

How to Power Tune a Device Running on a Linux Kernel for Better Suspend Battery Life

Mark.gross@intel.com

ELC2011

Introduction

- Based on a true story!
- Started with 2 weeks of focused effort hardening suspend to ram and use cases.
- Ended with 1 extra week at the ODM re-working hardware and doing more measurements
- Reduced suspend power from “way-too-high” to “not-too-bad”
- Identified HW issues to be fixed
- ASK QUESTIONS!!!
 - ~30min of content

Acknowledgments

- Veneesh K K (“destroyer of boards”)
 - Experimental power measurement to isolate what device and driver is causing higher power loads
- German Monroy (“designated adult”)
 - Drove priorities:
 - Disciplined tracking of issues and power.
- A number of others:
 - Vishwish, Youvdeep, Yanmin, Fengwei, Jack, Aswhar, and others participated in the activity

Outline

- The power BOM
- Suspend to Ram Hardening
- Measuring
- Locating Easy leakages
- Locating Harder leakages
- Getting line of site to spec power

Power BOM-1

- Start with schematic
- ID all peripherals
 - Spec sheets
 - Driver source code
 - Data bus
 - Voltage rails
 - Connectors
 - High, typical, low power ratings

Power BOM-2

- Understand the SOC
 - IP blocks
 - Voltage rails
 - Data buses
- Understand clock and power controls
 - What rails can be gated and how
 - What clocks can be gated
- Understand what platform states are available

Power BOM-3

- Derive board level power target from BOM
 - for this board it was 25mW
- Have an idea what rails have what power going through them
 - Know what rails can be gated from SW.
- Estimate power per connector
- Look for obvious leakages
 - Locate extra LDO's
 - LED's
 - Strong pull up / down resistors

Suspend to RAM Hardening

- Review drivers for suspend/resume functions
 - Check specs for going into and coming out of low power states.
- Turn on linux pm debugging
- Follow Documentation/power/basic-pm-debugging.txt
 - Use `sys/power/pm_test`
 - Use `sys/power/pm_trace`
- First get suspend and resume working.
- Second make sure use cases still work after resume.
 - This is hard.
- Stress test

Measuring

- At battery connectors
 - Battery simulators
 - Monsoon device
- Special instrumented measurement boards
 - Only available for reference hardware.
 - Didn't have one for this device so we had to “rough it”

First result

- ~685mW !
 - Yipes!

Started unplugging connectors to isolate power hogs

- Sensor hub (245mW)
- Cameras (30mW)
- Display (20mW)
- Touch Screen (10mW)
- Touch pad(20mW)
- Extra display(30mW)

Power gating rails to isolate more:

- Modem (100mW)
- BT (25mW)

Start fixing:

- New Sensor hub FW and driver changes
- Camera driver updates
- Display update
- Back light PWM gaiting
- Touch screen now held in reset
- Extra display driver off/on
- Power gated modem and BT

Second measurement

- 120mW board level
 - Still looking for 25-ish
- Found devices not going to low power
 - Touch screen was still hot
 - Few more driver updates
 - FW update
 - Found a few more clocks needing gating

How to know when done?

- For each connector that can be unplugged while in suspend state:
- For each power rail that can be gated from SW:
- For each peripheral you can pull out while in suspend:
 - power meter shows expected delta after.

3rd measurement

- 80mW from unmodified hardware
 - Pretty good for an early spin of a device that has a bright power LED always on
- Still have higher than expected leakage from 2 connectors (10mw)
- LED (20mw)
- 25mW of unexpected leakage yet to find in board.

Getting Hard Core

- Reworking device to measure and find floor.
 - Do this because we don't have an instrumented platform
 - Need to keep the device booting.
 - Yank parts off device
 - Cut power rails with exact-o-knife
 - Keep in mind what you end up with is no longer the gadget you started with and results are in some ways “academic”
- At the end we saw 40mW
 - We expected <10.
 - Filed a bug for the 30mW with team responsible for SOC
 - Filed bug with board layout team for the 30 MW

Random data

- We had 8 working devices going into power camp.
 - There was a variance of $\sim 20\text{mW}$ suspend power measured across all of them.
 - All no longer fully functional
- Don't know what expected variances should be.
 - Plans to measure units coming off the line are in place
- Numbers most times never added up all the way.
 - Source of frustration when looking for those last few mW
 - Contributed by per-device variances confusing measurements.
 - Maybe board level power is not as simple as I want to think.

Concluding remarks

- Went from 685mW to < 80mW through SW and FW updates alone
- At the end we had a small set of issues to get fixed for the next HW spin:
 - Identified HW mods needed to save 20mW (LED)
 - Identified 10 mW of leakage with touch pad and prox light sensor
 - Identified 25 mW higher than expected use from SOC + board
- Line of sight to power target believed achieved
- I had fun and I'd do it again.
- Need to do another pass on next spin of hardware.

Questions?