

ECE 571 – Advanced Microprocessor-Based Design Lecture 21

Vince Weaver

<http://www.eece.maine.edu/~vweaver>

vincent.weaver@maine.edu

9 April 2013

Project/HW Reminder

- Homework #4 comments
 - Good job finding references, found some new ones
 - The best battery life question was too vague, as is the case with most Energy/Power questions there are a lot of variables that need to be considered.
 - The source of the thermal error on the laptop was most likely due to the fact that the Linux version being run does not have a proper fan driver, so the processor



was overheating due to the fan not kicking in.

- Project Notes

- If you are willing to present your project on the earlier presentation date (April 30) let me know, otherwise I'll have to randomly assign people to go first.
- Homework 5 will be coming soon. Will be like HW4 and include Energy question and some project work.
- Cortex A9 L2 Prefetching enable seems to work now
- Cortex A9 L2 cache performance counters turn out to be available, but are in the cache controller so would



- require special Linux support to access.
- An Intel x86 machine is available for use (a mac-mini). I will set up accounts for anyone who needs one.
 - If you happen to enjoy this type of low-level research (or know someone that does) and you're looking for a professor to work for, let me know. I do have some startup money to use up.

- WattsUpPro

- A few groups want to use the WattsUpPro so we will have to schedule access.



- Using the meter is straight-forward but can be a bit tricky. Sometimes the program that reads the values can lose sync and it takes a few tries to get going.
- On a Linux machine, the output is just text, printed once a second, with various value (Watts is first).



Video Card Digression

- There were some questions last class about exactly how video cards work today, and some of their power implications.
- This is important; embedded systems like phones come with GPUs these days.



Old CRT Days

- Electron gun
- Horizontal Blank, Vertical Blank
- Atari 2600 – only enough RAM to do one scanline at a time
- Apple II – video on alternate cycles, refresh RAM for free
- Bandwidth key issue. SNES / NES, tiles. Double



buffering vs only updating during refresh



Old 2D Video Cards

- Framebuffer (possibly multi-plane), Palette
- Dual-ported RAM, RAMDAC (Digital-Analog Converter)
- Interface (on PC) various io ports and a 64kB RAM window
- Mode 13h
- Often commands for drawing lines, rectangles, blitting sprites, mouse cursors, video overlay



GPUs

- Display memory often broken up into tiles (improves cache locality)
- Massively parallel matrix-processing CPUs that write to the frame buffer (or can be used for calculation)
- Texture control, 3d state, vectors
- Front-buffer (written out), Back Buffer (being rendered)
Z-buffer (depth)



- Originally just did lighting and triangle calculations. Now shader languages and fully generic processing



Video RAM

- VRAM – dual ported. Could read out full 1024Bit line and latch for drawing, previously most would be discarded (cache line read)
- GDDR3/4/5 – traditional one-port RAM. More overhead, but things are fast enough these days it is worth it.
- Confusing naming, GDDR3 is equivalent of DDR2 but with some speed optimization and lower voltage (so



higher frequency)



Busses

- DDC – i2c bus connection to monitor, giving screen size, timing info, etc.
- PCIe (PCI-Express) – most common bus used in x86 systems
 - Original PCI and PCI-X was 32/64-bit parallel bus
 - PCIe is a serial bus, sends packets
 - Can power 25W, additional power connectors to supply can have 75W, 150W and more
 - Can transfer 8GT/s (giga-transfers) a second



In general PCIe is the main limiting factor to getting data to GPU.



Connectors

CRTC (CRT Controller) Can point to same part of memory (mirror) or different.

- RCA – composite/analog TV
- VGA – 15 pin, analog
- DVI – digital and/or analog. DVI-D, DVD-I, DVD-A
- HDMI – compatible with DVI (though content restrictions). Also audio. HDMI 1.0 – 165MHz, 1080p



or 1920x1200 at 60Hz. TMDS differential signalling. Packets. Audio sent during blanking.

- Display Port – similar but not the same as HDMI
- Thunderbolt – combines PCIe and DisplayPort. Intel/Apple. Originally optical, but also Copper. Can send 10W of power.
- LVDS – Low Voltage Differential Signaling – used to connect laptop LCD



LCD Displays


- Crystals twist in presence of electric field
- Asymmetric on/off times
- Passive (crossing wires) vs Active (Transistor at each pixel)
- Passive have to be refreshed constantly
- Use only 10% of power of equivalent CRT



- Circuitry inside to scale image and other post-processing
- Need to be refreshed periodically to keep their image
- New “bistable” display under development, requires not power to hold state



Power Saving Strategies



big.LITTLE / Heterogeneous Computing

- ARM
- big = Cortex A15 = power hungry, fast, high-leakage
- little = Cortex A7 = low power, slow
- “big.LITTLE switcher” by Pitre. have 1:1, move from slow to fast when need the speed
- have all procs visible to Linux, schedule them with intelligent scheduler



- Can use cpufreq interface, “big” just seen as higher frequency operating point



Race to Idle

- Good strategy on high-leakage chips (Intel?)
- Depends on how CPU bound process is
- Example 1:
 - If 34W full speed, 24W half speed, 1W idle, total time 1s
 - 1s at half speed, $24W * 1s = 24J$
 - 0.5s at full speed, 0.5s at idle: $34W * 0.5 + 1W * 0.5 = 17.5J$



- Example 2:

- Instead, 34W full speed, 24W half, 20W idle

- 1s at half speed, $24W * 1s = 24J$

- 0.5s at full speed, 0.5s at idle: $34W * 0.5s + 20W * 0.5s = 27J$



Adagio

- “Adagio: Making DVS Practical for Complex HPC Applications”, Rountree et al. ICS 2009.
- For HPC MPI workloads
- Predictor predicts when to DVS in realtime. Unlike other static predictors that base decision on old traces.
- 8-20% Energy reduction with 1% overhead
- Critical Path Analysis, want jobs to finish “just in time”



- Use hardware performance counters. Run at full speed first time, predict subsequent calls will be identical
- Approximate intermediate frequencies by running part at higher and part at lower
- 117-180W range, so best energy improvement = 39%
- “When *Not* to Race to Idle: When to Use (and Avoid) Dynamic Voltage Frequency Scaling”, Rountree 2011.
- If task is CPU bound, slowdown linear with frequency. If task is memory or I/O bound, no slowdown with



frequency.

