Introduction to CMOS VLSI Design

Power

Peter Kogge
University of Notre Dame
Fall 2012
Based on lecture slides by Jay Brockman & David Harris
http://www.cmosvlsi.com/coursematerials.html

Outline

- Motivation
- Thermal Resistance
- Energy 101
- Dynamic Power
- Static Power
- Low Power Design
- Historical Trends

Power and Energy (Chap 5)

- ***Energy*** = ability to do work – in ***Joules –*** watts*seconds
- ***Power*** = rate at which work performed or energy converted
  - CMOS: drawn from V_{DD} pin(s) of a chip, & converted to ***heat***
  - in units of ***Watts*** – volts*amps
- ***Instantaneous Power***: $P(t) = i_{DD}(t)V_{DD}$
- ***Energy***: $E = \int_0^T P(t)dt = \int_0^T i_{DD}(t)V_{DD}dt$
- **Average Power**: $P_{\text{avg}} = \frac{E}{T} = \frac{1}{T}\int_0^T i_{DD}(t)V_{DD}dt$

Getting A Feel for the Units

- ***Joules*** (J) are "Watt-seconds" = "Volts*Amps"
  - Relevant CMOS Unit: 1 ***picoJoule*** (**pJ**) = **10^{-12} J**
- Consider a typical cell phone (Motorola RAZR V3)
  - Battery: 3.6 V, 680mA Hours: Stored Energy:
    - = 3.6V * 0.68A Hr * 3600sec/Hr
    - = 8812.8 Watt-secs = 8812.8 Joules = 8.8×10^{15} pJ
  - Continuous talk time: 200-430 minutes
    - 430 min = 25,800 seconds
- So while talking, cell draws (8812.8/25800) = 0.34 Watts
  - or 0.34 J to talk 1 second
  - or 1 pJ can run cell phone for 2.94 picoseconds

Why *Else* do We Care About Power?

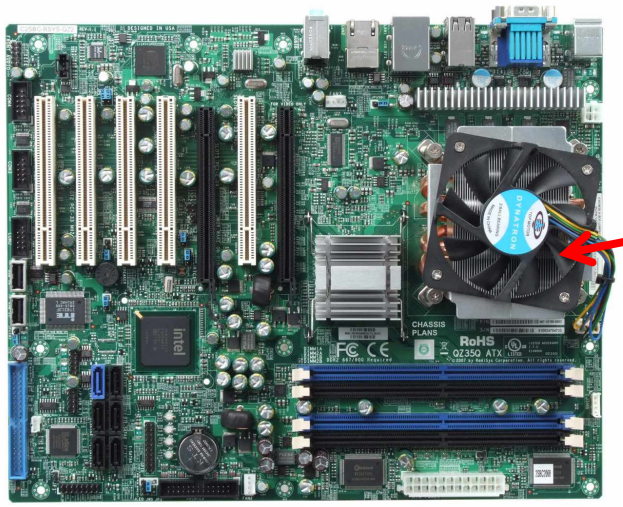

Question?: How big is the chip below this ~ 75x75mm fan?

Answer: About 15x15mm^2. (1/25th of the area)

Why *Else* Do We Care About Power?

- What goes in as electrical power-
- Goes out as ***HEAT!!!***
- And Heat ***KILLS*** electronics!
  - Leakage goes up
  - Dissipating more power
  - Causing temperature to go up even further
- Typical operational range of transistors: 40-125^{o}C
- Typical design points: <85^{o}C

Thermal Resistance

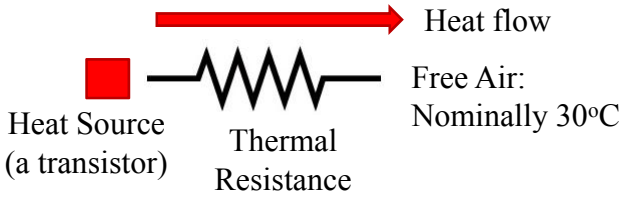

- **Thermal Resistance**: Steady state temperature difference between two endpoints per watt dissipated in source.
  - Depends on material and shape of object conducting heat
  - Units: °C/W
- For complex thermal path, obeys laws similar to resistance
  - Series paths increase thermal resistance
  - Parallel paths reduce thermal resistance
- In electronics, a function of
  - How chip is packaged
  - **Heat Sink** to air

Key Definitions

- **Q**: power dissipated by device
- **T_J**: temperature of a transistor "junction" in the device
- **T_{JMAX}**: maximum operating temperature of a transistor "junction" in the device
- **T_C**: temperature of the case
- **T_H**: temperature of the heatsink
- **T_{AMB}**: temperature of the air
- **ΔT_{ij}**: difference in temperature between points I and j
- **$R_{\theta JC}$**: thermal resistance from junction to case
- **$R_{\theta CH}$**: thermal resistance from case to heatsink
- **$R_{\theta HA}$**: thermal resistance from heatsink to air

Term	Typical
T_{JMAX}	125°C
T_{AMB}	70°C
$R_{\theta JC}$	1.5°K/W
$R_{\theta HA}$	4°K/W
$R_{\theta B}$ (heat transfer Pad)	0.1°K/W

http://en.wikipedia.org/wiki/Thermal_resistance

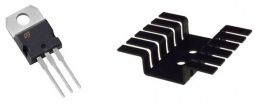

Equations

- $\Delta T = Q \times R_\theta$
- $R_\theta = R_{\theta JC} + R_{\theta CH} + R_{\theta HA}$
- $\Delta T = T_{JMAX} - T_{AMB}$

- Solving for Q:
  - $Q_{max} = (T_{JMAX} - T_{AMB})/(R_{\theta JC} + R_{\theta CH} + R_{\theta HA})$

- Using prior numbers for a TO-220 case:
  - $Q_{max} = (125 - 70) /(1.5 +0.1 +4) = 9.8W$

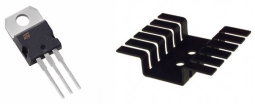

How Do We Do Better?

- Place chip as close as possible to metal shell
- Adding more surface area to Heat Sink adds more paths for heat escape
  - Add "fins"
- Don't let air sit over heat sink
  - Blow cool air to reduce "ambient"
- Net Result: we can cheaply air cool ~ 130W per chip
- More requires **liquid cooling**

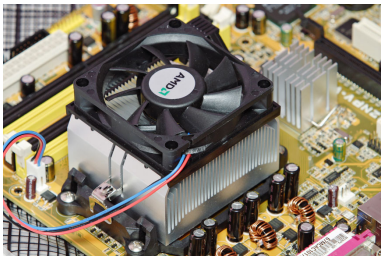

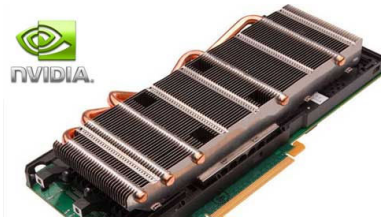

How About *REAL* Compute Power?

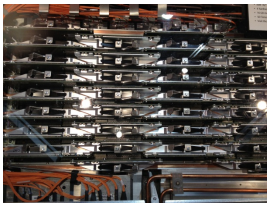

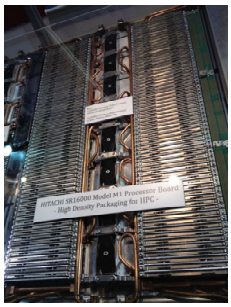

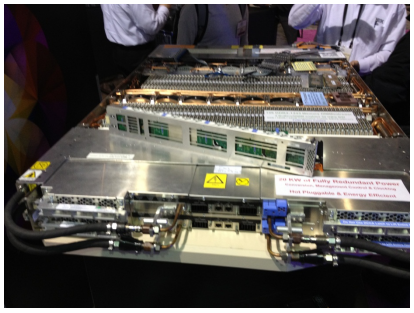

Moore's Law

- 1965, Gordon Moore: "the number of transistors that can be integrated on a die would double every 18 to 14 months."
  - Equivalent to Feature Size decreasing by factor of √2
  - Memory density 4X every 3 years

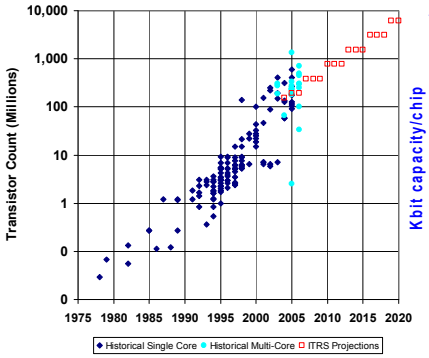

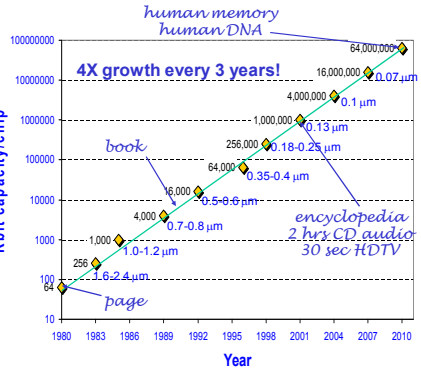

Smaller Feature Sizes Also Enable Higher Clock Rates

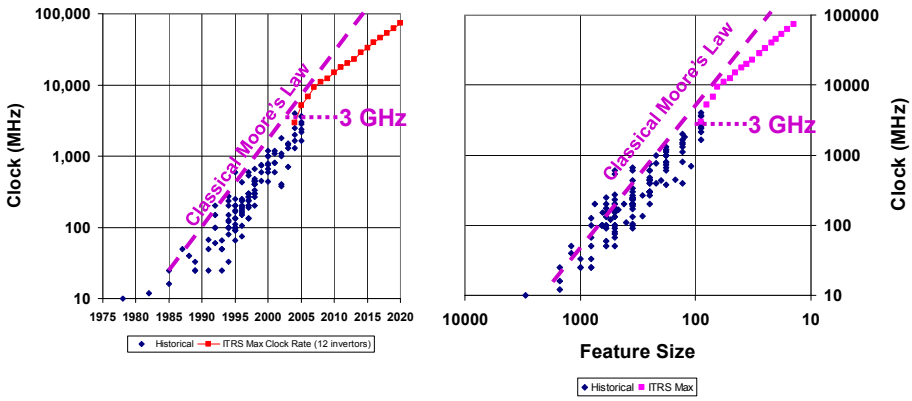

2005 projection was for 5.2 GHz – and we didn't make it.
Further, we're still stuck at 3+GHz in production.

Why the Clock Flattening?

POWER

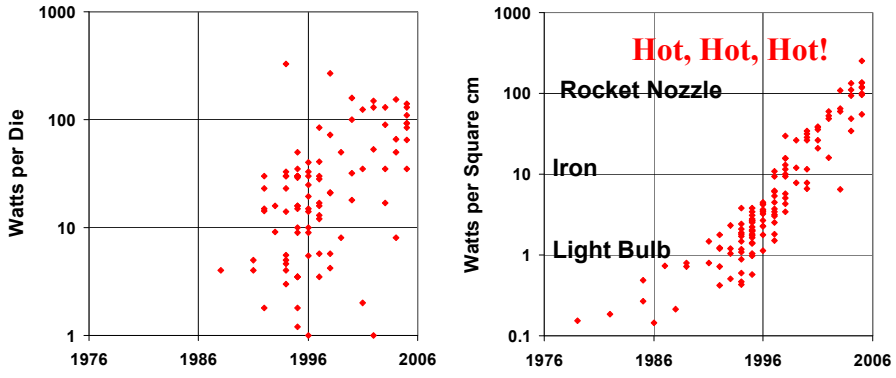

CMOS Energy 101

CMOS Energy 101

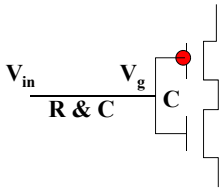

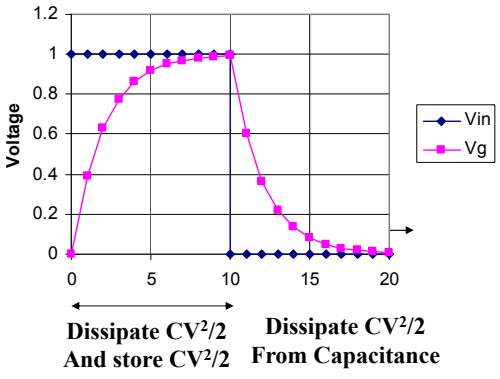

One clock cycle dissipates $C*V^2$

How Did CV^2 Improve With Time?

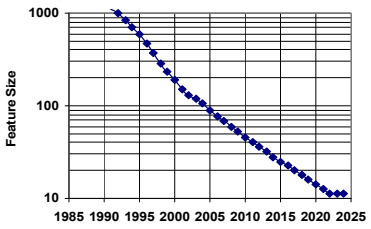

Assume capacitance of a circuit scales as feature size

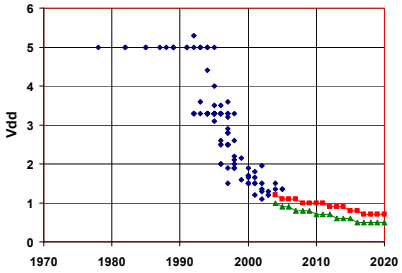

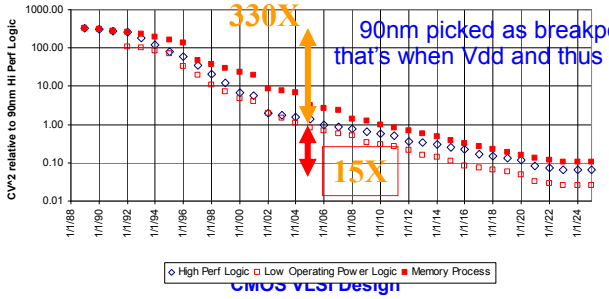

90nm picked as breakpoint because that's when Vdd and thus clocks flattened

Total CMOS Power

Dynamic Dissipation

- Power dissipated only at specific times
- Charging & Discharging of load capacitances
- "Short circuit" when both p and n types partially on

Static Dissipation

- "Always there"
- Subthreshold conduction thru OFF transistors
- Tunneling thru oxide
- Leakage thru reverse biased diodes
- Contention current in ratioed logic

$$P_{total} = P_{dynamic} + P_{static}$$

Dynamic Power

Dynamic Power I

- Charge & discharge load capacitances when transistors switch.

- One cycle involves a rising and falling output.
- On rising output, charge Q = CV_{DD} is required
- On falling output, charge is dumped to GND
- This repeats Tf_{sw} times over an interval of T

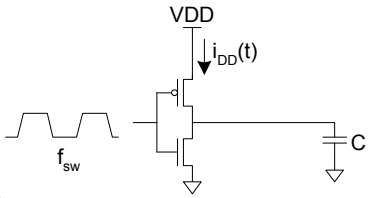

Remember: Current = charge/time

Charging a Capacitor

- When the gate output rises
  - Energy stored in capacitor is

    $E_C = \frac{1}{2} C_L V_{DD}^2$
  - But energy drawn from the supply is

    $$E_{VDD} = \int_0^\infty I(t) V_{DD} dt = \int_0^\infty C_L \frac{dV}{dt} V_{DD} dt$$
    $$= C_L V_{DD} \int_0^{V_{DD}} dV = C_L V_{DD}^2$$
  - Half the energy from V_{DD} is dissipated in the pMOS transistor as heat, other half stored in capacitor
- When the gate output falls
  - Energy in capacitor is dumped to GND
  - Dissipated as heat in the nMOS transistor

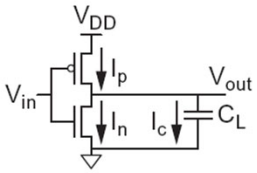

Dynamic Power II

- When transistors switch, both nMOS and pMOS networks may be momentarily ON at once
- Leads to a blip of "**short circuit**" current.
  - Also called ***Crowbar current***
- < 10% of dynamic power if rise/fall times are comparable for input and output

Switching Waveforms

- Example: V_{DD} = 1.0 V, C_L = 150 fF, f = 1 GHz

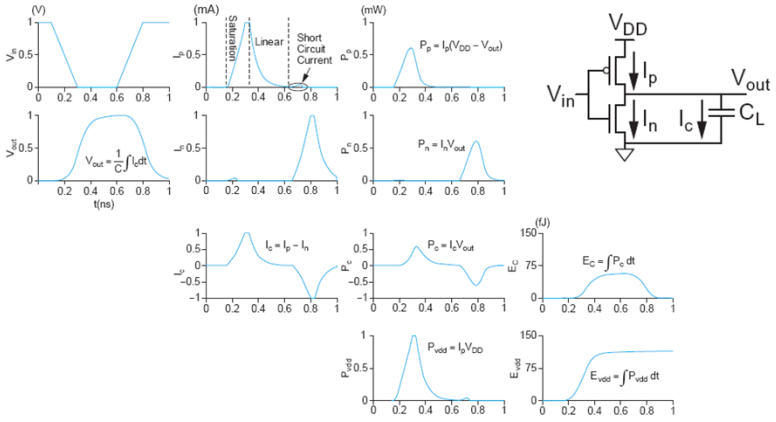

Dynamic Power Cont.

$$P_{\text{dynamic}} = \frac{1}{T}\int_0^T i_{DD}(t)V_{DD}\,dt$$

$$= \frac{V_{DD}}{T}\int_0^T i_{DD}(t)\,dt$$

$$= \frac{V_{DD}}{T}\left[Tf_{sw}CV_{DD}\right]$$

$$\boxed{= CV_{DD}^{\ 2} f_{sw}}$$

Integral of i is total charge converted

Charge in 1 cycle

# of cycles in T time

Remember this Equation

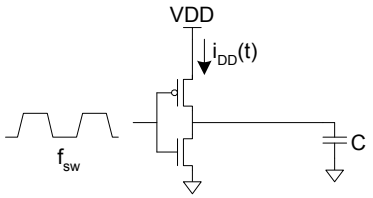

Activity Factor

- Not all gates switch every cycle
- Suppose the system clock frequency = f
- Let $f_{sw} = \alpha f$, where **α = activity factor**
  - If the signal is a clock, $\alpha = 1$
  - If the signal switches once per cycle, $\alpha = ½$
  - Dynamic gates:
    - Switch either 0 or 2 times per cycle, $\alpha = ½$
  - Static gates:
    - Depends on design, but typically $\alpha = 0.1$
- **Revised Dynamic power**: $P_{\text{dynamic}} = \alpha C V_{DD}^2 f$

Activity Factor Estimation

- Let P_i = Prob(node i = 1)
  - $\bar{P}_i = 1-P_i$
- $\alpha_i = \bar{P}_i * P_i$
- Completely random data has P = 0.5 and $\alpha = 0.25$
- Data is often not completely random
  - e.g. upper bits of 64-bit words representing bank account balances are usually 0
- Data propagating through ANDs and ORs has lower activity factor
  - Depends on design, but typically $\alpha \approx 0.1$

Switching Probability

Gate	P_Y
AND2	$P_A P_B$
AND3	$P_A P_B P_C$
OR2	$1 - \overline{P}_A \overline{P}_B$
NAND2	$1 - P_A P_B$
NOR2	$\overline{P}_A \overline{P}_B$
XOR2	$P_A \overline{P}_B + \overline{P}_A P_B$

Example

- A 4-input AND is built out of two levels of gates
- Estimate the activity factor at each node if the inputs have P = 0.5

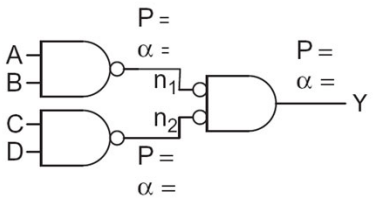

Dynamic Power Example

- 1 billion transistor chip
  - 50M logic transistors
    - Average width: 12 λ
    - Activity factor = 0.1
  - 950M memory transistors
    - Average width: 4 λ
    - Activity factor = 0.02
  - 1.0 V 65 nm process (with $\lambda = 25\text{nm}$)
  - C = 1 fF/μm (gate) + 0.8 fF/μm (diffusion)
- Estimate dynamic power consumption @ 1 GHz. Neglect wire capacitance and short-circuit current.

Solution

$$C_{\text{logic}} = \left(50\times10^{6}\right)\left(12\lambda\right)\left(0.025\mu m/\lambda\right)\left(1.8fF/\mu m\right) = 27\text{ nF}$$

$$C_{\text{mem}} = \left(950\times10^{6}\right)\left(4\lambda\right)\left(0.025\mu m/\lambda\right)\left(1.8fF/\mu m\right) = 171\text{ nF}$$

$$P_{\text{dynamic}} = \left[0.1C_{\text{logic}} + 0.02C_{\text{mem}}\right]\left(1.0\right)^{2}\left(1.0\text{ GHz}\right) = 6.1\text{ W}$$

Static Power

Static Power (5.3)

- **Static power** is consumed even when chip is quiescent.
- Leakage draws power from "nominally" OFF devices:
  - $P_{static} = I_{static} V_{dd}$
- Two forms of leakage current:
  - **Subthreshold Leakage**: conduction thru OFF transistors
  - **Gate Leakage**: conduction across transistor gates when transistor is on
- Both proportional to total widths of transistors in on and off states, AND thresholds of transistors
  - **Low Threshold Transistors**: turn on fast but leak more
  - **High Threshold Transistors**: slower, but leak much less

Leakage Example

- 200 Million transistor chip
  - 20M logic transistors ave width: 12 λ
  - 180M memory transistors ave width: 4 λ
- 1.2 V 100 nm process
  - λ = 0.05μm
- On average 50% of Transistors OFF
- 2 types of transistors
  - "**Low-threshold**"/Hi Leakage
    - Fast but leaky
    - Used in 20% of logic
    - Gate leakage = 3nA/μm
    - Subthreshold leakage for OFF = 20nA/μm
  - "**High threshold**"/Low Leakage
    - Slow but less leakage
    - In All Memory & 80% logic
    - Gate leakage = 0.002nA/μm
    - Subthreshold leakage for OFF = 0.02nA/μm

Estimate static power

- Compute μm of hi leakage transistors
  - 20Mx20% x 12λ x 0.05μm
  - = 2.4×10^6 μm
- Compute μm of low leakage transistors
  - 20Mx80% x 12λ x 0.05μm
  - + 180M x 4λ x 0.05μm
  - = 45.6×10^6 μm
- Gate leakage
  - Hi: 2.4×10^6 μm 50% x 3nA/μm = 3.6mA
  - Low: 45.6×10^6 μm x 50% 0.002nA/μm = 0.045mA
- Subthreshold leakage
  - Hi: 2.4×10^6 μm x 50% x 20nA/μm = 24mA
  - Low: 45.6×10^6 μm x 50% x 0.02nA/μm = 0.456mA
- Total Current = 28.1 mA
- At 1.2V this is **34mW**

Redo with No Low Leakage Devices

- 200 Million transistor chip
  - 20M logic transistors ave width: 12 λ
  - 180M memory transistors ave width: 4 λ
- 1.2 V 100 nm process
  - λ = 0.05μm
- 2 types of transistors
  - "Low-threshold"/Hi Leakage
    - Used in 20% of logic
    - Gate leakage = 3nA/μm
    - Subthreshold leakage for OFF = 20nA/μm
  - ~~"High threshold"/Low Leakage~~
    - ~~In Memory & 80% logic~~
    - ~~Gate leakage = 0.002nA/μm~~
    - ~~Subthreshold leakage for OFF = 0.02nA/μm~~
- All transistors have gate leakage
- On average 50% of Transistors OFF

Estimate static power

- Total μm of transistors
  - 2.4×10^6 μm + 45.6×10^6 μm
  - = 48×10^6 μm
- Gate leakage
  - 48×10^6 μm x 50% 3nA/μm = 72mA
- Subthreshold leakage
  - 48×10^6 μm x 50% x 20nA/μm = 480mA
- Total Current = 552mA
- At 1.2V this is **662 mW**
- This is ***20X MORE*** power

Power Reduction Techniques

Modern Terminology

- **Energy Delay Product**
  - Notionally:
    - "Energy used to perform some function"
    - times "Time required to do function"
  - Many variants
- **Dynamic Clock and Voltage Scaling**
  - Under periods of light loads, reduce clock
  - And if possible reduce V_{dd} to point where clock is "max possible"
- **Clock Gating**
  - In pipelined systems gating the clocks to latches off when circuits are not in use

Low Power Design

- Reduce dynamic power
  - α: clock gating, sleep mode
  - C: small transistors (esp. on clock), short wires
  - V_{DD}: lowest suitable voltage
  - f: lowest suitable frequency
- Reduce static power
  - Selectively use low V_t devices
  - Leakage reduction:
    stacked devices, body bias, low temperature

Example Technique: Intel "Speedstep"

- As we decrease clock rate, we can also reduce voltage
- Significant example: Pentium M
  - cf: ftp://download.intel.com/design/network/papers/30117401.pdf

F	Vdd	fV^2	Ref f	Rel fV^2
1.6	1.484	3.52	1.00	1.00
1.4	1.420	2.82	0.88	0.80
1.2	1.276	1.95	0.75	0.55
1.0	1.164	1.35	0.63	0.38
0.8	1.036	0.86	0.50	0.24
0.6	0.956	0.55	0.38	0.16

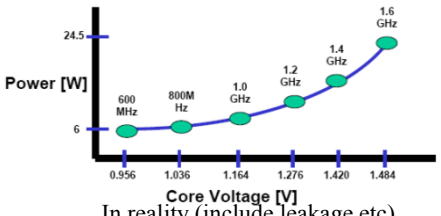

In reality (include leakage etc)

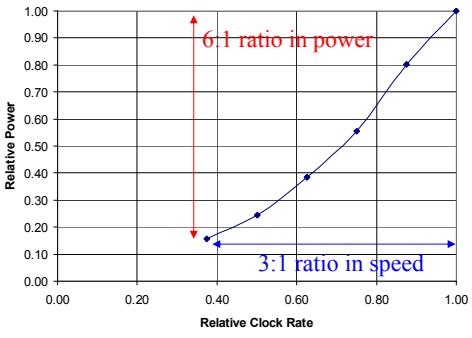

If capacitance remained Constant

Leakage Control

- Leakage and delay trade off
  - Aim for low leakage in sleep and low delay in active mode
- To reduce leakage:
  - Increase V_t: *multiple V_t*
    - Use low V_t only in critical circuits
  - Increase V_s: *stack effect*
    - *Input vector control* in sleep
  - Decrease V_b
    - *Reverse body bias* in sleep
    - Or forward body bias in active mode

Also: for series transistors, keep normally off signals at bottom

NAND3 Leakage Example

- 100 nm process

I_{gn} = 6.3 nA I_{gp} = 0

I_{offn} = 5.63 nA I_{offp} = 9.3 nA

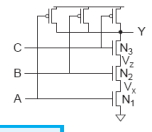

Input State (ABC)	I_{sub}	I_{gate}	I_{total}	V_x	V_z
000	0.4	0	0.4	stack effect	stack effect
001	0.7	0	0.7	stack effect	$V_{DD} - V_t$
010	0	1.3	1.3	intermediate	intermediate
011	3.8	0	10.1	$V_{DD} - V_t$	$V_{DD} - V_t$
100	0.7	6.3	7.0	0	stack effect
101	3.8	6.3	10.1	0	$V_{DD} - V_t$
110	5.6	12.6	18.2	0	0
111	28	18.9	46.9	0	0

Data from [Lee03]

Power Gating

- Turn OFF power to blocks when they are idle to save leakage
  - Use virtual V_{DD} (V_{DDV})
  - Gate outputs to prevent invalid logic levels to next block

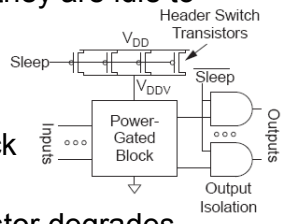

- Voltage drop across sleep transistor degrades performance during normal operation
  - Size the transistor wide enough to minimize impact
- Switching wide sleep transistor costs dynamic power
  - Only justified when circuit sleeps long enough

What About "Sleep States?"

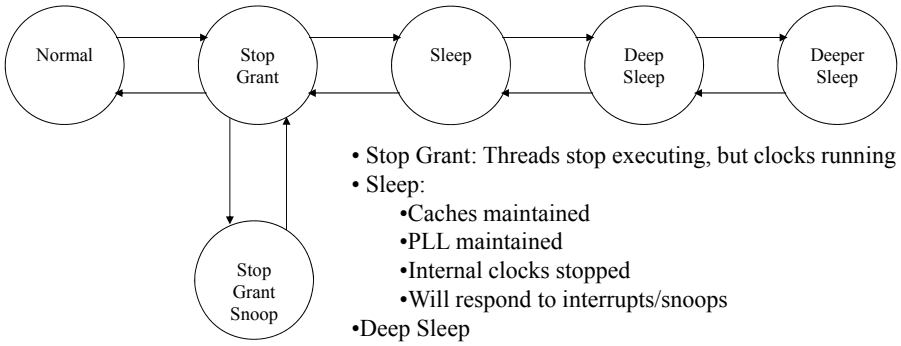

- Stop Grant: Threads stop executing, but clocks running
- Sleep:
  - Caches maintained
  - PLL maintained
  - Internal clocks stopped
  - Will respond to interrupts/snoops
- Deep Sleep
  - Will NOT respond to interrupts/snoops
- Deeper Sleep
  - L2 cache emptied
  - Vdd lowered even further

from ATOM chip documentation page
http://download.intel.com/design/processor/datashts/319535.pdf