ShiftMask: Dynamic OLED Power Shifting Based on Visual Acuity for Interactive Mobile Applications

Han-Yi Lin[1], Pi-Cheng Hsiu[2], Tei-Wei Kuo[1,2]
[1]Department of Computer Science and Information Engineering, National Taiwan University, Taiwan
[2]Center for Information Technology Innovation, Academia Sinica, Taiwan
E-mail: d03922006@csie.ntu.edu.tw, pchsiu@citi.sinica.edu.tw, ktw@csie.ntu.edu.tw

***Abstract*—OLED power management on mobile devices is very challenging due to the dynamic nature of human-screen interaction. This paper presents the design, algorithms, and implementation of a lightweight mobile app called ShiftMask, which allows the user to dynamically shift OLED power to the portion of interest, while dimming the remainder of the screen based on visual acuity. To adapt to the user's focus of attention, we propose efficient algorithms that consider visual fixation in static scenes, as well as changes in focus and screen scrolling. The results of experiments conducted on a commercial smartphone with popular interactive apps demonstrate that ShiftMask can achieve substantial energy savings, while preserving acceptable readability.**

***Index Terms*—OLED power, user interaction, visual acuity, mobile systems**

I. INTRODUCTION

Future mobile devices, like smartphones and tablets, will increasingly feature emerging *organic light-emitting diode* (OLED) displays, whose power consumption comes primarily from a large number of individual light emissive diodes [1]. Meanwhile, to provide higher visual realism, the display resolution of mobile devices has continued to increase, from HD to Full HD, and recently to QuaD HD and even 4K Ultra HD [2]. Although high density displays provide sharper images, they also require a larger number of light emissive diodes, thus imposing a heavier burden on the limited battery capacity of mobile devices, especially with the strong demand for larger, higher-resolution screens.

Numerous low-power techniques have been proposed in recent years. In particular, Shin et al. [3] developed *OLED dynamic voltage scaling*, a technique that reduces the supply voltage of each pixel's diode to the minimum needed to sustain the pixel's brightness. Based on the hardware-assisted technique, Chen et al. [4] divided a display panel into multiple rectangular regions and tried to minimize the voltage supply of each region while retaining the visual quality of the displayed image. To leverage the flexibility provided by OLED technology, Lin et al. [5] designed a technique called *image pixel scaling*, which allows the pixel values of any-shaped regions to be scaled down by different magnitudes based on the level of visual attention. As the contrast is much more important than the absolute brightness to the image quality perceived by the human visual system, Kang et al. [1] attempted to increase the contrast while decreasing the brightness in order to enhance the image quality and save OLED power simultaneously. They also implemented a win-win camera that allows users to take quality-enhanced power-saving pictures on their OLED smartphones. Subsequently, Chen et al. [6] utilized some basic configuration parameters computed during camera recording to improve OLED power savings in future playback. Essentially, the above schemes deal with (a sequence of) *still images* and do not consider dynamic user behavior. Moreover, when applied to applications with frequent content rendering, the "content-aware" schemes would induce significant computational overhead.

Many popular mobile apps, including those for social networking (e.g., Facebook), instant messaging (e.g., LINE), and news providers (e.g., BBC), are *interactive applications* in which content displayed on the screen changes dynamically based on the user's behavior. Motivated by the observation that much of the web content contains bright colors that are not power-efficient on OLED displays, Dong and Zhong [7] developed a color adaptive browser that renders web pages with power-optimized color schemes under user-supplied constraints. The OLED-friendly browser exploits the *color transformation* technique [8], which changes the original colors into harmonious colors that consume less power, and allows users to set their color and fidelity preferences to ensure that the transformed web pages are well accepted by users.

The most effective way to reduce OLED power consumption is to darken portions of the displayed content, referred to as *partial dimming* [9]. Unfortunately, some content will become dark or even invisible to the user. The challenge thus lies in preserving acceptable readability. Chen et al. [10] proposed dimming the screen areas covered by the user's fingers, thus reducing the energy consumption when the user is interacting with a smartphone via the touch screen. Actually, when a user "reads" rather than watches the displayed content, the focus is normally concentrated on only a small portion of the screen. Tan et al. [11] observed that most apps require the user to focus on just half of the screen. This inspired them to develop a scheme called Focus, which allows the user to create a dimming profile for each app and dim the screen's top and/or bottom portions that are of less interest. A user satisfaction survey of 30 participants showed that users are willing to trade off screen brightness for longer usage, especially when the battery is low, if acceptable readability can be preserved in the focused region. However, because the dimming profiles are created offline, the clear region determined for each individual app is *stationary*. Such an immutable scheme, which is not dynamically adaptable to the user's focus of attention, would adversely affect the user experience of the used application.

For touch-screen device usage, human-screen interaction is essential and, to a certain extent, nondeterministic. Even on a static scene, the focus of a user's attention may change over time, and individual apps may have different movement patterns. Furthermore, mobile users often interact with their devices by scrolling to search for content of interest, and sometimes switch between applications quickly. Yu et al. [12] observed that when an interactive app is used, about one-third of the time is spent scrolling, and it could be more than half the time for interaction-intensive apps like web browsing. Although partial dimming generally achieves more substantial power savings than other low-power techniques that maintain the entire screen's fidelity, such frequent changes in the displayed content, along with the dynamic nature of human-screen interaction, make OLED power management based on partial dimming very challenging.

In this paper, we introduce a lightweight mobile app called ShiftMask, which dynamically shifts OLED power to the portion of the screen that is currently of interest to the user and simultaneously dims the remainder of the screen to achieve substantial energy savings. To ensure that ShiftMask is of practical use, several requirements

978-1-5090-6023-8/17/$31.00 ©2017 IEEE

must be met. First, the energy consumption must be reduced without compromising the readability in an unacceptable way. This is the most important principle in aspect of making ShiftMask "usable". Our design exploits *visual acuity* to comply with the perception of the human eye while partially dimming the display. Second, ShiftMask should be "lightweight" and "adaptive" so that it can instantly fit the focus of the user's attention; otherwise, the reduction in display energy consumption would be offset by extra computational overhead. This is particularly difficult when the distribution of partial dimming over the screen varies frequently due to a high level of human-screen interaction and not addressed in prior studies like Focus [11]. To resolve the difficulty, we propose three efficient algorithms that perform partial dimming based on visual acuity for static scenes, changes in focus, and screen scrolling respectively. The third requirement is that ShiftMask should be "generic" to most (if not all) interactive applications, thereby supporting unmodified apps and commercial smartphones/tablets. To this end, ShiftMask is implemented as a stand-alone Android app that applies a transparent bitmap to the entire screen and uses the *alpha blending* technique provided in nearly all computer graphics libraries to realize partial dimming. Moreover, for the sake of convenience, ShiftMask provides a user-friendly interface for the user to easily indicate his focus and manipulate the focused region. Finally, to evaluate the efficacy of ShiftMask, we conducted experiments on a Samsung Galaxy Note 4 smartphone using popular mobile apps with different interaction patterns. With ShiftMask installed, the smartphone can achieve energy savings of between 42.5 and 62.1% for the OLED display, while preserving an acceptable level of readability. This accounts for 13.4 to 39.3% of the device's total energy consumption when the extra computational overhead is taken into account. The results of experiments performed to compare ShiftMask with Focus [11] also provide some valuable insights into dynamic OLED power management for interactive mobile applications.

The remainder of this paper is organized as follows. Section 2 provides some background information and the design rationale. We then present the three underlying algorithms in Section 3, describe the implementation issues in Section 4, and discuss the experiment results in Section 5. Section 6 contains some concluding remarks.

II. BACKGROUND AND MOTIVATION

A. OLED Power Savings

An OLED display panel comprises a two-dimension matrix of pixels, each of which contains three diodes that emit red, green, and blue light respectively. A diode's light intensity depends on the electric current that flows through it, and the amount of current is determined by the diode's color and pixel value [7]. In the RGB color space, each pixel is defined by three values, in the range 0 to 255, whose combinations can create various colors. Consequently, the power required to display a scene on an OLED screen depends primarily on the scene's content. As modeled in [4, 13], the required power is the sum of each pixel's power consumption, and the power consumed by each pixel is the sum of the power consumed by its RGB diodes. Therefore, the power model of an OLED display can be approximated by a strictly increasing function of the pixel value, and the increasing slopes may vary between OLED displays [5].

OLED power consumption increases dramatically with the pixel's RGB values, which can be converted into the $\mathrm{YC_bCr}$ coordinate space with Y representing the brightness and $\mathrm{C_bCr}$ representing the chrominance. The human visual system is much more sensitive to variations in brightness than color. Thus, decreasing each pixel's brightness by a *dimming ratio* while maintaining the chrominance of the pixel is a very effective way to reduce the power required to display a scene on an OLED screen [1]. However, the reduced power consumption is achieved at the cost of a dimmer scene, and this in turn will interfere with the readability of the scene if it is not dimmed appropriately. The technical challenge thus lies in determining an appropriate dimming ratio for each pixel so as to achieve substantial OLED power savings, while retaining acceptable readability.

B. Human Visual Acuity

Visual acuity refers to the clarity of vision and defines the ability of the eye to discern fine details. When a user looks at a scene, the visual acuity is not uniform across the visual field. Instead, the acuity degrades considerably as the *eccentricity*, denoted by ϕ, increases from the fovea located in the center of the eye [14], as illustrated in Figure 1(a). Vision within the fovea is called central vision, while that outside the fovea is generally called peripheral vision. The former only covers about $\theta = 5^\circ$ of the visual field [15], but it allows us to read details clearly when our eyes are focused straight ahead. On a mobile screen, the corresponding region is called the *central region* in which the content is currently of interest to the user. The size and location of the central region depend on the application used. Let the screen contain H rows of pixels. If the region's upper and lower boundaries are the rows indexed by s and t respectively, the distance d between the eyes and the screen such that the central region remains readable can be derived by

$$\tan(\frac{\theta}{2}) = \frac{|s-t|/2}{d}, \tag{1}$$

where the tangent function relates an angle of a triangle to the lengths of its adjacent and opposite sides. The remaining regions corresponding to peripheral vision are called *peripheral regions*. Their acuity drops to about 50% when $\phi = 2^\circ$ and to less than 30% when $\phi = 7^\circ$. Figure 1(b) shows the acuity function, denoted as $a()$, derived based on the clinical study in [16]. The function, which gives the same row of pixels the same acuity to comply with the horizontal reading pattern of the human eye, provides a foundation for partial dimming.

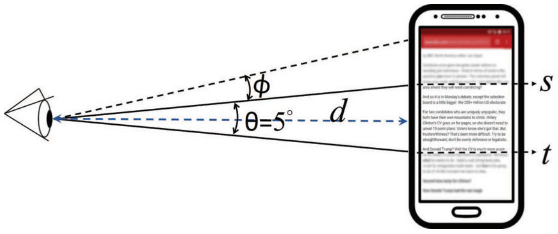

(a) Central and peripheral regions

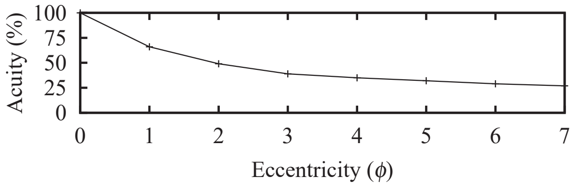

(b) An acuity function derived in [16]

Fig. 1. An illustration of human visual acuity

When the user changes (or adjusts) the central region with his visual focus, the distribution of visual acuity over the screen also changes; accordingly, the dimming ratios of all the pixels should be readjusted to adapt to the new distribution. However, if the dimming ratios of a significant proportion of pixels change abruptly, flickering may occur and interfere with the user's visual experience. Thus, the magnitude of changes applied to dimming ratios should be limited, so that the changes are too subtle to be discerned by the human eye. In the human visual system, the *just noticeable difference* (JND) is the minimum amount by which the stimulus intensity must be changed to produce a noticeable variation in sensory experience. For many sensory modalities, including brightness, the ratio of JND to the stimulus intensity is a constant, and the constant for brightness perceived by the rods of the eye is $\rho = 0.14$ [17]. In other words, a dimming ratio can be increased or decreased by at most ρ per change. Furthermore, when the user scrolls the screen, a relative movement speed occurs between the eyes and the scene. That speed, denoted as v, has a direct impact on the *useful field of view* (UFOV), whose size shrinks as the speed increases. As measured in [18], the accepted maximal speed at which perception can occur is $\bar{v} = 30°/sec$. This implies that the human eye cannot perform visual cognitive tasks correctly when the scrolling speed is greater than $\bar{v}$. As a result, the UFOV will shrink to a straight line in the middle of the central region.

III. ACUITY-AWARE DYNAMIC PARTIAL DIMMING

We exploit visual acuity to achieve substantial OLED power savings, while ensuring acceptable readability. In Section III-A, we consider visual fixation in a static scene. Then, we deal with focus changes and screen scrolling in Sections III-B and III-C respectively.

A. Static Scenes

When the user focuses on a region in a static scene, the visual acuity decreases as the eccentricity increases. Here, we present an algorithm to perform partial dimming according to the distribution of visual acuity.

Algorithm 1 Static Scenes

```
Input: A central region with boundaries $s$ and $t$ on a screen
    comprised of $H$ rows, the angle of central vision $\theta = 5°$, and
    an acuity function $a()$
Output: An array of dimming ratios $\sigma[]$
1: $d \leftarrow \frac{|s-t|/2}{\tan(\frac{\theta}{2})}$
2: for $i \leftarrow 1$ to $H$ do
3:     if $s \le i \le t$ then
4:         $\sigma[i] \leftarrow 1$
5:     else
6:         $\phi \leftarrow \tan^{-1}\left(\frac{|(s+t)/2\ -\ i|}{d}\right) - \frac{\theta}{2}$
7:         $\sigma[i] \leftarrow a(\phi)$
8: return $\sigma[]$;
```

Given a central region whose upper and lower boundaries are respectively the s_{th} and t_{th} rows on a screen comprised of H rows, with the angle of central vision $\theta = 5°$ and an acuity function $a()$, Algorithm 1 returns an array $\sigma[]$, in which a dimming ratio is determined for each row of pixels. At the beginning of the algorithm, the distance d between the eyes and the screen is computed based on the tangent function given in Equation (1), where the central vision covers two identical right-angled triangles (Line 1). Then, we determine a dimming ratio for each row $i = 1, ..., H$ (Line 2). If $s \le i \le t$, then row i is located within the central region, and we leave its pixels unchanged to retain the readability (Lines 3-4). Otherwise, if row i is in a peripheral region, we decrease the brightness of its pixels based on the acuity function to achieve OLED power savings (Lines 5-7). To this end, the eccentricity of row i is computed via the inverse tangent function according to Equation (1), where the triangle has an angle $\phi + \frac{\theta}{2}$ and two sides with lengths d and $|(s+t)/2 - i|$ respectively, as illustrated in Figure 1(a). We then substitute the eccentricity ϕ into the acuity function $a()$ to obtain the corresponding acuity level, as shown in Figure 1(b), and set the dimming ratio for the row proportional to the obtained acuity level. At the end, the array of dimming rations $\sigma[]$ is returned (Line 8).

B. Focus Changes

While reading a static scene, the user may change his visual focus. As the distribution of visual acuity will change with the focus, the dimming ratios of all pixels should be readjusted to adapt to the new distribution. To prevent flickering, we propose an algorithm that gradually adjusts the dimming ratios of all pixels so that the transition of partial dimming can be made smoothly from the original central region to the new one.

Algorithm 2 Focus Changes

```
Input: A new central region with boundaries $\hat{s}$ and $\hat{t}$ on a screen
    comprised of $H$ rows, the angle of central vision $\theta = 5°$, an
    acuity function $a()$, the original dimming ratio array $\sigma[]$, and the
    JND constant $\rho = 0.14$
1: $\hat{\sigma}[] \leftarrow$ Algorithm 1$(\hat{s}, \hat{t}, H, \theta, a())$
2: while $\sigma[] \neq \hat{\sigma}[]$ do
3:     for $i \leftarrow 1$ to $H$ do
4:         if $\sigma[i] < \hat{\sigma}[i]$ then
5:             $\sigma[i] \leftarrow \min(\sigma[i] \times (1+\rho), \hat{\sigma}[i])$
6:         else
7:             $\sigma[i] \leftarrow \max(\sigma[i] \times (1-\rho), \hat{\sigma}[i])$
8:     Perform partial dimming based on $\sigma[]$
```

Given a new central region with two boundaries $\hat{s}$ and $\hat{t}$ on a screen comprised of H rows, the angle of central vision $\theta = 5°$, an acuity function $a()$, the array of the original dimming ratios $\sigma[]$, and the JND constant $\rho = 0.14$, Algorithm 2 determines a sequence of arrays of dimming ratios. Meanwhile, it applies the intermediate ratios to the scene whenever an array is determined. The process is repeated until all the dimming ratios have been adapted to the new acuity distribution. Initially, Algorithm 1 is invoked to obtain another array of dimming ratios, stored in $\hat{\sigma}[]$, with respect to the new acuity distribution (Line 1). Next, each of the dimming ratios stored in the original array $\sigma[]$ gradually moves towards its counterpart in the new array $\hat{\sigma}[]$ until the two arrays become identical (Line 2). For each round of changes, we increase or decrease the dimming ratio in $\sigma[i]$, $\forall\ 1 \le i \le H$, depending on whether $\sigma[i] < \hat{\sigma}[i]$ (Lines 3-7). The increment or decrement is limited by the JND constant ρ to ensure that the change is too subtle to be discerned by the human eye. At the end of the round, we apply partial dimming to the scene based on $\sigma[]$ so that the readability of the new central region will be improved smoothly, while that of the original region will be degraded gradually to save power (Line 8).

C. Screen Scrolling

When a user interacts with a mobile application, he may browse to look for some content of interest by scrolling the screen. As a consequence, a relative speed occurs between the eyes and the scene. The size of UFOV shrinks as the speed increases. Here, we present an algorithm that dynamically shrinks or enlarges the central region when a scrolling speed is given.

Given the original central region with boundaries s and t on a screen comprised of H rows, the angle of central vision $\theta = 5°$, an acuity function $a()$, the array of the original dimming ratios $\sigma[]$, the JND constant $\rho = 0.14$, the maximal UFOV speed $\bar{v} = 30°/sec$, and a screen scrolling speed v, Algorithm 3 determines a new central region based on the speed v. It then shrinks or enlarges the original central region to fit the new one. First, we determine the boundaries $\hat{s}$ and $\hat{t}$ of the new central region with respect to speed v (Lines 1-2). Specifically, if v is faster than the maximal UFOV speed $\bar{v}$, the central region is shrunk to two adjacent rows (i.e., $|s - t| = 1$) in the middle of the central region. In contrast, when v is equal to zero, the

Algorithm 3 Screen Scrolling

```
Input: The original central region with boundaries $s$ and $t$ on a screen comprised of $H$ rows, the angle of central vision $\theta = 5^\circ$, an acuity function $a()$, the array of the original dimming ratios $\sigma[]$, the JND constant $\rho = 0.14$, the maximal UFOV speed $\bar{v} = 30^\circ/sec$, and a screen scrolling speed $v$
1: $\hat{s} \leftarrow s + \frac{|s-t|}{2} \times \frac{\min(v,\bar{v})}{\bar{v}}$
2: $\hat{t} \leftarrow t - \frac{|s-t|}{2} \times \frac{\min(v,\bar{v})}{\bar{v}}$
3: $\hat{\theta} \leftarrow 2 \times \tan^{-1}\left(\tan(\frac{\theta}{2}) \times \frac{|\hat{s}-\hat{t}|}{|s-t|}\right)$
4: Algorithm 2$(\hat{s}, \hat{t}, H, \hat{\theta}, a(), \sigma[], \rho)$
```

scrolling stops and the central region returns to the original size. For a speed v within $\bar{v}$, the size of the central region is set proportional to v. Moreover, the angle of central vision is affected by the scrolling speed, while the distance d between the eyes and the screen remains the same. Then, the new angle can be computed based on the tangent equation $d = \frac{|s-t|/2}{\tan(\frac{\theta}{2})} = \frac{|\hat{s}-\hat{t}|/2}{\tan(\frac{\hat{\theta}}{2})}$ (Line 3). Finally, Algorithm 2, along with $\hat{s}$, $\hat{t}$, and $\hat{\theta}$, is invoked to achieve a smooth transition between the two central regions of different sizes (Line 4).

IV. SHIFTMASK IMPLEMENTATION

Next, we describe the ShiftMask app, which allows users to dynamically shift OLED power to the portion of interest by dimming the remainder of the screen online. The app is especially useful in low battery situations when readability is less important than energy savings. To support existing interactive apps, we developed ShiftMask as a stand-alone app that can be installed on any Android smartphone or tablet without any modifications of the underlying framework. Basically, it uses an RGBA bitmap (in which each pixel has a transparency channel in addition to the red, green, and blue color channels) to mask the entire screen, where the RGBA values of all pixels are initialized at 0 to display a transparent black image. When the dimming ratios are determined by one of our algorithms, the app sets the transparency of each row of pixels in proportion to the corresponding ratio. It then exploits the alpha blending technique (provided in the computer graphics library) to combine the translucent image with the displayed content of the used app in the background to perform partial dimming. Note that the bitmap's penetrability is enabled so that the user can still interact with the background app directly as usual.

Fig. 2. Snapshots of the ShiftMask app

To enable ShiftMask to be conveniently used much in the way that users interact with mobile apps, it allows the user to manipulate the central region via the control bars[1] on the right-hand side of the screen, as shown in Figure 2, and supports four intuitive operations as follows. First, the central region can be enlarged or shrunk by adjusting the distance between the two red bars to fit the needs of the interactive app used. Second, the blue bar allows the central region to be scrolled up and down to follow the user's focus. Third, the user can change the central region directly by touching the green bar when his focus jumps to another portion of the screen. The row touched will become the center of the new central region. Fourth, if the user scrolls the screen, or the app used is zoomed in or out, the central region will be resized automatically based on the scrolling speed.

The first two operations can be realized by invoking Algorithm 1 repeatedly. Each invocation requires less than 0.5 ms on a Samsung Galaxy Note 4 smartphone. To ensure the central region is changed smoothly without incurring a significant computational overhead, instead of invoking Algorithm 1 whenever the central region is resized or shifted by one row, we invoke it immediately before any row with a dimming ratio smaller than 1-$\frac{\rho}{2}$ becomes the central region's boundary. The objective is to minimize the number of invocations while satisfying the JND constraint. Intuitively, the third operation can be implemented by simply invoking Algorithm 2. However, it takes a maximum of 6.5 ms, which is too short to transit the central region "gradually" from one portion to another on a display with a frame rate of 60 frames per second. Thus, we deliberately update the dimming ratios once every 17 ms to ensure that each round of change during the execution of Algorithm 2 can be displayed successfully. Finally, the fourth operation can be performed by invoking Algorithm 3 with a given scrolling speed. When the user slides his finger on the screen, a sequence of speeds will be detected. We use the speed immediately before the finger is moved away from the screen for shrinking the central region. When the screen scrolling stops, Algorithm 3 is invoked again with the scrolling speed set at zero to enlarge the central region back to the original size. Each invocation of the algorithm takes approximately 7 ms, which is also too short. Thus, we exploit the lazy display approach used for Algorithm 2 to allow the central region to be resized "gradually". Based on the above measurement, the algorithms are very computationally efficient and could be implemented in different software stacks for other online scenarios, such as a screen power-saving mode for graphical user interfaces on OLED tablets [5].

V. PERFORMANCE EVALUATION

A. Experiment Setup

We installed ShiftMask on a Samsung Galaxy Note 4 smartphone and conducted a series of experiments with some popular mobile apps. The smartphone features an Octa-core big.LITTLE processor and a 5.7-inch AMOLED display with a Quad HD resolution of 2560×1440 pixels. Table I details the specifications of the related hardware and software. If necessary, the smartphone can access the Internet via a D-Link 802.11ac access point dedicated for our experiments. We used the power monitor produced by Monsoon Solutions to measure the smartphone's transient power and energy consumption. In addition, to reduce the potential influence of human intervention, we used a touch event recorder called RepetiTouch Pro to record and replay all the screen touch events for each investigated app. This ensured that the same interaction patterns were reproducible in the experiments.

To cover various interaction patterns, we chose four representative apps, namely BBC News, Google Maps, Facebook, and LINE, which have different frequencies of focus changes and screen scrolling. Figure 3 shows some snapshots of the apps when ShiftMask was applied. The central region of BBC News was initialized in the top one-fourth of the screen, and the user read each quarter every 15 seconds from top to bottom sequentially. For Google Maps, the central region was set as one-third of the screen in the middle. The

---

[1] Eye tracking, if available on the smartphone, could make the implementation less intrusive with the proposed algorithms unmodified.

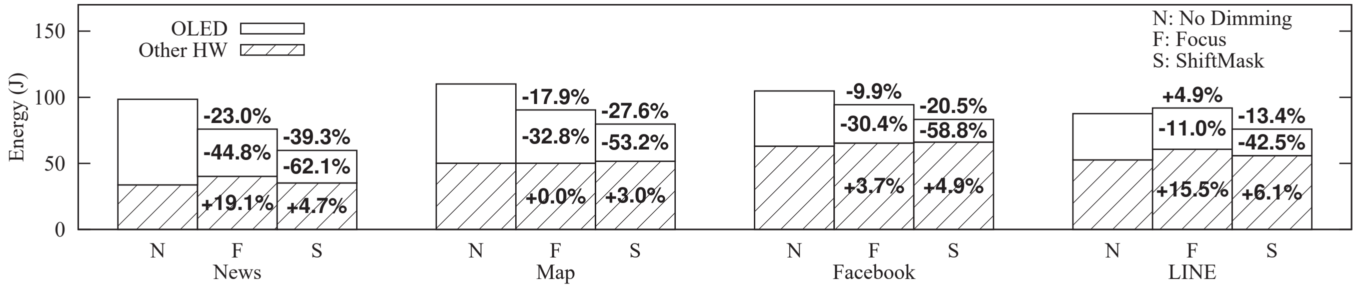

Fig. 4. Energy consumption

TABLE I
SPECIFICATIONS OF SAMSUNG GALAXY NOTE 4

Hardware	
SoC Model	Samsung Exynos 7 Octa 5433
CPU	ARM Quad-core Cortex-A57@1.9GHz ARM Quad-core Cortex-A53@1.3GHz
Memory	3GB LPDDR3 RAM
Display	5.7” Quad HD Super AMOLED 2560×1440 pixels, 518 ppi
GPU	ARM Mali-T760
Network	IEEE 802.11 a/b/g/n/ac Wi-Fi
Storage	32GB eMMC 5.0
Battery	Li-Ion 3220 mAh
Software	
OS	Android 6.0.1 Marshmallow Linux Kernel 3.10.9

user took 10 seconds to zoom in on a location by swiping the screen three times and looked at the map for 20 seconds; then, the process was repeated except that the map was zoomed out to the original scale this time. Similarly, the central region of Facebook was initialized as one-third in the middle, and the user took 10 seconds to find some content of interest by scrolling the screen continuously. Then, he read the content for 10 seconds and changed his focus to the bottom one-third of the screen and read that for 10 seconds. The process was repeated twice in one minute. To accommodate the keyboard, the central region of LINE was initialized in the bottom half of the screen. The user typed for 7 seconds and then changed his focus to read the latest message in the top half of the screen for 3 seconds. The process was repeated six times. Such settings allowed us to observe some typical yet diverse interaction patterns of mobile apps.

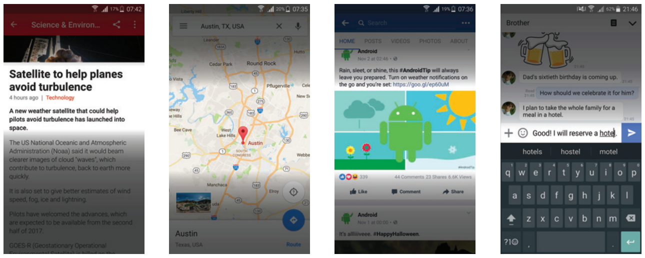

(a) News (b) Map (c) Facebook (d) LINE

Fig. 3. Snapshots of the apps studied with ShiftMask

To gain further insight, we compared ShiftMask with Focus [11], which requires the user to create a dimming profile for each app. The profile must specify the size and location of the clear region, as well as the intensities of the top and bottom of the screen. Based on the profile, Focus keeps the clear region undimmed and dims the remaining portions from 100% intensity *linearly* to the specified intensities at the screen’s edges. To ensure a fair comparison, the clear regions of Focus were set exactly the same as the central regions of ShiftMask, and the intensities at the screen’s edges were set at the dimming ratios determined by ShiftMask for the edges. Note that the dimming profiles were created offline, so the clear region for each app was stationary. As a result, the user’s interaction patterns must change to accommodate the clear region if necessary.

Both ShiftMask and Focus trade off readability in unfocused regions for energy savings. In [11], a user study with 30 participants has shown that dimmed apps can save significant energy and remain quite readable, when compared to the original undimmed apps. Here, we did not repeat the user satisfaction study but investigated the impact of user interaction patterns on the energy consumption of representative apps with ShiftMask or Focus applied. Moreover, to validate whether the readability is preserved in accordance with the visual acuity, we defined a metric called the *acuity distance*. The metric quantifies the difference between the array of dimming ratios $\sigma[]$ applied to a scene and the array $\hat{\sigma}[]$ in accordance with the visual acuity, i.e., $\frac{1}{H}\sum_{i=1}^{H}|\sigma[i]-\hat{\sigma}[i]|$. A shorter acuity distance implies that the distribution of partial dimming over the screen coincides better with the focus of the user’s attention and the content that is currently of interest to the user. To compute the acuity distance, we used the AZ Screen Recorder to capture 30 screenshots per second as each studied app was used.

B. *Experiment Results*

Figure 4 shows the individual energy consumption of the four apps without dimming, as well as with ShiftMask and Focus applied. Both ShiftMask and Focus can achieve considerable OLED energy savings, but they incur different degrees of computational overhead. The amount of OLED energy saved depends primarily on the app’s displayed content, regardless of the application. For all the apps, ShiftMask requires significantly less OLED energy than Focus because it performs partial dimming based on the acuity function, which decreases more quickly than the linear function used by Focus. On the other hand, the amount of extra computational energy used depends on the invoked algorithms, as well as on the “extra” human-screen interaction caused by ShiftMask or Focus. ShiftMask incurs a small overhead when the user manipulates it, and the overhead is similar for all the apps. In contrast, Focus requires nearly zero overhead for Map, but it incurs a significant overhead for News and LINE. The reason is that the clear and dimmed regions will never be changed after they are set according to the dimming profile created offline. As a result, for the News app, some extra scrolling operations are needed to align different paragraphs with the clear region. In the LINE app, the keyboard needs to occupy the clear region when the user is typing and must be folded up so that the user can read the message. Thus, the amount of overhead incurred is related to the app’s typical interaction pattern. Interestingly, instead of saving energy, Focus could increase the energy consumption if the OLED energy saved is offset by extra overhead, as in the case of LINE. The results show that ShiftMask can save 42.5 to 62.1% of OLED energy, which accounts for 13.4

to 39.3% of the total energy consumption when the extra overhead is taken into account. ShiftMask could obtain similar benefits on devices with different sizes of screens because the ratio of OLED energy saved is unrelated to the screen size and the computational overhead, which would increase or decrease slightly with the screen size, will not outweigh the energy consumption reduction.

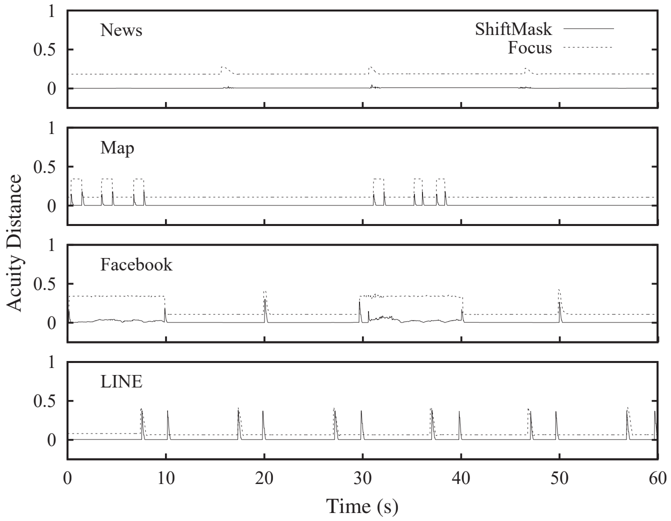

Fig. 5. Acuity distance

Figure 5 shows the variations in the acuity distance when ShiftMask and Focus are applied to each app respectively. In general, the acuity distance achieved by ShiftMask is shorter than that achieved by Focus. This result is as expected because ShiftMask dims a static scene based on the distribution of the acuity, while Focus simply utilizes a linear function. The difference between the distances achieved by the two schemes is more manifest when the focused region is smaller, as in the case of News. This is because a larger number of pixels can actually be dimmer than when they are dimmed by Focus, provided that their details cannot be clearly discerned by the human eye. When the user interacts with the screen, the acuity distance varies and an abrupt increase may occur during a change in focus or screen scrolling. The reason is that, instead of adapting immediately to the new distribution of acuity, ShiftMask gradually adjusts the dimming ratios to prevent flickering. As a result, during a short period, the dimming ratios applied to a scene may not be completely in accordance with the visual acuity. By contrast, Focus requires the user to scroll the screen up or fold up the keyboard to align the content of interest with the clear region. If the content does not coincide with the region, an abrupt increase in the acuity distance occurs. This also explains why the abrupt increases in Map and Facebook last much longer with Focus than with ShiftMask. As the displayed content is of less interest during screen scrolling, ShiftMask shrinks the central region, but Focus employs a stationary clear region. There is one exception where the acuity distances achieved by Focus at some points in LINE are shorter than those derived by ShiftMask. The reason is that when the user is ready to use the keyboard, it is “fortunate” to pop up within the clear region fixed in the bottom half of the screen. The results show that ShiftMask can instantly fit the focus of the user’s attention so that the content of interest is always located within the central region and thereby preserve an acceptable level of readability.

VI. CONCLUDING REMARKS

We have presented ShiftMask, which allows users to dynamically shift OLED power on mobile devices. In contrast to existing schemes intended for still images or stationary visual attention, we consider changes in focus and screen scrolling. ShiftMask can adapt to the dynamic nature of human-screen interaction because it reacts instantly to the change of the visual acuity and readjusts the distribution of partial dimming to reflect the change. We conducted a number of experiments on a commercial smartphone with four popular apps. The results show that ShiftMask, even with some computational overhead incurred by its extra human-screen interaction, can save a significant amount of energy while preserving the readability in accordance with the visual acuity.

The energy saved by ShiftMask already accounts for a considerable portion of the OLED energy consumption. In our future work, to further reduce the display energy used by interaction-intensive apps, we will extend ShiftMask to incorporate *dynamic resolution scaling* [2] to jointly consider the OLED and GPU energy consumption for frequent rendering of the displayed content.

ACKNOWLEDGMENT

This work was partially supported by the Ministry of Science and Technology under Grant 103-2218-E-003-002-MY3 and Grant 105-3111-Y-001-041.

REFERENCES

[1] C.-K. Kang, C.-H. Lin, and P.-C. Hsiu, “A Win-Win Camera: Quality-Enhanced Power-Saving Images on Mobile OLED Displays,” in *Proc. of IEEE/ACM ISLPED*, 2015, pp. 267–272.
[2] S. He, Y. Liu, and H. Zhou, “Optimizing Smartphone Power Consumption through Dynamic Resolution Scaling,” in *Proc. of ACM MobiCom*, 2015, pp. 27–39.
[3] D. Shin, Y. Kim, N. Chang, and M. Pedram, “Dynamic Voltage Scaling of OLED Displays,” in *Proc. of IEEE/ACM DAC*, 2011, pp. 53–58.
[4] 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*, 2012, pp. 1000–1005.
[5] C.-H. Lin, C.-K. Kang, and P.-C. Hsiu, “Catch Your Attention: Quality-Retaining Power Saving on Mobile OLED Displays,” in *Proc. of IEEE/ACM DAC*, 2014, pp. 1–6.
[6] X. Chen, J. Mao, J. Gao, K. W. Nixon, and Y. Chen, “MORPh: Mobile OLED-friendly Recording and Playback System for Low Power Video Streaming,” in *Proc. of IEEE/ACM DAC*, 2016, pp. 153:1–153:6.
[7] M. Dong and L. Zhong, “Chameleon: A Color-Adaptive Web Browser for Mobile OLED Displays,” in *Proc. of ACM MobiSys*, 2011, pp. 85–98.
[8] M. Dong, Y. K. Choi, and L. Zhong, “Power-Saving Color Transformation of Mobile Graphical User Interfaces on OLED-based Displays,” in *Proc. of IEEE/ACM ISLPED*, 2009, pp. 339–342.
[9] J. Betts-LaCroix, “Selective Dimming of OLED Displays,” US Patent 0149223 A1, 2010.
[10] X. Chen, K. W. Nixon, H. Zhou, Y. Liu, and Y. Chen, “FingerShadow: An OLED Power Optimization based on Smartphone Touch Interactions,” in *Proc. of USENIX HotPower*, 2014, pp. 6:1–6:5.
[11] 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*, 2013, pp. 573–582.
[12] J. Yu, H. Han, H. Zhu, Y. Chen, J. Yang, Y. Zhu, G. Xue, and M. Li, “Sensing Human-Screen Interaction for Energy-Efficient Frame Rate Adaptation on Smartphones,” *IEEE Trans. on Mobile Computing*, vol. 14, no. 8, pp. 1698–1711, 2015.
[13] M. Dong and L. Zhong, “Power Modeling and Optimization for OLED Displays,” *IEEE Trans. on Mobile Computing*, vol. 11, no. 9, pp. 1587–1599, 2012.
[14] M. Boucart, C. Moroni, M. Thibaut, S. Szaffarczyk, and M. Greene, “Scene Categorization at Large Visual Eccentricities,” *Vision Research*, vol. 86, pp. 35–42, 2013.
[15] H. Strasburger, I. Rentschler, and M. Jüttner, “Peripheral vision and pattern recognition: A review,” *Journal of Vision*, vol. 11, no. 4, p. 13, 2011.
[16] M. Millodot, *Dictionary of Optometry and Visual Science*, 7th ed. Butterworth Heinemann - Elsevier, 2009.
[17] K. M. and L. C., “Light and dark adaptation,” in *Webvision: The Organization of the Retina and Visual System*, 2005.
[18] K. Yamanaka and A. Minochi, “Measurement of Useful Field of View during Ocular Following Response,” in *Proc. of HCI*, 2013, pp. 455–459.