ESE534
Computer Organization

Day 8: February 10, 2010
Energy, Power, Reliability

Today

- Energy Tradeoffs?
- Voltage limits and leakage?
- Variations
- Transients
- Thermodynamics meets Information Theory (brief, if we get to it)

At Issue

- Many now argue **power** will be the ultimate scaling limit
  - (not lithography, costs, …)
- Proliferation of portable and handheld devices
  - …battery size and life biggest issues
- Cooling, energy costs may dominate cost of electronics
  - Even server room applications

Preclass 1

- 1GHz case
  - Voltage?
  - Energy per Operation?
  - Power required for 2 processors?
- 2GHz case
  - Voltage?
  - Energy per Operation?
  - Power required for 1 processor?

Energy and Delay

$$E = \frac{1}{2}CV^2$$

$$\tau_{gd}=Q/I=(CV)/I$$

$$I_{d,sat}=(\mu C_{OX}/2)(W/L)(V_{gs}-V_{TH})^2$$

Tradeoff

- $E \approx V^2$
- $\tau_{gd} \approx 1/V$

$$E = \frac{1}{2}CV^2$$

$$T_{gd}=(CV)/I$$

$$I_{d,sat} \propto (V_{gs}-V_{TH})^2$$

- We can trade speed for energy
- $E \times (\tau_{gd})^2 \approx$ constant

Martin *et al*. *Power-Aware Computing*, Kluwer 2001
http://caltechcstr.library.caltech.edu/308/

Parallelism

- We have Area-Time tradeoffs
- Compensate slowdown with additional parallelism

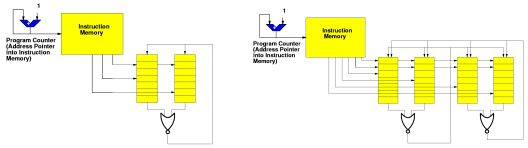

- …trade Area for Energy ➔ Architectural Option

Question

- How far can this go?
- What limits us?

Challenge: Power

Origin of Power Challenge

- Limited capacity to remove heat
  - ~100W/cm^2 force air
  - 1-10W/cm^2 ambient
- Transistors per chip grow at Moore's Law rate = $(1/F)^2$
- Energy/transistor must decrease at this rate to keep constant power density
- E/tr $\propto CV^2$
  - …but V scaling more slowly than F

Energy per Operation

$$E = \frac{1}{2} CV^2$$

C_{total} = # transistors × C_{tr}

C_{tr} scales (down) as F
transistors scales as F^{-2}
…ok if V scales as F…

ITRS Vdd Scaling: V Scaling more slowly than F

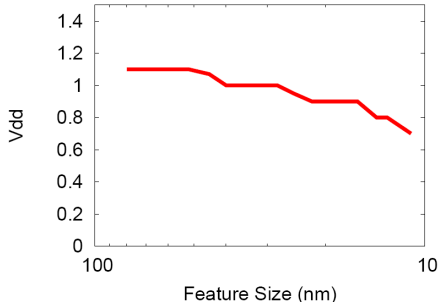

CV^2 scaling from ITRS: More slowly than $(1/F)^2$

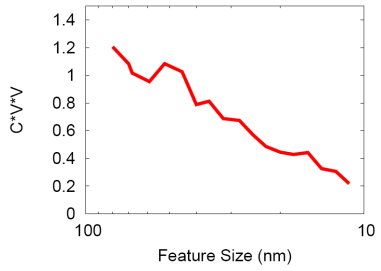

Origin of Power Challenge

- Transistors per chip grow at Moore's Law rate = $(1/F)^2$
- Energy/tr must decrease at this rate to keep constant
- $E/tr \propto CV^2f$

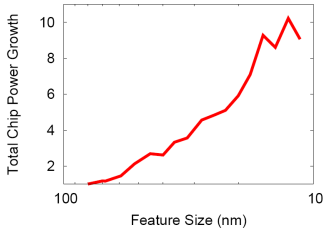

Historical Power Scaling

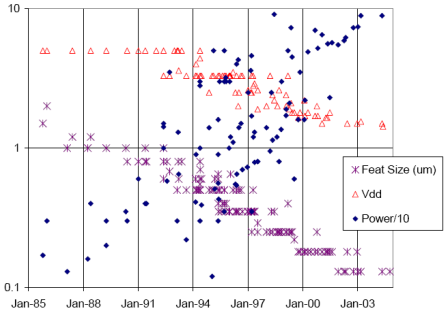

[Horowitz et al. / IEDM 2005]

Impact

- Can already place more transistors on a chip than we can afford to turn on.
- Power potential challenge/limiter for all future chips.

Impact

Power Limits Integration

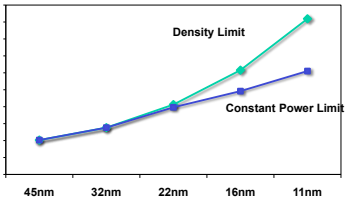

Source: Carter/Intel

How far can we reduce voltage?

Limits

- Ability to turn off the transistor
- Noise
- Parameter Variations

MOSFET Conduction

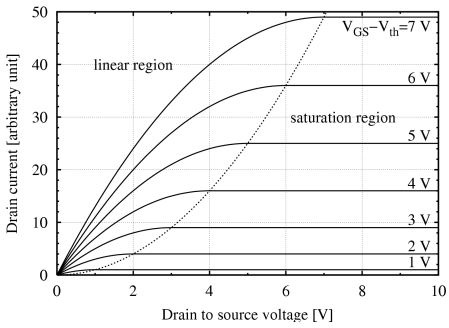

From: http://en.wikipedia.org/wiki/File:IvsV_mosfet.png

Transistor Conduction

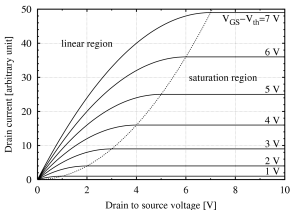

- Three regions
  - Subthreshold ($V_{gs}<V_{TH}$)
  - Linear ($V_{gs}>V_{TH}$) and ($V_{ds} < (V_{gs}-V_{TH})$)
  - Saturation ($V_{gs}>V_{TH}$) and ($V_{ds} > (V_{gs}-V_{TH})$)

Saturation Region

- Saturation Region
- ($V_{gs}>V_{TH}$)
- ($V_{ds} > (V_{gs}-V_{TH})$)

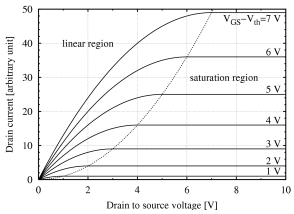

$$I_{d,sat}=(\mu C_{OX}/2)(W/L)(V_{gs}-V_{TH})^2$$

Linear Region

- ($V_{gs}>V_{TH}$)
- ($V_{ds} < (V_{gs}-V_{TH})$)

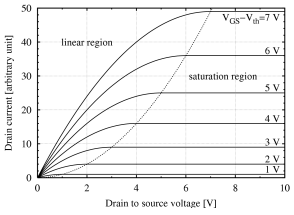

$$I_{d,lin}=(\mu C_{OX})(W/L)(V_{gs}-V_{TH})V_{ds}-(V_{ds})^2/2$$

Subthreshold Region

- ($V_{gs}<V_{TH}$)

$$I_{sub} = I_{VT} \times 10^{\left((V_{gs}-V_{TH})/S\right)}$$

$$S = (\ln(10))\eta kT/e$$

[Frank, IBM J. R&D v46n2/3p235]

Operating a Transistor

- Concerned about I_{on} and I_{off}
- I_{on} drive current for charging
  - Determines speed: $T_{gd} = CV/I$
- I_{off} leakage current
  - Determines leakage power
    - $E_{leak} = V \times I_{leak} \times T_{cycle}$

Leakage

- To avoid leakage want I_{off} very small
- Switch V from 0 to V_{dd}
- V_{gs} in off state is 0 → $V_{gs} < V_{TH}$

$$I_{sub} = I_{VT} \times 10^{\left((V_{gs} - V_{TH})/S\right)}$$

$$I_{off} = I_{VT} \times 10^{-\left((V_{TH})/S\right)}$$

Leakage

$$I_{off} = I_{VT} \times 10^{-\left((V_{TH})/S\right)}$$

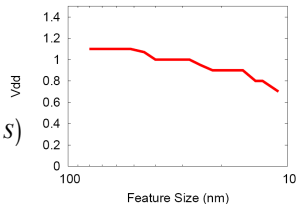

- S≈90mV for single gate
- S≈70mV for double gate
- 4 orders of magnitude I_{VT}/I_{off} → V_{TH}>280mV

Leakage limits V_{TH} in turn limits V_{dd}

How maximize I_{on}/I_{off} ?

- Maximize I_{on}/I_{off} – for given V_{dd} ? $E_{sw} \propto CV_2$
- Get to pick V_{TH}, V_{dd}

$$I_{d,sat} = (\mu C_{OX}/2)(W/L)(V_{gs} - V_{TH})^2$$

$$I_{d,lin} = (\mu C_{OX})(W/L)(V_{gs} - V_{TH})V_{ds} - (V_{ds})^2/2$$

$$I_{sub} = I_{VT} \times 10^{\left((V_{gs} - V_{TH})/S\right)}$$

Preclass 2

- $E = E_{sw} + E_{leak}$
- $E_{leak} = V \times I_{leak} \times T_{cycle}$
- $E_{sw} \propto CV_2$
- $I_{chip\text{-}leak} = N_{devices} \times I_{tr\text{-}leak}$

$$I_{sub} = I_{VT} \times 10^{\left((V_{gs} - V_{TH})/S\right)}$$

Preclass 2

- $E_{leak}(V)$?
- $T_{cycle}(V)$?

In Class

- Assign calculations
- Collect results

Values

V	T(v)	Esw(V)	Eleak(V)	E(V)
0.36	3.6E-09	1.296E-09	1.296E-11	1.30896E-09
0.27	0.000000027	7.29E-10	7.29E-11	8.019E-10
0.24	5.17064E-08	5.76E-10	1.24095E-10	7.00095E-10
0.21	9.74734E-08	4.41E-10	2.04694E-10	6.45694E-10
0.205	1.08137E-07	4.2025E-10	2.21682E-10	6.41932E-10
0.2	1.19897E-07	4E-10	2.39794E-10	6.39794E-10
0.19	1.4711E-07	3.61E-10	2.79509E-10	6.40509E-10
0.18	0.00000018	3.24E-10	3.24E-10	6.48E-10
0.15	3.23165E-07	2.25E-10	4.84748E-10	7.09748E-10
0.12	5.56991E-07	1.44E-10	6.68389E-10	8.12389E-10
0.09	0.0000009	8.1E-11	8.1E-10	8.91E-10

Graph for In Class

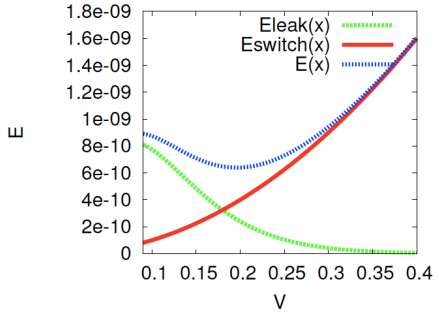

Impact

- Subthreshold slope prevents us from scaling voltage down arbitrarily.
- Induces a minimum operating energy.

Challenge: Variation

(This section was a little rushed)

Statistical Dopant Count and Placement

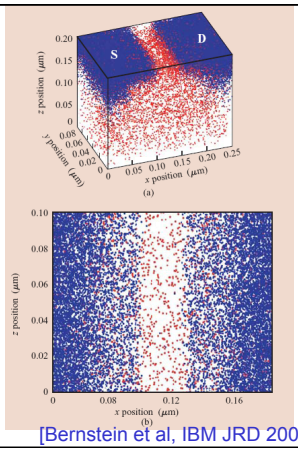

[Bernstein et al, IBM JRD 2006]

V_{th} Variability @ 65nm

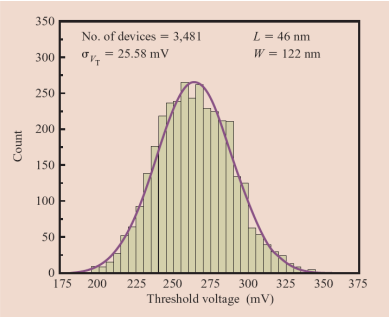

[Bernstein et al, IBM JRD 2006]

Variation

- Fewer dopants, atoms → increasing Variation
- How do we deal with variation?

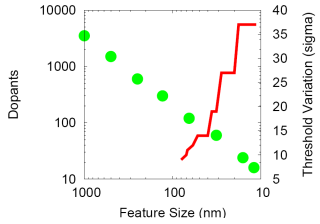

% variation in V_{TH}
(From ITRS prediction)

Impact of Variation?

- Higher V_{TH}?
  - Not drive as strongly
  - $I_{d,sat} \propto (V_{gs}-V_{TH})^2$
- Lower V_{TH}?
  - Not turn off as well → leaks more

$$I_{off} = I_{VT} \times 10^{-((V_{TH})/S)}$$

Variation

- Margin for expected variation
- Must assume V_{TH} can be any value in range

$$I_{on,min}=I_{on}(V_{th,max})$$
$$I_{d,sat} \propto (V_{gs}-V_{TH})^2$$

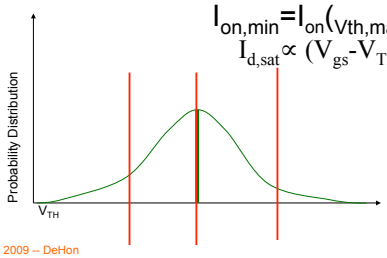

Margining

- Must raise V_{dd} to accommodate worst-case value
- → increase energy

$$I_{on,min}=I_{on}(V_{th,max})$$
$$I_{d,sat} \propto (V_{gs}-V_{TH})^2$$

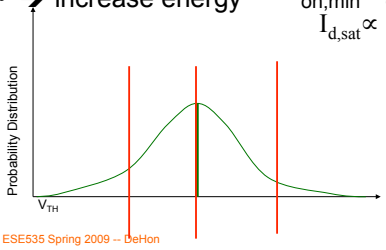

Variation

- Increasing variation forces higher voltages
  - On top of our leakage limits

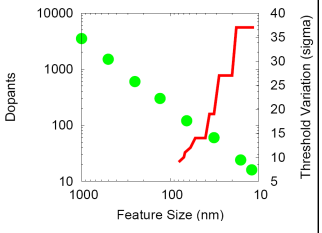

Variations

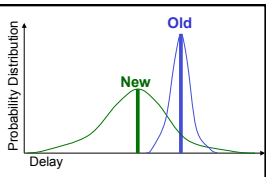

- Margins growing due to increasing variation
- Margined value may be worse than older technology?

End of Energy Scaling?

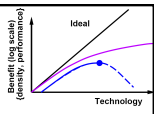

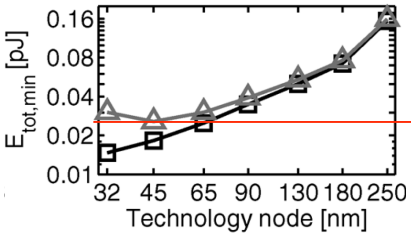

Black nominal
Grey with variation

[Bol *et al.*, IEEE TR VLSI Sys 17(10):1508—1519]

Chips Growing

- Larger chips → sample further out on distribution curve

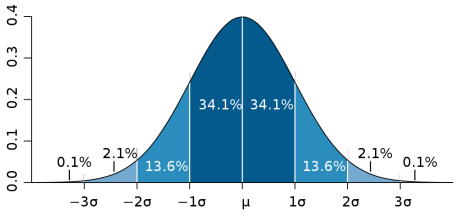

From: http://en.wikipedia.org/wiki/File:Standard_deviation_diagram.svg

Lecture Ended Here

(Didn't really cover material in transient and thermodynamics sections)

Challenge Transients

Transient Sources

- Effects
  - Thermal noise
  - Timing
  - Ionizing particles
    - α particle 10^5 to 10^6 electrons
    - Calculated gates with 15--30 electrons last time
      - Even if CMOS restores, takes time

Voltage and Error Rate

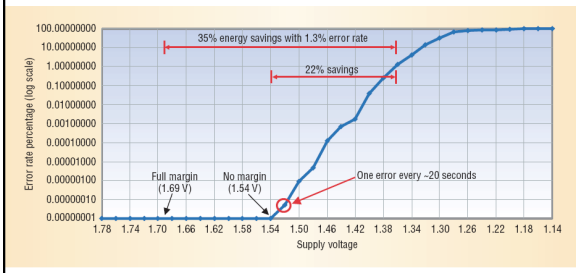

[Austin *et al.--IEEE Computer*, March 2004]

Errors versus Frequency

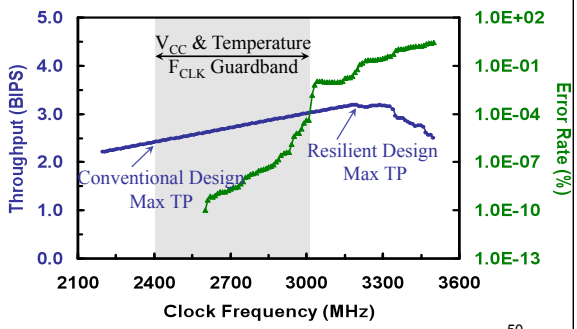

[Bowman, ISSCC 2008]

Scaling and Error Rates

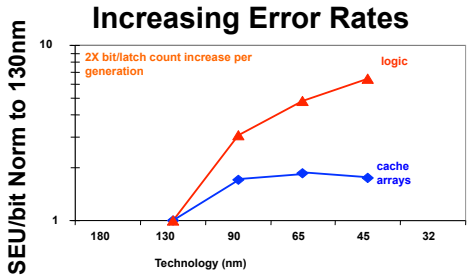

Source: Carter/Intel

Power and Reliability

- Intersection is the challenge
- Push V_{dd} in opposite directions
- Both reach inflection points
  - From doesn't matter
  - To major concern

Thermodynamics

Lower Bound?

- Reducing entropy costs energy
- Single bit gate output
  - Set from previous value to 0 or 1
  - Reduce state space by factor of 2
  - Entropy: $\Delta S = k \times \ln(\text{before/after}) = k \times \ln 2$
  - Energy=T $\Delta S = kT \times \ln(2)$
- Naively: setting a bit costs at least $kT \times \ln(2)$

Numbers (ITRS 2005)

- kT×ln(2) = 2.87×10^{-21}J (at R.T. K=300)

Year in Production	2014	2015	2016	2017	2018	2019	2020
DRAM ½ Pitch (nm) (contacted)	28	25	22	20	18	16	14
MPU/ASIC Metal 1 (M1) ½ Pitch (nm)(contacted)	28	25	22	20	18	16	14
MPU Physical Gate Length (nm)	11	10	9	8	7	6	6
L_g Physical gate length for LOP (nm) [1]	14	13	11	10	9	8	7
V_{dd} Power Supply Voltage (V) [6]	0.6	0.6	0.5	0.5	0.5	0.5	0.5
$C_{g,total}$ Total gate capacitance for calculation of CV/I [14]							
Extended Planar Bulk (F/μm)							
UTB FD (F/μm)	5.83E-16	5.44E-16	5.05E-16				
DG (F/μm)	6.43E-16	6.14E-16	5.55E-16	5.24E-16	4.82E-16	4.41E-16	4.00E-16

W/L=3 ➔ W=21nm=0.021μm

C≈8×10^{-18}F ≈10^{-17}F

Table 41d

E_{op}=CV^2=2.5×10^{-18}F

Recycling…

- Thermodynamics only says we have to dissipate energy if we discard information
- Can we compute without discarding information?
- Will that help us?

Three Reversible Primitives

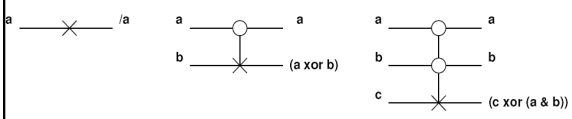

Universal Primitives

- These primitives
  - Are universal
  - Are all reversible
- If keep all the intermediates they produce
  - Discard no information
  - Can run computation in reverse

Thermodynamics

- In theory, at least, thermodynamics does not demand that we dissipate any energy (power) in order to compute.

Admin

- Assignment grades, feedback on blackboard for HW1 and HW2
- Class Wed.
- No class next Monday (2/22)

Big Ideas

- Can trade time for energy
  - …area for energy
- Noise and leakage limit voltage scaling
- Power major limiter going forward
  - Can put more transistors on a chip than can switch
- Continued scaling demands
  - Deal with noisier components
    - High variation and high transient upsets
- Thermodynamically admissible to compute without dissipating energy