ECE 571 – Advanced Microprocessor-Based Design Lecture 5

Vince Weaver

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

vincent.weaver@maine.edu

16 September 2014

Announcements

- HW#1 was posted
- Question about Supercomputing OSes.
  Linux is currently 97% (not 95%)
  Two Windows machines
  Rest "UNIX". Power/IBM so assume they are AIX.
- MSRs question. Use to set what happens on sysenter.

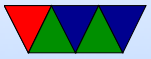

Power and Energy

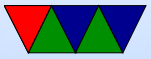

Definitions and Units

People often say Power when they mean Energy

- Energy – Joules, kWH (3.6MJ), Therm (105.5MJ), 1 Ton TNT (4.2GJ), eV (1.6×10^{-19} J), BTU (1055 J), horsepower-hour (2.68 MJ), calorie (4.184 J)
- Power – Energy/Time – Watts (1 J/s), Horsepower (746W), Ton of Refrigeration (12,000 Btu/h)
- Volt-Amps (for A/C) – same units as Watts, but not same thing
- Charge – mAh (batteries) – need voltage to convert to

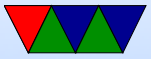

Energy

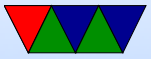

Power and Energy in a Computer System

Power Consumption Breakdown on a Modern Laptop, A. Mahersi and V. Vardhan, PACS'04.

- Thinkpad Laptop, 1.3GHz Pentium M, 256M, 14" display
- Oscilloscope with voltage probe and clamp-on current probe.
- Measured V and Current. P=IIR. V=IR P=IV, subtractive for things without wires
- Hard Drive 0.5-2W (Flash/SSD less)

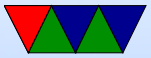

- LCD 1W (slightly more black than white)
- Backlight Inverter (this is before LED) 1-4W depending on brightness
- Total System Power 14-30W
- CPU 2-15W (with scaling)
- GPU 1-5W
- Memory 0.45 - 1.5W
- Power Supply Loss - 0.65W
- Wireless 0.1 - 3W (wifi on cellphones)
- CDROM 3-5W
- (USB 2.0 – 5V, can draw 5 units of 100mA each, 2.5W)

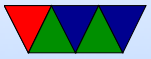

CPU Power and Energy

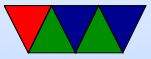

CMOS Transistors

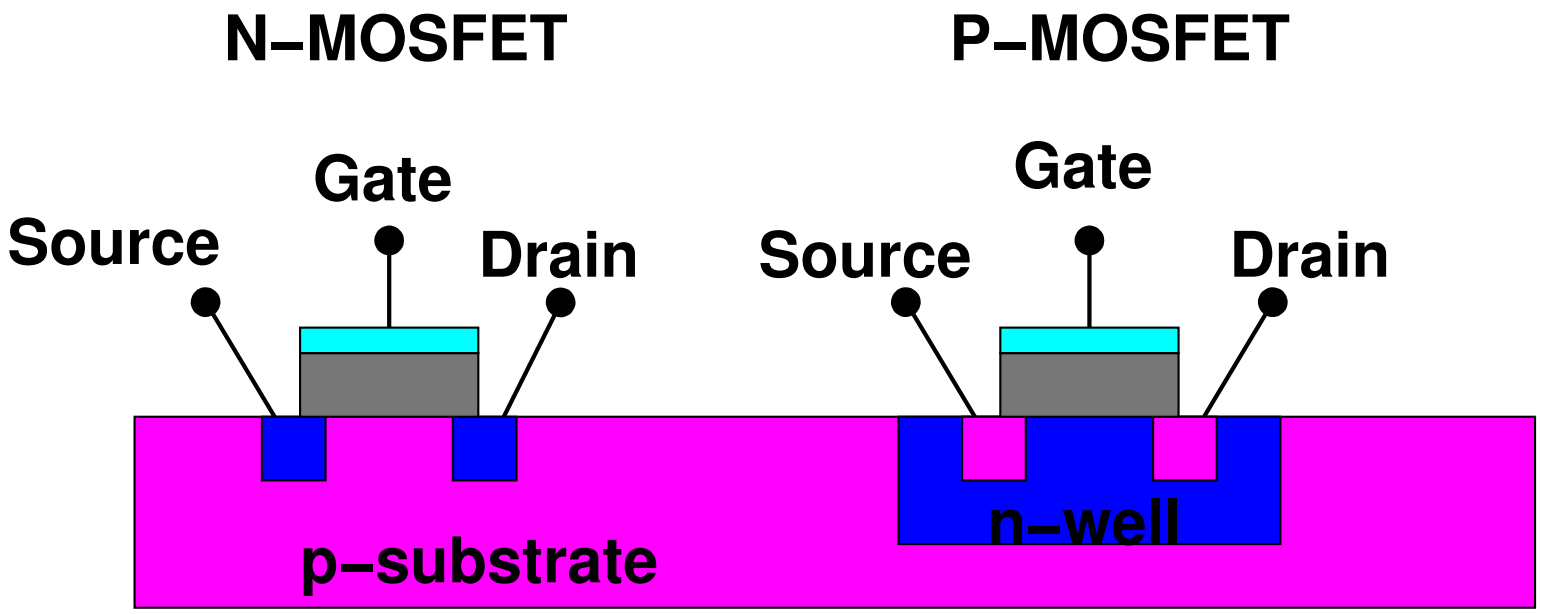

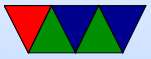

CMOS Dynamic Power

- $P = C\Delta V V_{dd} \alpha f$
  Charging and discharging capacitors big factor ($C\Delta V V_{dd}$) from V_{dd} to ground
  α is activity factor, transitions per clock cycle
  F is frequency

- α often approximated as $\frac{1}{2}$, $\Delta V V_{dd}$ as V_{dd}^2 leading to $P \approx \frac{1}{2} C V_{dd}^2 f$

- Some pass-through loss (V momentarily shorted)

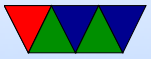

CMOS Dynamic Power Reduction

How can you reduce Dynamic Power?

- Reduce C – scaling
- Reduce V_{dd} – eventually hit transistor limit
- Reduce α (design level)
- Reduce f – makes processor slower

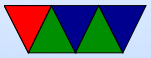

CMOS Static Power

- Leakage Current – bigger issue as scaling smaller. Forecast at one point to be 20-50% of all chip power before mitigations were taken.
- Various kinds of leakage (Substrate, Gate, etc)
- Linear with Voltage: $P_{static} = I_{leakage} V_{dd}$

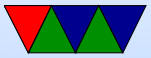

Leakage Mitigation

- SOI – Silicon on Insulator (AMD, IBM but not Intel)
- High-k dielectric – instead of SO2 use some other material for gate oxide (Hafnium)
- Transistor sizing – make only the critical transistors fast; non-critical ones can be made slower and less leakage prone
- Body-biasing
- Sleep transistors

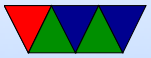

Total Energy

- $E_{tot} = [P_{dyanmic} + P_{static}]t$
- $E_{tot} = [(C_{tot}V_{dd}^2\alpha f) + (N_{tot}I_{leakage}V_{dd})]t$

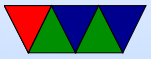

Delay

- $T_d = \frac{C_L V_{dd}}{\mu C_{ox}(\frac{W}{L})(V_{dd}-V_t)}$
- Simplifies to $f_{MAX} \sim \frac{(V_{dd}-V_t)^2}{V_{dd}}$
- If you lower f, you can lower V_{dd}

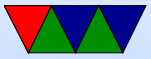

Thermal Issues

- Temperature and Heat Dissipation are closely related to Power
- If thermal issues, need heatsinks, fans, cooling

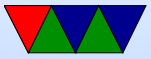

Metrics to Optimize

- Power
- Energy
- MIPS/W, FLOPS/W (don't handle quadratic V well)
- $Energy * Delay$
- $Energy * Delay^2$

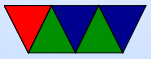

Power Optimization

- Does not take into account time. Lowering power does no good if it increases runtime.

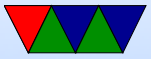

Energy Optimization

- Lowering energy can affect time too, as parts can run slower at lower voltages

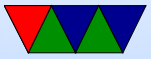

Energy Delay – Watt/t*t

- Horowitz, Indermaur, Gonzalez (Low Power Electronics, 1994)
- Need to account for delay, so that lowering Energy does not made delay (time) worse
- Voltage Scaling – in general scaling low makes transistors slower
- Transistor Sizing – reduces Capacitance, also makes transistors slower

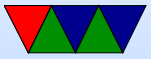

- Technology Scaling – reduces V and power.
- Transition Reduction – better logic design, have fewer transitions
  Get rid of clocks? Asynchronous? Clock-gating?
- Example with inverse ED2 (higher better):

Alpha 21064	SPEC=155	Power=30W	SPEC*SPEC/W=800
PPC603	SPEC=80	Power=3W	SPEC*SPEC/W=2100

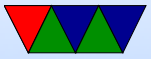

Energy Delay Squared– E*t*t

- Martin, Nyström, Pénzes – Power Aware Computing, 2002
- Independent of Voltage in CMOS
- Et can be misleading
  Ea=2Eb, ta=tB/2
  Reduce voltage by half, Ea=Ea/4, ta=2ta, Ea=Eb/2, ta=tb
- Can have arbitrary large number of delay terms in Energy

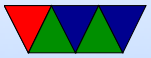

product, squared seems to be good enough

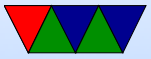

Energy Delay / Energy Delay Squared

Lower is better.

Energy	Delay	ED	ED^2
5J	2s	$10Js$	$20Js^2$
5J	3s	$15Js$	$45Js^2$

Same ED, Different ED^2

Energy	Delay	ED	ED^2
5J	2s	$10Js$	$20Js^2$
2J	5s	$10Js$	$50Js^2$

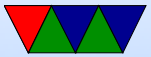

Energy Example