



(51) International Patent Classification:

G06F 1/32 (2006.01) *G09G* 3/20 (2006.01)
H04W 4/02 (2009.01)

(21) International Application Number:

PCT/US2012/039905

(22) International Filing Date:

30 May 2012 (30.05.2012)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

13/151,401 2 June 2011 (02.06.2011) US

(71) Applicant (for all designated States except US): **INTEL CORPORATION** [US/US]; 2200 Mission College Boulevard, Santa Clara, California 95052 (US).

(72) Inventor; and

(75) Inventor/Applicant (for US only): **PARK, Hee-jun** [KR/US]; 3613 NW Talamore Ter., Portland, Oregon 97229 (US).(74) Agents: **RYDER, Douglas** et al.; RYDER IP LAW, c/o PortfolioIP, P.O. Box 52050, Minneapolis, Minnesota 55402 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished upon receipt of that report (Rule 48.2(g))

(54) Title: RESCHEDULING ACTIVE DISPLAY TASKS TO MINIMIZE OVERLAPPING WITH ACTIVE PLATFORM TASKS




FIG. 1

(57) Abstract: In general, in one aspect, a mobile device display includes panel electronics, a backlight driver and a rescheduler. The panel electronics is to generate images on an optical stack of the display based on input from a processing platform of the mobile device. The backlight driver is to control operation of a backlight used to illuminate the optical stack so that the user can see the images generated on the display. The rescheduler is to determine when a timing critical task of the processing platform overlaps with a non-timing critical task of the panel electronics or the backlight driver and reschedule the non-timing critical task until the timing critical task is inactive or a visual tolerance limit has been reached. The rescheduling minimizes overlap between the timing critical tasks and non-timing critical tasks and accordingly reduces power consumption without effecting performance or impacting a user's visual experience.

**RESCHEDULING ACTIVE DISPLAY TASKS TO
MINIMIZE OVERLAPPING WITH ACTIVE PLATFORM TASKS**

BACKGROUND

Mobile devices are continuing to increase in popularity and the processing capability continues to increase. Accordingly, more and more tasks are performed on mobile devices. In order to ensure the continued proliferation of mobile devices extending the life of the batteries used to power the mobile devices is a necessity. High current peaks (wide current fluctuations) for the mobile device significantly increases the power loss of the mobile device.

The significant power loss (P) is due to a squared relationship with current (I), where the power loss equals the product of the current squared and the resistance (R), such that $P = I^2R$. Both the internal resistance (cell impedance) of the battery and the resistance of the power delivery path from the battery to the mother board (e.g., circuits, connectors, circuit board) effect the power loss of the device. By way of example, Lithium (Li)-ion battery cells may have a cell impedance of approximately 70-350 mili-ohms (mΩ) and older batteries with many charge/discharge cycles have even higher cell impedance. The power delivery path may have a resistance of about 30 mΩ. The power loss reduces the actual amount of usable energy from the battery for the mobile device (actual energy from the battery is less than the specified battery capacity).

FIG. 1 illustrates an example graph of the relationship between system power consumption and power loss. The power loss is associated with a battery powering the system and the power delivery path. The relationship is not linear due to the squared relationship between power loss and current. As illustrated, when the system power increases three fold from 5 watts (W) to 15W the power loss increases ten fold from 0.3W to 3W.

High current peaks may also temporarily cause deeper battery voltage drops, and if the temporary voltage drops go below a defined cut-off voltage, the system may shut-down prematurely. Furthermore, the trend for some processors is to enter an idle state quicker to reduce power consumption. This trend may decrease platform power but may increase the peak currents further which will result in an increased power loss and reduction in the actual amount of usable energy from the battery.

BRIEF DESCRIPTION OF THE DRAWINGS

The features and advantages of the various embodiments will become apparent from the following detailed description in which:

FIG. 1 illustrates an example graph of the relationship between system power
5 consumption and battery power loss;

FIG. 2 illustrates a high level block diagram of an example mobile device;

FIG. 3 illustrates an example timing diagram showing overlapping active
periods between the platform and the display;

FIG. 4 illustrates an example flow chart of operations for rescheduling the
10 display device active periods when the platform is active, according to one
embodiment;

FIG. 5 illustrates an example timing diagram showing the rescheduling of
display device active periods (delay in the overall display task cycle), according to
one embodiment;

15 FIG. 6 illustrates example timing diagrams showing the rescheduling of a
display device active period (no impact on overall display task cycle), according to
one embodiment;

FIG. 7 illustrates an example flow chart of operations for halting the display
device active periods when a platform active period is detected, according to one
20 embodiment;

FIG. 8 illustrates example timing diagrams showing the halting of a display
device active period, according to one embodiment;

FIG. 9 illustrates an example timing diagram showing rescheduling of
backlight active and image refresh active periods to minimize conflicts with the
25 platform active periods, according to one embodiment; and

FIG. 10 illustrates an example timing diagram showing rescheduling of
backlight active and image refresh active periods to minimize conflicts with the
platform active periods, according to one embodiment.

DETAILED DESCRIPTION

30 **FIG. 2** illustrates a high level block diagram of an example mobile device
200. The mobile device **200** may include a processing platform **210**, a display **240**
and a battery **260**. The processing platform **210** may include a main processor (CPU)
215, a graphics processor **220**, memory **225** and input/output (I/O) **230**. The display

240 may include panel electronics **245** and a backlight **250**. The battery **260** provides power to both the processing platform **210** and the display **240**. The power loss in the device **200** may be based on the cell impedance **270** in the battery **260**, the resistance in the connectors **275** (battery **260**, platform **210**, and display **240**), the
5 resistance in the power delivery means **280**, and the resistance in the power rails on the printed circuit boards **285** (platform **210**, display **240**).

The platform **210** and the display **240** may be active at the same time during video playback on the mobile device **200**. The platform **210** (CPU **215**, graphics **220**, memory **225** and/or I/O **230**) may receive and process data related to a video in order
10 to present the video to the display **240**. The panel electronics **245** may generate images on the optical stack (not illustrated) of the display **240** based on input from the platform **210**. A backlight driver (not illustrated) may control operation of the backlight **250** to illuminate the optical stack so that the user can see the images generated on the display **240**.

15 During video playback, the platform **210** (CPU **215**, graphics **220**, memory **225** and/or I/O **230**) may be active approximately 50 to 100 times per second, the panel electronics **245** may be active to refresh the image presented on the display **240** approximately 50 to 60 times/sec, and the backlight **250** may be active to illuminate the image presented on the display approximately 200 times/sec. When both the
20 platform **210** and the display **240** are active at the same time they are both drawing current and the increased current increases the power loss. The increased power loss may result in the battery needing to be recharged more often, the increased recharging may increase the cell impedance of the battery **260**, and the increased cell impedance may reduce the life of the battery **260**.

25 The activities performed by the device **200** can be divided into timing critical activities and visual experience (non-timing) critical activities. Timing critical activities are those that can not be delayed without impacting the operation of the device **200**. Visual experience critical activities are those that may be delayed up to several milliseconds (msec) without impacting a user's visual experience as human
30 eyes are not sensitive enough to detect such a delay. The timing critical activities are those performed by the platform **210** (CPU **215**, graphics **220**, memory **225** and/or I/O **230**) while the visual experience critical activities are those performed by the

display **240** (panel electronics **245**, backlight **250**). Accordingly, there may be an opportunity to reschedule some of the visual experience critical activities (display activities) to reduce the device's power peak.

FIG. 3 illustrates an example timing diagram showing overlapping active periods between the platform and the display (backlight, image refresh). The platform is active when processing is required and the duration of the active periods and the time between active periods may vary. The backlight may alternate between on and off at a fast rate, pulse width modulation (PWM). The flickering on/off of the backlight can not be detected by the human eye. Rather, the amount of time that the backlight is on controls the brightness of the display. By way of example, in **FIG. 3** it is assumed that the backlight is on for 1 msec and off for 2 msec every cycle. The panel electronics may be active when an image is being generated on the display (image refresh) and then may be inactive for a certain period, vertical blanking interval (VBI), while the image is presented before a new image is generated. By way of example, in **FIG. 3** it is assumed that the panel electronics are on for 3 msec and off for 6 msec every cycle.

As illustrated, there are three overlapping active periods **300, 320, 340** (illustrated as shaded regions) between the platform and the backlight and/or panel electronics. The first overlapping period **300** occurs during the first platform active period **305**. The platform is active for the start of both a backlight active period **310** and an image refresh period **315**. The first platform active period **305** lasts the entire backlight active period **310** and a portion of the image refresh active period **315**. The second overlapping period **320** occurs during the second platform active period **325**. The start of the platform active period **325** is aligned with the start of both a backlight active period **330** and a panel electronics active period **335**. The second platform active period **325** lasts the entire backlight active period **330** and the entire image refresh active period **335**. The third overlapping period **340** occurs during the fourth platform active period **345**. The platform is active for the start of a backlight active period **350** and remains active through the entire backlight active period **350** and a portion of a next backlight active period **355**. A panel electronics active period **360** is occurring when the fourth platform active period **345** begins and the fourth active period **345** remains active after the image refresh active period **360** ends.

As the display active periods (backlight, panel electronics) are not timing critical it may be possible to adjust the timing of the display active periods to minimize (reduce and possibly avoid) the overlap with the performance critical platform active periods. If the adjustments to the display active periods are limited to adjustments not detected by the human eye (several msec) the adjustments will not adversely affect the display quality and the user will be unaware of any changes. By intelligently rescheduling the display device active periods I^2R power loss may be reduced and battery life may be preserved.

FIG. 4 illustrates an example flow chart of operations for rescheduling the non-timing critical display device active periods (backlight, image refresh) around the timing critical platform active periods. The operations delay the initiation of the non-timing critical display device active periods when the timing critical platform operations are active. Initially the display prepares for initiation of a display task **400**. The display task may be turning on the backlight or may be processing an image refresh. Once the task is initiated a visual (non-timing critical) tolerance timer is reset and started **410**. The visual tolerance timer is the amount of time between when a display event is initiated (or scheduled) and when it is actually processed. A determination is then made as to whether the platform is active at that point **420**. If the platform is not active (**420 No**), the display task is processed **440**.

If the platform is determined to be active (**420 Yes**), a determination is made as to whether the timer has reached a visual tolerance limit **430**. The visual tolerance limit is amount of time that a display event can be delayed without impacting the display quality and/or users visual experience. The visual tolerance limit may be the same for the backlight and the image refresh or may be different. The visual tolerance limit may vary based on the quality of what is being displayed. The visual tolerance limit may be programmable. If the timer has reached the limit (**430 Yes**), the display task is processed **440**. If the timer has not yet reached the limit (**430 No**), the operations return to determining if the platform is active **420**. This loop **420, 430** results in the processing of the display task being delayed while waiting for the processor to become inactive (**420 No**) or the limit to be reached (**430 Yes**). Once the display task has been processed **440** it will again wait for the initiation (scheduling) of the next display task **400**.

The rescheduling of display device active periods is not limited to the operations or the order of operations described and illustrated in FIG. 4. Rather, the order of operations could be modified, various operations may be modified, additional operations may be added, operations may be removed, operations may be combined and/or operations may be separated without departing from the current scope. For example, the visual tolerance limit may be the amount of time between the completion of one display device active period and the start of the next display device active period and the timer may be reset and started at the conclusion of each active period.

The rescheduling (delay) of a display device active period to avoid overlapping with a platform active period may delay the overall cycle for the display task (backlight, image refresh). **FIG. 5** illustrates an example timing diagram showing the rescheduling of display device active periods (delay in the overall display task cycle). The timing diagram illustrated is for a backlight having a PWM where the backlight is on for 1 msec and then off for 2 msec. The top diagram illustrates the PWM of the backlight prior to any rescheduling. The diagram starts at time zero and the backlight is on for the first (from 0 to 1), fourth, seventh and tenth msec. The bottom diagram assumes that prior to initiation of the first backlight active period it was determined that the platform was active. In order to limit (and possibly avoid) an overlap with the platform active period the initiation of the first backlight active period as well as the overall backlight cycle and all the active periods may be delayed by 2 msec. Accordingly, the backlight is active during the third, sixth, ninth and twelfth msec.

Alternatively, the rescheduling (delay) of a display device active period (or several display device active periods) to avoid overlapping with a platform active period may not impact the overall display task cycle. If the overall backlight on-time ratio is maintained with the rescheduling, the device user may not notice any brightness change as the human eye is not sensitive enough to detect msec periods of brightness.

FIG. 6 illustrates example timing diagrams showing the rescheduling of a display device active period (no impact on overall display task cycle). The timing diagram illustrated is for a backlight having a PWM where the backlight is on for 1 msec and then off for 2 msec. The top diagram illustrates the PWM of the backlight

prior to any rescheduling. The second diagram shows a delay of 1 msec in the first active period so the first active period occurs during the second msec and all other active periods remain the same.

The third diagram shows a delay of 2 msec in the first active period so the first active period occurs during the third msec and is combined with the original second active period that occurs during the fourth msec so that the backlight is active during the third and fourth msec. The bottom diagram shows a delay of 3 msec in the first active period so the first active period would occur during the fourth msec which is when an active period is already scheduled for. Accordingly, the first active period may occur during the fourth msec and the originally scheduled second active period may be delayed 1 msec and occur during the fifth msec. Alternatively, the first active period may be delayed an additional 1 msec and occur during the fifth msec. In either event, the backlight may be active during the fourth and fifth msec.

FIG. 7 illustrates an example flow chart of operations for rescheduling the non-timing critical display device active periods (backlight, image refresh) around the timing critical platform active periods. The operations halt the non-timing critical display device active periods when the timing critical platform operations become active. Initially the display begins processing a display task **700** (backlight, image refresh). While the display task is being processed (during the display device active period), a determination is made as to whether the platform is active **710**. If the platform is not active (**710 No**), the display task continues to be processed **700**. If the platform is determined to be active (**710 Yes**), the processing of the display task is halted **720** and a visual tolerance timer is reset and started **730**. A determination is then made as to whether the platform is active **740**. If the platform is not active (**740 No**), the processing of the display task is restarted **700**. It should be noted that the display task may be restarted from the beginning (e.g., begin the image refresh again) or it may begin where it left off (e.g., image refresh continues from where it left off, may require register to record position left off).

If the platform is determined to be active (**740 Yes**), a determination is made as to whether the timer has reached a visual tolerance limit **750**. If the timer has reached the limit (**750 Yes**), the processing of the display task is restarted **700**. If the timer has not yet reached the limit (**750 No**), the operations return to determining if

the platform is active **740**. This loop **740, 750** results in the processing of the display task being delayed while waiting for the processor to become inactive (**740 No**) or the limit to be reached (**750 Yes**).

The halting of the display tasks is not limited to the operations or the order of
5 operations described and illustrated in FIG. 7. Rather, the order of operations could be modified, various operations may be modified, additional operations may be added, operations may be removed, operations may be combined and/or operations may be separated without departing from the current scope.

FIG. 8 illustrates example timing diagrams showing the halting of a display
10 device active period. The timing diagram illustrated is for an image refresh where the display is active for 3 msec while the image is generated and then is off for 6 msec before the next image is generated. The top diagram illustrates the image refresh cycle prior to any rescheduling. The diagram starts at time zero and the image refresh is active for the first through third (from 0 to 3) and ninth through
15 twelfth msec. The bottom two diagrams assume that during the active refresh a platform active is detected and interrupts the image refresh active for 2 msec (either that is the duration of the platform active or that is the visual tolerance limit).

In the second diagram the active image refresh process is halted after 2 msec, and then after a 2 msec delay the whole active image refresh is processed over the
20 fifth through seventh msec. In the third diagram the active image refresh process is halted after 2 msec, and then after a 2 msec delay the image refresh process continues where it left off during the fifth msec. It should be noted that as illustrated the overall image refresh cycle was not effected, that the second image refresh active portion still occurs between the tenth and twelfth msec. This was done for ease of
25 illustration and in no way is intended to limit the scope thereto. Rather, the overall display task (image refresh) cycle could be delayed each time there is a need to reschedule a display task to avoid an overlap with platform active periods or certain times that may be determined based on various parameters (the parameters may be programmable).

FIG. 9 illustrates an example timing diagram showing rescheduling of display
30 device (backlight and image refresh) active periods to minimize conflicts with the platform active periods. As illustrated, when a display device active period is delayed it delays the entire display task cycle from that point (as illustrated in FIG.

5). As the display device active periods are being shifted each platform active period is highlighted **900, 920, 940, 960** to determine if shifting caused any new overlaps.

As illustrated when a display device active period is halted, it continues from the point it left off at when processing is resumed (as illustrated in bottom diagram of

5 FIG. 8).

As illustrated, the backlight active period **310** and the entire backlight cycle from that point forward have been shifted (delayed) 2 msec. The shifted backlight active period **910** does not occur during the first platform active period **900**.

Backlight active period **930A** (which is delayed 2 msec from backlight active period **330** of FIG. 3) and the entire backlight cycle from that point forward have been
10 delayed 1 msec so the shifted backlight active period **930B** does not occur during the second platform active period **920**. Backlight active period **975A** (which is aligned with backlight active period **350** of FIG. 3 because there was a total delay of a 3 msec cycle) and the entire backlight cycle from that point forward have been delayed
15 3 msec (assuming 3 msec is the maximum visual tolerance limit) so the shifted backlight active period **975B** only slightly overlaps with the fourth platform active period **970**.

As illustrated, the image refresh active period **315** and the entire image refresh cycle from that point forward have been shifted (delayed) 2 msec. The
20 shifted image refresh active period **915** does not occur during the first platform active period **900**. Image refresh active period **935A** (which is delayed 2 msec from image refresh active period **335** of FIG. 3) and the entire image refresh cycle from that point forward have been delayed 1 msec so the shifted image refresh active period **935B** does not occur during the second platform active period **920**. Image
25 refresh active period **945A** (which is delayed 3 msec from the third image refresh active period, which was not numbered, of FIG. 3) is occurring when the third platform active period **940** begins. The image refresh active period **945A** is halted for 2 msec so as not to overlap the third platform active period **940**. When processing is restarted the processing of the image refresh continues from where it
30 left off (last 2 msec) **945B**. The result is that the entire image refresh cycle from that point forward has been delayed 2 msec. Image refresh active period **980A** (which is delayed 5 msec from the image refresh active period **360** of FIG. 3) and the entire

image refresh cycle from that point forward have been delayed 1.5 msec so the shifted image refresh active period **980B** does not overlap with the fourth platform active period **970**.

FIG. 10 illustrates an example timing diagram showing rescheduling of display device (backlight and image refresh) active periods to minimize conflicts with the platform active periods. As illustrated, when a display device active period is delayed it delays only that active period and not the entire display task cycle (as illustrated in FIG. 6). As illustrated when a display device active period is halted, it starts over when processing is resumed (as illustrated in middle diagram of FIG. 8). The platform active periods, the non-modified display event active periods, and the overlapping active periods are the same as those illustrated in FIG. 3 and are accordingly numbered the same.

As illustrated, the backlight active period **310** has been shifted (delayed) 2 msec to avoid the first overlapping active period **300**. The shifting 2 msec means that two backlight active periods will abut to produce one 2 msec backlight active period **1010**. The backlight active period **330** has been delayed 3 msec (or possibly 4 msec if delayed until after current backlight display period) to avoid the second overlapping active period **320**. The shifting 3 msec means that two backlight active periods will abut to produce one 2 msec backlight active period **1030**. The backlight active period **350** have been delayed 3 msec (assuming 3 msec is the maximum visual tolerance limit) to limit the overlap with the platform active period **345** (minimize the third overlapping active period **340**). The shifting 3 msec means that two backlight active periods will abut to produce one 2 msec backlight active period **1050**.

As illustrated, the image refresh active period **315** has been shifted (delayed) 2 msec to avoid the first overlapping active period **300**. The image refresh active period **335** has been delayed 3 msec to avoid the second overlapping active period **320**. The image refresh active period **360** is occurring when the platform active period **345** begins. The image refresh active period **360** is halted for 4 msec (assuming 4 msec is the maximum visual tolerance limit) to limit the overlap with the platform active period **345** (minimize the third overlapping active period **340**). When processing is restarted the processing of the image refresh starts at the

beginning (active for all 3 msec even though 2 msec were processed prior to halting) **1060**. The result of the delay of 4 msec and the reprocessing of 2 msec is that two image refresh active periods will abut to produce one 6 msec image refresh active period **1060**.

- 5 Referring back to FIG. 2, the rescheduling operations described in the flow charts of FIGs. 4 and 7 may be implemented in one or more of the devices of the platform **210** (CPU **215**, graphics **220**, memory **225**, I/O **230**), a discrete component on the platform **210**, one or more devices of the display **240** (panel electronics **245**, backlight diver), a discrete component on the display **240**, or some combination thereof.
- 10 Furthermore, the operations may be implemented in hardware, software, firmware or some combination thereof. The various devices (CPU **215**, graphics **220**, memory **225**, I/O **230**, panel electronics **245**, backlight diver) may have access to device (e.g., processor) readable storage (on the device, off the device, or some combination thereof) that contains instructions that when executed by the device causes the device
- 15 to perform at least a subset of the operations described above in FIGs. 4 and 7.

The devices may include registers that track whether the device is active or inactive. The status maintained in each register may be provided to and/or accessible by each of the other devices. Accordingly, the various devices may know the active status of the other devices and may take necessary actions to minimize the

20 overlapping of platform (timing critical) and display (visual experience critical) active periods.

- For example, the CPU **215** may determine that the backlight driver is active at the same time as the CPU **215** (and/or one or more of the other devices of the platform **210**) and may instruct the backlight driver to halt the active period. The
- 25 CPU **215** may instruct the backlight driver upon completion of the platform active period so that the backlight driver can continue operations. The CPU **215** may maintain the timer and determine when the visual tolerance limit has been met and instruct the backlight driver. Alternatively, the backlight driver may maintain the timer and determine when the visual tolerance limit has been met and restart
- 30 operations at that point even though the CPU has not indicated the platform active period is complete.

By way of another example, the panel electronics **245** may determine that one or more of the devices on the platform **210** is active at the same time it is active and may halt the image refresh active period and start the timer. The panel electronics **245** may restart the image refresh either after the platform active period is concluded
5 or after visual tolerance limit has been met. The panel electronics may write where the image refresh was when it was halted so that upon re-initiation it may continue from the same place.

The various embodiments described above were described with respect to being incorporated within a mobile device where reducing power loss and thus increasing
10 amount of usable energy from the battery is vital. However, the various embodiments are not limited to mobile devices. Rather, they could be implemented in any electronic device having a display as the implementation would reduce power loss and thus overall power consumed (conserve energy). For example, the various embodiments may be incorporated into a personal computer (PC), laptop computer, ultra-laptop computer, tablet, touch pad,
15 portable computer, handheld computer, palmtop computer, personal digital assistant (PDA), cellular telephone, combination cellular telephone/PDA, television, smart device (e.g., smart phone, smart tablet or smart television), mobile internet device (MID), messaging device, data communication device, and so forth.

Although the disclosure has been illustrated by reference to specific embodiments,
20 it will be apparent that the disclosure is not limited thereto as various changes and modifications may be made thereto without departing from the scope. Reference to “one embodiment” or “an embodiment” means that a particular feature, structure or characteristic described therein is included in at least one embodiment. Thus, the appearances of the phrase “in one embodiment” or “in an embodiment” appearing in
25 various places throughout the specification are not necessarily all referring to the same embodiment.

The various embodiments are intended to be protected broadly within the spirit and scope of the appended claims.

CLAIMS

What is claimed:

1. A method comprising
determining, in a mobile device, when a timing critical task is active at same
5 time as a non-timing critical task is active; and
rescheduling the non-timing critical task until the timing critical task is
inactive or a threshold time associated with the rescheduling of the non-timing
critical task has been reached, wherein the rescheduling of the non-timing critical
task minimizes overlap between the timing critical tasks and non-timing critical tasks
10 without effecting performance of the mobile device, and wherein the minimizing of
the overlap reduces power consumption and power loss of the mobile device.
2. The method of claim 1, further comprising refreshing and starting a
timer and comparing the timer to the threshold time.
3. The method of claim 2, wherein the timer is refreshed and started when
15 a previous non-timing critical activity is completed.
4. The method of claim 2, wherein the timer is refreshed and started when
the non-timing critical activity is initiated.
5. The method of claim 2, wherein the timer is refreshed and started when
the non-timing critical activity is rescheduled.
- 20 6. The method of claim 1, wherein the rescheduling includes delaying
start of the non-timing critical task if the determining determines that the timing
critical task is active when the non-timing critical task is initiated.
7. The method of claim 1, wherein the rescheduling includes halting the
non-timing critical task if the determining determines that the timing critical task has
25 become active during active operation of the non-timing critical task.
8. The method of claim 7, wherein the rescheduling includes resuming
the non-timing critical task from where it was halted.
9. The method of claim 7, wherein the rescheduling includes resuming
the non-timing critical task from its beginning.
- 30 10. The method of claim 1, wherein the non-timing critical tasks are visual
experience critical tasks and the threshold time is time that the visual experience
critical tasks may be delayed without impacting a user's visual experience.
11. A mobile computing device comprising

one or more processors to perform timing critical tasks;
a display to present information to a user;
one or more circuits to control operation of the display, wherein tasks performed by the one or more circuits are non-timing critical and may be delayed up
5 to a visual tolerance limit without impacting a user's visual experience; and
a rescheduler to
determine when a timing critical task of the one or more processors overlaps with a non-timing critical task of the one or more circuits; and
reschedule the non-timing critical task until the timing critical task is
10 inactive or the visual tolerance limit has been reached, wherein the rescheduling of the non-timing critical task minimizes overlap between the timing critical tasks and non-timing critical tasks without effecting performance of the mobile device, and wherein the minimizing of the overlap reduces power consumption and power loss of the mobile device.

15 12. The mobile computing device of claim 11, wherein the rescheduler is further to refresh and start a timer and compare the timer to the visual tolerance limit.

13. The mobile computing device of claim 11, wherein the rescheduler is to delay start of the non-timing critical task if the timing critical task is active when
20 the non-timing critical task is initiated.

14. The mobile computing device of claim 11, wherein the rescheduler is to temporarily halt the non-timing critical task if the timing critical task has become active during active operation of the non-timing critical task.

15. The mobile computing device of claim 11, wherein the rescheduler is
25 implemented in the one or more processors.

16. The mobile computing device of claim 11, wherein the rescheduler is implemented in the one or more circuits.

17. The mobile computing device of claim 11, wherein the rescheduler is implemented in a discrete component.

30 18. A mobile device display comprising
panel electronics to generate images on an optical stack of the display based on input from a processing platform of the mobile device;

a backlight driver to control operation of a backlight used to illuminate the optical stack so that the user can see the images generated on the display; and

a rescheduler to

determine when a timing critical task of the processing platform

5 overlaps with a non-timing critical task of the panel electronics or the backlight driver; and



reschedule the non-timing critical task until the timing critical task is inactive or a visual tolerance limit has been reached, wherein the rescheduling of the non-timing critical task minimizes overlap between the timing critical tasks and non-
10 timing critical tasks without effecting performance of the mobile device or impacting a user's visual experience, and wherein the minimizing of the overlap reduces power consumption and power loss of the mobile device.

19. The mobile device display of claim 18, wherein the rescheduler is implemented in the panel electronics and the backlight driver.

15 20. The mobile computing device of claim 18, wherein the rescheduler is a discrete component.

1/5

Non-Linear Relationship Between Power Loss (battery and power delivery) and System Power

**FIG. 1****FIG. 2**

2/5




FIG. 3




FIG. 4

3/5




FIG. 5




FIG. 6

4/5




FIG. 7




FIG. 8

5/5




FIG. 9




FIG. 10