Catch Your Attention: Quality-retaining Power Saving on Mobile OLED Displays

Chun-Han Lin, Chih-Kai Kang, and Pi-Cheng Hsiu
Research Center for Information Technology Innovation, Academia Sinica, Taipei 115, Taiwan, R.O.C.
{chunhan, akaikang, pchsiu}@citi.sinica.edu.tw

ABSTRACT

Organic light-emitting diode (OLED) technology is considered as a promising alternative to mobile displays. This paper explores how to reduce the OLED power consumption by exploiting visual attention. First, we model the problem of OLED image scaling optimization, with the objective of minimizing the power required to display an image without adversely impacting the user's visual experience. Then, we propose an algorithm to solve the fundamental problem, and prove its optimality even without the accurate power model. Finally, based on the algorithm, we consider implementation issues and realize two application scenarios on a commercial OLED mobile tablet. The results of experiments conducted on the tablet with real images demonstrate that the proposed methodology can achieve significant power savings while retaining the visual quality.

Categories and Subject Descriptors

C.3 [**Special-Purpose and Application-Based Systems**]: Real-time and embedded systems

General Terms

Algorithms, Design, Experimentation, Human Factors

Keywords

OLED displays, power saving, visual attention, mobile devices

1. INTRODUCTION

Organic light-emitting diode (OLED) technology is deemed a promising display alternative for the emerging genre of mobile devices as it provides brighter colors, wider viewing angles, and faster response times, as well as allows fabrication on flexible plastic substrates, compared with conventional *liquid crystal display* (LCD) technology [14]. Unlike LCD displays whose power consumption is dominated by the external lighting, an OLED display is self-emissive and its power consumption is highly dependent on the image content (more precisely, the pixel values). While OLED displays consume nearly zero power when presenting a black image, they could consume power more than twice that consumed by LCD displays when presenting a white image [7]. As the OLED power consumption increases dramatically with the pixel values of the displayed image, such a property brings new opportunities and challenges to the optimization of display power consumption.

Displays account for a large proportion of the total power consumption in mobile devices [8]. Many approaches based on *backlight scaling*, some of which incorporate *image compensation*, have been proposed to dim the LCD backlight while limiting the distortion and/or maintain the fidelity of the presented image [9, 13]. However, they are not suitable seamlessly for OLED displays that obviate external lighting due to their emissive nature. In recent years, researchers and vendors have explored various low-power techniques for OLED displays [3, 6, 14]. An intuitive way is to darken the contents that are not of interest to the user [3], referred to as *partial display disabling/dimming*. The challenge thus lies in identifying those unimportant contents [15]. Motivated by an observation that different colors require drastically different power, the *color remapping* technique changes the original colors into colors that consume less power [6]. This technique is applicable to graphics user interfaces and applications not dealing with natural images [7]. To retain the original visual experience, another technique called *OLED dynamic voltage scaling* was developed [14]. It tries to decrease the supply voltage of each pixel's circuit to the minimum sufficient to sustain the pixel's (luminance) value. The technique requires hardware support and extra costs; thus, a display is normally partitioned into a limited number of rectangular regions [5].

In this paper, we present an alternative technique, called *image pixel scaling*, which leverages the flexibility provided by OLED technology to scale down the pixel values in regions of any shapes. With the new technique, we introduce *visual attention* into the quality-retaining power-saving design on mobile OLED displays. The rationale behind the notion is based on some interesting findings of psychological experiments [12]: 1) regions in an image receive varying degrees of attention, and thus 2) they can tolerate different degrees of image distortion. In addition to the above notion, the contributions of this study are as follows.

First, we model the fundamental problem of carrying out the notion as an *OLED image scaling optimization problem*. The objective is to minimize the power required to display a given image without adversely impacting the user's visual experience. Then, we propose an efficient algorithm to solve the fundamental problem. Importantly, we prove that the algorithm is optimal in terms of power savings even if the OLED display's accurate power model is unknown, in the sense that the optimal solution derived is applicable to any OLED devices. Next, we discuss some technical issues that arise when implementing the proposed algorithm on commercial mobile devices. We have also realized two application scenarios, namely an image converter and a power-saving mode, on the Samsung Galaxy Tab 7.7. Finally, to obtain further insights, we conduct extensive experiments on the Samsung tablet with some images of different characteristics, and compare the proposed methodology with a gird-based approach revised based on that proposed in [5]. The results demonstrate that the proposed methodology can achieve substantial OLED power savings between 38% and 42%, depending on the image's characteristics, while retaining satisfied visual quality.

The remainder of this paper is organized as follows. Section 2 describes the system model and formulates the problem. In

DAC'14, June 01-05, 2014, San Francisco, CA, USA.

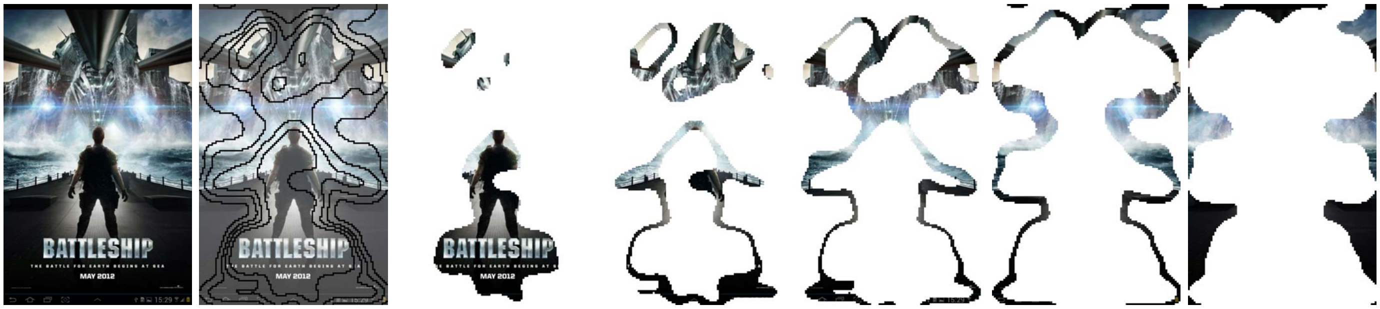

(a) Original (b) Segmented (c) Attention Regions

Figure 2: Five attention regions of an image segmented based on visual attention

Section 3, we propose an optimal algorithm to solve the problem and discuss implementation issues. The experiment results are reported in Section 4. Section 5 reviews related work, and Section 6 contains some concluding remarks.

2. RATIONALE AND PROBLEM DEFINITION

This section presents the rationale behind the image pixel scaling technique and how we exploit visual attention to reduce OLED power consumption while retaining image quality. Then, the OLED image scaling optimization problem is formulated accordingly.

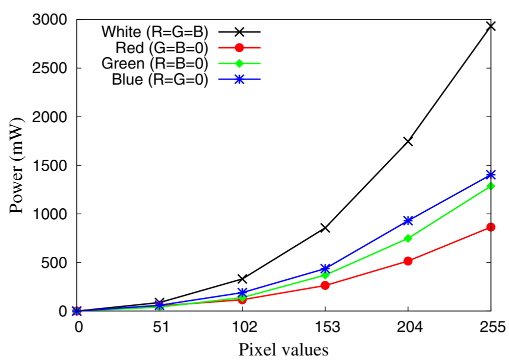

Figure 1: An OLED power model

An OLED display comprises a matrix of pixels, each of which consists of three subpixels emitting red, green, and blue light respectively. Unlike an LCD display that relies on a light source to illuminate from behind, each subpixel on the OLED display is self-emissive and emits light individually in response to an electric current that corresponds to the given pixel value [8]. Thus, the amount of power that a subpixel consumes varies significantly depending on its value ranging from 0 to 255. As modeled in [5, 7, 8], the power consumption of an image on an OLED display is the sum of each individual pixel's power consumption, which is the sum of the power consumed by its three subpixels. Figure 1 shows the power model measured from the Samsung Galaxy Tab 7.7 in practice, where the power increases dramatically with the pixel value. Although the power model, denoted as $P(x)$, of any OLED display is a strictly increasing function of the pixel value x, the increasing slopes may be different between OLED displays. In this paper, the proposed algorithm is proved to be optimal in terms of power savings even without the accurate power model measured in advance.

Lowering the pixel values of an image is an effective way to reduce its power consumption on an OLED display; however, this can also have an adverse impact on visual experience if the pixel values are scaled down inappropriately. Note that any change in a pixel value will lead to a change to the *chrominance* and/or the *brightness* of the pixel, but not every change is noticeable to the *human visual system* [12]. As the same magnitude of change is much more noticeable in chrominance than in brightness, to maintain the chrominance of a pixel when its value is scaled down, its red, green, and blue subpixels must be applied with the same scaling ratio. In contrast, a single scaling ratio should not be applied to all the pixels so as to save more power while retaining perceptual visual quality to a tolerable extent. This is because scaling down the values of different pixels can have varying degrees of impact on *image distortion*, which is normally defined as the resemblance between the original image and the pixel-scaled image [16].

Psychological experiments have demonstrated that not every pixel or region in an image receives the same attention level [12]. This is due to the human visual system's selectivity to first respond to the most attractive parts in a visual scene, with subsequent eye movements from one fixation location to another, the so-called *focus of attention*. In computer vision, many efforts have been made in modeling the mechanism of human attention [4]. In particular, the saliency map indicates a saliency value for each pixel or region in an image [10]. Thus, an image can be segmented based on its saliency map into a set of N disjoint attention regions, $R = \{r_1, r_2, r_3, ..., r_N\}$, associated with an $N \times N$ adjacency matrix A, where $A[i, j]$ is 1 if r_i and r_j are adjacent regions, and 0 otherwise. Note that an attention region can comprise several non-adjacent subregions. Figure 2 shows an image and its five attention regions segmented according to its saliency map in practice.

As regions in an image receive varying degrees of attention, they can tolerate different degrees of image distortion. A number of *visual quality metrics* intended for the human visual system have been proposed to quantify image distortion [12]. In other words, attention regions should be given tolerable distortion in inverse proportion to their attention levels. Thus, each region $r_i \in R$ is associated with *a critical scaling ratio*, $c(i)$, representing the lowest ratio that will not violate the tolerable distortion (quantified by a certain visual quality metric). This *distortion constraint* limits the lowest scaling ratio of each attention region. However, lowering the pixel values by simply applying the critical scaling ratio to each region may result in sharp edges between adjacent regions, and these sharp edges will severely interfere with visual experience [5]. Therefore, the difference between the scaling ratios applied to two adjacent regions should be limited, so that the region boundaries are too indistinct to be discerned by the human eye. In human visual systems, the *just noticeable difference* is the minimum amount by which the stimulus intensity must be changed in order to produce a noticeable variation in sensory experience [13]. Weber's Law states that the just noticeable difference between two stimuli is proportional to the magnitude of the stimuli [1]. Based on the law, we define the *differential constraint* that the difference between the scaling ratios applied to two adjacent regions is not greater than either ratio multiplied by a differential constant d.

Accordingly, our objective is to determine an appropriate scaling ratio for each region such that the power required for a given image on any OLED display is minimized. The determination of a scaling ratio $\sigma(i)$, in the range 0 to 1, for every region $r_i \in R$ is called a *scaling assignment*. A scaling assignment is *feasible* if both the distortion constraint, i.e., $\sigma(i) \geq c(i)$, and the differential constraint, i.e., $|\sigma(i) - \sigma(j)| \leq d \times \sigma(j)$ if $A[i,j] = 1$, are satisfied, $\forall r_i, r_j \in R$. Next, we formally define the OLED image scaling optimization problem as follows.

Instance: A set of attention regions $R = \{r_1, r_2, r_3, ..., r_N\}$ and its adjacency matrix A, where each region $r_i \in R$ is associated with a critical scaling ratio $c(i)$; and a differential constant d for any two adjacent regions.

Objective: A feasible scaling assignment σ such that the total power consumption, $\sum_{r_i \in R} \sum_{x_k \in r_i} P(\lceil x_k \sigma(i) \rceil)$, is minimized, where $P(x)$ is any OLED power model.

3. OLED IMAGE SCALING OPTIMIZATION

In this section, we present an optimal algorithm to solve the OLED image scaling optimization problem. We also discuss some implementation issues, especially image segmentation and distortion estimation, which arise when applying the algorithm to two application scenarios for mobile OLED displays.

3.1 Algorithm Description

Given an image represented by a region set R with an adjacency matrix A, as well as critical scaling ratios $c()$ and a differential constant d, Algorithm 1 determines a feasible scaling assignment σ such that the power required for the image on any OLED display is minimized. To begin with, each region is initially assigned with its critical scaling ratio (Line 1). Throughout the algorithm, we maintain a priority queue Q that initially contains all the regions in R, keyed by their current scaling ratios stored in an array σ (Line 2). We repeatedly remove from Q a region whose key is maximum until Q is empty (Lines 3-4). Whenever a region, say r_i, is extracted, its scaling ratio is determined and will never change. Then, all the remaining regions in Q are examined whether their scaling ratios should be updated so as to satisfy the differential constraint (Line 5). Let r_j be any region that remains in Q and is adjacent to r_i (Line 5). Since the difference between $\sigma[i]$ and $\sigma[j]$ cannot be greater than either ratio (or, equivalently, the smaller ratio $\sigma[j]$) multiplied by the differential constant d, region r_j should be assigned a scaling ratio at least $\frac{\sigma[i]}{1+d}$. The current scaling ratio assigned to r_j should be updated if $\sigma[j]$ is smaller than $\frac{\sigma[i]}{1+d}$ (Line 7). At the end, the scaling assignment σ is returned (Line 8).

Algorithm 1

Input: A region set R with an adjacency matrix A, as well as crtical sacling ratios $c()$ and a differential constant d
Output: A feasible assignment σ
1: $\sigma \leftarrow c$;
2: $Q \leftarrow R$;
3: **while** $Q \neq \emptyset$ **do**
4:     $r_i \leftarrow$ remove from Q a region whose key is maximum
5:     **for all** $r_j \in Q$ **do**
6:         **if** $A[i,j] = 1$ **then**
7:             $\sigma[j] \leftarrow \max(\sigma[j], \frac{\sigma[i]}{1+d})$
8: **return** σ

3.2 Properties

Next, we analyze the time complexity of Algorithm 1 and prove that it solves the OLED image scaling optimization problem optimally without the information about the power model.

LEMMA 1. *The time complexity of Algorithm 1 is* $O(N^2 \ln N)$.

Proof. This analysis assumes that the priority queue Q is implemented by a binary heap. The initialization of array σ and the building of heap Q based on R in Lines 1-2 can be done in $O(N)$ time. The **while** loop in Lines 3-7 is exactly executed N times. Within the loop, removing a region r_i from Q in Line 4 costs time $O(\ln N)$. Because r_i has at most $N-1$ adjacent regions and each update to σ (if necessary) in Line 7 takes time $O(\ln N)$, each loop contributes $O(N \ln N)$ to the running time. Note that Line 7 involves an implicit operation on the heap to maintain the heap property. Thus, the time complexity of Algorithm 1 is $O(N^2 \ln N)$. □

THEOREM 1. *Algorithm 1 is an optimal algorithm.*

Proof. Each region's scaling ratio is initialized as its critical scaling ratio and will never decrease, so σ returned by Algorithm 1 satisfies the distortion constraint. Moreover, the scaling ratios of two adjacent regions are subjected to the differential constraint, because a region's scaling ratio can be updated only when any adjacent region r_i is removed from Q, and the updated ratio is not lower than $\frac{\sigma[i]}{1+d}$ and will never decrease thereafter. Thus, σ is a feasible scaling assignment.

Next, we prove that the power consumption under σ is minimum by contradiction. Suppose there exists a feasible assignment σ' with less power consumption. Since $P()$ is a strictly increasing function, there must be at least one region whose scaling ratio in σ' is lower than that in σ. Let r_j be the first region with $\sigma'[j] < \sigma[j]$ during the execution of Algorithm 1. At the time immediately before r_j is removed from Q, we should have $\sigma'[j] < \sigma[j]$ and $\sigma'[i] \geq \sigma[i], \forall r_i \notin Q$. We delineate two possible cases, depending on whether any adjacent regions of r_j have been removed from Q.

1. If no adjacent region of r_j has been removed from Q, its scaling ratio must have not been updated yet. Thus, $\sigma'[j] < \sigma[j] = c(j)$.

2. Otherwise, some adjacent regions of r_j have been removed from Q. Let r_i be an adjacent region whose scaling ratio is the maximum among the ratios of all the adjacent regions that have been removed from Q. Then, $\sigma'[j] < \sigma[j] = \max(c(j), \frac{\sigma[i]}{1+d}) \leq \max(c(j), \frac{\sigma'[i]}{1+d})$. In other words, $\sigma'[j] < c(j)$ or $\sigma'[i] - \sigma'[j] > d \times \sigma'[j]$.

To conclude, $\sigma'[j]$ violates either the distortion constraint or the differential constraint, which contradicts the assumption that σ' is a feasible assignment. Thus, the theorem follows. □

3.3 Implementation Issues

Next, we explain how to segment an image into a set R of attention regions with an adjacent matrix A, as well as how to derive the critical scaling ratios $c()$, as our algorithm's inputs. Not every pixel in an image receives the same attention level. Visual saliency is the distinct subjective perceptual quality which makes some pixels in an image stand out from their surroundings and immediately grab human attention. Many visual attention models have been proposed [4]. In our implementation, we adopt *Itti's model* [10], which was specially designed for still images. Given an image, Itti computes a saliency map which quantitatively evaluates each pixel' conspicuousness in the range 0 to 1, with 1 representing the highest attention level. We divide the range evenly into N subranges, and classify all the pixels into N corresponding attention regions according to their saliency values. Thus, each region's shape is formed by the pixels with saliency values in the same subrange. In addition, all the entries of the adjacency matrix A are initialized as 0, and the entry $A[i,j]$ is updated to 1 if two adjacent pixels in the image belong to different regions r_i and r_j. To this end, for each pixel, we check the two respective pixels right to and below the pixel by scanning the image from left to right and top to bottom. Note that R and A can be constructed simultaneously by scanning the saliency map once sequentially. Thus, the time required to segment an image based on its saliency map is linear to the number of pixels.

The regions in R receive varying attention levels and thus can tolerate different degrees of image distortion. A number of visual quality metrics have been proposed to estimate the distortion of an image [12]. Our implementation utilizes the *structural similarity* (SSIM) index [16], a metric specially designed to comply with the perception of the human eye and widely used in related studies [5, 9]. The resultant SSIM index is a decimal value between -1 and 1, where 1 is only achievable in the case of two identical images. Let the SSIM index required for the region with the lowest attention level be set at s. We assign each of the N regions an index requirement uniformly distributed over the range s to 1 based on their attention levels, i.e., the SSIM index required for the region with the ith highest attention level is set at $1 - \frac{i(1-s)}{N}$. Then, for each region $r_i \in R$, we find its critical scaling ratio $c(i)$, which represents the lowest scaling ratio satisfying the SSIM index requirement assigned to r_i, by a binary search on all the 256 possible scaling ratios. It implies that eight SSIM indexes need to be computed to derive a region's critical scaling ratio. Moreover, to preserve the structural information in the original image, our implementation derives a region's critical scaling ratio by applying SSIM to the minimum rectangle that covers all the pixels of the region.

Now, we delineate two possible application scenarios. First, the proposed methodology (along with the underlying algorithm) can be implemented as an image converter (or integrated into a photo editor) to generate the power-saving versions of images for mobile OLED displays. Note that an image is only converted once (which can be done on a personal computer), and the power-saving version can be displayed on any OLED mobile devices. In the second scenario, a mobile device can implement a screen power-saving mode using the methodology. Specifically, each scene to be displayed on the screen is written into a frame buffer in the kernel space. When the device is switched to the power-saving mode, the scene stored in the frame buffer is deemed an image and analyzed by our methodology to determine its scaling assignment. Then, the scene in the buffer is updated based on the scaling assignment so that its power-saving version is displayed on the screen instead. We have implemented both scenarios on the Samsung Galaxy Tab 7.7, equipped with a 1.4 GHz dual-core processor and a 1280×800 AMOLED display. A major difference between the two scenarios is that the computational overhead should be justifiable in the second scenario. Because both Itti and SSIM involve pixel-level image processing, our methodology requires a few seconds to convert a 1280×800 image on the Samsung tablet. This is unacceptable for an online scenario that requires timely responses. Thus, in the second scenario, *Lanczos resampling* [11] is used to scale down the resolution of a scene to 320×200; then, each grid of 16 pixels in the scene is applied with the same scaling ratio determined based on the scaled resolution. As a result, a scene can be processed within a few hundred milliseconds and replaced with its power-saving version when it remains unchanged for one second, at the cost of increased OLED power consumption compared to that consumed in the first scenario. The proposed methodology could be further accelerated by exempting some routine computations from the CPU to the GPU on the graphics card.

4. PERFORMANCE EVALUATION

4.1 Experiment Setup

To evaluate the efficacy of the proposed methodology, we conducted extensive experiments on the Samsung Galaxy Tab 7.7 with four snapshot images of popular mobile apps, as shown in Figure 7. The four images, namely, Weather, Twitter, Facebook, and CNN, have different characteristics in terms of the average *luminance*[1] and the average saliency value. The respective luminance of Weather and Twitter (i.e., 120 and 112) is higher than that of Facebook and CNN (i.e., 87 and 79). The respective saliency values of Twitter and CNN (i.e., 0.31 and 0.29) are larger than those of Weather and Facebook (i.e., 0.21 and 0.17). Note that a larger luminance value implies higher power consumption on OLED displays, while a larger saliency value implies more pixels that attract user attention. Based on our measurements, Weather, Twitter, Facebook, and CNN require amounts of power 990, 805, 591, and 474 mW respectively; moreover, the pixels whose saliency values above 0.5 account for 11%, 18%, 9%, and 15% of all the pixels in each image.

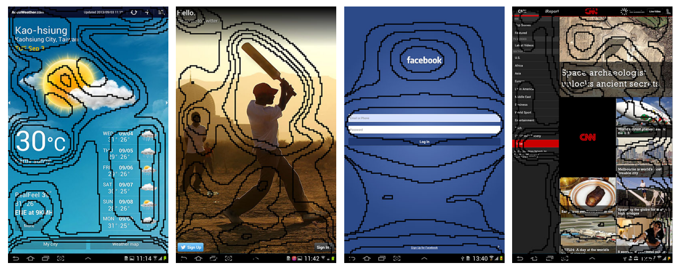

(a) Weather (b) Twitter (c) Facebook (d) CNN

Figure 3: Investigated images with attention regions

We investigated the impact of each of the three parameters on the algorithm's efficacy: 1) the impact of the SSIM index in the range 1 to 0.88, 2) the impact of the number of regions, N, in the range 1 to 9, and 3) the impact of the differential constant, d, in the range 0% to 4%. The default values of the SSIM index, the number of regions N, and the differential constant d were set respectively at 0.94, 5, and 2%, unless otherwise stated. The settings were based on the following observations. The power consumption was saturated when an image was segmented into five or more regions; thus, we set N at 5 to allow a trade-off between the computational overhead and visual quality. The SSIM index was set at 0.94 so that all regions still maintained acceptable quality [2]. Moreover, it was difficult to discern the region boundaries when d was not larger than 2%.

In addition to the above studies, we compared the proposed methodology, denoted as CURA, with a grid-based approach [5], denoted as GRID. In GRID, an image is partitioned into multiple rectangles of the same size. Then, each rectangular region is given an initial threshold between 0 and 255 based on its pixel values and a predefined *sacrificed luminance ratio*. Any pixel value that exceeds the threshold is simply truncated. Next, GRID computes the SSIM indexes of every rectangular region and the whole image. If the image's SSIM index is higher (resp. lower) than the required index, the threshold of the region with the highest (resp. lowest) index is decreased (resp. increased) by 1. The process is repeated until the image's SSIM index satisfies the required index. Following [5], the sacrificed luminance ratio was set at 0.1. For a fair comparison, the number of rectangles was set at 5, and the required SSIM index was set at 0.94. Note that GRID does not defines a differential constant to smooth the sharp edges between adjacent rectangles. We measured the execution time (seconds) that includes everything required to process an image, as well as the power (watts) required to display the processed image, on the Samsung tablet.

[1]A pixel's luminance can be derived by converting its RGB values to its grayscale representation between 0 and 255.

4.2 Experiment Results

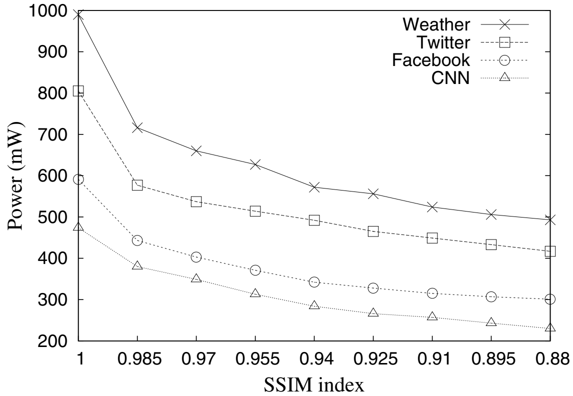

Figure 4: The impact of SSIM

Figure 4 shows the impact of the SSIM index on the power consumption of each image processed by CURA. As expected, the power consumption decreases as the SSIM index decreases. The reason is that a lower SSIM index leads to lower critical scaling ratios. This in turn implies lower visual quality but less power consumption. We notice that the power consumption drops abruptly when the SSIM index decreases from 1 to 0.985. This is because the power consumed by an OLED pixel decreases dramatically as its value decreases, as shown in Figure 1. Thus, OLED image scaling (if applied appropriately) can achieve a significant reduction in power consumption at a cost of a slight degradation in visual quality. In addition, the power consumption is higher when an image's luminance is larger, such as the Weather image whose luminance is the largest among that of the four images, as mentioned previously. The results show that, when SSIM $= 0.94$, CURA can achieve power savings between 38% and 42%, depending on the image's characteristics.

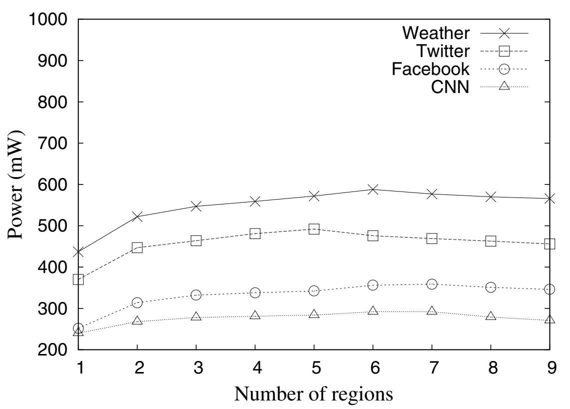

Figure 5: The impacts of the number of regions (N)

Figure 5 shows the impact of the number of regions on the power consumption achieved by CURA. The power consumption increases as the number of regions increases from 1 to 5. This is because the average SSIM index requirement increases with the number of regions, which implies higher critical scaling ratios. Moreover, the more the regions, the denser the adjacent matrix in general. As a result, the scaling ratios determined by CURA are relatively high. However, the power consumption is saturated when $N = 5$, because the increments in the average SSIM index requirement and the adjacent matrix's density are slowing down and become negligible when $N = 5$. Therefore, we set N at 5 in our implementations to trade off the visual quality against the computational overhead. We observe that the power consumption decreases slightly as an image is segmented into 9 regions. The reason for this interesting phenomenon is that the influence of the differential constraint on adjacent regions decreases as the number of regions increases. Consequentially, a lower scaling ratio could be applied to the least-attention region, which usually accounts for a large proportion of an image.

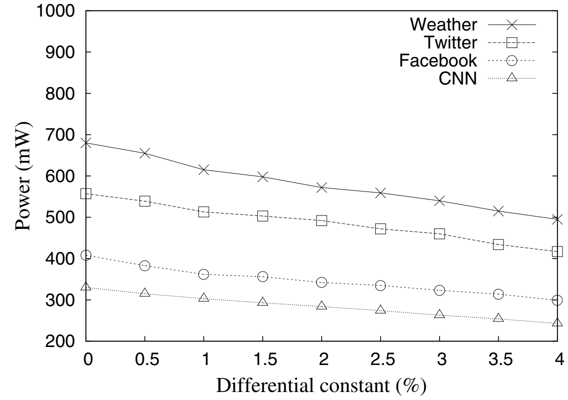

Figure 6: The impacts of the differential constant (d)

Figure 6 shows the impact of the differential constant on the power consumption achieved by CURA. The power consumption decreases as the differential constant increases. This result is as expected because a larger differential constant allows the pixel values of a region to be scaled down by a larger magnitude. Thus, a region's power consumption decreases until the scaling ratio determined by CURA reaches its critical scaling ratio. The characteristic is specially beneficial for images whose average saliency values are low to save power. As mentioned previously, Weather and Facebook contain more pixels with low saliency values than Twitter and CNN. This explains why the decrease in power consumption is more manifest for Weather and Facebook than for Twitter and CNN. Finally, to conclude the impacts of the three parameters based on the above experiments, it appears that the impact is more evident by varying the differential constant than by varying the number of regions, but neither of the impacts is as significant as the impact of varying the SSIM index.

Table 1: Power (mW) achieved by GRID and CURA

Scenario 1	Weather	Twitter	Facebook	CNN
GRID	648	340	237	343
CURA	572	492	342	284
Scenario 2	Weather	Twitter	Facebook	CNN
GRID	797	451	362	357
CURA	595	503	378	305

Table 2: Time (seconds) required by GRID and CURA

Scenario 1	Weather	Twitter	Facebook	CNN
GRID	28	219	193	27
CURA	8.8	8	7.6	8
Scenario 2	Weather	Twitter	Facebook	CNN
GRID	4.2	4.77	0.97	3.4
CURA	0.73	0.811	0.731	0.72

Table 1 shows the power consumption achieved by GRID and CURA under the two scenarios, i.e., the image converter and the power-saving mode, respectively. In both scenarios, CURA requires less power for Weather and CNN, while GRID requires less power for Twitter and Facebook. The reason is as follows: CURA scales down pixel values based on human attention, so it usually saves more power when the low-attention regions contain a large number of bright pixels. In contrast, GRID treats all pixels equally and may save more power by darkening bright pixels in the high-attention regions. However, we observe that the difference in power consumption between the two scenarios are more significant under GRID than under CURA. This phenomena implies that GRID is more sensitive to image resizing (i.e., Lanczos resampling used in the power-saving mode scenario). The reason is as follows. To maintain the required visual quality, the pixel values scaled down in Scenario 1 have to be increased slightly in Scenario 2 because the latter is based on low-resolution versions. CURA distributes the increment over all the pixels, while GRID only increases the pixel values that exceed the threshold. As a result, the power achieved by GRID in Scenario 2 increases significantly due to the OLED power model' characteristic. Table 2 shows the time required by GRID and CURA under the two scenarios. CURA outperforms GRID in terms of the processing time for all the cases because GRID has to compute much more SSIM indexes. Moreover, the time required by CURA is very stable in the same scenario, but that required by GRID varies significantly. This is because GRID's processing time highly depends on whether the initial threshold is predicted accurately. Importantly, we observe that the images processed by GRID often suffer from obvious chrominance changes and sharp edges, as shown in Figure 7.

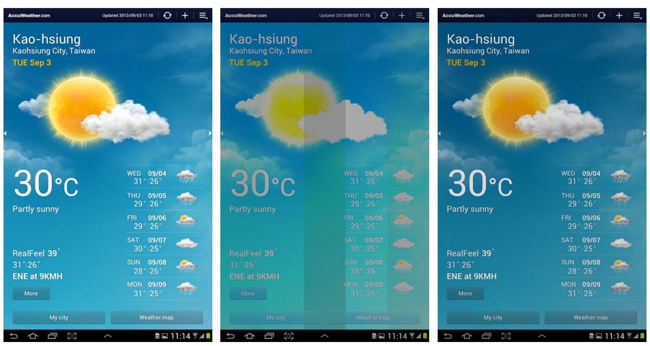

(a) Original (b) GRID (c) CURA

Figure 7: Images processed by GRID and CURA

5. RELATED WORK

Much of the early research on low-power display techniques was intended primarily for LCD displays [9, 13]. However, the design of low-power techniques dealing with OLED displays has attracted more attention recently. Comprehensive measurements of OLED power modeling helped understand the characteristics of OLED power consumption [8], and have facilitated the empirical design of various low-power display techniques, such as partial display disabling/dimming [3], color remapping [6], and OLED dynamic voltage scaling [14], as well as the technique, called image pixel scaling, introduced in this work.

Based on the above techniques, a number of researchers have proposed effective power-saving methodologies and algorithms that focus on a variety of mobile applications on OLED displays. In particular, inspired by an observation that the user usually focuses on just one half of the screen for most (but not all) interactive applications, an approach based on partial panel dimming was developed to dim the top or bottom portions of the screen that contain content relatively unimportant to the user [15]. In [7], a color-remapping methodology that renders web pages with power-optimized color schemes was realized on mobile devices to make web browsing more power-efficient on OLED displays. Moreover, an algorithm based on OLED dynamic voltage scaling was proposed in [5] to reduce supply voltage while retaining visual quality in video streaming applications. It divides the OLED panel into multiple rectangular regions and optimizes the voltage of each region under a given quality requirement. The algorithm is a closely related approach that also aims at retaining image quality; thus, it was revised to comply with image pixel scaling for comparison, as described in Section 4.

6. CONCLUDING REMARKS

This paper presents a methodology, called CURA, for quality-retaining power savings on mobile OLED displays. Specifically, CURA exploits visual attention, and models a fundamental optimization problem with scaling constraints to limit the image distortion of regions according to their attention levels and to avoid sharp edges between adjacent regions. To solve the problem, CURA relies on an underlying algorithm that can derive optimal scaling assignments without the information about the OLED display's power model. In order to evaluate the improvement in power savings, based on CURA, we have realized two application scenarios on a commercial Samsung Galaxy Tab 7.7 tablet, and conducted a series of experiments with some snapshot images of popular mobile apps. The results show that CURA can achieve a significant reduction in power consumption, while maintaining the high visual quality of the regions that dominate the user's attention. The reason for the reduction is that CURA scales down pixel values by different magnitudes appropriately according to their tolerable distortion from the perspective of visual experience. As a result, the reduction is more obvious when an image has a large luminance value but a small saliency value.

Acknowledgement

This work was supported in part by the National Science Council under Grants No. 101-2219-E-002-002 and No. 102-2221-E-001-007-MY2.

REFERENCES

[1] T. Acharya and A. K. Ray. *Image Processing: Principles and Applications*, pages 38–39. John Wiley & Sons, Inc., 2005.

[2] A. Bartolini, M. Ruggiero, and L. Benini. Visual Quality Analysis for Dynamic Backlight Scaling in LCD Systems. In *Proc. of IEEE/ACM DATE*, pages 1428–1433, 2009.

[3] J. Betts-LaCroix. Selective Dimming of Oled Displays. US Patent 0149223 A1, 2010.

[4] A. Borji and L. Itti. State-of-the-Art in Visual Attention Modeling. *IEEE Trans. on PAMI*, 35(1):185–207, 2013.

[5] X. Chen, J. Zheng, Y. Chen, M. Zhao, and C. J. Xue. Quality-retaining OLED Dynamic Voltage Scaling for Video Streaming Applications on Mobile Devices. In *Proc. of IEEE/ACM DAC*, pages 1000–1005, 2012.

[6] M. Dong, Y.-S. K. Choi, and L. Zhong. Power-Saving Color Transformation of Mobile Graphical User Interfaces on OLED-based Displays. In *Proc. of IEEE/ACM ISLPED*, pages 339–342, 2009.

[7] M. Dong and L. Zhong. Chameleon: A Color-Adaptive Web Browser for Mobile OLED Displays. In *Proc. of ACM MobiSys*, pages 85–98, 2011.

[8] M. Dong and L. Zhong. Power Modeling and Optimization for OLED Displays. *IEEE Trans. on Mobile Computing*, 11(9):1587–1599, 2012.

[9] P.-C. Hsiu, C.-H. Lin, and C.-K. Hsieh. Dynamic Backlight Scaling Optimization for Mobile Streaming Applications. In *Proc. of IEEE/ACM ISLPED*, pages 309–314, 2011.

[10] L. Itti, C. Koch, and E. Niebur. A Model of Saliency-based Visual Attention for Rapid Scene Analysis. *IEEE Trans. on PAMI*, 20(11):1254–1259, 1998.

[11] C. Lanczos. *The Variational Principles of Mechanics*. Univ. of Toronto Press, 1970.

[12] W. Lin and C. J. Kuo. Perceptual Visual Quality Metrics: A Survey. *Jour. of Visual Communication and Image Representation*, 22(4):297–312, 2011.

[13] A. lranli and M. Pedram. DTM: Dynamic Tone Mapping for Backlight Scaling. In *Proc. of IEEE/ACM DAC*, pages 612–616, 2005.

[14] D. Shin, Y. Kim, N. Chang, and M. Pedram. Dynamic Voltage Scaling of OLED Displays. In *Proc. of IEEE/ACM DAC*, pages 53–58, 2011.

[15] K. W. Tan, T. Okoshi, A. Misra, and R. K. Balan. FOCUS: A Usable & Effective Approach to OLED Display Power Management. In *Proc. of ACM UbiComp*, pages 573–582, 2013.

[16] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image Quality Assessment: From Error Visibility to Structural Similarity. *IEEE Trans. on Image Processing*, 13(4):600–612, 2004.