MobiSLIC: Content-Aware Energy Saving for Educational Videos on Mobile Devices

Qiyam Tung, Maximiliano Korp, Chris Gniady(✉),
Alon Efrat, and Kobus Barnard

University of Arizona, Tucson, Arizona
{qtung,mkorp,gniady,alon,kobus}@cs.arizona.edu

Abstract. We present a context-aware system that simultaneously increases energy-efficiency and readability for educational videos on smartphones with OLED displays. Our system analyzes the content of each frame of the video and intelligently modifies the colors and presentations of specific regions of the frame to drastically reduce display energy consumption while retaining relevant content of the lecture video. We achieve this by leveraging the mapping between frames and electronic versions of slides used in the lecture. This enables separate manipulation of the slide area and the background. Further, since the slides can themselves be analyzed for content (e.g. text and images within) this approach provides substantive control over energy use and user experience. We evaluate the system using extensive energy measurements performed on phones using two different display technologies. Our method was able to reduce energy usage up to 59.2 % of the energy used by the display which amounts to 27 % of the total energy used by the device.

Keywords: Energy-efficiency · OLED · Smartphones · Lecture videos

1 Introduction

We are becoming more dependent on smartphones and expect long battery life, despite steady increases in features and performance. As more smartphone users are spending their time on multimedia applications [1], such as watching movies, the demand for brighter displays that can be visible even in bright light is growing. Among the videos that are watched on the mobile devices are educational videos available on-line [2–7]. Large and high resolution displays of modern smartphones have made it easier to access and watch educational videos lectures in situations where using a regular desktop or laptop is inconvenient, such as commuting. However, being able to watch long videos depends on a reasonable battery life.

Displays in smartphones often account for a significant amount of the total energy consumption, making them one of the primary targets for energy optimizations. In particular, organic light-emitting diodes (*OLED*) displays, whose

© Institute for Computer Sciences, Social Informatics and Telecommunications Engineering 2014
I. Stojmenovic et al. (Eds.): MOBIQUITOUS 2013, LNICST 131, pp. 396–408, 2014.
DOI: 10.1007/978-3-319-11569-6_31

energy consumption is directly related to the color and brightness of the pixels [8], present opportunities for energy savings, such as modifying the colors and intensity of pixels in regions that are less important to the viewer.

This paper develops a system for modifying educational videos with consideration of the *content* appearing in each frame. Our techniques enable us to make selective adjustments that minimize the impact on the viewing experience while at the same time significantly reducing energy consumption. In addition, our technique has the ability to improve the readability of the slides in the video. We present two main approaches for reducing energy consumption for lecture videos. First, we utilize *backprojection* to replace portions of each frame occupied by the slide by the original slide used in the presentation. The slide color scheme is also revised to reduce energy consumption and increases readability. Second, we alter background regions of the video frame to improve energy efficiency, while only impacting the visibility of the region that is less important to the context of the video.

Hints to automatically understanding the content are obtained from analyzing slides used in the lecture and cross correlate them with the video itself. Our techniques work for recorded videos and hence are applicable to any of the discussions, lectures, colloquiums, etc., recorded already, as long as one can access the electronic slides being used in the lecture.

2 Related Work

Energy conservation has been a challenging focus of mobile research due to the ever-growing demand for more features. In an extensive study, Mahesri and Vardhan identified that the CPU and display consume the most amount of power in a laptop in [9]. The LCD screen dominates during idle times, using up to 29 % of the total power draw from the backlight alone. For mobile phones, even finding the optimal schedule for switching an active display to an inactive display can decrease total energy usage by 60 % [10]. For this reason, much research has been focused on minimizing the energy consumed by displays.

Flinn and Satyanarayanan propose a method of reducing the amount of energy consumed without dimming the backlight [11]. Specifically, they use different compression settings, such as reducing the frame rate and decreasing the dimensions of the video. Similarly, Park *et al.* [12] show that by degrading video quality by 13 %, they were able to get achieve savings of 42 %, measured on an ARM/Linux-based platform. While those techniques can offer energy savings, they degrade video quality. Degrading video quality is not desirable in our scenario as lecture videos often need to be large and sharp enough for text to be legible.

OLED displays can have a significant range of power demand from 0.25 W to 2 W as compared to the CPU's range of 50 mW to 600 mW [13]. Although the maximum luminance of OLED displays rarely occur in the usage of a phone, the phone usually supplies enough voltage to do so. Chen *et al.* [14] show that by dynamically adjusting the voltage after analyzing the display contents can save from 19.1 % to 49.1 % of the OLED power. Considering that video playback on a Samsung Galaxy S2 uses 35.6 % of the phone's total power [15], finding ways to optimize display energy usage is an important task.

In fact, substantial savings can be achieved by dimming parts of the display alone. Iyer *et al.* [16] found that reducing the intensity for the background and inactive windows in a desktop environment resulted in energy savings of up to 20 %. Similar optimizations have been applied to reduce the intensity of non-active portions of an iPAQ display [17]. A user study showed that these changes unobtrusive and sometimes even preferable. A similar technique of dimming regions that are not the locus of attention in video games can also achieve savings of 11 % of the display energy [18].

In addition, Dong *et al.* show that significant energy savings can also be obtained by modifying the color scheme of the graphical user interface [19]. In the same manner, for visualization of data, such as weather data or grouping voxels by color, Chuang *et al.* [20] show that one can optimize energy by selecting energy-aware but perceptually distinguishable colors. These techniques are not applicable for lecture videos as the color selection was not designed for natural scenes. Nevertheless, selectively dimming colors is the basis of reducing energy for our system.

3 OLED Displays

While dimming the backlight was the main method of decreasing energy consumption for liquid crystal displays (LCD), OLED displays offer greater potential for energy savings as the diodes are emissive and do not depend on a global backlight. The energy in an OLED display is directly correlated to how many pixels are on and what intensity level and color they are displaying. However, different display technologies may alter the energy consumption behavior of the display. In particular, phone displays may have different subpixel arrangement of the three primary colors, as is the case with the devices we used (Fig. 1). The Samsung Galaxy S Vibrant uses the PenTile matrix array (PMA) technology [21] and the red, blue, and green OLEDs have different dimensions. Furthermore, the pixels are laid out so that each one contains only two OLEDs. In order to generate the full range of colors, PMA displays utilize subpixel sharing. For example, in order for a pixel that does not contain a blue OLED to display blue, its adjacent pixel's blue OLED would need to be activated.

Samsung Galaxy S2, on the other hand, uses a standard RGB stripe pixel array. The layout is simpler and the three OLEDs are arranged in vertical stripes of red, green, and blue and hence do not require subpixel sharing. The specifications of the two phones used in our experiments are shown in Table 1. Despite the differences, the blue channel consistently consumes the most energy in both devices, as seen in Fig. 2. Therefore, in terms of color manipulation, our methods either dim all colors equally or dim out the blue channel. The choice depends on the viewer's preference and content of the lecture.

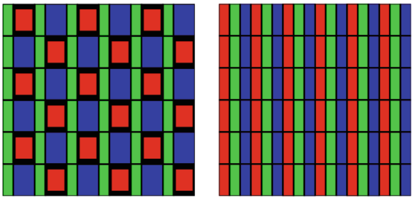

Fig. 1. Layout of pixels of different channels for the (left) Pentile matrix and (right) RGB stripe matrix.

Table 1. Smartphone specifications

Phone	Galaxy S	Galaxy S2
OS	Android 2.2	Android 2.3.5
Color depth	24 bit	24 bit
Display size	4 inches	4.5 inches
Display type	PenTile	RGB Stripe
Resolution	480 × 800	480 × 800
CPU	1 GHZ	2 × 1.5 GHZ
Memory	512 MB	1 GB

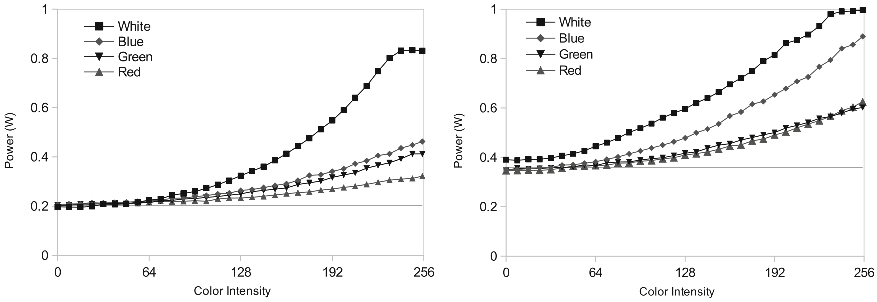

(a) Galaxy S (b) Galaxy S2

Fig. 2. Measurements of the power draw of the three primary colors for all sRGB values

4 Improving Energy Efficiency

Energy optimizations we are proposing are tightly related to the layout of a video frame. A typical frame consists of several regions (see Fig. 3):

1. The *non-slide background* consists of the audience, walls around the screen, and any other regions outside of the projector screen.
2. The *slide region* contains elements that carry vital information, such as text, figures, pie charts, graphs, etc.
3. The region of the slide not containing the text or images is the *slide background.*

The challenge are in identifying these regions in a deck of electronic slides, as well as in the video containing these slides. The energy optimizations can be vary from recoloring of the whole frame to recoloring the individual regions as identified above.

4.1 Altering the Non-Slide Background

The slide background may contain speaker's gestures, audience reactions or interactions, etc. While a speaker's gestures may be important, the absence of a color

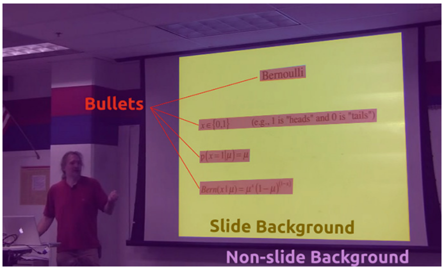

Fig. 3. The three kinds of regions in a video frame.

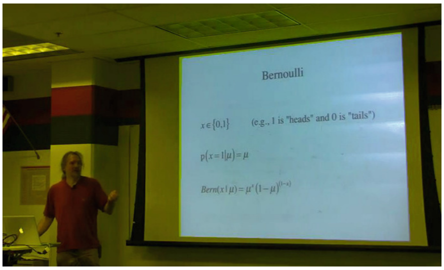

Fig. 4. The blue channel is dimmed from the background (color figure online).

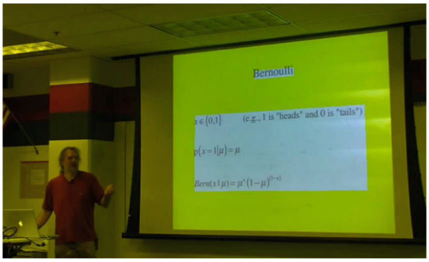

Fig. 5. The blue channel is dimmed from the non-bullets regions (color figure online).

or background all together is unlikely to prevent a user from understanding the lecture content. However, the desirability of the background depends on the person, and is orthogonal to the rest of the discussion.

To detect the background and slide regions, we first identify the slide used in each frame, along with the geometric mapping between slide and frames. We use the automatic slide detection already developed in the SLIC project [22], which gives us both the geometric transformation, the *homography*, as well the *temporal* information, which tells us which slide in the slide deck is shown in each video frame. Once the video frame layout is identified, we can dim or alter color by removing the most power hungry color, which is blue, from the background frame, as seen in Fig. 4. Surprisingly, this approach is rather ineffective for saving energy, as the non-slide background area is often small and much dimmer than the slide portion of the frame. Subsequently, the slide portion of the frame is the most power hungry portion and critical for successful energy optimizations.

4.2 Altering the Slide Background

To further take advantage of dimming nonessential regions we need to consider slide area that is not covered by text or images as shown in Fig. 5. Slide text and images tend to occupy a relatively small portion of the slide and the video

Vegetables
- Vegetables contain many vitamins and minerals
- However, different vegetables contain different spreads
- For example
  - Green vegetables typically contain vitamin A.
  - Dark orange and dark green vegetables contain vitamin C.

Vegetables
- Vegetables contain many vitamins and minerals
- However, different vegetables contain different spreads
- For example
  - Green vegetables typically contain vitamin A.
  - Dark orange and dark green vegetables contain vitamin C.

Vegetables

Fig. 6. Finding the bounding box by diffing two images.

frame under consideration. Subsequently, by dimming or removing color from such regions not occupied by text or images, we can get significant savings.

To identify text, embedded equations, and figure regions in the slide, we again rely on techniques developed as a part of the SLIC project [23]. Each presentation, which could be in the Keynote or the PowerPoint format, is converted to the latest PowerPoint format for processing. The latest PowerPoint format is XML based and allows simple alteration of text color. We use such color manipulations to identify the bounding boxes of the text region. We set a particular text region to different color in the slide. Subsequently, both slides, the original and the altered text color one, are exported as images. Since the only difference between the two slides is the color of the modified text, the difference reveals every pixel of the text. This process is illustrated in Fig. 6. We repeat this sequence for each text region and identify the entire text layout on the slides in the video frame.

We use an analogous method to find the bounding boxes of images in slides. The PowerPoint format stores all of a slide's images as individual image files. We extract these images and we replace them with uniquely colored monochromatic images, and export the slides to images. We proceed by taking the pixel-difference between the images to find the maximum x and y coordinates that bound the differing pixels. This method requires minimal knowledge about the format of the presentation, and can potentially extend to other presentation formats such as Google presentations, Adobe PDFs, etc.

Once we find the text and image locations in the slide we can frame them as shown in Fig. 5 and remove the power hungry colors from or dim the regions that not text or image areas. While this seems a rather straightforward process once we know image location on the slides, in reality, it can be complicated as slides may have different sizes or angles in the video frames and the text location needs to be translated to such layout.

In order to modify the color of the video frames, we need to know the temporal information (which slide is shown in each frame) and the location of the slide, which is given by the homography. To find them, we follow Fan *et al.*'s algorithm for finding slides in a video [24]. Finding the correct pixels to color can be computed with the aforementioned homography H, a 3×3 matrix. A homography describes how to map points on a plane from one camera's view to another, which makes it ideal for describing where the slide is in the frame (see Fig. 7). The mapping is only valid, however, if the points are described in homogeneous coordinates.

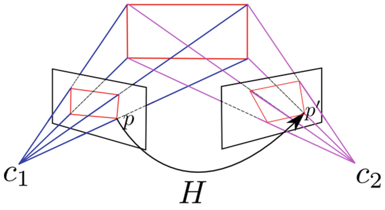

Fig. 7. A homography H describes how to map planes from one camera's viewpoint to another.

The homogeneous coordinates of a pixel in the slide, P_s, is related to its corresponding pixel coordinate in the frame P_f by the following equivalence: $HP_f = P_s$. H was derived by matching slide pixels in the frame to the slide, so to find the coordinates P_f given P_s, we take the inverse, giving us $P_f = H^{-1}P_s$. With this, dimming the color channel of the pixel is finally straightforward. We scan every pixel in the slide and reduce the color value of the channel if its corresponding frame pixel is outside the bounding box.

4.3 Altering and Replacing Slides

To achieve slide background modification we went through color alteration of text and images to detect text and image regions as well as slide mapping to the video frame. We can combine those techniques and replace the entire slide in the video frame with the slide optimized for energy efficiency. To generate energy efficient slides we can either do them by hand or utilize previously described techniques to automatically change text and slide background while keeping images unaltered.

To maximize energy efficiency and readability, each slide has its background modified to black and the text to white. Figures, pie charts, etc., appear in their original colors and intensities since automatic processing is not able to recolor images and guarantee that they contain original information content. Once modified, the slide is backprojected into each frame it was used [25], after making the geometric manipulations necessary to account for size and geometry of the slide in the video frame. The comparison of before and after optimization is in Fig. 8.

While backprojection can significantly improve energy efficiency as well as readability by improving image quality, it can hide events that happen on the projector screen, such as the lecturer pointing (with hands or laser pointer), slide animations, and videos embedded in the slides. This is due to the backprojected slide overlaying the original slide in its proper location in the video frame. For example, if the speaker walks into the region of the slide, he or she would be hidden by the slide due to the fact that it is drawn on top of the video frame.

Fig. 8. A comparison of the original (right) with the backprojected slide (left).

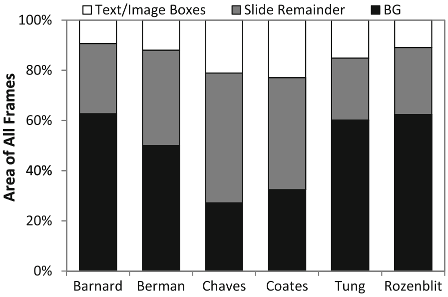

Fig. 9. Average area distributed among a slide, slide text and images, and the background areas.

Table 2. Video statistics

Video	Length [min]	Average intensity of		
		Red	Green	Blue
Barnard	4:00	86	101	105
Berman	3:15	111	139	132
Chaves	4:00	24	119	105
Coates	4:00	89	93	145
Tung	5:27	50	80	94
Rozenblit	3:15	120	114	114

5 Results

We evaluated the energy consumption of the proposed mechanisms using a National Instruments PCI-6230 Data Acquisition card (NI PCI-6230) and measuring software. The phone was connected directly to the power supply, instead of battery, for current measurements. To make the measurements reproducible, we turn off all wireless signals as well as automatic brightness adjustment. We evaluated our proposed mechanisms on the two phones, Samsung Galaxy S and Samsung Galaxy S2, over six lecture videos. We selected videos with varying lighting conditions, camera angles, and the amount of screen space that the lecture slide occupies versus background, as well as color mixes and Table 2 shows a basic statistics such selected video length and the average color composition in the video.

5.1 Efficiency of Altering the Non-Slide Background

Figure 9 shows the frame distribution between non-slide background (BG), background of the slide (Slide Remainder), and the text or image areas of the slide (Text/Image Boxes). The slide covers roughly 50 % of the total area in the video, which is not surprising as the slide has to occupy a significant portion of the screen so that the text is large enough to be legible. The figure also shows that the text and images only take up 15 % of the screen real estate on average, which means that only a fraction of the energy needed to play the original video is necessary to display the text and images.

Blacking out the non-slide background corresponds to removing the energy consumption from 50 % of the video frame, on average. However, the resulting energy savings are only 23.4 % and 19.9 % on the Galaxy S and S2, respectively. The modest savings can be accounted from the fact that the non-slide background is typically much darker that the slide area since the slides are displayed with a high power projector on a white screen and the room is dimmed to provide better slide readability.

5.2 Efficiency of Altering the Slide Background

Figure 11 shows progressive stages of dimming the slide background by reducing the RGB values by 20 %, 40 %, 60 %, and 80 %. The corresponding energy savings on Galaxy S2 are 20.0 %, 36.2 %, 44.2 %, and 50.1 % of the display energy usage. We observe that the energy savings come at the expense of alteration of the original image. While removing colors can make a lecture video seem unnatural, this particular transformation preserves the chromaticity of the original background but reduces the intensity, putting focus on the relevant part of the lecture, similar to [16–18].

Figures 10(a) and 10(b) show the average power demand of the entire phone for the combination of the display dimming techniques: video with all frames completely black (*All Black*); the video where only text and image regions are displayed (*Non-Text Black*); the video where we only remove blue, the most power hungry color, from the entire frame except the text and image regions (*Non-Text Red-Green*); and the original video without any alterations (*Full Color*). The upper bound on energy savings is attained by setting the color of the entire screen to black for the duration of the video. Playing such a video with all frames being black on the Galaxy S and S2 show the average upper bound on savings of 43.4 % and 46.4 % of the total energy needed to play the

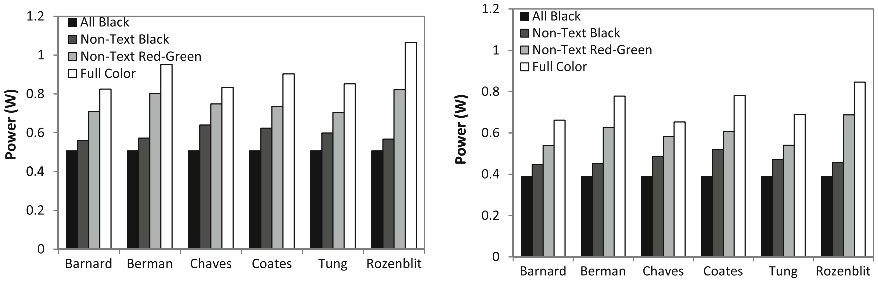

(a) Galaxy S (b) Galaxy S2

Fig. 10. Energy consumption after dimming less relevant regions.

Important Facts

The Beatles were an English rock band formed in Liverpool in 1960

- Influential albums include
  - *Rubber Soul* (1965),
  - *Revolver* (1966),
  - *Sgt Pepper's Lonely Hearts Club Band* (1967),
  - *The Beatles* (1968),
  - *Abbey Road* (1969).

Pythagoras' theorem

$a^2+b^2=c^2$

Fig. 11. Images whose slide background is dimmed by 20 %, 40 %, 60 %, and 80 % from left to right.

original video, respectively. For the remaining results, we will cite the savings in terms of the display energy, which is calculated by subtracting the energy to display a black video from the original video.

Non-Text Red-Green saves 37.1 % to 39.8 % for the Galaxy S and S2, respectively. The savings are significant, especially when considering the diversity of the color distribution in the videos. Table 2 shows the average RGB values of the 6 videos. In cases where the slide background is white, such as in *Rozenblit*, savings can increase to 19 % on the Galaxy S and 22 % on the Galaxy S2. The other factor in such significant savings is also due to the fact that some videos, such as *Rozenblit* and *Berman*, have long periods where no slide was shown.

While the blue channel requires the most power, red and green combined still constitute a significant portion of energy. The only exception to this is the *Coates* talk due to the fact that the slides shown have a blue background (Table 2). Dimming out all colors (*Non-Text Black*) from the slide and non-slide background can reduce energy consumption to levels close to an entirely black video (*All Black*). In general, dimming out colors from non-bounding boxes tends to save a lot of energy as bounding boxes take up only a small percentage of the screen real estate.

5.3 Efficiency of Altering and Replacing Slides

Figure 12 shows the average power demand for the Galaxy S2 and the combination of the video alteration techniques: the video with all frames completely black (*All Black*); the video where the slide region was replaced by the entirely black side and the non-slide background was unaltered (*Black Slide*); the video where the slide region was replaced by the energy optimized side and the non-slide background was unaltered (*Backprojected*); and the original video without any alterations (*Full Color*).

The *Black Slide* is the upper bounds on what can be saved by recoloring only the slide, leaving the non-slide background unaltered, and can offer 80.9 % on average of screen power demand reduction. Backprojection is very efficient offering 59.2 % reduction in average display power which translates to 27.4 % reduction in power demand of the entire phone. Further power reduction may be possible at the expense of lower color intensity of the text and images in the slide.

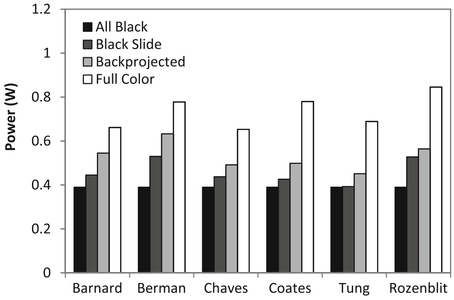

Fig. 12. Energy consumption of videos after slide backprojection on the Galaxy S2.

The efficiency of backprojecting white-on-black slides can be seen in the backprojected *Tung* video; which, despite the fact that the slide takes up only 41 % of the screen, uses only 9 % more energy than a totally black video. The savings are significant because the total area of the screen taken up by bullets is less than the 15 % that is taken up by the bounding boxes. Since the text was recolored rather than just keeping the original color of the whole region containing the text, the total area taken up by the text will have decreased as well. In other words, more pixels are blacked out and which increases the savings.

6 Conclusion

We have described a system that intelligently reduces energy consumption by replacing colors from less important regions of a lecture video. A key contribution is using slide-to-frame matching and slide semantic extraction to provide fine-grained understanding of content that can be exploited to enhance viewing and save energy. Subsequently, we have shown that removing colors from lecture videos of multiple lighting conditions is a viable method for saving a significant amount of energy consumed in mobile devices during playback. In addition, we have presented several methods to selectively remove varying degrees of different colors from portions of video frames. The resulting optimizations provided significant power reduction of displaying educational videos while minimizing the disruption of video quality by utilizing information about which areas of a video frame are the most informationally important.

Acknowledgments. This research was funded by the National Science Foundation under Grant No. 0844569 and the Microsoft Research Faculty grant.

References

1. Smith, A.: Mobile access 2010. Pew Internet & American Life Project, Washington, DC (2010)
2. MIT OpenCourseWare (2013). https://ocw.mit.edu/index.htm

3. UC Berkeley Extension Online (2013) https://learn.berkeley.edu
4. Open Yale Courses (2013). https://oyc.yale.edu
5. Coursera (2013). http://www.coursera.org
6. Udacity (2013). http://www.udacity.com
7. edX (2013). http://www.edx.org
8. Dong, M., Choi, Y., Zhong, L.: Power modeling of graphical user interfaces on oled displays. In: 46th ACM/IEEE Design Automation Conference, DAC'09, pp. 652–657 (2009)
9. Mahesri, A., Vardhan, V.: Power consumption breakdown on a modern laptop. In: Falsafi, B., VijayKumar, T.N. (eds.) PACS 2004. LNCS, vol. 3471, pp. 165–180. Springer, Heidelberg (2005)
10. Falaki, H., Govindan, R., Estrin, D.: Smart screen management on mobile phones (2009)
11. Flinn, J., Satyanarayanan, M.: Managing battery lifetime with energy-aware adaptation. ACM Trans. Comput. Syst. **22**(2), 137–179 (2004)
12. Park, S., Lee, Y., Lee, J., Shin, H.: Quality-adaptive requantization for low-energy mpeg-4 video decoding in mobile devices. IEEE Trans. Consum. Electron. **51**(3), 999–1005 (2005)
13. Kennedy, M., Venkataraman, H., Muntean, G.-M.: Energy consumption analysis and adaptive energy saving solutions for mobile device applications. In: Kim, J.H., Lee, M.J. (eds.) Green IT: Technologies and Applications, pp. 173–189. Springer, Heidelberg (2011)
14. Chen, X., Zheng, J., Chen, Y., Zhao, M., Xue, C.J.: Quality-retaining oled dynamic voltage scaling for video streaming applications on mobile devices. In: 2012 49th ACM/EDAC/IEEE Design Automation Conference, pp. 1000–1005. IEEE (2012)
15. Duan, L.-T., Guo, B., Shen, Y., Wang, Y., Zhang, W.L.: Energy analysis and prediction for applications on smartphones. J. Syst. Archit. **59**(10), 1375–1382 (2013)
16. Iyer, S., Luo, L., Mayo, R., Ranganathan, P.: Energy-adaptive display system designs for future mobile environments. In: MobiSys, pp. 245–258 (2003)
17. Harter, T., Vroegindeweij, S., Geelhoed, E., Manahan, M., Ranganathan, P.: Energy-aware user interfaces: an evaluation of user acceptance. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, pp. 199–206 (2004)
18. Wee, T.K., Balan, R.K.: Adaptive display power management for oled displays. In: Proceedings of the First ACM International Workshop on Mobile Gaming, pp. 25–30. ACM (2012)
19. Dong, M., Choi, Y., Zhong, L.: Power-saving color transformation of mobile graphical user interfaces on oled-based displays. In: Proceedings of the 14th ACM/IEEE International Symposium on Low Power Electronics and Design, pp. 339–342 (2009)
20. Chuang, J., Weiskopf, D., Möller, T.: Energy aware color sets. Comput. Graph. Forum **28**(2), 203–211 (2009). (Wiley Online Library)
21. Oled info (2011). www.oled-info.com/pentile
22. The SLIC browsing system (2013), http://slic.arizona.edu
23. Tung, Q., Swaminathan, R., Efrat, A., Barnard, K.: Expanding the point–automatic enlargement of presentation video elements. In: Assocation of Computing Machinery Multimedia, ACM MM (2011)

24. Fan, Q., Amir, A., Barnard, K., Swaminathan, R., Efrat, A.: Temporal modeling of slide change in presentation videos. In: IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2007, vol. 1, pp. I-989–I-992 (2007)
25. Winslow, A., Tung, Q., Fan, Q., Torkkola, J., Swaminathan, R., Barnard, K., Amir, A., Efrat, A., Gniady, C.: Studying on the move: enriched presentation video for mobile devices. In: 2nd IEEE Workshop on Mobile Video Delivery (2009)