ESE680-002 (ESE534): Computer Organization

Day 7: January 31, 2007
Energy and Power

Today

- Energy Tradeoffs?
- Voltage limits and leakage?
- Thermodynamics meets Information Theory
- Adiabatic Switching

- [This is an ambitious lecture]

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

What can we do about it?

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

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

$$I_d=(\mu C_{OX}/2)(W/L)(V_{gs}-V_{TH})^2$$

Tradeoff

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

- 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/

Questions

- How far can this go?
  - (return to later in lecture)
- What do we do about slowdown?

Parallelism

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

- …trade Area for Energy ➔ Architectural Option

Ideal Example

- Perhaps: 1nJ/32b Op, 10ns cycle
- Cut voltage in half
- 0.25nJ/32b Op, 20ns cycle
- Two in parallel to complete 2ops/20ns
- 75% energy reduction
  - Also 75% power reduction

Power Density Constrained Example

- Logic Density: 1 foo-op/mm^2
- Energy cost: 10nJ/foo-op @ 10GHz
- Cooling limit: 100W/cm^2
- How many foo-ops/cm^2/s?
  - 10nJ/mm^2 x 100mm^2/cm^2=1000nJ/cm^2
  - ➔ top speed 100MHz
  - 100M x 100 foo-ops = 10^{10} foo-ops/cm^2/s

Response

- How many foo-ops/cm^2/s?
  - 10nJ/mm^2 x 100mm^2/cm^2=1000nJ/cm^2
  - ➔ top speed 100MHz
  - 100M x 100 foo-ops = 10^{10} foo-ops/cm^2/s
- Power constraint won't let us run at 10GHz
  - might as well lower voltage, save energy

What can we support?

$E\times(t_{gd})^2 \approx \text{constant} \rightarrow 10nJ\times(100ps)^2 = E\times(t_{cycle})^2$

$$100W/cm^2 = \left(\frac{10nJ}{\left(\frac{t_{cycle}}{100\,ps}\right)^2}\right)\times 100\times\left(\frac{1}{t_{cycle}}\right)$$

(Pushing through the Math)

$$\left(t_{cycle}\right)^3 = \frac{10nJ\times 100\times\left(100\,ps\right)^2}{100\,J/s}$$

$$t_{cycle} = \sqrt[3]{10^{-8}\times\left(10^{-10}\right)^2 s^3}$$

$$t_{cycle} = 4.64\times10^{-10}\,s \approx 500\,ps$$

Improved Power

- How many foo-ops/cm^2/s?
  - 2GHz x 100 foo-ops = 2 ×10^{11} foo-ops/cm^2/s
  - At 5× lower voltage
  - [vs. 100M x 100 foo-ops = 10^{10} foo-ops/cm^2/s]

How far?

Limits

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

Sub Threshold Conduction

- To avoid leakage want I_{off} very small
- Use I_{on} for logic – determines speed
- Want I_{on}/I_{off} large

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

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

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

Sub Threshold Conduction

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

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

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

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

ITRS2005 – High Performance

Year of Production	2005	2006	2007	2008	2009	2010	2011	2012	2013
DRAM ½ Pitch (nm) (contacted)	80	70	65	57	50	45	40	36	32
MPU/ASIC Metal 1 (M1) ½ Pitch (nm)(contacted)	90	78	68	59	52	45	40	36	32
MPU Physical Gate Length (nm)	32	28	25	22	20	18	16	14	13
L_g Physical L_{gate} for High Performance logic (nm) [1]	32	28	25	22	20	18	16	14	13
V_{tsat} Saturation Threshold Voltage [7]									
Extended Planar Bulk (mV)	195	168	165	160	159	151	146	148	
UTB FD (mV)				169	168	167	170	166	167
DG (mV)							181	184	185
$I_{sd,leak}$ Source/Drain Subthreshold Off-State Leakage Current [8]									
Extended Planar Bulk (μA/μm)	0.06	0.15	0.2	0.2	0.22	0.28	0.32	0.34	
UTB FD (μA/μm)				0.17	0.19	0.22	0.22	0.29	0.29
DG (μA/μm)							0.1	0.11	0.11
I_{dsat} effective NMOS Drive Current [9]									
Extended Planar Bulk (μA/μm)	1020	1130	1200	1570	1810	2050	2490	2300	
UTB FD (μA/μm)				1486	1625	1815	2015	2037	2198
DG (μA/μm)							1899	1932	2220

Table 40a

ITRS2005 – Low Power

Year in Production	2005	2006	2007	2008	2009	2010	2011	2012	2013
DRAM ½ Pitch (nm) (contacted)	80	70	65	57	50	45	40	36	32
MPU/ASIC Metal 1 (M1) ½ Pitch (nm)(contacted)	90	78	68	59	52	45	40	36	32
MPU Physical Gate Length (nm)	32	28	25	22	20	18	16	14	13
L_g: Physical gate length for LOP (nm) [1]	45	37	32	28	25	22	20	18	16
V_{dsat}: Saturation Threshold Voltage [7]									
Extended Planar Bulk (mV)	288	303	285	274	275	226	233	231	
UTB FD (mV)							273	266	272
DG (mV)							261	255	257
$I_{off,nmk}$: Source/Drain Subthreshold Off-State Leakage Current [8]									
Extended Planar Bulk (µA/µm)	3.0E-03	3.0E-03	5.0E-03	5.0E-03	5.0E-03	5.0E-03	1.8E-02	2.5E-02	
UTB FD (µA/µm)							8.0E-03	1.0E-02	1.0E-02
DG (µA/µm)							5.0E-03	7.0E-03	7.0E-03
$I_{d,sat}$: effective NMOS Drive Current [9]									
Extended Planar Bulk (µA/µm)	589	607	573	712	775	749	749	774	
UTB FD (µA/µm)							740	765	718
DG (µA/µm)							783	822	709

Table 41c

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)

- $kT \times \ln(2) = 2.87 \times 10^{-21}$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

$C \approx 8 \times 10^{-18}$F $\approx 10^{-17}$F

Table 41d

$E_{op} = CV^2 = 2.5 \times 10^{-18}$F

Sanity Check

- V=0.5V
- $Q = CV = 0.5 \times 10^{-17}$ columbs
- $e = 1.6 \times 10^{-19}$ columbs
- Q≈30 electrons?

- Energy in α particle?
  - 10^5—10^6 electrons?

Hmm…

- $CV^2 = 2.5 \times 10^{-18}$J
- 18 Billion Transistors in 2.5cm^2
  - Generous, assumes no interconnect capacitance
- 4.5×10^{-8}J/2.5cm^2 ≈2×10^{-8}J/cm^2
- Cooling limit of @100W/cm^2
- Maximum operating frequency?
- 5GHz

Recycling…

- Thermodynamics only says we have to dissipate energy if we discard information
- Can we compute without discarding information?
- Can we use this?

Three Reversible Primitives

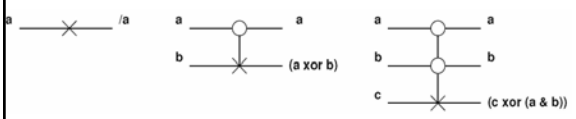

Universal Primitives

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

Cleaning Up

- Can keep "erase" unwanted intermediates with reverse circuit

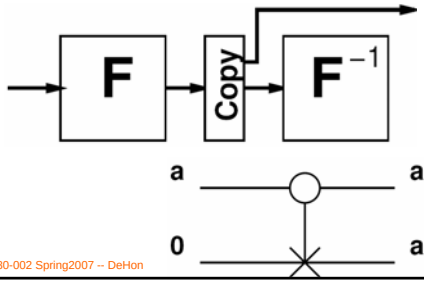

Thermodynamics

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

Adiabatic Switching

Two Observations

1. Dissipate power through on-transistor charging capacitance
2. Discard capacitor charge at end of cycle

Charge Cycle

- Charging capacitor
  - Q=CV
  - E=QV
  - $E=CV^2$
    - Half in capacitor, half dissipated in pullup

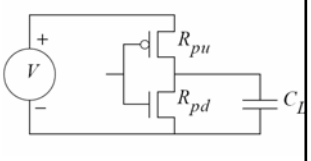

[Athas/Koller/Svensoon, USC/ISI ACMOS-TR-2 1993]

Adiabatic Switching

- Current source charging:
  - Ramp supplies slowly so supply constant current
    - $P=I^2R$
    - $E_{total}=P*T$
    - $Q=IT=CV$
    - $I=CV/T$
    - $E_{total}=I^2R*T=(CV/T)^2R*T$
    - $E_{total}=I^2R*T=(RC/T)\,CV^2$

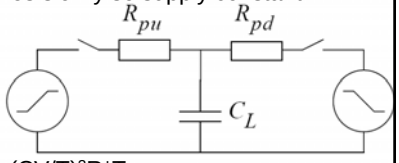

Ignores leakage …
May require large V_t

Impact of Adiabatic Switching

- $E_{total}=I^2R*T=(RC/T)\,CV^2$
- $RC=\tau_{gd}$
- $E_{total} \propto (\tau_{gd}/T)$
- Without reducing V
  - Can trade energy and time
- E×T=constant

Adiabatic Discipline

- Never turn on a device with a large voltage differential across it.
- $P=\Delta V^2/R$

SCRL Inverter

- Φ's, nodes, at Vdd/2
- P1 at ground
- Slowly turn on P1
- Slow split Φ's
- Slow turn off P1's
- Slow return Φ's to Vdd/2

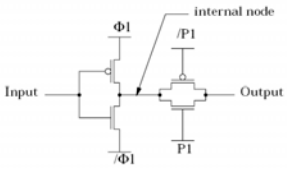

[Younis/Knight ISLPED(?) 1994]

SCRL Inverter

- Basic operation
  - Set inputs
  - Split rails to compute output adiabatically
  - Isolate output
  - Bring rails back together
- Have transferred logic to output
- Still need to worry about resetting output adiabatically

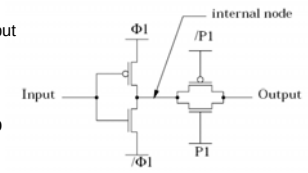

SCRL NAND

- Same basic idea works for nand gate
  - Set inputs
  - Adiabatically switch output
  - Isolate output
  - Reset power rails

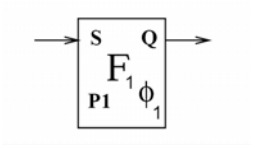

SCRL Cascade

- Cascade like domino logic
  - Compute phase 1
  - Compute phase 2 from phase 1…
- How do we restore the output?

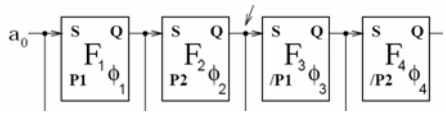

SCRL Pipeline

- We must **uncompute** the logic
  - Forward gates compute output
  - Reverse gate restore to Vdd/2

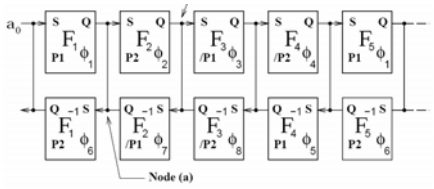

SCRL Pipeline

- P1 high (F1 on; F1 inverse off)
- Φ1 split: a=F1(a0)
- Φ2 split: b=F2(F1(a0))
- $F2^{-1}(F2(F1(a0))=a$
- P1 low – now $F2^{-1}$ drives a
- F1 restore by Φ1 converge
- …restore F2
- Use $F2^{-1}$ to restore a to Vdd/2 adiabatically

SCRL Rail Timing

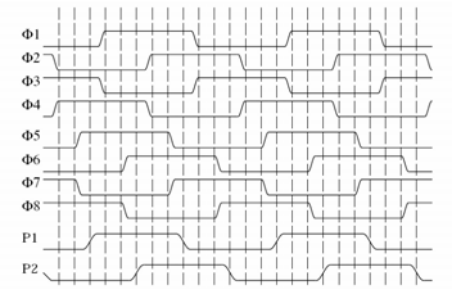

SCRL

- Requires Reversible Gates to uncompute each intermediate
- All switching (except IO) is adiabatic
- Can, in principle, compute at any energy

Trickiness

- Generating the ramped clock rails
- Use LC circuits
- Need high-Q resonators
- Making this efficient is key to **practical** implementation
  - Some claim not possible in practice

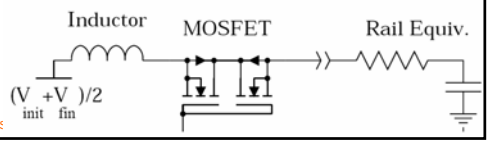

Big Ideas

- Can trade time for energy
  - …area for energy
- Noise and subthreshold conduction limit voltage scaling
- Thermodynamically admissible to compute without dissipating energy
- Adiabatic switching alternative to voltage scaling
- Can base CMOS logic on these observations