixel map for correction [6]. Our bank protection technique keeps the luma channels pristine avoiding the problem all together.

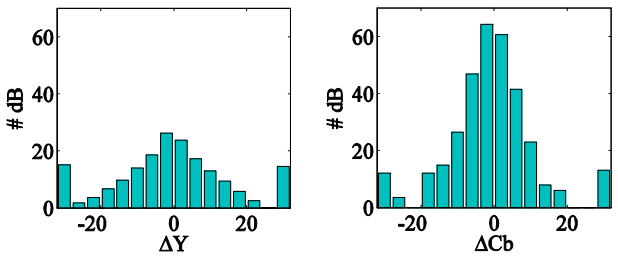

Figure 3. Representative neighboring luma and chroma difference distribution for the benchmark video "bus." The chroma channel has already been subsampled at 4:2:0 rate. Note the histogram y axis is count in decibel (log scale) to emphasize the trailing delta magnitude of Cb and Cr.

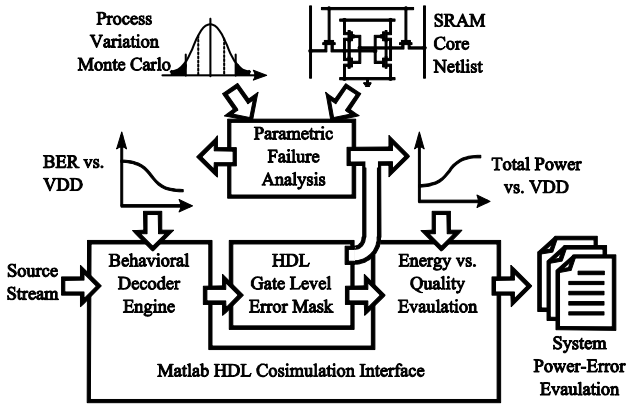

Figure 4. The Matlab-HDL framework for energy to image quality trade-off evaluation.

IV. Simulation Results

The evaluation framework is illustrated in Fig. 4. We first perform parametric error analysis on the target SRAM. We sweep the SRAM bank power addition to the error masking filter to acquire the total power versus supply voltage. This analysis gives us the bit error rate versus supply voltage trade-off. Second, we supply these parametric failures to the Matlab based frame-buffer. The co-simulation interface between Matlab and Hardware Description Language (HDL) drive the HDL filter for accurate power estimation generated during error filtering. Third, the image quality metrics are evaluated against the energy dissipation of the frame-buffer. The Matlab-HDL co-simulation interface glues the Verilog gate level netlist and high level behavioral simulation together. The Matlab wrapper feed consecutive pixels to the gate level filter. Then it collects the filtered vectors for quality evaluation. The input vectors are a set of video test media (akiyo, bus, foreman, and sign_irene) listed on [17]. The sequence "akiyo" and "sign_irene" includes pure red and blue forms and shapes to evaluate chroma specific corruptions. The "bus" sequence contains defined edges, fine shadow and lighting details, textured foliage background, and pure tone logo at the corner. It is suitable to test filter artifacts and detail preserving features. The "foreman" sequence was selected to match the sequence from the prior work [6]. The bit error is injected through bit XOR with the error bit stream before sending to the Verilog HDL filter.

A. SRAM Error-Energy Tradeoff

The parametric analysis uses commercial 130 nm CMOS technology considering traditional 6-T cell parameters. The following metrics are the parametric stabilities under inspection: (1) *Static Noise Margin:* The read static noise margin (SNM) is estimated as the largest square fits within the smaller lobe of voltage transfer characteristic (VTC) curves while the word-line is on. The read SNM is used to predict the read disturb failure. The hold SNM may be estimated similarly while the word-line is off. The hold SNM is used to predict the hold failure. (2) *Read/Write Time:* The read time is measured as the time required developing 100mV voltage differential across the bit-lines. The read time is related to the read failure. Note that the commercial SRAM read sense amp might work

below this given voltage; this metric might be pessimistic. The write time is measured as the time-difference between the wordline tuned on and the cell nodes cross each other. The write time is related to the write failure. The detail Monte-Carlo analysis of the SRAM is performed considering the internal process variation parameters of the commercial 130 nm technology. The voltage and temperature variations assume all cells have same supply noise (10% max Vdd), clock jitter (10% clock frequency), and temperature (125 ^{o}C). Fig. 5a includes the failure rates for the cell's metrics as a function of the supply voltage. The mean read SNM and read time together form the outer couture of the bit error rate at the given supply voltage.

B. *Frame Buffer Specification*

The embedded framebuffer is sized to hold 786kb of pixel information. The size chosen to store few decoded QCIF frames. The buffer is arranged as three 262k SRAM banks using commercial 130 nm SRAM core generator, two of which are capable running at low voltage. Each bank is arranged as 512 × 512 cells for balanced row decoder and column decoder stages. The word size is 32 bits to match one row of 4 × 4 macroblock partition. The SRAM parametric failure simulation is performed using statistical simulations at different supply voltages. The read static noise margin (SNM), standby SNM, and bit line write margin are collected to determine the cell's read disturb, hold failure, and write disturb failures. The read time and write time also contribute to these failures when timing constraints are included. We modeled these factors accordingly. The voltage scaled power profile considers the power from row decoder simulation, column decoder simulation, and column precharge/sense simulation. The power estimation was performed using HSPICE based circuit simulation. The combined read disturb error was used for the error-energy tradeoff analysis for the framebuffer as shown in Fig. 5b It is important to note that all the banks are operating at the same frequency (~100MHz). This implies a higher read-time error for low-voltage banks. Also, the bank-wise voltage scaling allows the decoder to operate at the reduced voltage, while in the bit-level protection the decoder operates at the nominal voltage or requires extra level shifter circuit.

C. *Different Frame-buffer Configurations*

1) *Baseline HOB Protection*

HOB protection technique protects the higher order bits from erroneous bit flips. The brightness channel (Y) and color channel (CbCr) are treated equally. We split the bit protection equally as 4 bits, 4 bits, and 4 bits. That leaves 4 bits each channel for voltage scaling. The energy reduction per pixel access is defined as follows:

$$\left(1 - r_{conf}\right) \times \frac{B_{low}}{B_{total}} \times \frac{E_{prot} - E_{low}}{E_{prot}}. \qquad (1)$$

Where r_{conf} is the additional power associated with the decoder at the nominal voltage, B_{low} is the available low voltage bits, B_{total} is the total pixel bits, E_{prot} is the active energy to perform a pair of read and write operation on each protected bit, and E_{low} is the energy to perform the same operation on a low voltage bit.

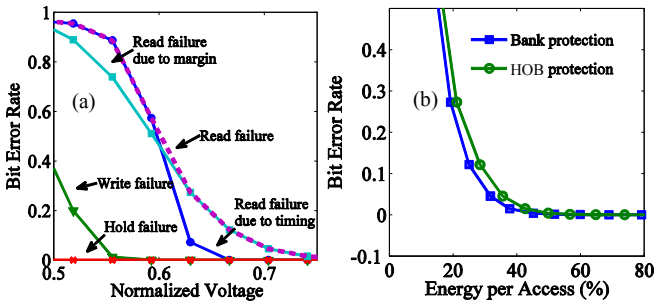

Figure 5. Shows (a) the bit error rate versus normalized voltage. The maximum error curve is bounded by the error probability of the read time and read SNM; (b) is the energy versus bit error rate for the modified and unmodified banks.

2) *HOB Protection with Filtering*

Because our proposed technique utilizes filters to mask erroneous bit error. We compare our technique with HOB protection with an identical filter. This comparison differentiates our technique's effectiveness over generic HOB protection; showing filtering may not be applied arbitrarily to corrupted image to improve quality. The equation is largely the same as Eq (1), but includes extra term for filter energy calculation.

3) *Proposed Channel Protection with Filtering*

The channel protection technique protects the Y channel bank and voltage scale the Cb and Cr color banks. This provides 8 bits protection within the same bank and 16 bits in the voltage scaled banks. The energy saving is computed as:

$$\frac{B_{low}}{B_{total}} \times \frac{E_{prot} - E_{low}}{E_{prot}} - \frac{n_{filter} \times E_{filter}}{E_{prot}}. \qquad (2)$$

Note that B_{low} in channel protection has 4 more bits than the HOB protection due to our technique groups pixels channels in the same memory bank for overhead reduction.

D. *Simulation Results*

We evaluated the power to quality trade-off for energy efficient filtering through introducing additional noise in the decoding flow. Fig 6a shows the correlation between the median filter power and the additional bit errors in a noisy video sequence. We applied this additonal energy in our system to benchmark our system quality versus the baseline system quality through the structure disimilarity (DSSIM) metric [16]. In the origional SSIM metric, the author suggested to use desaturate image for quality assessment. We reported both desaturated DSSIM as well as the CbCr combined DSSIM to evaulate the color corectness of the image. In the experiment, 0.1 level of DSSIM was considered as the acceptable threshold for visual quality evaulation, see Fig. 7. The median filtering worked extremely well for the overall bank-wise protected YCbCr. Applying similar filtering in three channels in the bit-wise protection showed significant reduction in quality due to the smoothing in the Y channel (Fig 6b). Applying known-bad-cell filter to HOB protected memory also did not improve the quality beyond the unfiltered bit-wise protection. We therefore compares our technique with the unfiltered HOB protection

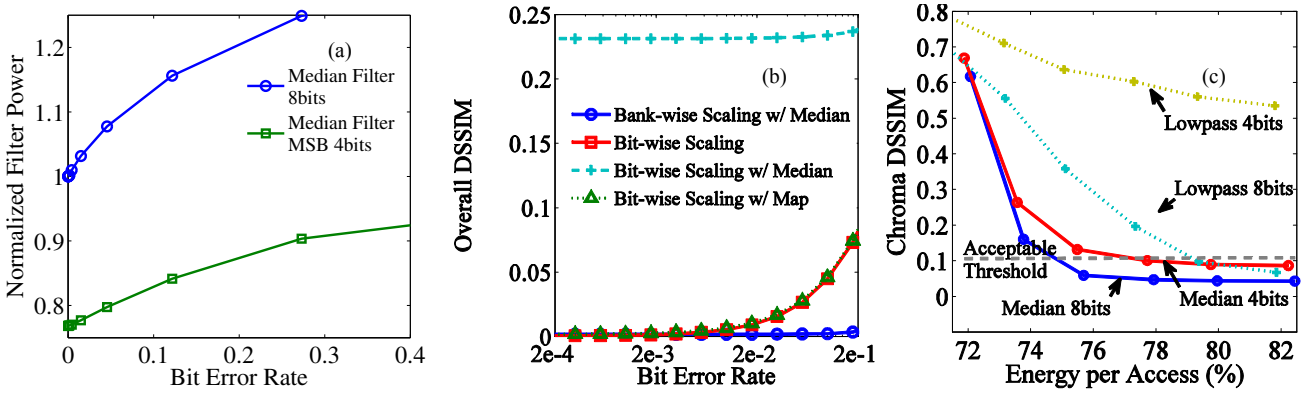

Figure 6. Filter related figures showing (a) normalized filter power versus input image sequences' bit error rate, (b) DSSIM metric after filtering, and (c) quality versus energy for various filters. In (b) the quality of the HOB protection with the known-bad-cell map does not improve beyond unfiltered image sequence.

technique. Note that the unfiltered HOB technique (the baseline) incurs no power overhead for the error masking filter. However, the proposed method requires additional filter power which was included in our results.

We compared the DSSIM of each filter for filter effectiveness in CbCr channels. The basic moving window lowpass filter and median filter were evaulated in our analysis. The corresponding error rate versus power is reported in Fig. 6c. The energy reduction in both filters improved with more bits processed due to the correction quality inprovement was higher than the filtering cost. The lowpass filter used less power and the quality decayed at a slower rate than the median filter but was not effective to protect against the higher-order-bit (salt and pepper) errors. Considering the filter power versus quality, median filter has better power at the acceptable video quality level. The median filter, which was more effective in error reduction versus energy tradeoff, was compared against the HOB protection technique. The benchmarking result is captured figures in Fig 8 . All the power savings are normalized according to the uncompressed video banks. The bank received ~10 % additional power reduction compared to HOB protection. Our technique showed ~45 % reduction and the HOB technique showed ~35 % reduction at 0.1 level. Note that beyond 35 % power reduction, our proposed technique produced better quality than the HOB technique in all channels.

E. Compression Assisted Power Reduction

Low power video decoders harness the redundancy within the data stream for error masking [6]. However, when compression is applied to the data stream, the redundancy reduced significantly. Many existing video framebuffers or caches are capable of compressing and decompressing data on the fly [8] . We must evaluate the robustness of our proposed technique within compressed video RAMs. We consider the following well-known compression techniques for evaluation:

Figure 7. Barely acceptable sequence quality for (a) bank protection after filtering (b) HOB protection. The filter artifact may be observed at the edges for bank protection. Snowy background may be observed in HOB protection due to lower bit corruptions.

1) Chroma Subsampling: The chroma subsampling technique is common in video decoders [11][13]. Sampling the color pixels every other row and column for each luma sample removes considerable amount of the spatial redundancy in the chroma channels and halve the required storage area. Since our technique utilizes similar property in the chroma channel, we should evaluate the additional gain from our technique on top of a compressed video RAM. We evaluated standard 4:4:4 uncompressed video sampling and 4:2:0 subsampled chroma stream (sampled every other row and column).

2) Truncation Coding: Truncation code is common in texture cache for graphic memories [9]. It reduces the storage bit width simply by discarding the lower order bits. We evaluated the full 888 bit stream and 555 bit compressed stream.

The compressed bank results are presented in Fig 9. We observed the proposed channel protection scheme required ~15 % higher energy than HOB protection schemes to achieve the same quality. This is a direct consquence of reducing spatial redundancy in the chroma channels. In the truncation coding bank, the channel protection allowed energy to scale down over ~20 % more than the HOB protection. This is because the HOB protection technique reduces efficiency due to removal of the protected lower order bits. For the final compressed framebuffer with both chroma sub-sampling and trucncation, the proposed approach allowed more energy reduction (~5 % at 0.1 DSSIM) than the HOB protection. Fig. 9d provides an overview of the overall power reduction normalized to the uncompressed bank. The compression provided a singificant power by itself and there was an additional ~5 % power improvement using the channel protection and a median error masking filter. This out perform the HOB protection by roughtly 2 %.

V. Conclusion

We have presented a low power video frambuffer, by separately storing Y and CbCr channels in SRAM banks with nominal and lower voltages. Further, the median filter is used to mask errors in the color (CbCr) channels. The architecture provides better quality-energy trade-off than the traditional higher-order-bit protection. The proposed framebuffer was evaluated for both uncompressed and compressed image storages. The channel protection provides significant advantage over bit protection for uncompressed buffers. The technique remains competitive when applied to compressed buffers providing additional power savings. The bank-wise voltage scaling does not need customized SRAM cells and peripheral circuits, and may take advantage of the conventional SRAM core generators. Therefore, the proposed memory shows potential in low-voltage low-power multimedia applications.

Acknowledgment

This material is based on the work supported by National Science Foundation (CCF# 0916083)

References

[1] T.M. Liu, et al., "A 125 µW , Fully Scalable MPEG-2 and H.264/AVC Video Decoder for Mobile Applications," Solid-State Circuits, IEEE Journal of , vol.42, no.1, pp.161-169, Jan. 2007.

[2] P. Tseng, et al., "Advances in Hardware Architectures for Image and Video Coding - A Survey," Proceedings of the IEEE , vol.93, no.1, pp.184-197, Jan. 2005.

[3] M. Ohashi, et al., "A 27 MHz 11.1 mW MPEG-4 video decoder LSI for mobile application ," Solid-State Circuits Conference, 2002. Digest of Technical Papers. ISSCC. 2002 IEEE International , vol.1, no., pp.366-474 vol.1, 2002.

[4] T.M. Liu, et al., "An 865-µW H.264/AVC Video Decoder for Mobile Applications," Asian Solid-State Circuits Conference, 2005 , vol., no., pp.301-304, 1-3 Nov. 2005.

[5] M. Cho, et al., "Reconfigurable SRAM Architecture With Spatial Voltage Scaling for Low Power Mobile Multimedia Applications," Very Large Scale Integration (VLSI) Systems, IEEE Transactions on , vol.19, no.1, pp.161-165, Jan. 2011.

[6] F.J. Kurdahi, et al., "Low-power multimedia system design by aggressive voltage scaling." IEEE Trans. Very Large Scale Integr. Syst. 18, 5 (May 2010), 852-856.

[7] K. Yi, et al., "A partial memory protection scheme for higher effective yield of embedded memory for video data," Computer Systems Architecture Conference, 2008. ACSAC 2008. 13th Asia-Pacific , vol., no., pp.1-6, 4-6 Aug. 2008.

[8] P.K. Krause and I. Polian, "Adaptive voltage over-scaling for resilient applications," Design, Automation & Test in Europe Conference & Exhibition (DATE), 2011 , vol., no., pp.1-6, 14-18 March 2011.

[9] M. Doggett, "Texture Caches," Micro, IEEE , vol.32, no.3, pp.136-141, May-June 2012.

[10] [J. Ström and T. Akenine-Mollör, "PACKMAN: texture compression for mobile phones." In ACM SIGGRAPH 2004 Sketches (SIGGRAPH '04), Ronen Barzel (Ed.). ACM, New York, NY, USA.

[11] A. Skodras, et al., "The JPEG 2000 still image compression standard," Signal Processing Magazine, IEEE , vol.18, no.5, pp.36-58, Sep 2001.

[12] J. Munkberg, et al. "High dynamic range texture compression for graphics hardware." In ACM SIGGRAPH 2006 Papers (SIGGRAPH '06). ACM, New York, NY, USA, 698-706.

[13] T. Sikora, "MPEG digital video-coding standards," Signal Processing Magazine, IEEE , vol.14, no.5, pp.82-100, Sep 1997.

[14] M.A. Breuer, et al., "Defect and error tolerance in the presence of massive numbers of defects," Design & Test of Computers, IEEE , vol.21, no.3, pp. 216- 227, May-June 2004.

[15] M. Shafique, et al., "Power-efficient error-resiliency for H.264/AVC Context-Adaptive Variable Length Coding," Design, Automation & Test in Europe Conference & Exhibition (DATE), 2012 , vol., no., pp.697-702, 12-16 March 2012A.

[16] Z. Wang, et al., "Image quality assessment: From error visibility to structural similarity," IEEE Transactions on Image Processing, vol. 13, no. 4, pp. 600-612, Apr. 2004.

[17] "Xiph.org Video Test Media. " Internet: http://media.xiph.org/video/derf/, Sep. 5, 2000.

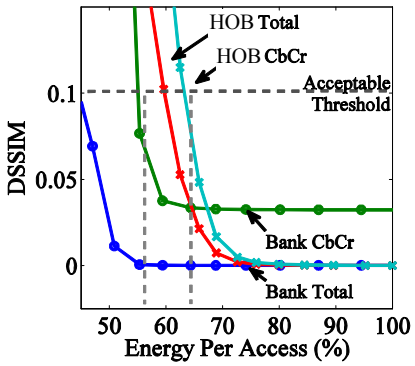

Figure 8. Combined DSSIM and chroma only DSSIM with no chroma subsampling and no truncation coding.

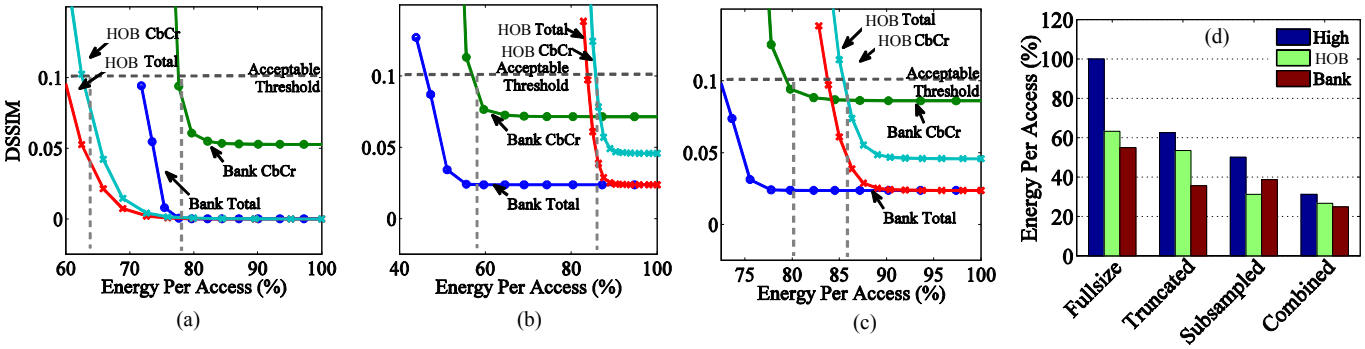

Figure 9. Combined DSSIM and chroma only DSSIM in (a) chroma subsampled every other row and column and no truncating code (b) no chroma subsampling and least 3 bits truncated (c) combine both chroma subsampling and truncating code. (d) shows the overall improvement after voltage scaling and compression.