
CS 150

Digital Design

Lecture 24 – Power and Energy

2012-4-17

Professor John Wawrzynek

today's lecture by John Lazzaro

TAs: Shaoyi Cheng, Daiwei Li, James Parker

www-inst.eecs.berkeley.edu/~cs150/



Sad fact: Computers turn
electrical energy into heat.
Computation is a byproduct.

Energy and Performance

**Air or water carries heat
away, or chip melts.**



The Joule: Unit of energy. Can also be expressed as **Watt-Seconds**. Burning 1 Watt for 100 seconds uses 100 Watt-Seconds of energy.

This is how electric tea pots work ...

1 Joule heats 1 gram of water
0.24 degree C


1 Joule of Heat Energy
per Second

1 Watt

The Watt: Unit of power. The amount of energy burned in the resistor in 1 second.

1 Ohm
Resistor

20 W rating: Maximum power the package is able to transfer to the air. Exceed rating and resistor **burns**.



Cooling an iPod nano ...



Like resistor on last slide, iPod relies on passive transfer of heat from case to the air.

Why? Users don't want fans in their pocket ...

To stay “cool to the touch”
via passive cooling,
power budget of 5 W.

If iPod nano used 5W all the time, its battery would last 15 minutes ...

Powering an iPod nano (2005 edition)



1.2 W-hour battery:
Can supply 1.2 watts
of power for 1 hour.

$$1.2 \text{ W} / 5 \text{ W} = 15 \text{ minutes.}$$

More W-hours require bigger battery
and thus bigger "form factor" --
it wouldn't be "nano" anymore :-).

Real specs for iPod nano :
14 hours for music,
4 hours for slide shows.

85 mW for music.
300 mW for slides.


Finding the (2005) iPod nano CPU ...



A close relative ...



digital media management system-on-chip



Two 80 MHz CPUs.
One CPU used for
audio, one for slides.

Low-power ARM
roughly 1mW per
MHz ... variable
clock, sleep modes.

85 mW **system**
power realistic ...

Year-to-year: continuous improvements



iPod nano
2005
14 hours
battery
life
(audio
playback)

iPod nano
2006
24 hours
battery
life
(audio
playback)



What
changed
inside?





iPod nano
2005 -
a C-shaped PC
board, with a
battery in the
"C" opening.


iPod nano
2006 -
battery lies
on top of PC
board.



How? Small IC packages, fewer parts

iPod nano 2006 →

iPod nano 2005 →



Aluminum permits thinner case ...



What's happened since 2006?

1.61"



nearly the same depth

2010 Nano



0.74 ounces

2010 Shuffle



0.44 ounces

2010 Nano:
"up to" **24 hours**
audio playback

2010 Shuffle:
"up to" 15 hours
audio playback



0.39 W Hr
(33% of 2005 Nano)



0.19 W Hr



	2007	2008	2009	2010	2011	Today
	iPhone	iPhone 3G	iPhone 3GS	iPhone 4	iPhone 4 (CDMA)	iPhone 4S
Display	3.5" TFT LCD, 480 x 320 pixels	3.5" TFT LCD, 480 x 320 pixels	3.5" TFT LCD, 480 x 320 pixels	3.5" IPS LCD, 960 x 640 Pixels	3.5" IPS LCD, 960 x 640 Pixels	3.5" IPS LCD, 960 x 640 Pixels



Desired screen size sets smartphone W x L
 Depth? : Thin body vs. battery life

	2007	2008	2009	2010	2011	Today
	iPhone	iPhone 3G	iPhone 3GS	iPhone 4	iPhone 4 (CDMA)	iPhone 4S
Battery	Li-Ion Polymer, 3.7V, 1170mAh	Li-Ion Polymer, 3.7V, 1150mAh	Li-Ion Polymer, 3.7V, 1220mAh	Li-Ion Polymer, 3.7V, 1420mAh		Li-Ion Polymer, 3.7V, 1430mAh

22% gain in battery energy over 5 iterations



iPhone (2007)



Mainboard

Battery

Antennas




iPhone
4{,S}

Battery

L-shape
Main Board

Metal frame acts as antenna



WARNING
Potential for fire or
burning. Do not
disassemble, puncture,
crush, heat, or burn.
Li-Ion Polymer Battery
3.7V 5.25Whr
APH: 616-0513
VPH: LIS1445APPC
Assembled in China

CE
RoHS
Li-ion

Authorized Service
Provider Only

	2007
	iPhone
Apps Processor	Samsung S5L8900B01 ARM Core
Process Geometry	90nm
Die Size	8.5 x 8.5 mm
Pin Count	424
ARM Core (Instruction Set)	ARM9 (ARMv5)
Clock Speed	~600MHz
GPU	PowerVR MBX
SDRAM	1Gb Mobile DDR


In 4 years:

6.8x increase in transistor count


33% max clock speed increase

Attached DRAM:
128 MB -> 512 MB

6.8x transistors:
Dual CPU and GPU,
and to save energy.



	2008	2009	2010	2011	Today
					iPhone 4S
Apps Processor					Apple A5 Dual ARM Core Processor
Process Geometry					45nm
Die Size					12.15 x 10.1 mm
Pin Count					1360
ARM Core (Instruction Set)					Cortex-A9 (ARMv7)
Clock Speed					~800MHz
GPU					PowerVR SGX (dual core)
SDRAM					4Gb Mobile DDR2



Notebooks ... as designed in 2006 ...

2006 Apple MacBook -- 5.2 lbs



* **Performance:** Must be “close enough” to desktop performance ... most people no longer used a desktop (even in 2006).

* **Size and Weight.** Ideal: paper notebook.

* **Heat:** No longer “laptops” -- top may get “warm”, bottom “hot”. Quiet fans OK.

Battery: Set by size and weight limits ...




Battery rating:
55 W-hour.

At 2.3 GHz, Intel Core Duo CPU consumes **31 W** running a heavy load - **under 2 hours battery life!** And, just for CPU!

46x more energy than iPod nano battery. And iPod lets you listen to music for 14 hours!

Almost full 1 inch depth.
Width and height set by available space, weight.

At 1 GHz, CPU consumes **13 Watts.** "Energy saver" option uses this mode ...



55 W-hour battery
stores the energy of
1/2 a stick of dynamite.

If battery short-circuits,
catastrophe is possible ...



MacBook Air ... design the laptop like an iPod




2011 Air: 11.8 in x 7.56 in x 0.68 in; 2.38 lbs



2006 Macbook: 12.8 in x 8.9 in x 1 in; 5.2 lbs

Mainboard: fills about 25% of the laptop



35 W-h battery: 63% of 2006 MacBook's 55 W-h




2011 Air: 35 W-h battery, 5 hour battery life*
2012 iPad: 42.5 W-h battery, 10 hour battery life*

*For a content-consumption workload.



Battery-Life-Hour/W-h: 1.7x iPad advantage

iPad: iPhone++



Li-Ion Polymer Battery
Model No.: A1399
3.7V 43.6Whr 11560mAh
APN: G16-05973
VPH: 9697A103H
Assembled in China

WARNING
Potential for fire or
burning. Do not
disassemble, puncture,
crush, heat, or burn.


苹果公司 型号: A1399
中国制造 锂电池
产品标准: GB4943-2000
GB/T 18287-2000
额定电压: 3.7V
额定容量: 11560mAh



iPad: iPhone++


A5X: 2 ARM Cortex9 Cores,
expanded PowerVR GPU

Top




Up to 64 GB Flash

Bottom



1 GB DRAM

Cellular RF



Cellular front-end chips
on separate board.

MacBook Air: Full PC




Thunderbolt I/O

Platform
Controller
Hub


Core i5
CPU/GPU

Top



Up to
4GB DRAM

Bottom



2011 Air: \$999 -- 64 GB SSD, 2 GB RAM, x86
2012 iPad: \$699 -- 64 GB SSD, 1 GB RAM, ARM
iPad 2012 CPU: iPhone 4S, with 2X GPU and RAM




"Content Creation vs. Content Consumption"

iPad 2011→2012,
battery W-hours
increased by 70%

Weight increase of
0.11 pound, thicker
by 0.03 inches.

Increase needed
to double display
resolution while
keeping battery
life @ 10 hours.




2011 IPAD

2012 iPad

The CPU is only part of power budget!

2004-era notebook
running a full workload.



“Amdahl’s Law for Power”

If our CPU took **no power** at all to run, that would **only double** battery life!

iPad 2011 → 2012 power++ comes from this side.

Servers: Total Cost of Ownership (TCO)




Reliability: running computers
hot makes them **fail more often.**

Machine rooms
are expensive.
Removing heat
dictates how
many servers to
put in a machine
room.

Electric bill adds
up! Powering the
servers +
powering the air
conditioners is a
big part of TCO.

Computations per W-h doubles every 1.6 years, going back to the first computer.

(Jonathan Koomey, Stanford).




Processors and Energy



Switching Energy: Fundamental Physics

Every logic transition dissipates energy.




Strong result: Independent of technology.


How can we limit switching energy?

- (1) Reduce # of clock transitions. But we have work to do ...
- (2) Reduce V_{dd} . But lowering V_{dd} limits the clock speed ...
- (3) Fewer circuits. But more transistors can do more work.
- (4) Reduce C per node. One reason why we scale processes.

Scaling switching energy per gate ...



IC process scaling
("Moore's Law")




Due to reducing V and C (length and width of C s decrease, but plate distance gets smaller).

Recent slope more shallow because V is being scaled less aggressively.

From: "Facing the Hot Chips Challenge Again", Bill Holt, Intel, presented at Hot Chips 17, 2005.

Second Factor: Leakage Currents

Even when a logic gate isn't switching, it burns power.




I_{sub} : Even when this nFET is off, it passes an I_{off} leakage current.

We can engineer any I_{off} we like, but a lower I_{off} also results in a lower I_{on} , and thus a lower maximum clock speed.

I_{gate} : Ideal capacitors have zero DC current. But modern transistor gates are a few atoms thick, and are not ideal.

Intel's 2006 processor designs, leakage vs switching power



65nm

■ Leakage



■ Dynamic

A lot of work was done to get a ratio this good ... 50/50 is common.

Bill Holt, Intel, Hot Chips 17

Engineering “On” Current at 25 nm ...

We can increase I_{on} by raising V_{dd} and/or lowering V_t .


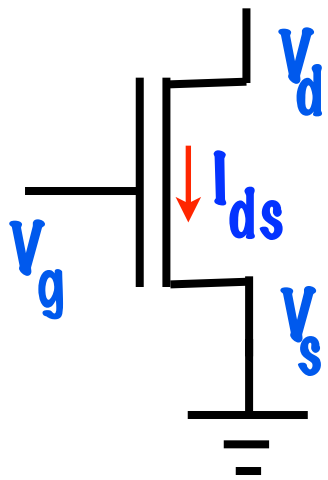


$1.2 \text{ mA} = I_{on}$

$I_{off} = 0$???

Plot on a “Log” Scale to See “Off” Current

We can decrease I_{off} by raising V_t - but that lowers I_{on} .



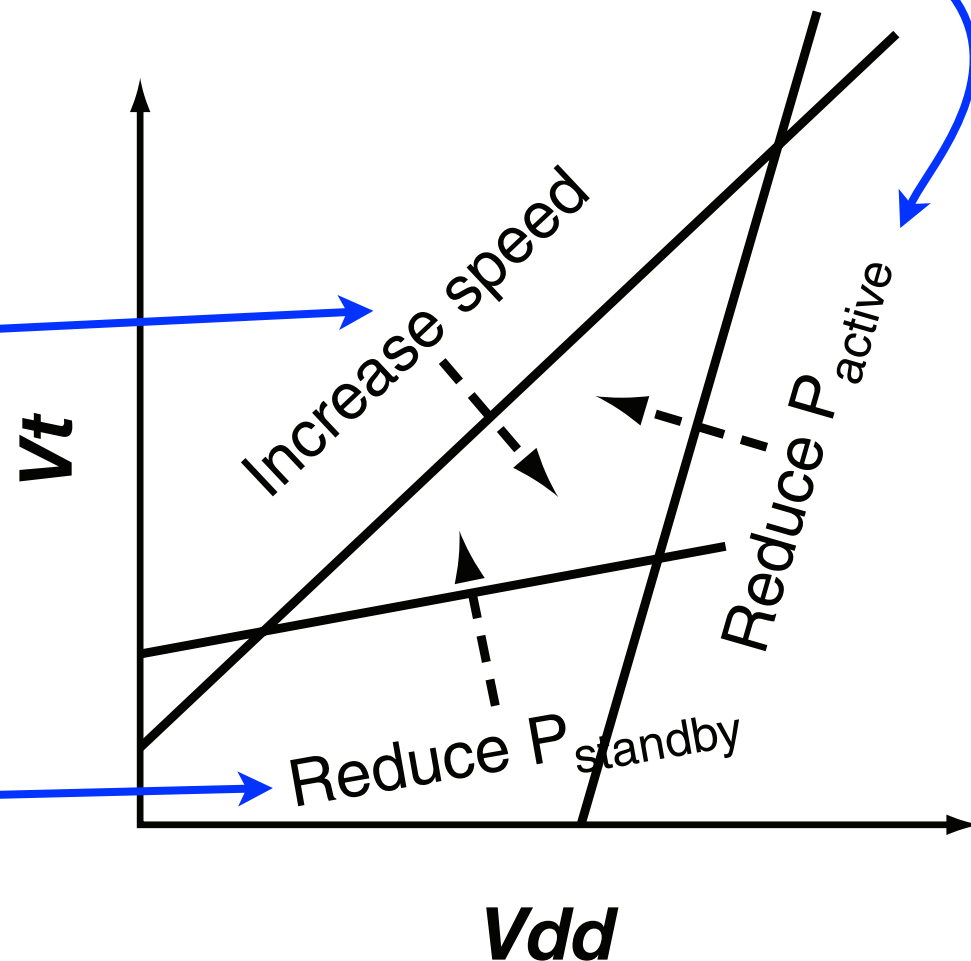
$I_{off} \approx 10 \text{ nA}$

Device engineers trade speed and power

We can reduce CV^2 (P_{active})
by lowering V_{dd} .

We can increase speed
by raising V_{dd} and
lowering V_{t} .

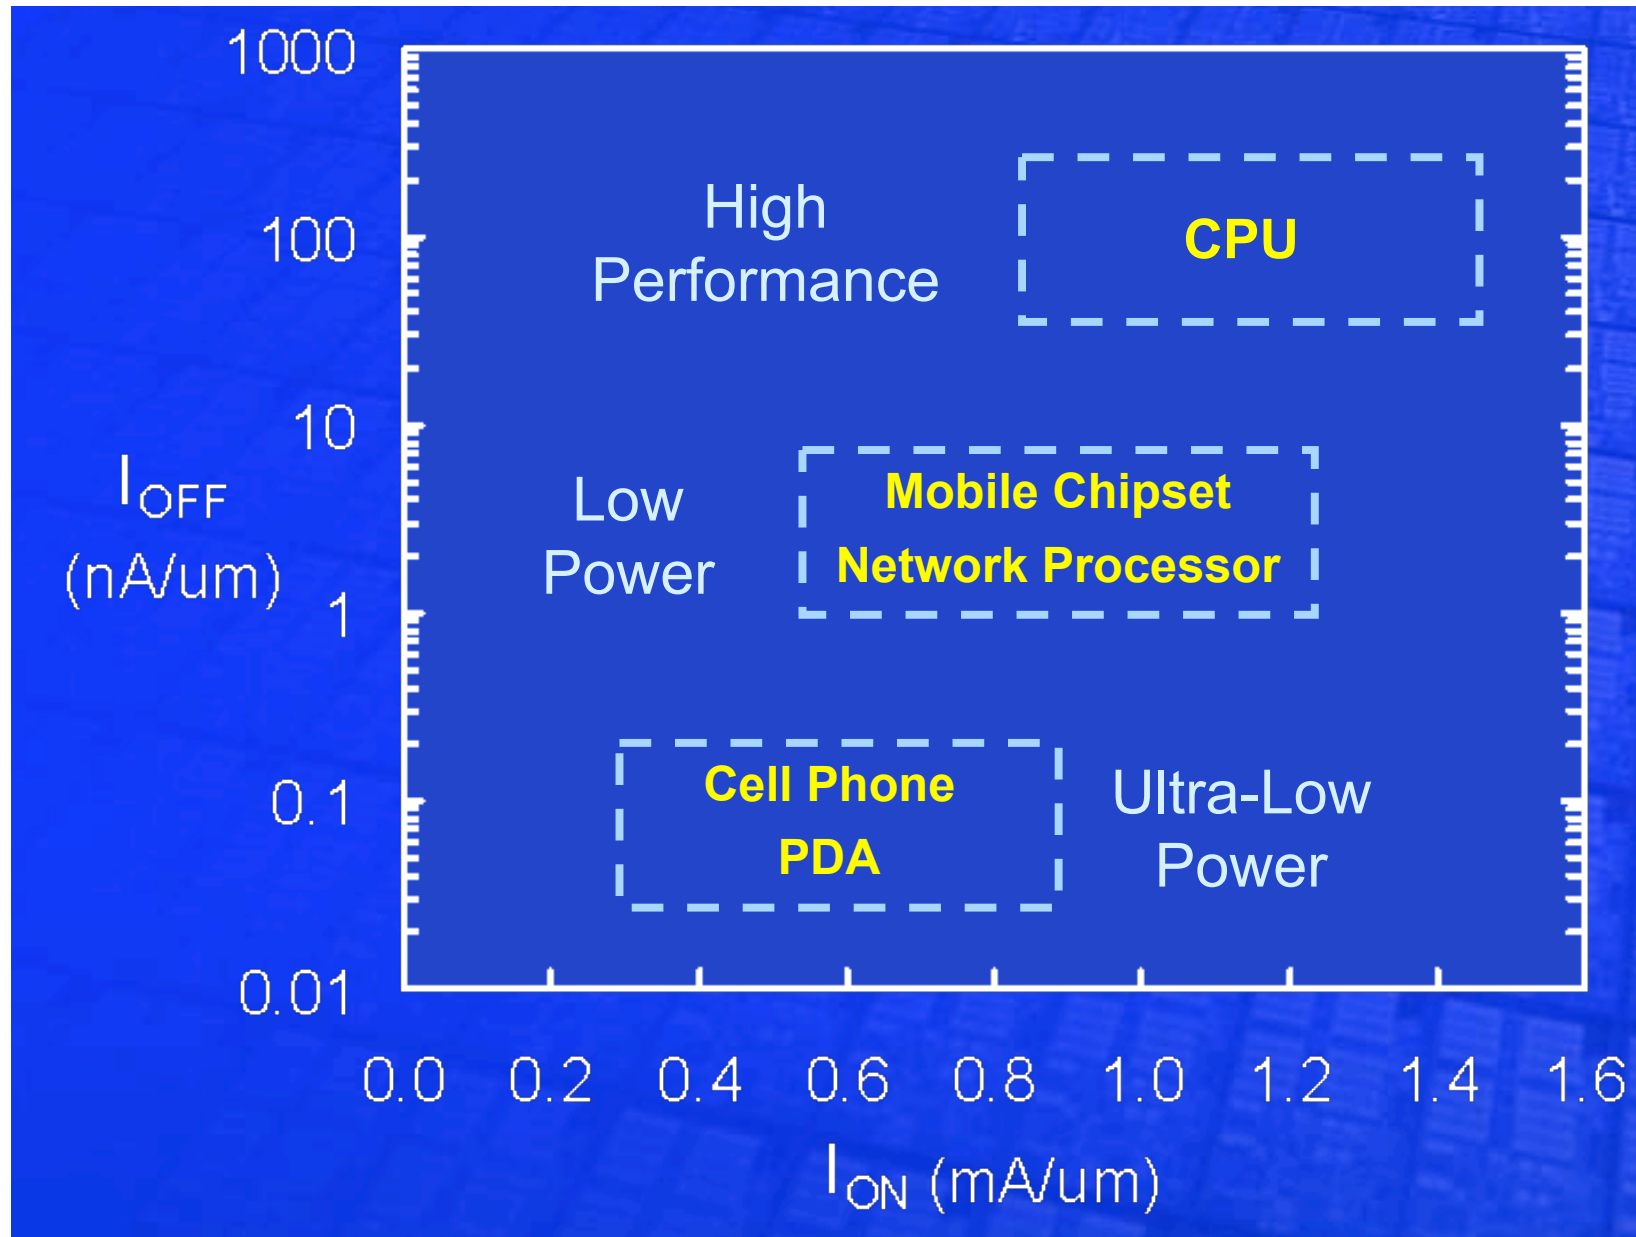
We can reduce leakage
(P_{standby}) by raising V_{t} .



From: Silicon Device Scaling to the Sub-10-nm Regime
Meikei Jeong,^{1*} Bruce Doris,² Jakub Kedzierski,¹ Ken Rim,¹ Min Yang¹

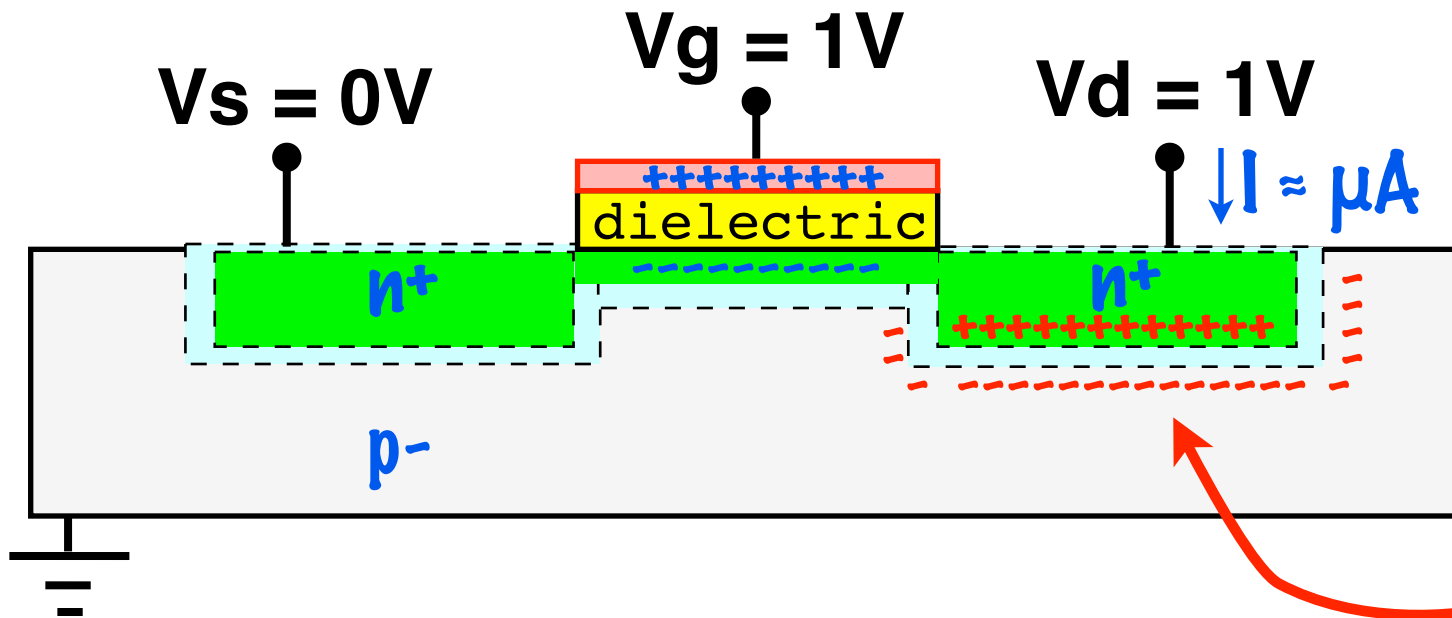


Customize processes for product types ...

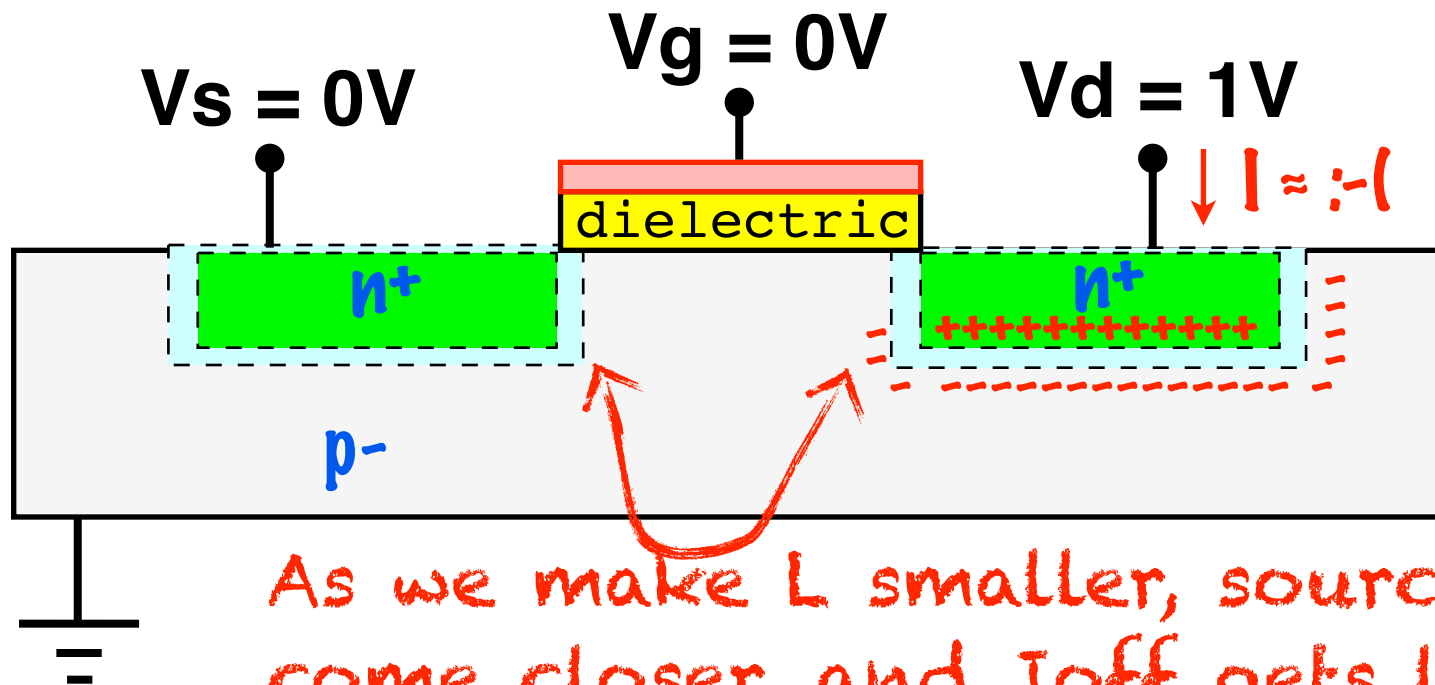


From: "Facing the Hot Chips Challenge Again", Bill Holt, Intel, presented at Hot Chips 17, 2005.

Transistor physics revisited ...



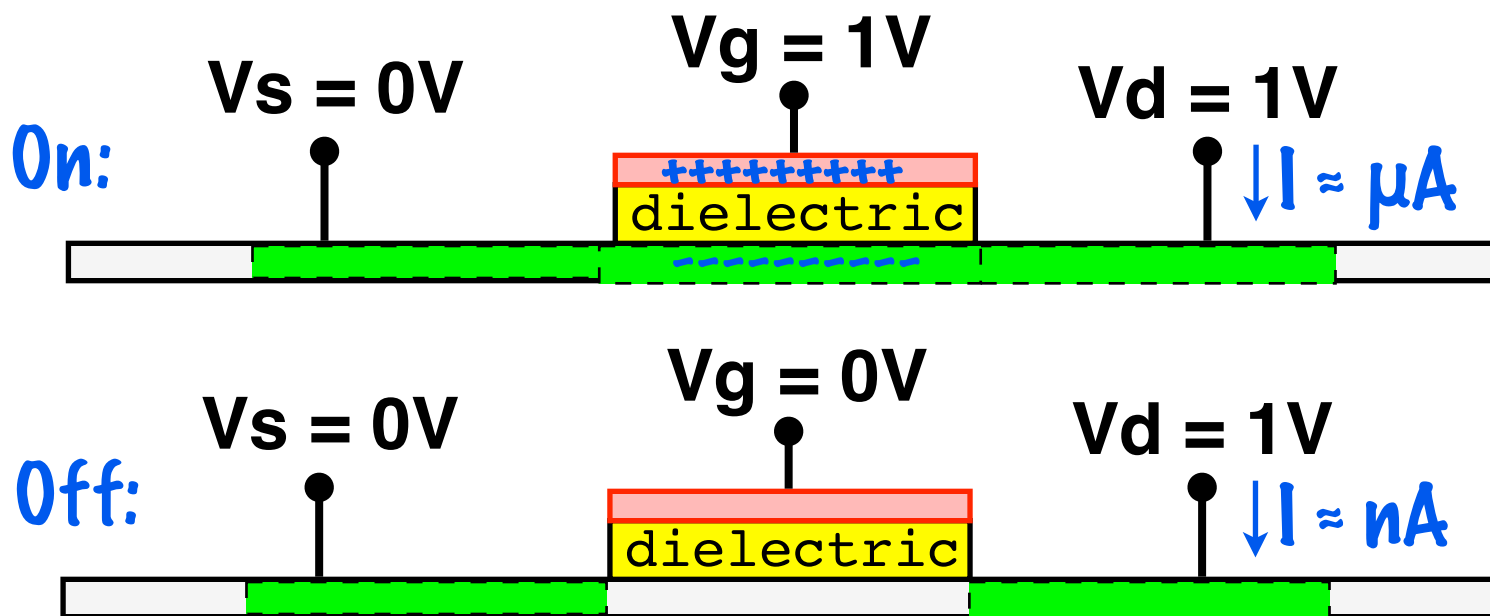
The drain junction is also a capacitor, and puts - charges in the substrate.



Away from the surface, the drain-induced charges remain even when the gate is off!

As we make L smaller, source and drain come closer, and I_{off} gets larger!

Solution concept: Fully-depleted channel




We limit the depth of the channel so that the gate voltage “wins” over the drain voltage.

Done as shown, **5 to 7 nm** depth for a 20 nm transistor.

Requires expensive wafers


“FD-SOI” -- Fully-Depleted Silicon-On-Insulator



Transistor channel is a raised fin.

Gate controls channel from sides and top.

Channel depth is fin width. 12-15nm for $L=22\text{nm}$.




(12) **United States Patent**
 Hu et al. Filed: Oct. 23, 2000

(54) **FINFET TRANSISTOR STRUCTURES HAVING A DOUBLE GATE CHANNEL EXTENDING VERTICALLY FROM A SUBSTRATE AND METHODS OF MANUFACTURE**

(75) Inventors: **Chenming Hu**, Alamo; **Tsu-Jae King**, Fremont; **Vivek Subramanian**, Redwood City; **Leland Chang**, Berkeley; **Xuejue Huang**; **Yang-Kyu Choi**, both of Albany; **Jakub Tadeusz Kedzierski**, Hayward; **Nick Lindert**, Berkeley; **Jeffrey Bokor**, Oakland, all of CA (US); **Wen-Chin Lee**, Beaverton, OR (US)

Intel "Ivy Bridge" 22nm CPUs, first production parts




3-D Tri-Gates

Source/Drain Contacts

Sandy
Bridge

32nm
planar

1.16B
transistors




"Less than half the power @ same performance"

Ivy Bridge


22nm
FinFet

1.4B
transistors





Long-term possibility: New devices

Electrostatic mechanical relays at the nanoscale



Working inverter at 500 kHz ... for a while.



- + 10 fA leakage current
- + Works at 500 degrees C
- Fails after 1-10 days of 500 kHz toggles.
- Switching requires 6V V_{dd}

Five low-power design techniques

- ✱ **Parallelism and pipelining**
- ✱ **Power-down idle transistors**
- ✱ **Slow down non-critical paths**
- ✱ **Clock gating**
- ✱ **Thermal management**

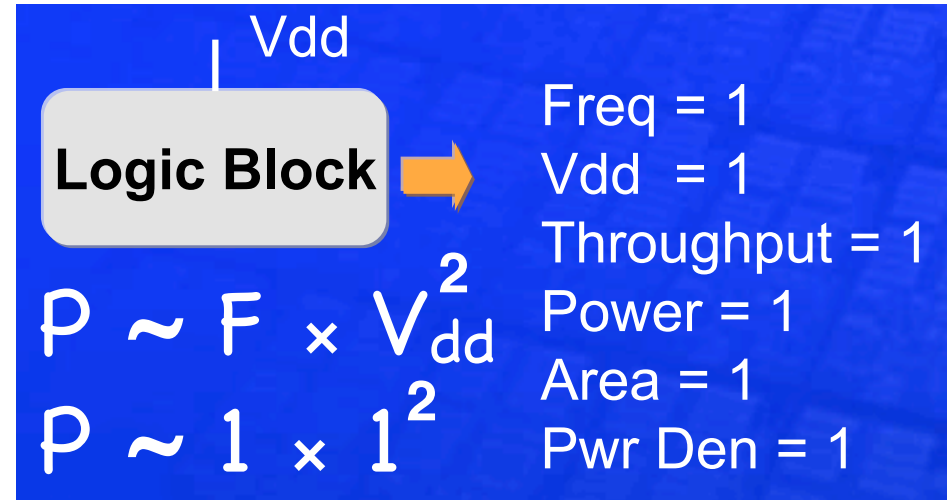
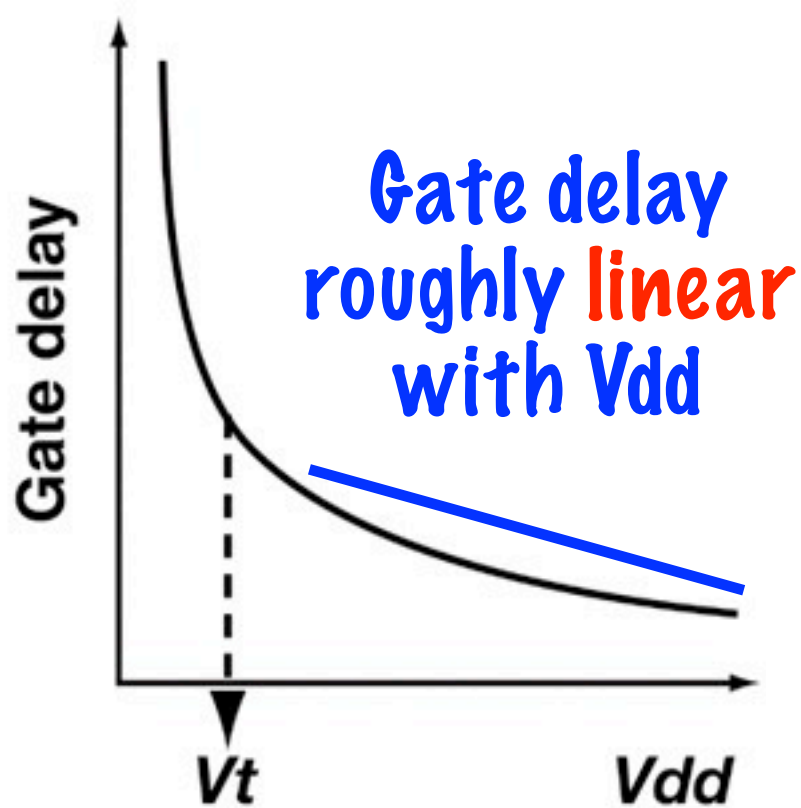


Trading Hardware for Power

via Parallelism and Pipelining ...



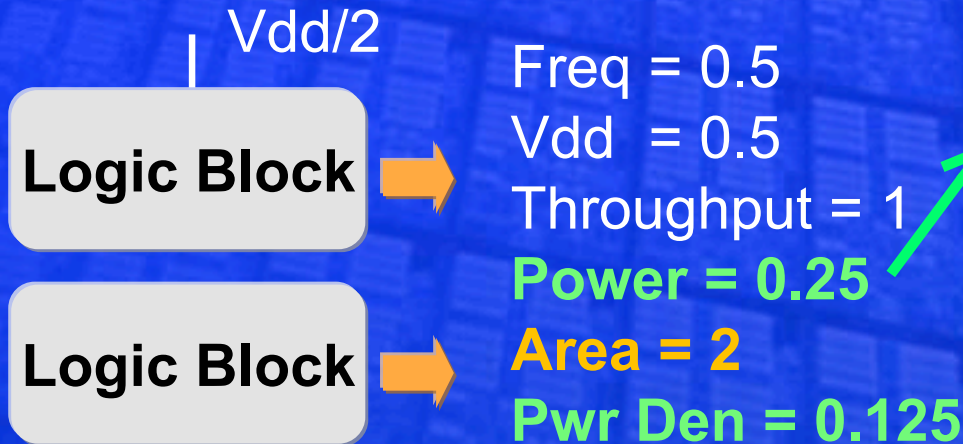
And so, we can transform this:



Block processes stereo audio. 1/2 of clocks for "left", 1/2 for "right".

Into this:

Top block processes "left", bottom "right".



$$P \sim \#blks \times F \times V_{dd}^2$$

$$P \sim 2 \times 1/2 \times 1/4 = 1/4$$

CV² power only

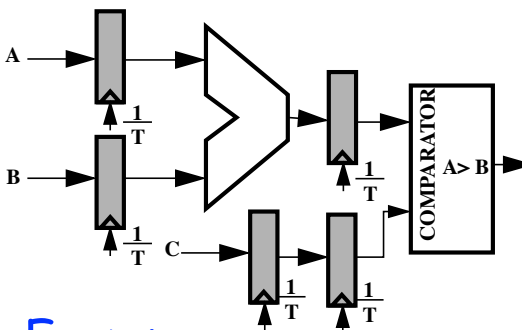


THIS MAGIC TRICK BROUGHT TO YOU BY CORY HALL ...

Chandrakasan & Brodersen (UCB, 1992)

Architecture	Power (normalized)
Simple	1
Parallel	0.36
Pipelined	0.39
Pipelined-Parallel	0.2


Architecture	Area (normalized)
Simple	1
Parallel	3.4
Pipelined	1.3
Pipelined-Parallel	3.7

Architecture	Voltage
Simple	5V
Parallel	2.9V
Pipelined	2.9V
Pipelined-Parallel	2.0




From:


Minimizing Power Consumption in CMOS Circuits



Simple



Parallel



Pipelined

Multiple Cores for Low Power

Trade hardware for power,
on a large scale ...



Cell: The PS3 chip




SONY



**COMPUTER
ENTERTAINMENT**

TOSHIBA




2006

Cell (PS3 Chip): 1 CPU + 8 “SPUs”


L2 Cache
512 KB

8
Synergistic
Processing
Units
(SPUs)

PowerPC



One Synergistic Processing Unit (SPU)



SPU issues 2 inst/cycle (in order) to 7 execution units

256 KB Local Store, 128 128-bit Registers

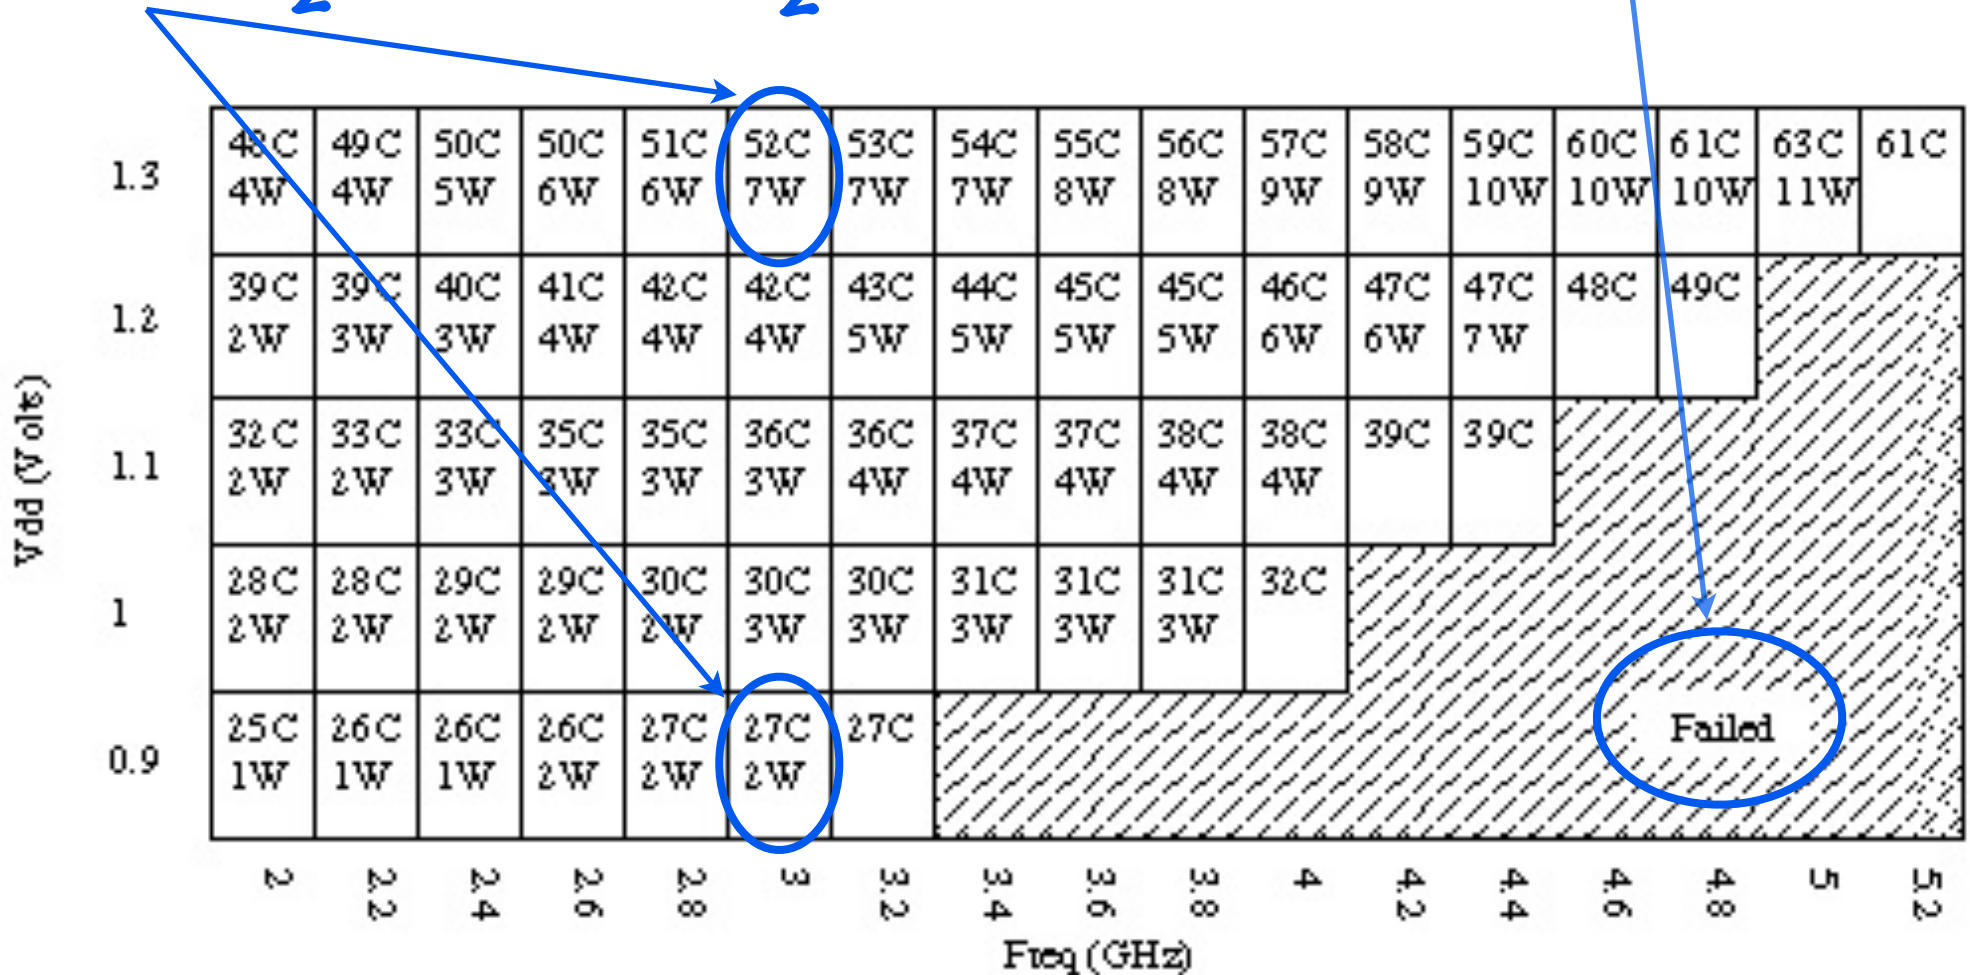
SPU fills Local Store using DMA to DRAM and network

A “Schmoo” plot for a Cell SPU ...

The lower V_{dd}, the less dynamic energy consumption.

The lower V_{dd}, the longer the maximum clock period, the slower the clock frequency.


$$E_{0 \rightarrow 1} = \frac{1}{2} C V_{dd}^2 \quad E_{1 \rightarrow 0} = \frac{1}{2} C V_{dd}^2$$



Clock speed alone doesn't help E/op ...

But, lowering clock frequency while keeping voltage constant **spreads** the same amount of work over a longer time, so chip stays **cooler** ...

$$E_{0 \rightarrow 1} = \frac{1}{2} C V_{dd}^2 \quad E_{1 \rightarrow 0} = \frac{1}{2} C V_{dd}^2$$




Scaling V and f does lower energy/op

1 W to get 2.2 GHz performance. 26 C die temp.

7W to reliably get 4.4 GHz performance. 47C die temp.


If a program that needs a 4.4 GHz CPU can be recoded to use two 2.2 GHz CPUs ... big win.



How iPod nano 2005 puts its 2 cores to use ...



PP5020 
digital media management system-on-chip



Dual ARM Processors

- Dual 32-bit ARM7TDMI processors
- Up to 80 MHz processor operation per core with independent clock-skipping feature on COP
- Efficient cross-bar implementation providing zero wait state access to internal RAM
- Integrated 96KB of SRAM
- 8KB of unified cache per processor
- Six DMA channels


Two 80 MHz CPUs.
Was used in several
nano generations,
with one CPU doing
audio decoding, the
other doing photos,
etc.

Powering down idle circuits



Add “sleep” transistors to logic ...

Sleep Transistor



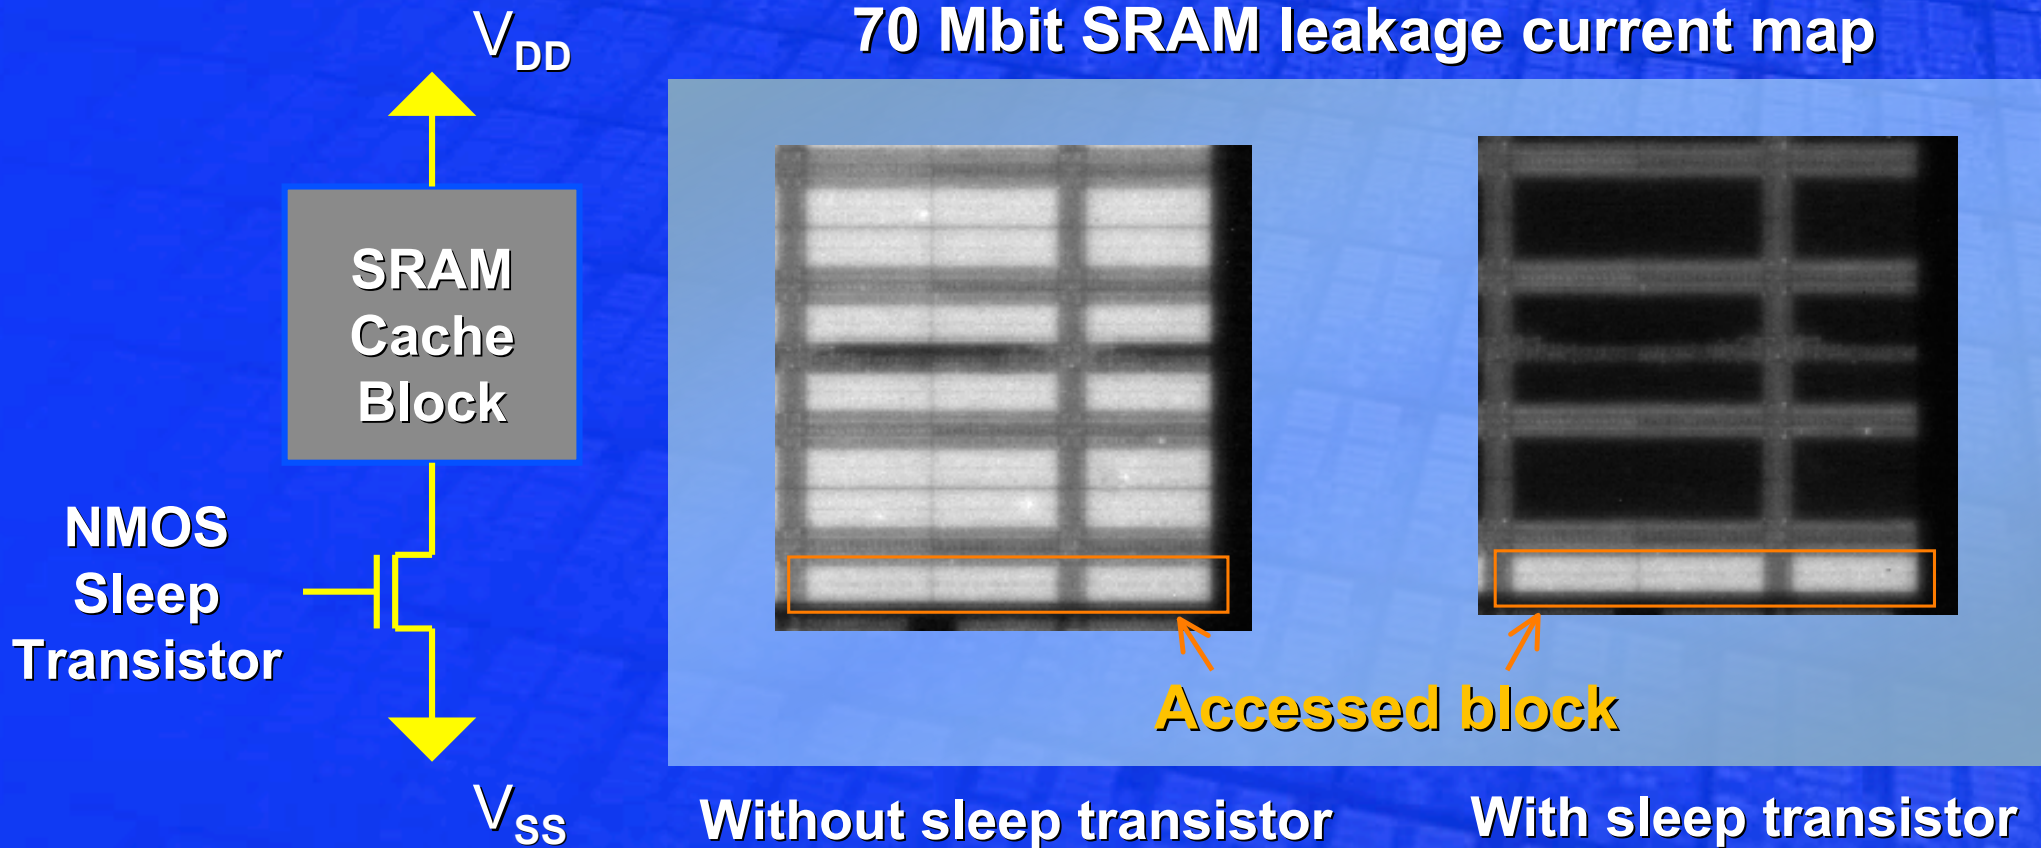
Example: Floating point unit logic.

When running fixed-point instructions, put logic “to sleep”.

+++ When “asleep”, leakage power is dramatically reduced.

--- Presence of sleep transistors slows down the clock rate when the logic block is in use.

Intel example: Sleeping cache blocks



>3x SRAM leakage reduction on inactive blocks


Slow down “slack paths”



Fact: Most logic on a chip is “too fast”

The critical path

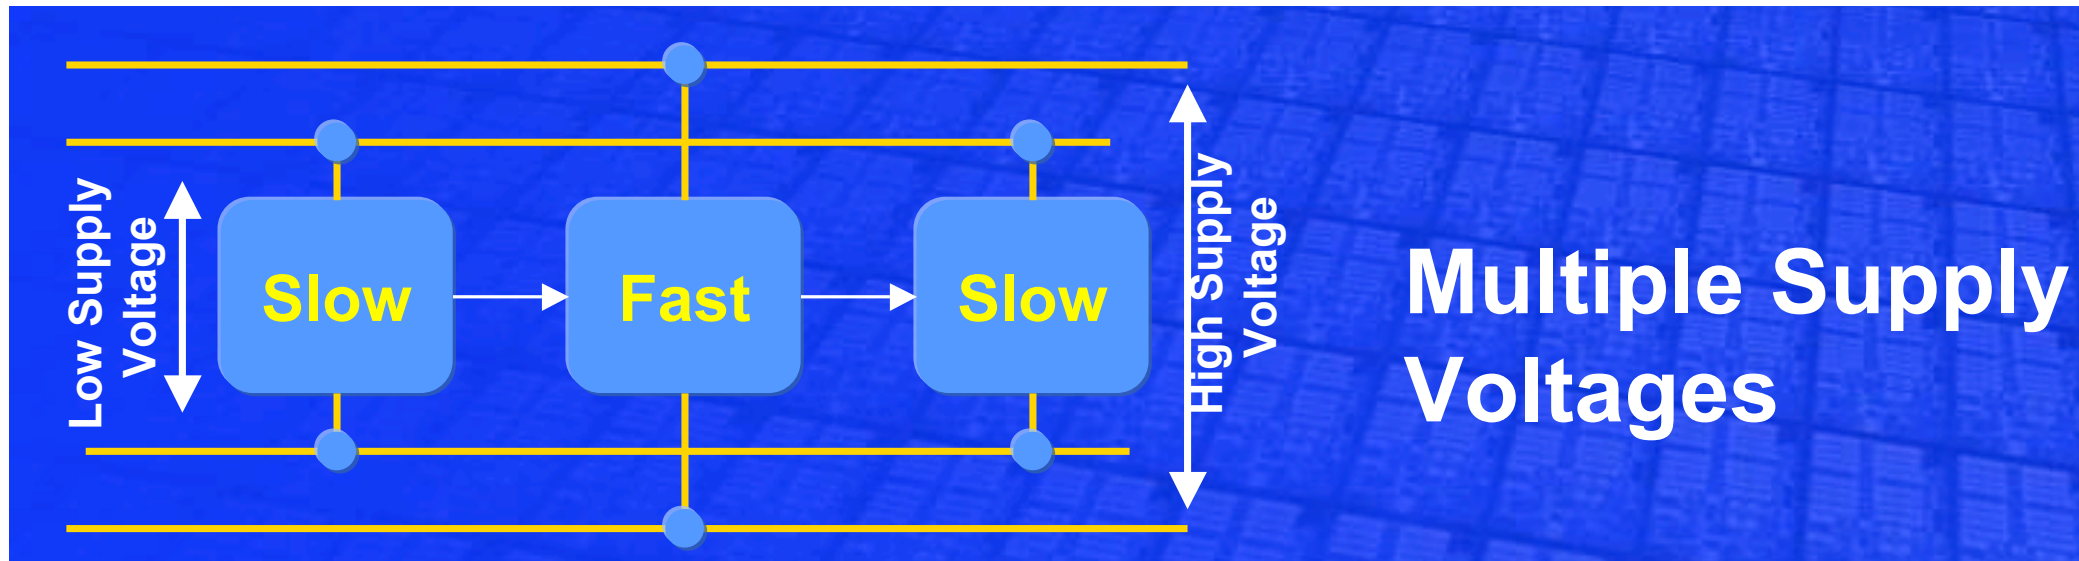
Most wires have hundreds of picoseconds to spare.



From “The circuit and physical design of the POWER4 microprocessor”, IBM J Res and Dev, 46:1, Jan 2002, J.D. Warnock et al.



Use several supply voltages on a chip ...



Why use multi-V_{dd}? We can reduce **dynamic** power by using low-power V_{dd} for logic off the critical path.

What if we can't do a multi-V_{dd} design?

In a multi-V_t process, we can reduce **leakage** power on the slow logic by using high-V_{th} transistors.

From: "Facing the Hot Chips Challenge Again", Bill Holt, Intel, presented at Hot Chips 17, 2005.

LOW POWER ARM 1136JF-S™ DESIGN


George Kuo, Anand Iyer
Cadence Design Systems, Inc.
San Jose, CA 95134, USA

Logical partition into
0.8V and 1.0V nets
done manually to meet
350 MHz spec (90nm).

Level-shifter insertion
and placement done
automatically.

Dynamic power in 0.8V
section cut 50% below
baseline.

Leakage power in 1.0V
section cut 70% below
baseline.




From a chapter from new book on ASIC design by Chinnery and Keutzer (UCB).

Gating clocks to save power



On a CPU, where does the power go?



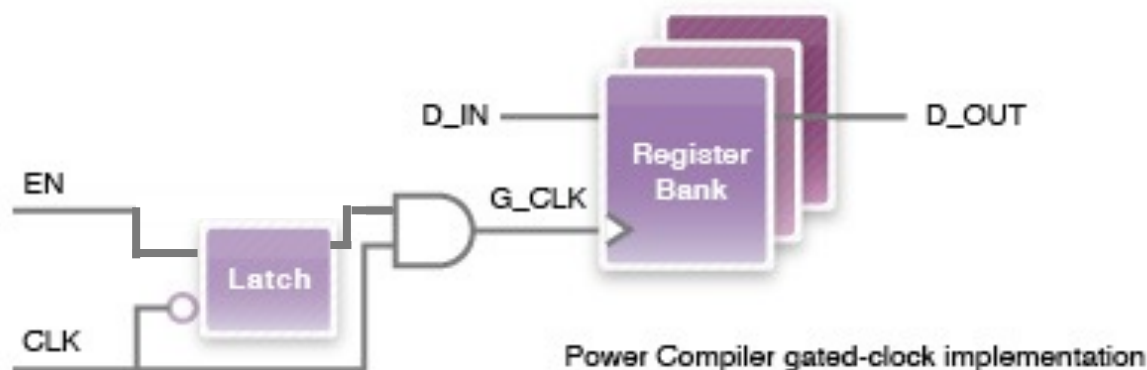
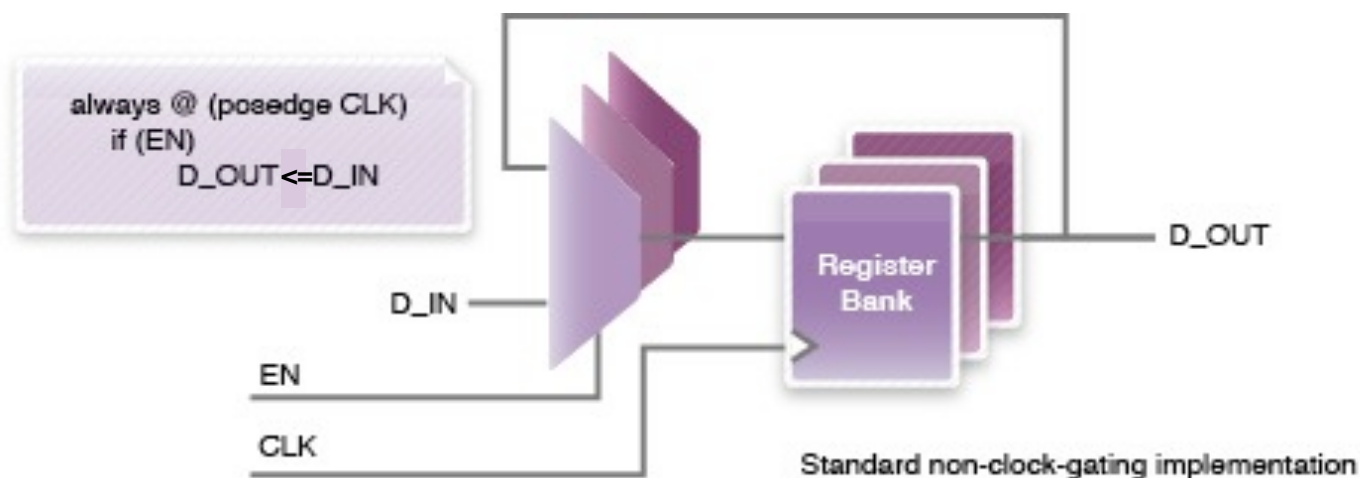
Half of the power goes to latches (Flip-Flops).

Most of the time, the latches don't change state.

So (gasp) gated clocks are a big win.
But, done with CAD tools in a disciplined way.



Synopsis Power Compiler can do this ...



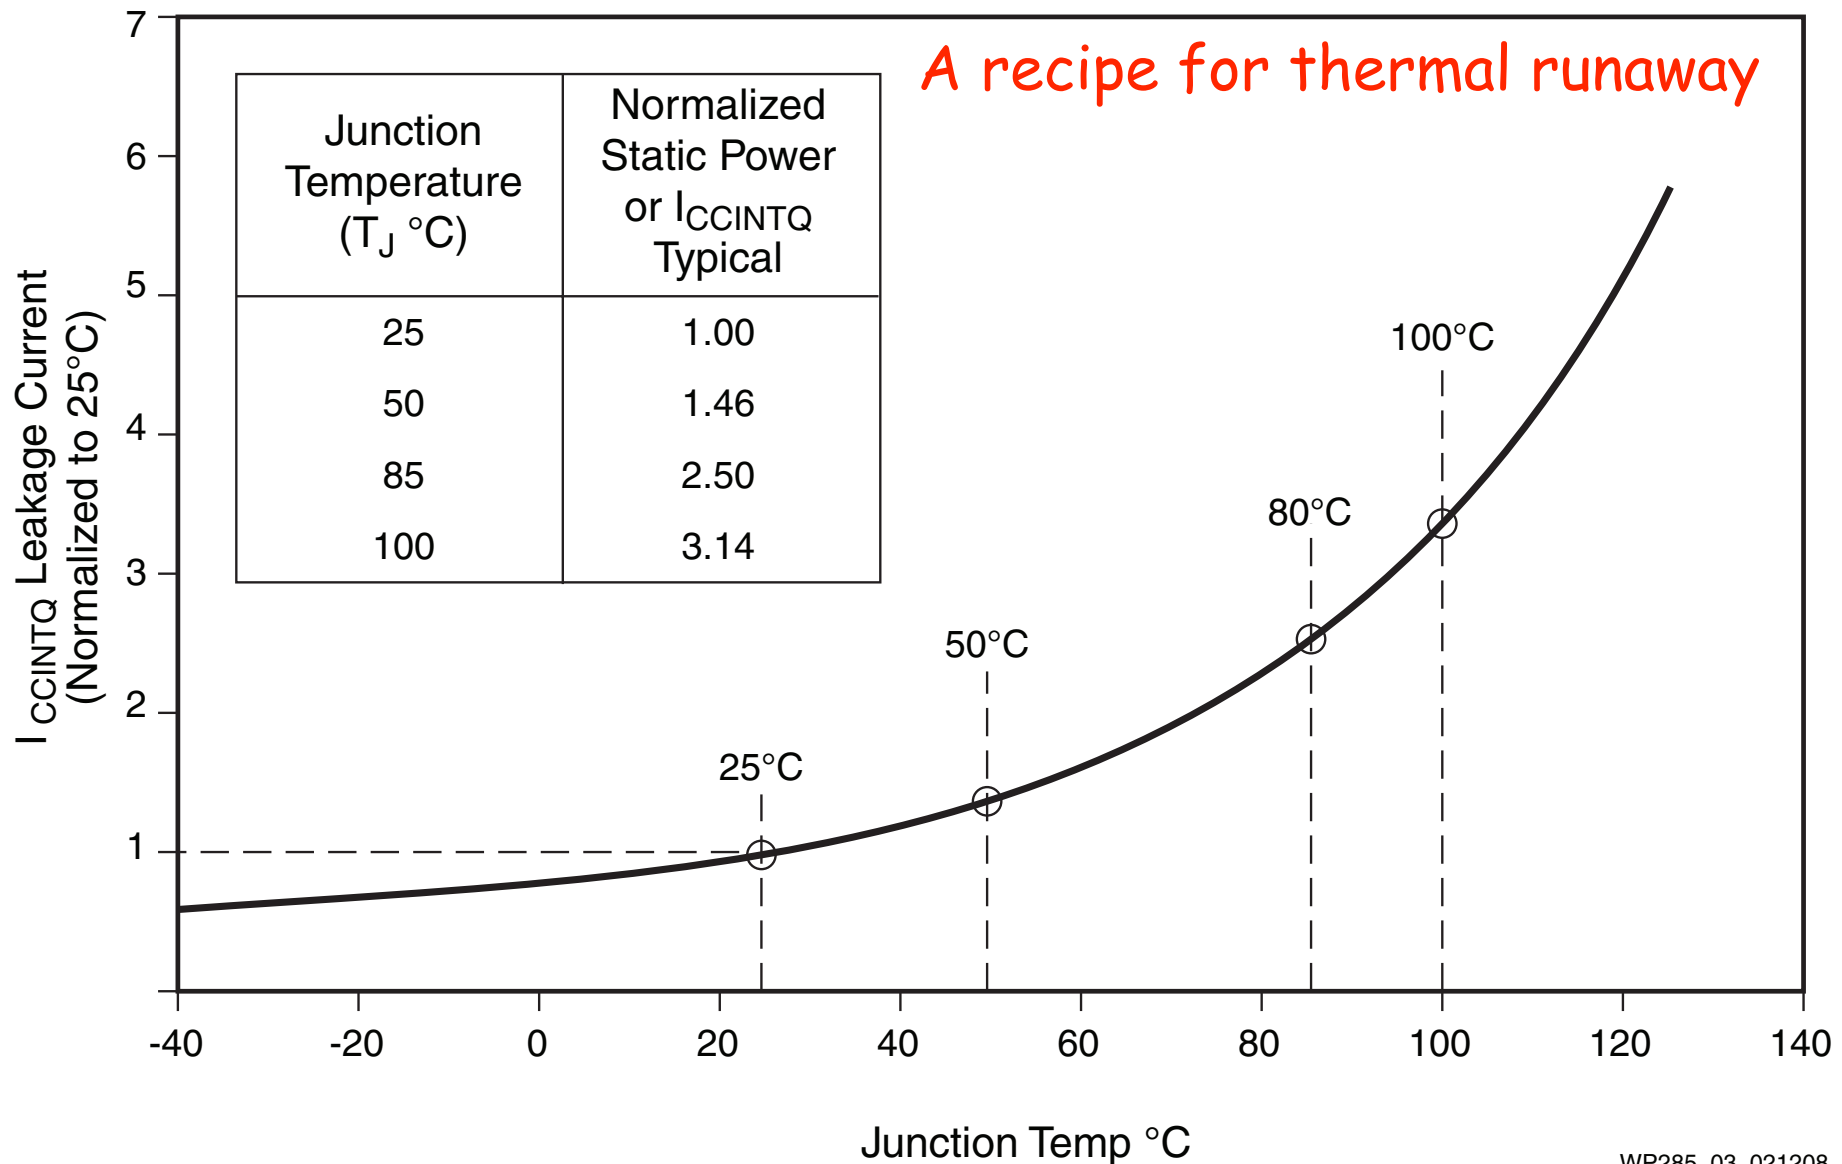
“Up to 70%
power savings
at the block
level, for
applicable
circuits”

Synopsis Data
Sheet

Thermal Management




Keep chip cool to minimize leakage power



WP285_03_021208


Figure 3: I_{CCINTQ} vs. Junction Temperature with Increase Relative to 25°C

IBM Power 4: How does die heat up?




↑
**4 dies on a
multi-chip
module**

**2 CPUs
per die** →




115 Watts: Concentrated in “hot spots”



66.8 C == 152 F

82 C == 179.6

Idea: Monitor temperature, servo clock speed



Thursday's lecture: Graphics ...



Thu
4/19

Lec #26: Graphics Processing Chips ([Slides](#))

