

# CS250

## VLSI Systems Design

### Lecture 5: Physical Realities: Beneath the Digital Abstraction, Part 2: Power & Energy

---

Spring 2016

John Wawrzynek  
with  
Chris Yarp (GSI)

**Thanks to John Lazaro for the slides**

**The Watt:**  
Unit of power.  
A rate of  
energy (J/s).  
A gas pump  
hose delivers  
**6 MW.**

**The Joule:** Unit of  
energy. A **1 Gallon**  
gas container holds  
**130 MJ** of energy.



**120 KW:** The power  
delivered by a  
**Tesla Supercharger.**  
Tesla Model S has a  
**306 MJ** battery  
(good for 265 miles).

$$1 \text{ J} = 1 \text{ W} * \text{s} \quad 1 \text{ W} = 1 \text{ J/s.}$$

**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.

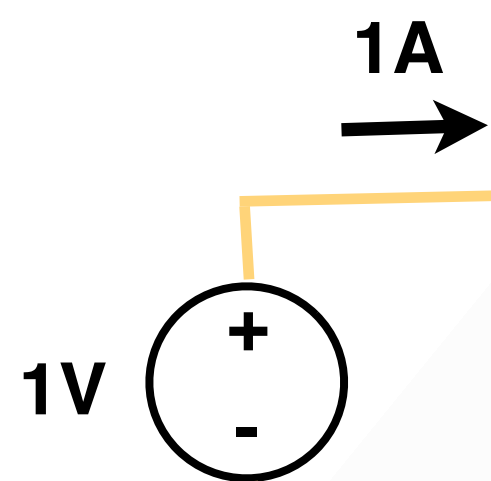
This is how electric tea pots work ...

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.

$1.2 \text{ W-hr} / 5 \text{ W} \approx 15 \text{ minutes.}$

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 ...

# What's happened since 2005?

2010 nano  
0.74 ounces  
(50% of  
2005 Nano)



"Up to" 24 hours  
audio playback.  
70% improvement  
from 2005 nano.



0.39 W Hr  
(33% of 2005 Nano)





2015  
**Apple**  
Watch




**3.8 V, 0.78 Wh lithium-ion battery on 38mm model. Apple claims the 205 mAh battery should provide up to 18 hours of use (which translates to 6.5 hours of audio playback, 3 hours of talk time, or 72 hours of Power Reserve mode.)**






Reality                      Layer

A clever prism projects a layer over reality light.



2.1 Wh battery - 5.3x as much energy as 2010 Nano.


Battery life very usage dependent.



640 x 360  
Liquid Crystal  
on Silicon  
(LCoS) prism  
projector.

1.76 ounces -  
4X the weight  
of iPod Shuffle

Logic Board



# iPhone6



## 4.7 inch iPhone6: 1,810mAh battery




## iPhone 5s: 1570mAh

- The front side of the logic board:

- Apple A8 APL1011 SoC + SK Hynix RAM as denoted by the markings H9CKNNN8KTMRWR-NTH (we presume it is 1 GB LPDDR3 RAM, the same as in the iPhone 6 Plus)
- Qualcomm MDM9625M LTE Modem
- Skyworks 77802-23 Low Band LTE PAD
- Avago A8020 High Band PAD
- Avago A8010 Ultra High Band PA + FBARs
- SkyWorks 77803-20 Mid Band LTE PAD
- InvenSense MP67B 6-axis Gyroscope and Accelerometer Combo

The A8 is manufactured on a 20 nm process by TSMC. It contains 2 billion transistors. Its physical size is 89 mm<sup>2</sup>.<sup>1</sup> It has 1 GB of LPDDR3 RAM included in the package. It is dual core, and has a frequency of 1.38 GHz.



# 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 ...



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.

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!

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






MacBook Air ... design the laptop like an iPod



Mainboard: fills about 25% of the laptop



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

# MacBook Air: Full PC




Thunderbolt I/O

Platform  
Controller  
Hub


Core i5  
CPU/GPU

Top



Up to  
4GB DRAM

Bottom



# Servers: Total Cost of Ownership (TCO)

---




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.

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

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

(Jonathan Koomey, Stanford).



# Processors and Energy

---



2.6 Billion

# Moore's Law

2,600,000,000

1,000,000,000

100,000,000

10,000,000

1 Million

1,000,000


100,000

10,000

2,300

Transistor count


curve shows transistor count doubling every two years



2 Thousand

# Switching Energy: Fundamental Physics

Every logic transition dissipates energy.




Strong result: Independent of technology.

How can we limit switching energy?

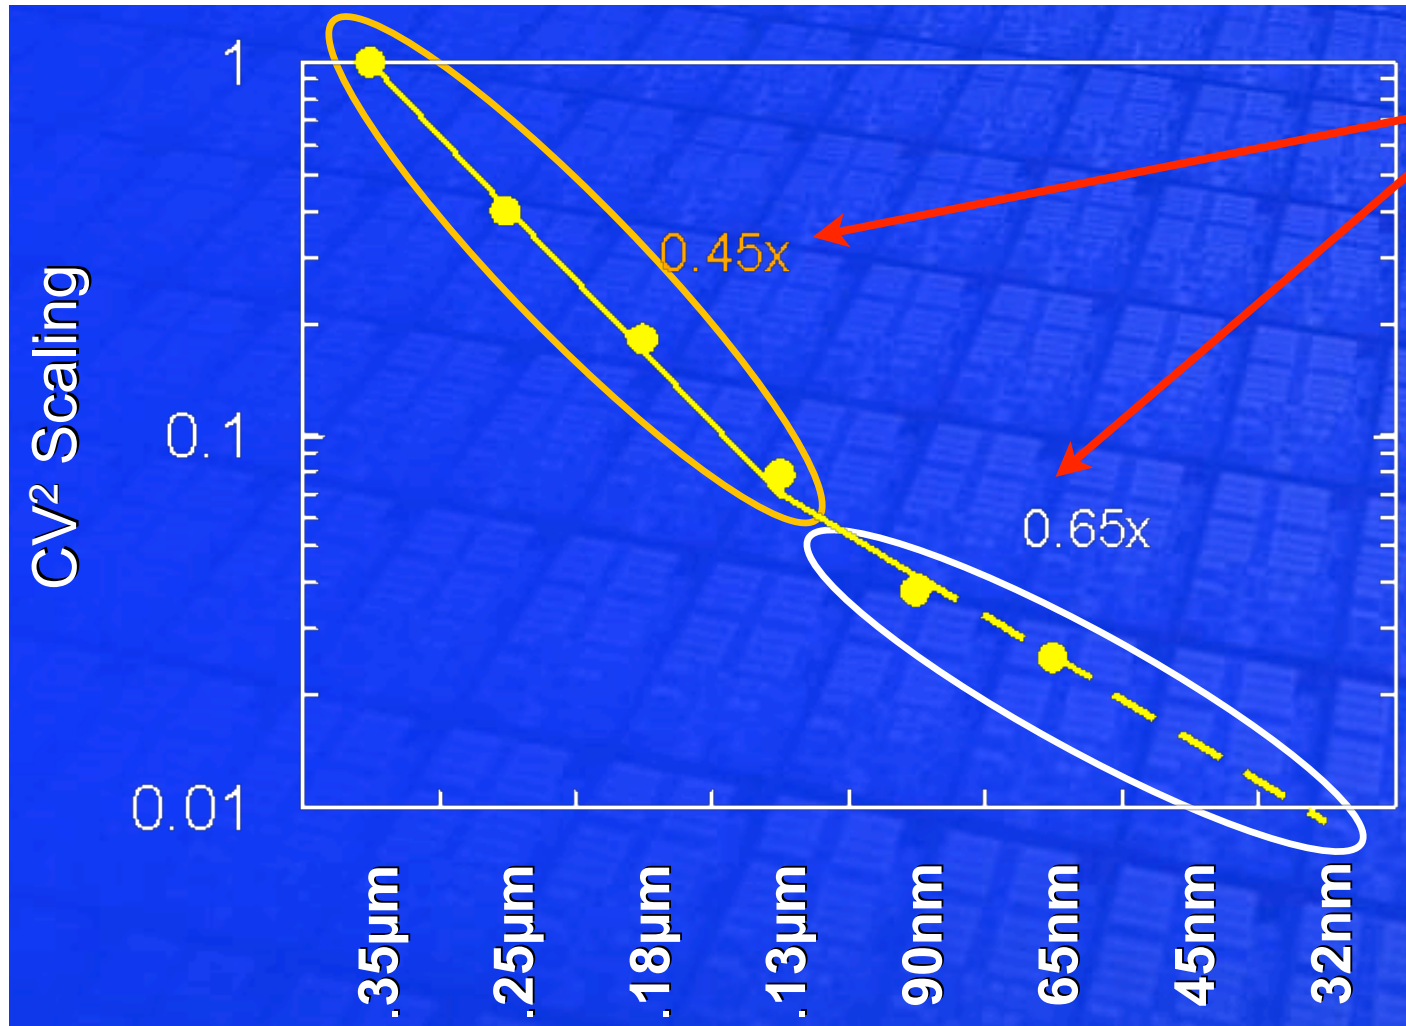
- (1) Reduce # of clock transitions. But we have work to do ...
- (2) Reduce  $V_{dd}$ . But lowering  $V_{dd}$  limits the clock speed ...
- (3) Fewer circuits. But more transistors can do more work.
- (4) Reduce  $C$  per node. One reason why we scale processes.



# 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>sub</sub>:** Even when this nfet is off, it passes an **I<sub>off</sub>** leakage current.

We can engineer any **I<sub>off</sub>** we like, but a **lower I<sub>off</sub>** also results in a **lower I<sub>on</sub>**, and thus a lower maximum clock speed.

**I<sub>gate</sub>:** 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

UC Regents Spring 2016 © UCB

# Engineering “On” Current at 25 nm ...

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


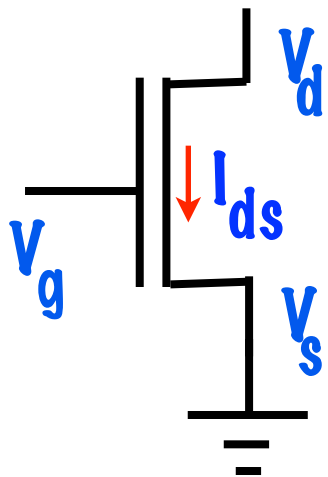


$I_{off} = 0$  ???

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

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

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



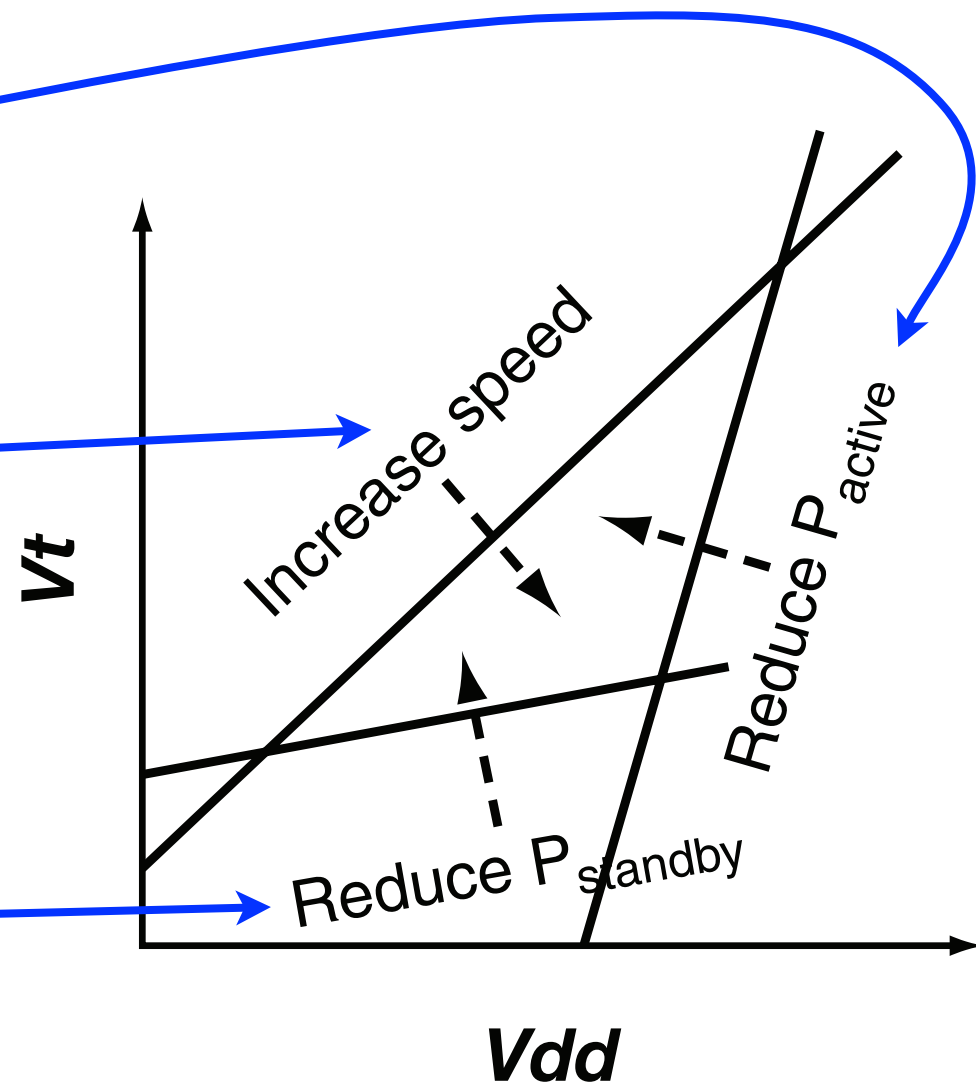
$I_{off} \approx 10$  nA

# Device engineers trade speed and power

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

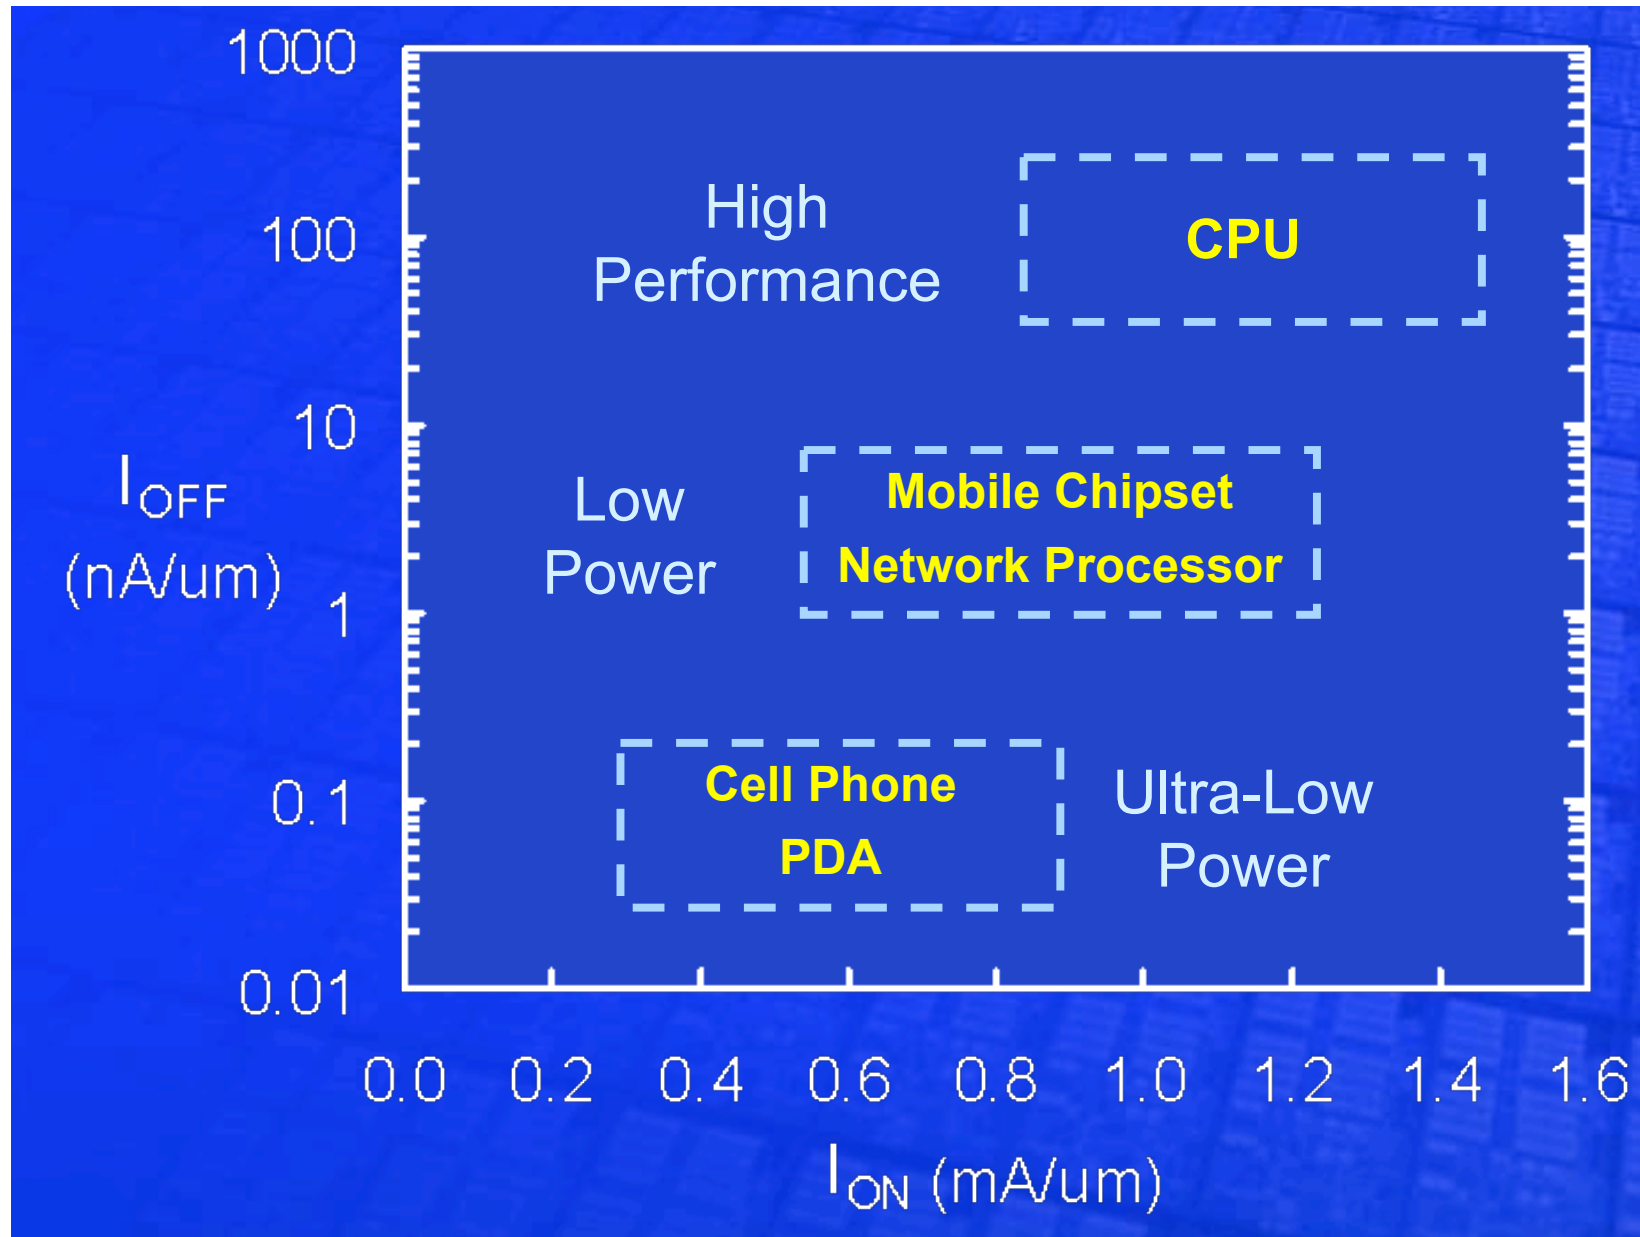
We can increase speed  
by raising  $V_{\text{dd}}$  and  
lowering  $V_{\text{t}}$ .

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



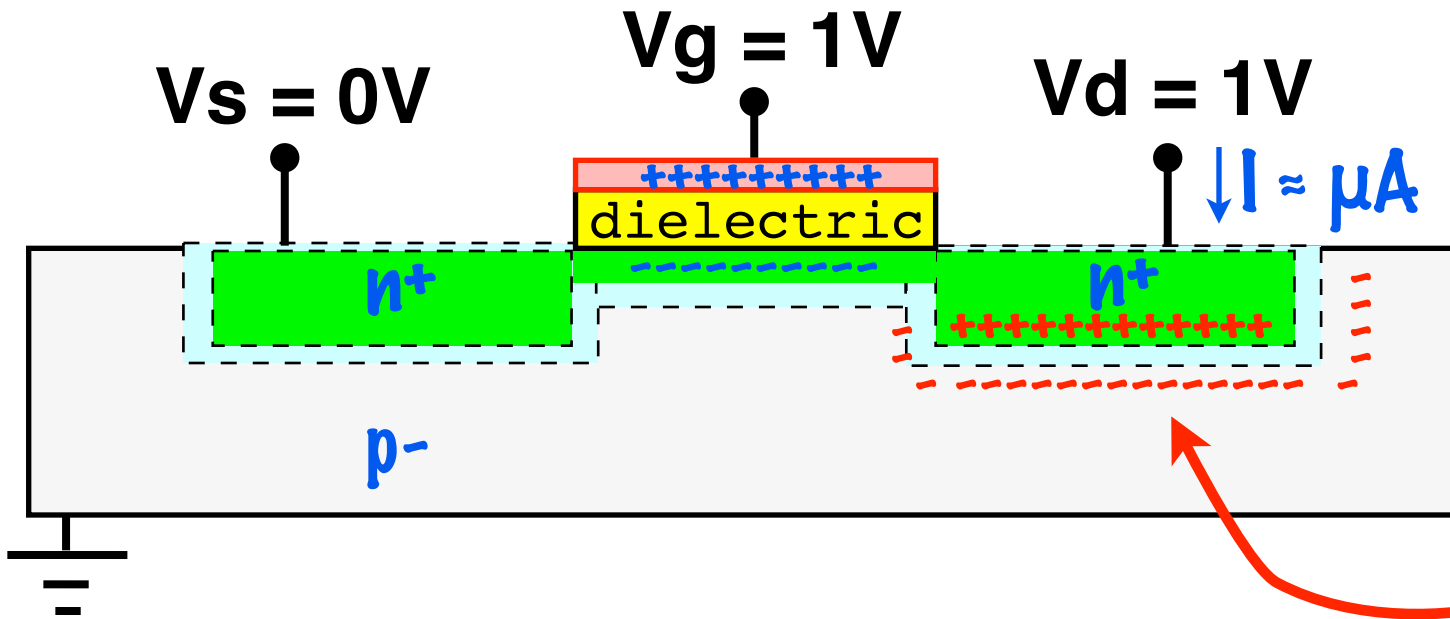
From: Silicon Device Scaling to the Sub-10-nm Regime  
Meikei Jeong,<sup>1\*</sup> Bruce Doris,<sup>2</sup> Jakub Kedzierski,<sup>1</sup> Ken Rim,<sup>1</sup> Min Yang<sup>1</sup>

# Customize processes for product types ...

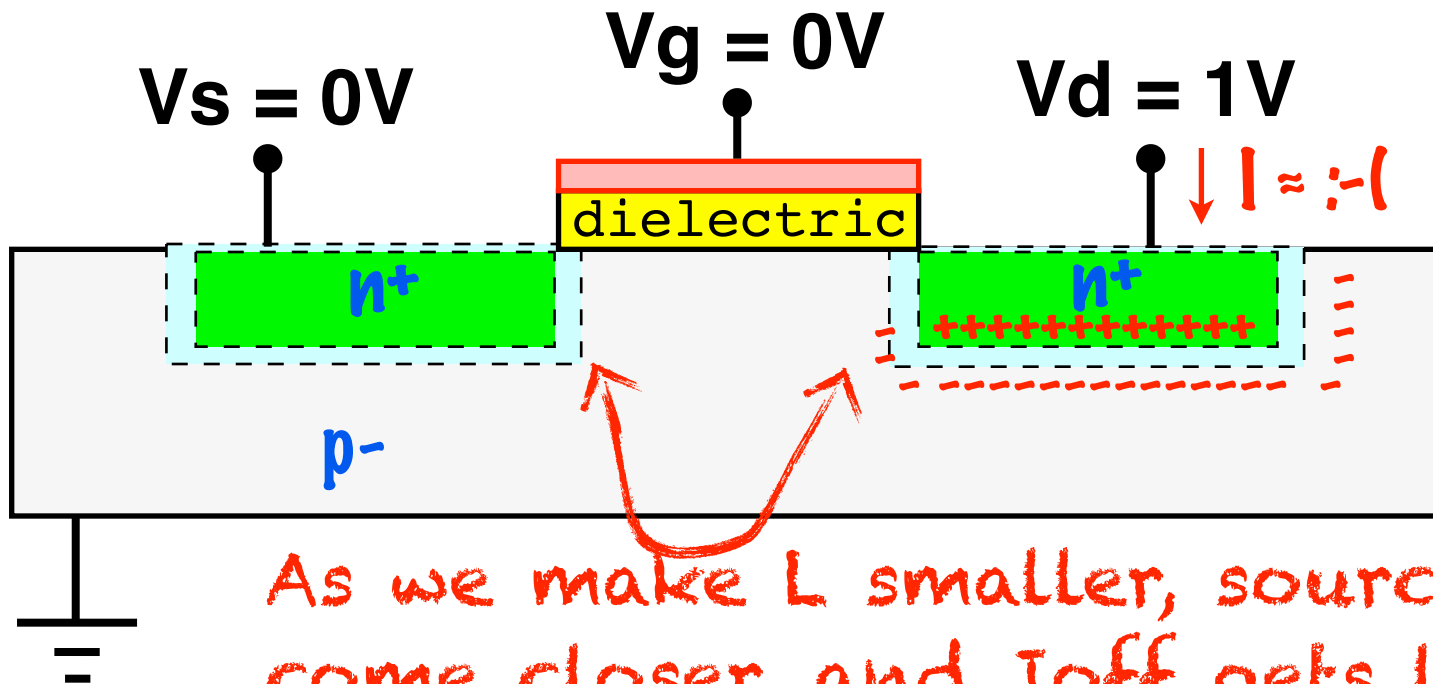


From: "Facing the Hot Chips Challenge Again", Bill Holt, Intel, presented at Hot Chips 17, 2005.

# Transistor physics revisited ...



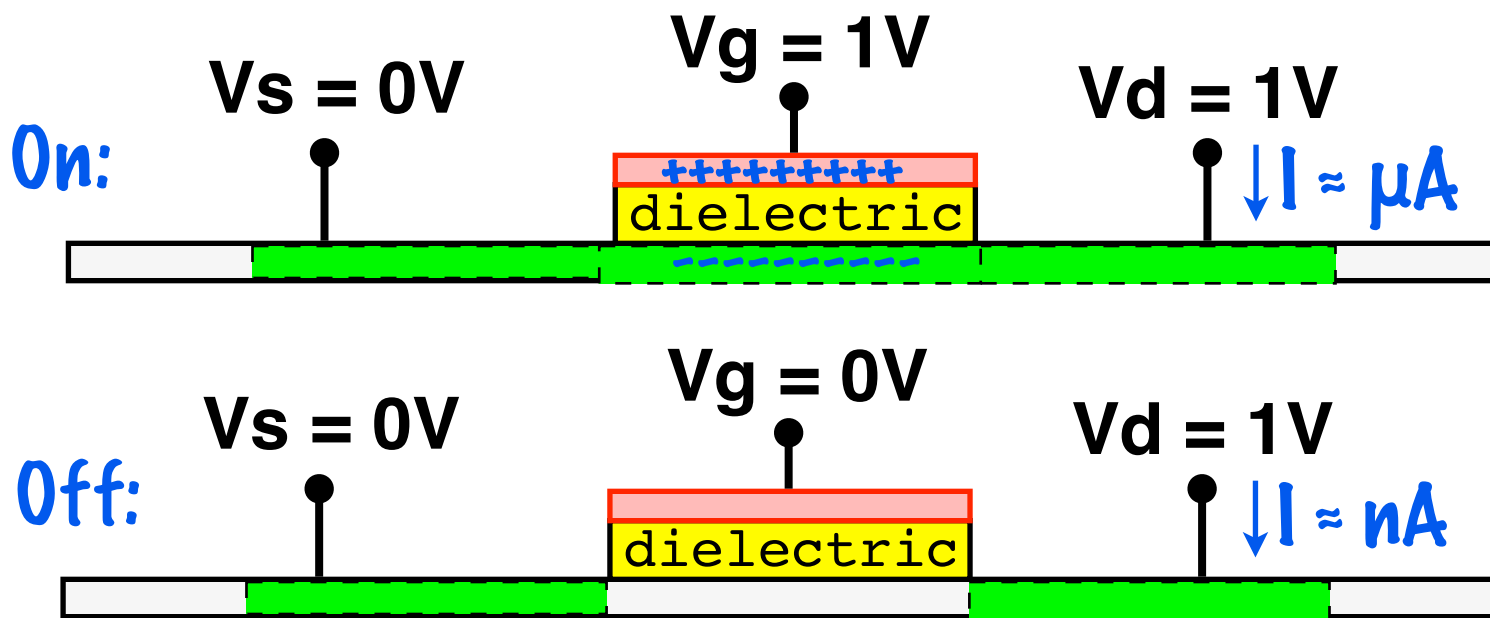
The drain junction is also a capacitor, and puts - charges in the substrate.



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  
22nm  
NMOS



Sandy  
Bridge

32nm  
planar

1.16B  
transistors




"Less than half the power @ same performance"


Ivy Bridge

22nm  
FinFet

1.4B  
transistors



# Leakage reduction in "Tock" 22nm Intel CPUs



—Cary Chin is director of marketing for low-power solutions at Synopsys.

# Break

---



# Six low-power design techniques

---


- \* **Parallelism and pipelining**
- \* **Power-down idle transistors**
- \* **Slow down non-critical paths**
- \* **Clock gating**
- \* **Data-dependent processing**
- \* **Thermal management**

# Trading Hardware for Power

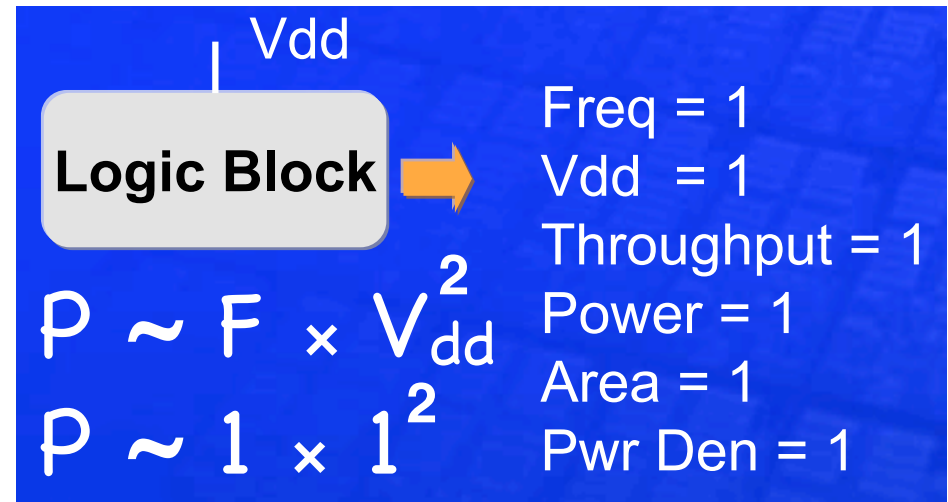
---

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".



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



Simple



Parallel



Pipelined

From:

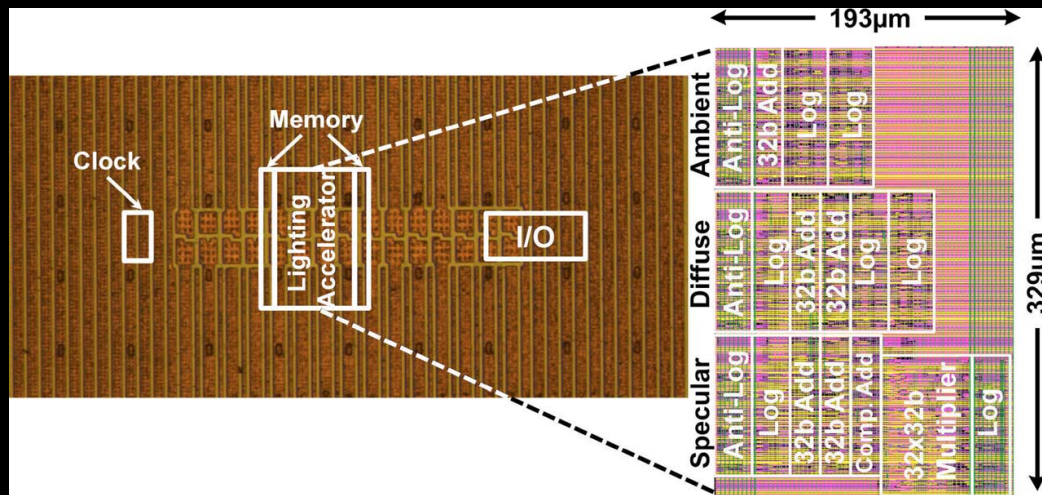
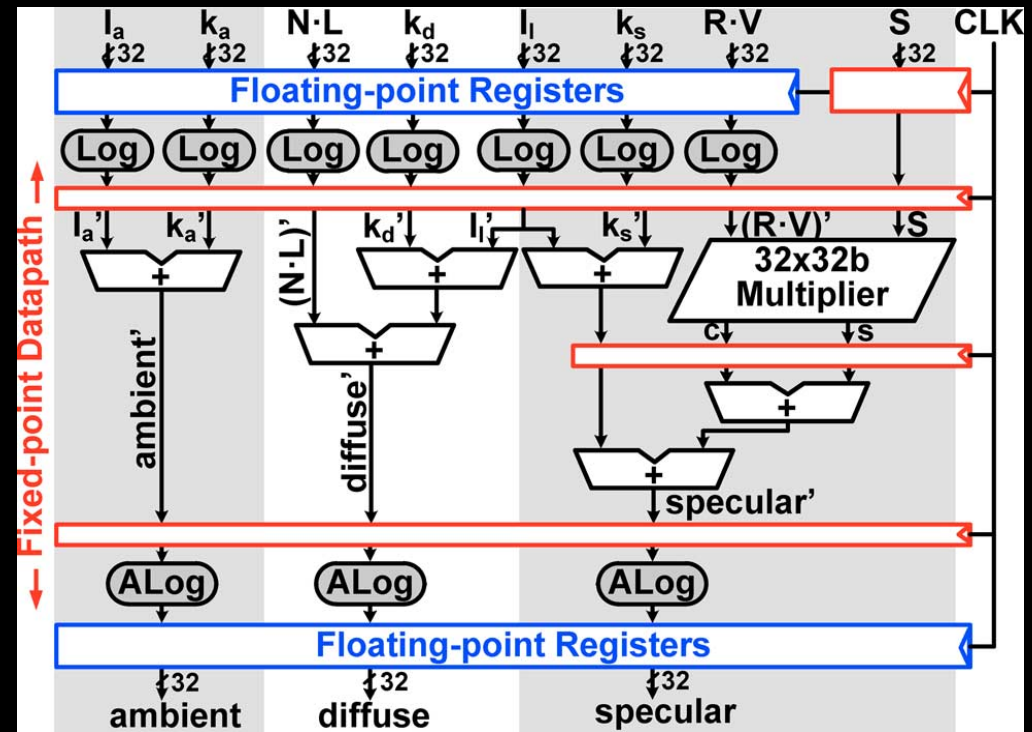

Minimizing Power Consumption in CMOS Circuits

Anantha P. Chandrakasan

Robert W. Brodersen

Regents Spring 2016 © UCB


# Example: Intel Graphics Pipeline IP



A 2.05 GVertices/s 151 mW Lighting Accelerator for 3D Graphics Vertex and Pixel Shading in 32 nm CMOS

Farhana Sheikh, *Member, IEEE*, Sanu K. Mathew, *Member, IEEE*, Mark A. Anders, *Member, IEEE*, Himanshu Kaul, *Member, IEEE*, Steven K. Hsu, *Member, IEEE*, Amit Agarwal, *Member, IEEE*, Ram K. Krishnamurthy, *Fellow, IEEE*, and Shekhar Borkar, *Fellow, IEEE*

# Clock Rate and Power vs Voltage



A 2.05 GVertices/s 151 mW Lighting Accelerator for 3D Graphics Vertex and Pixel Shading in 32 nm CMOS

Farhana Sheikh, *Member, IEEE*, Sanu K. Mathew, *Member, IEEE*, Mark A. Anders, *Member, IEEE*, Himanshu Kaul, *Member, IEEE*, Steven K. Hsu, *Member, IEEE*, Amit Agarwal, *Member, IEEE*, Ram K. Krishnamurthy, *Fellow, IEEE*, and Shekhar Borkar, *Fellow, IEEE*

# 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

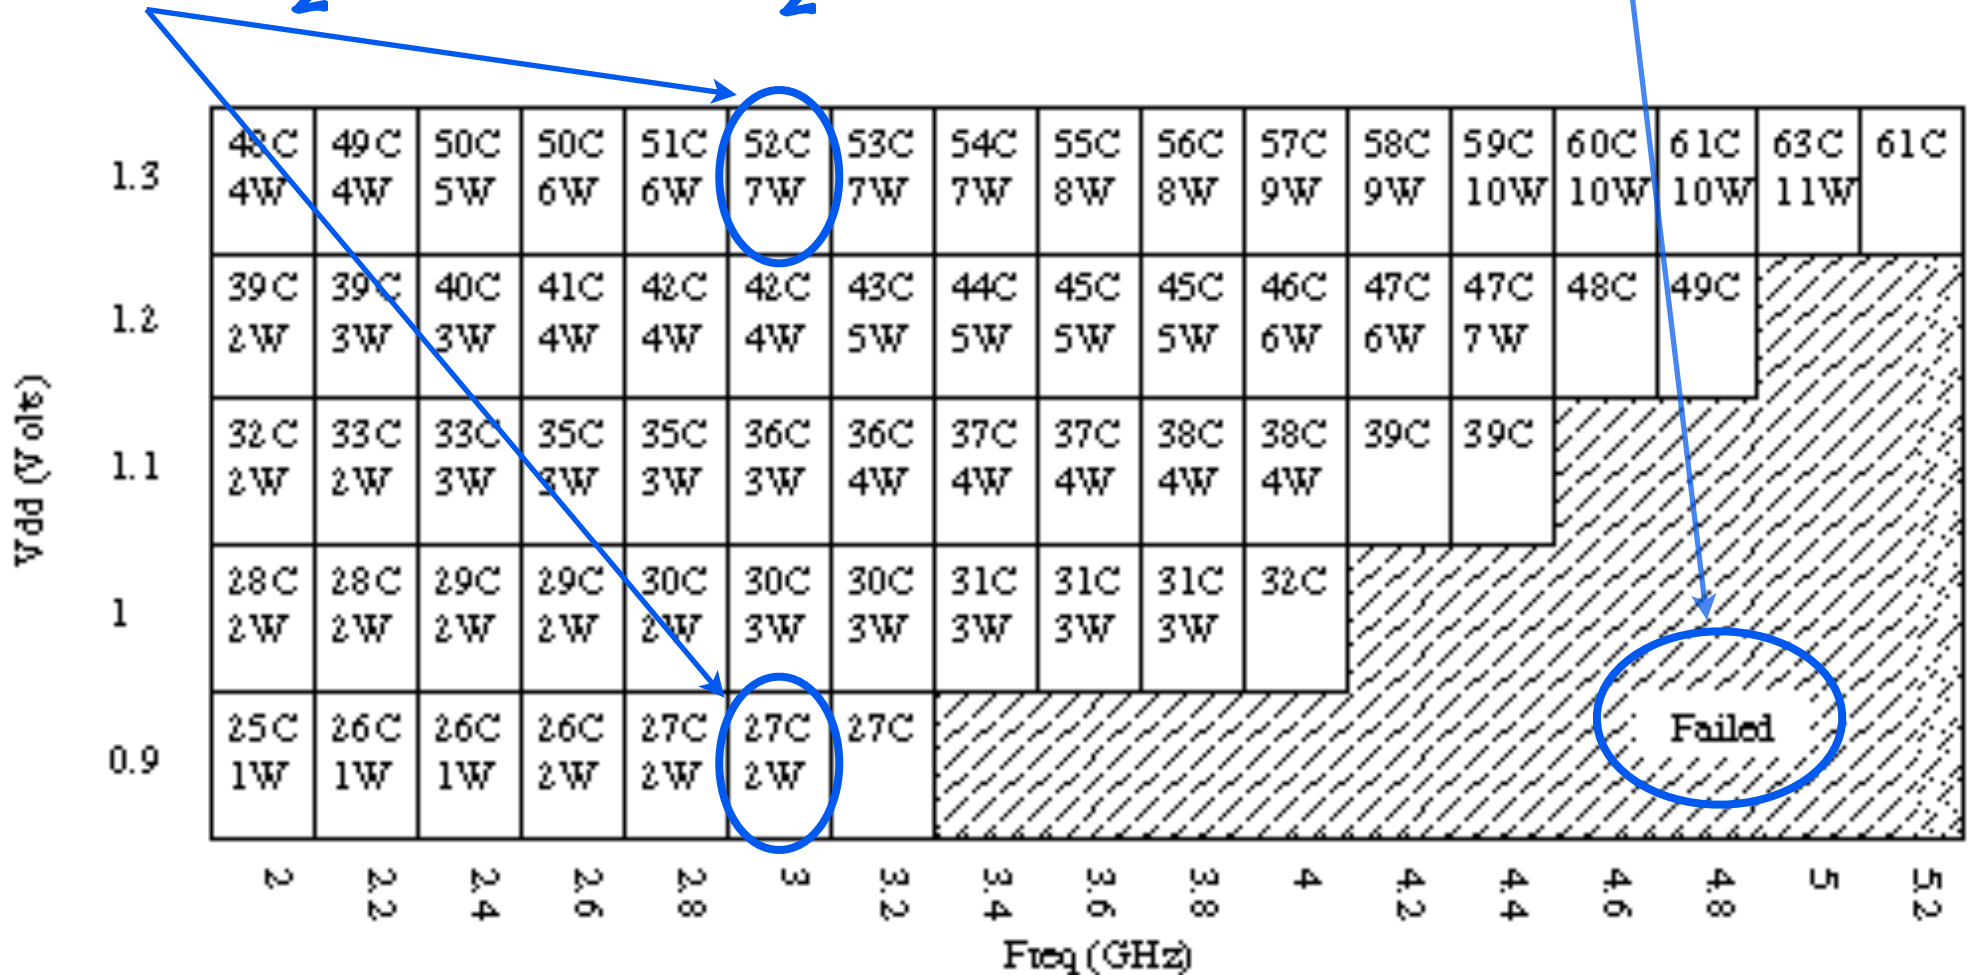


# A “Schmoo” plot for a Cell SPU ...

The lower V<sub>dd</sub>, the less dynamic energy consumption.

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


The lower V<sub>dd</sub>, the longer the maximum clock period, the slower the clock frequency.



# 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** ...

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



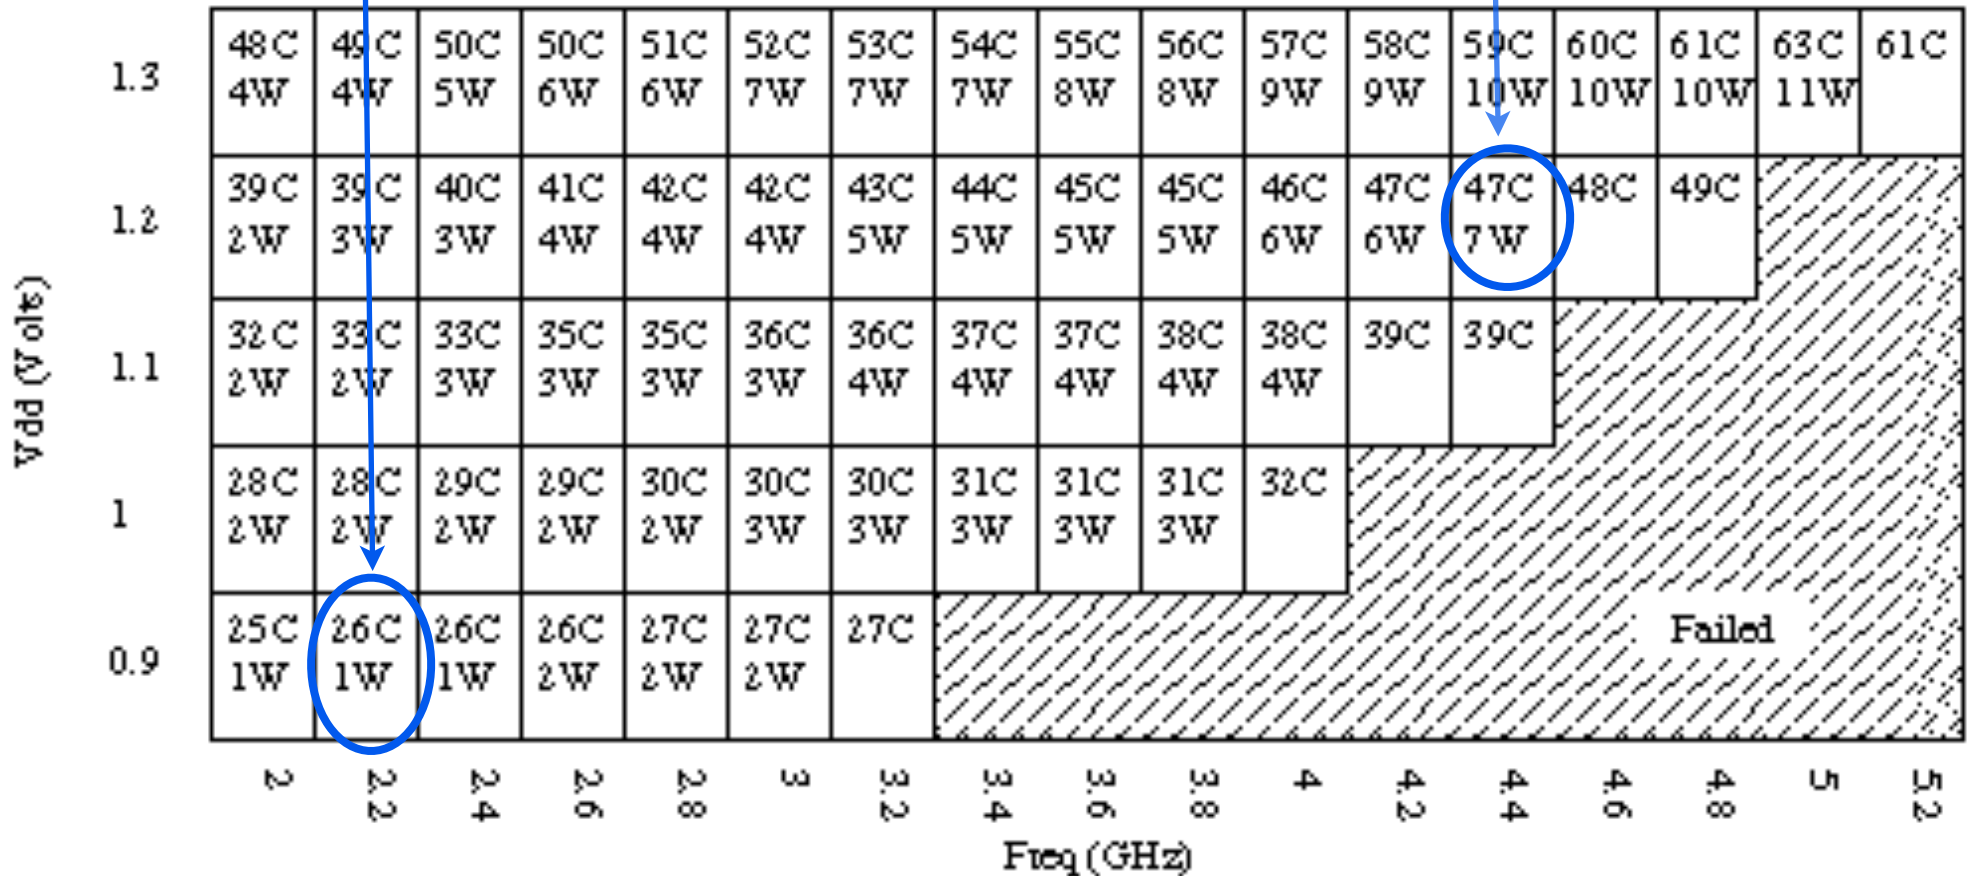


# Scaling V and f does lower energy/op

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



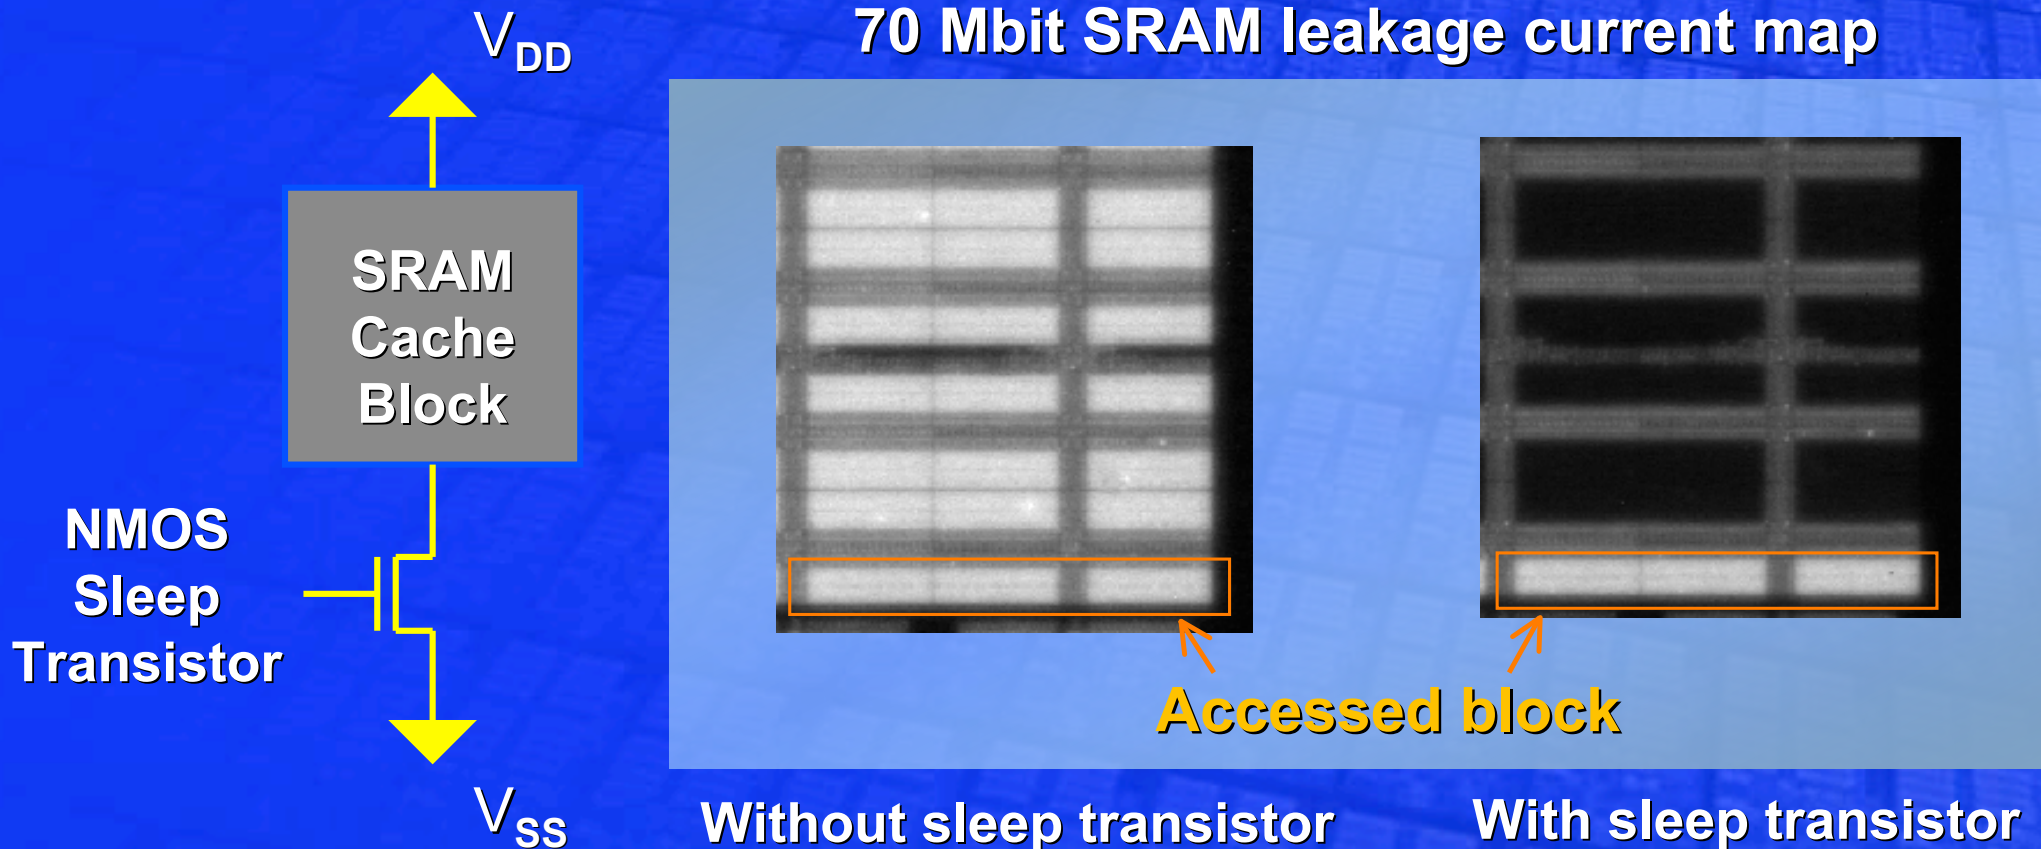
Example: Floating point unit logic.

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

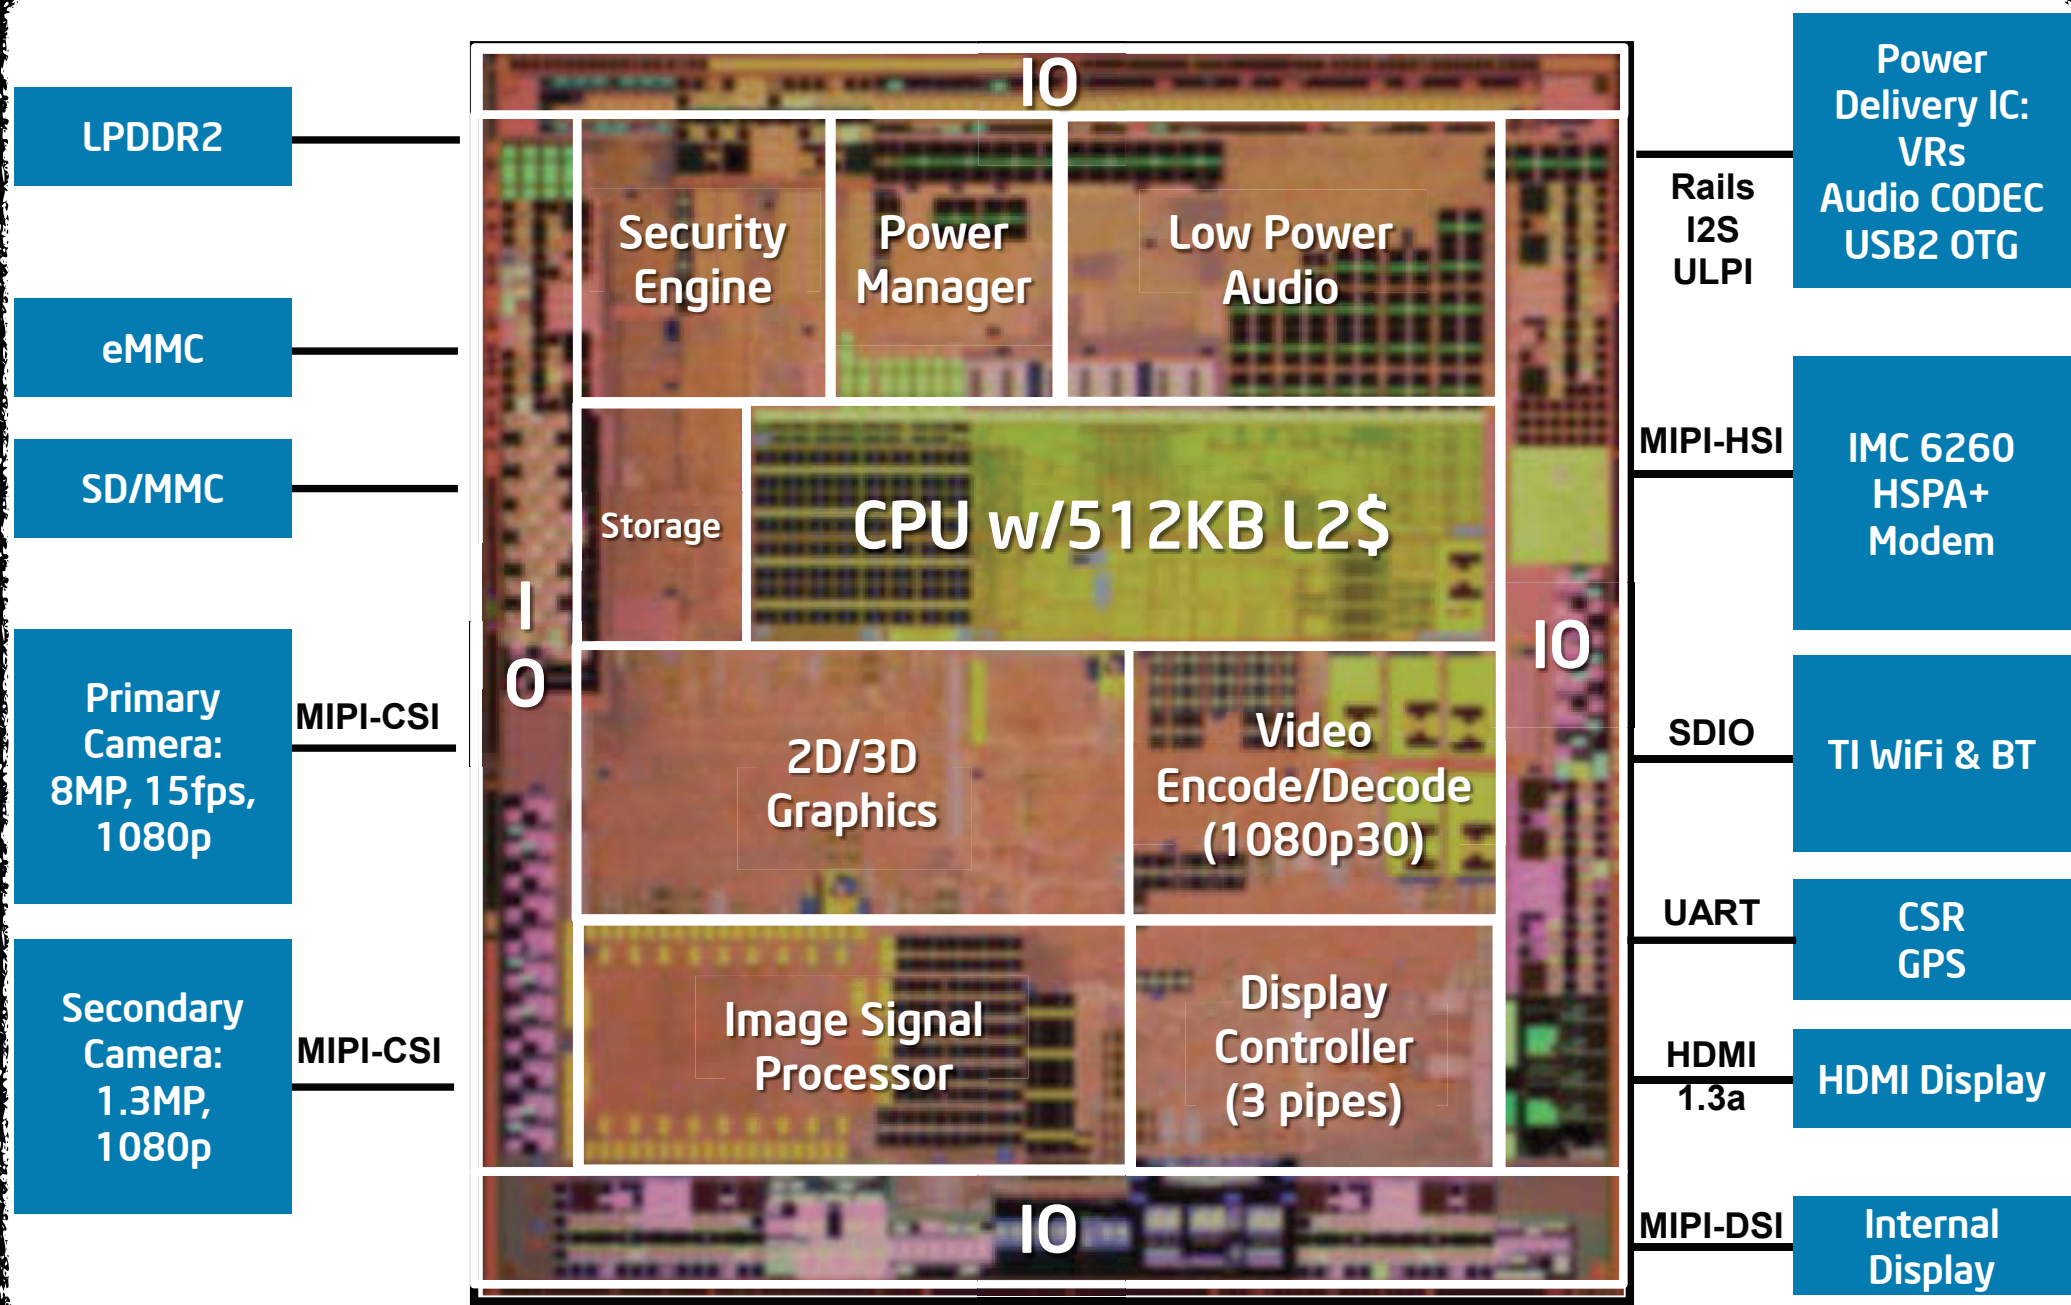


**>3x SRAM leakage reduction on inactive blocks**

**A tiny current supplied in "sleep" maintains SRAM state.**

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

# Intel Medfield



# Intel Medfield

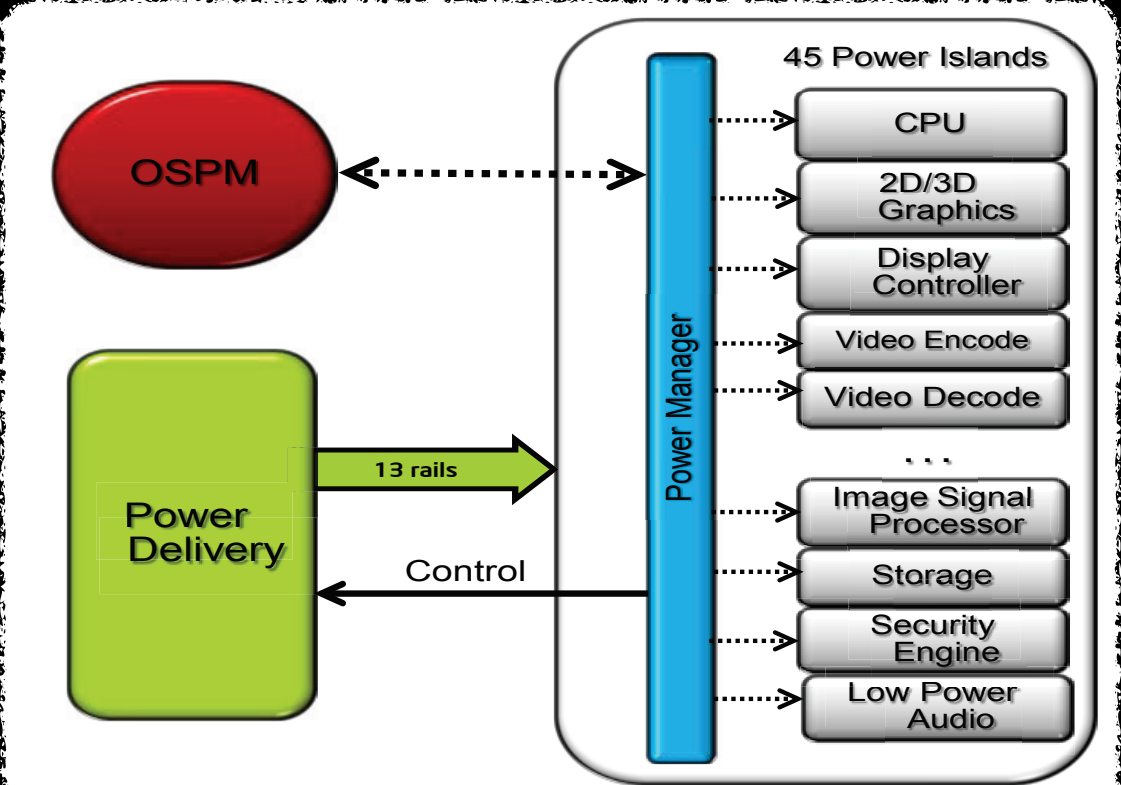
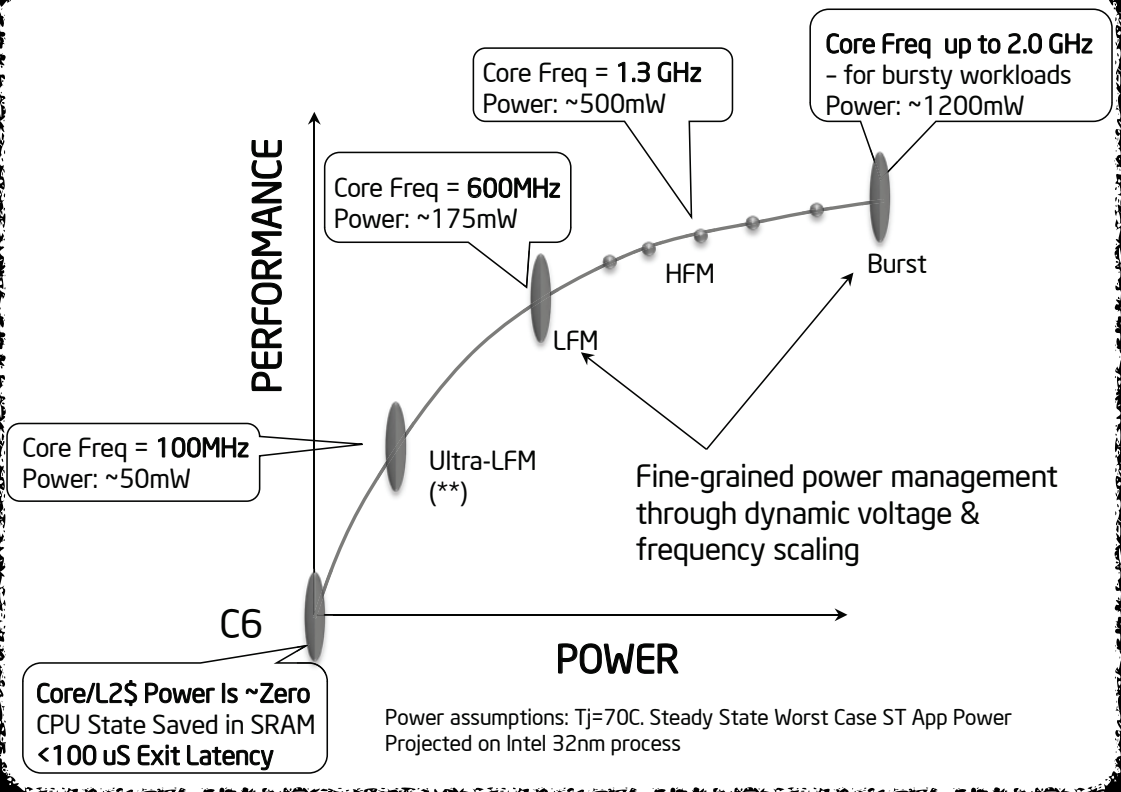
Switches 45 power "islands."

Fine-grained control of leakage power, to track user activity.


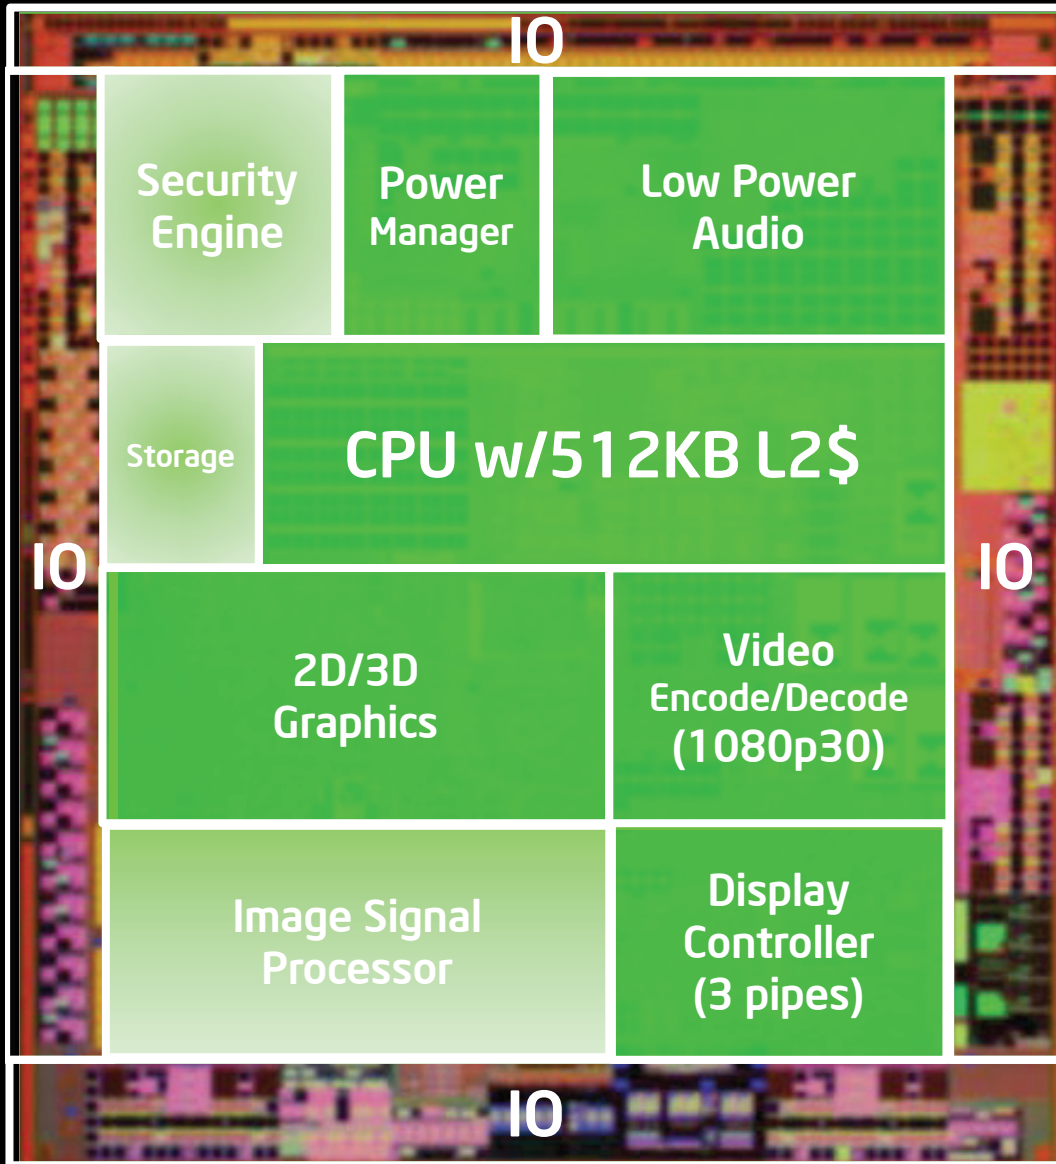
"Race to idle" strategy -- finish tasks quickly, to get to power down.



Intel® Smartphone Reference Design


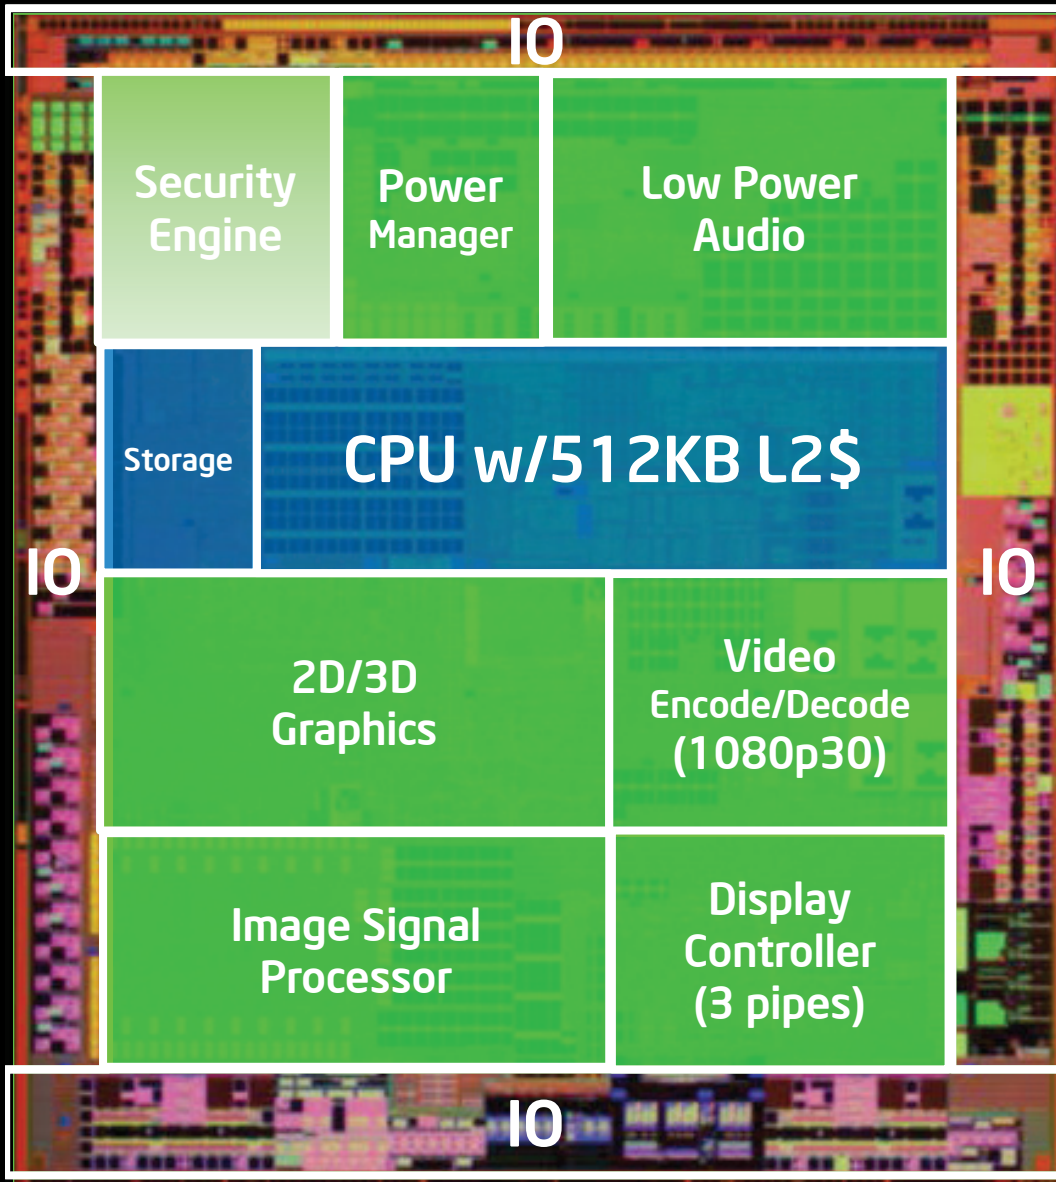


# Playing a game ...


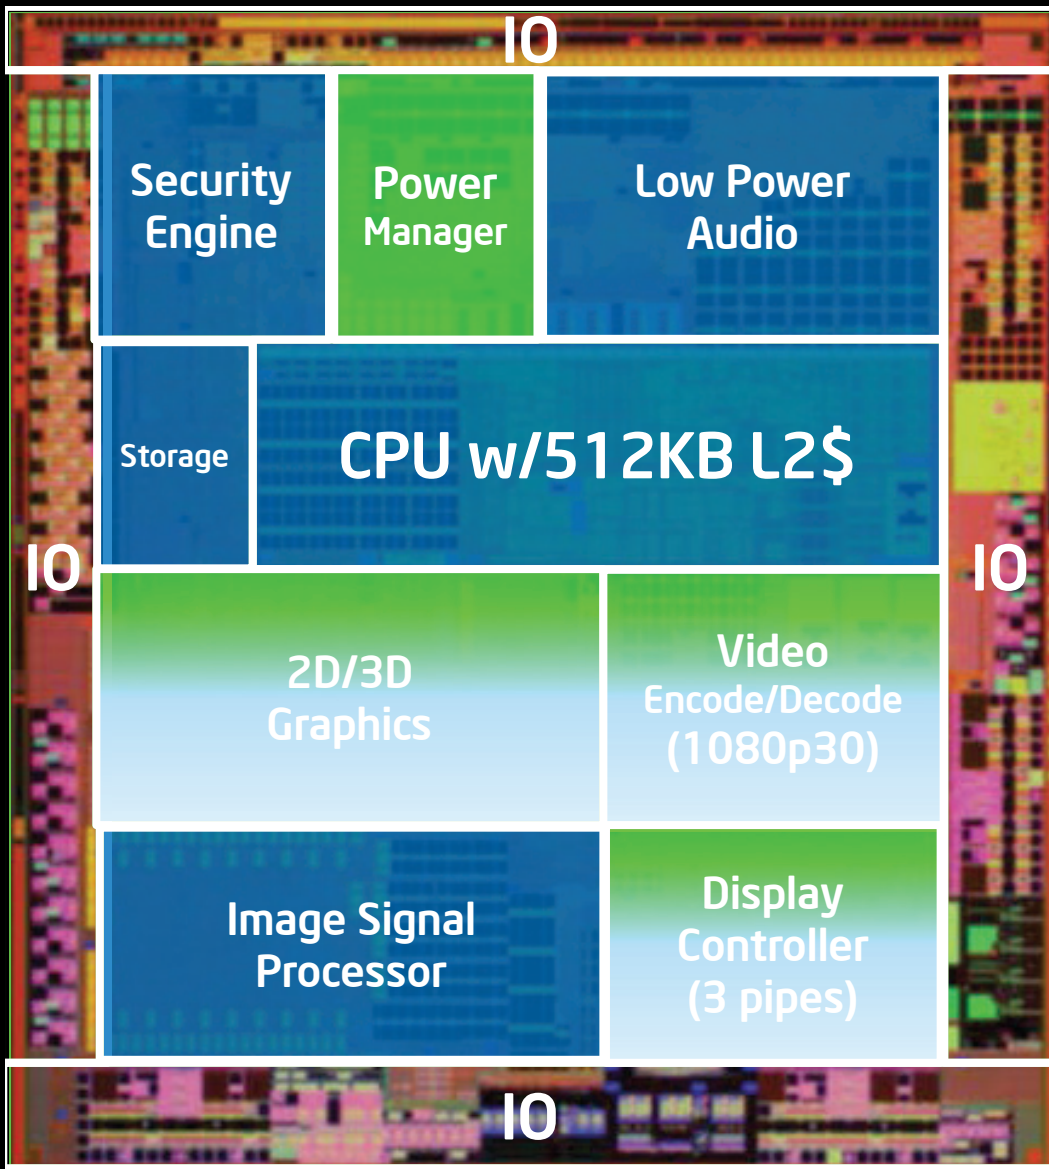





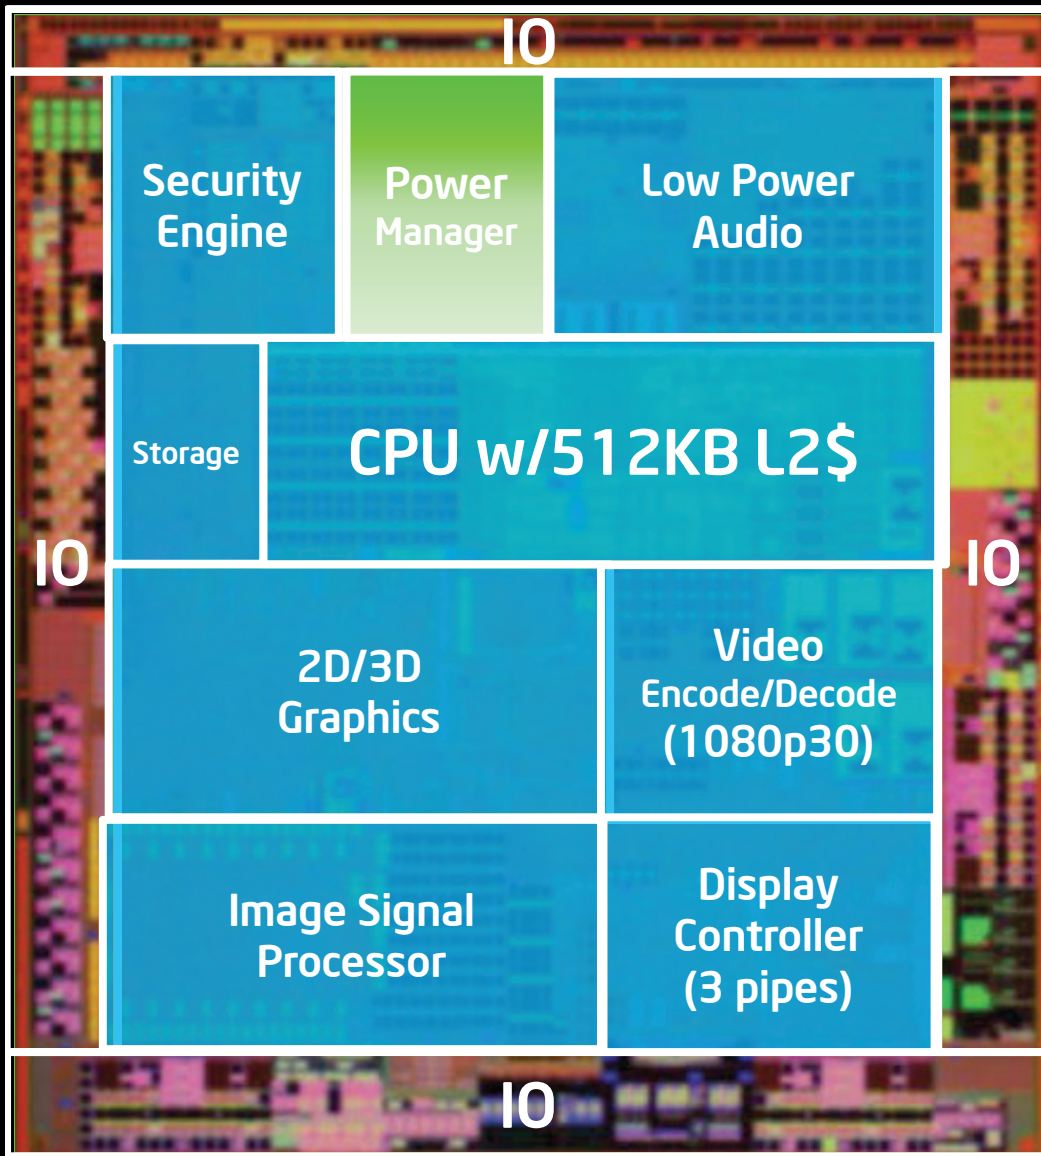
# Watching a video ...



# Looking at phone screen, not doing anything ...



# Phone in your pocket, waiting for a call ...



# Slow down “slack paths”


---



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

The critical path

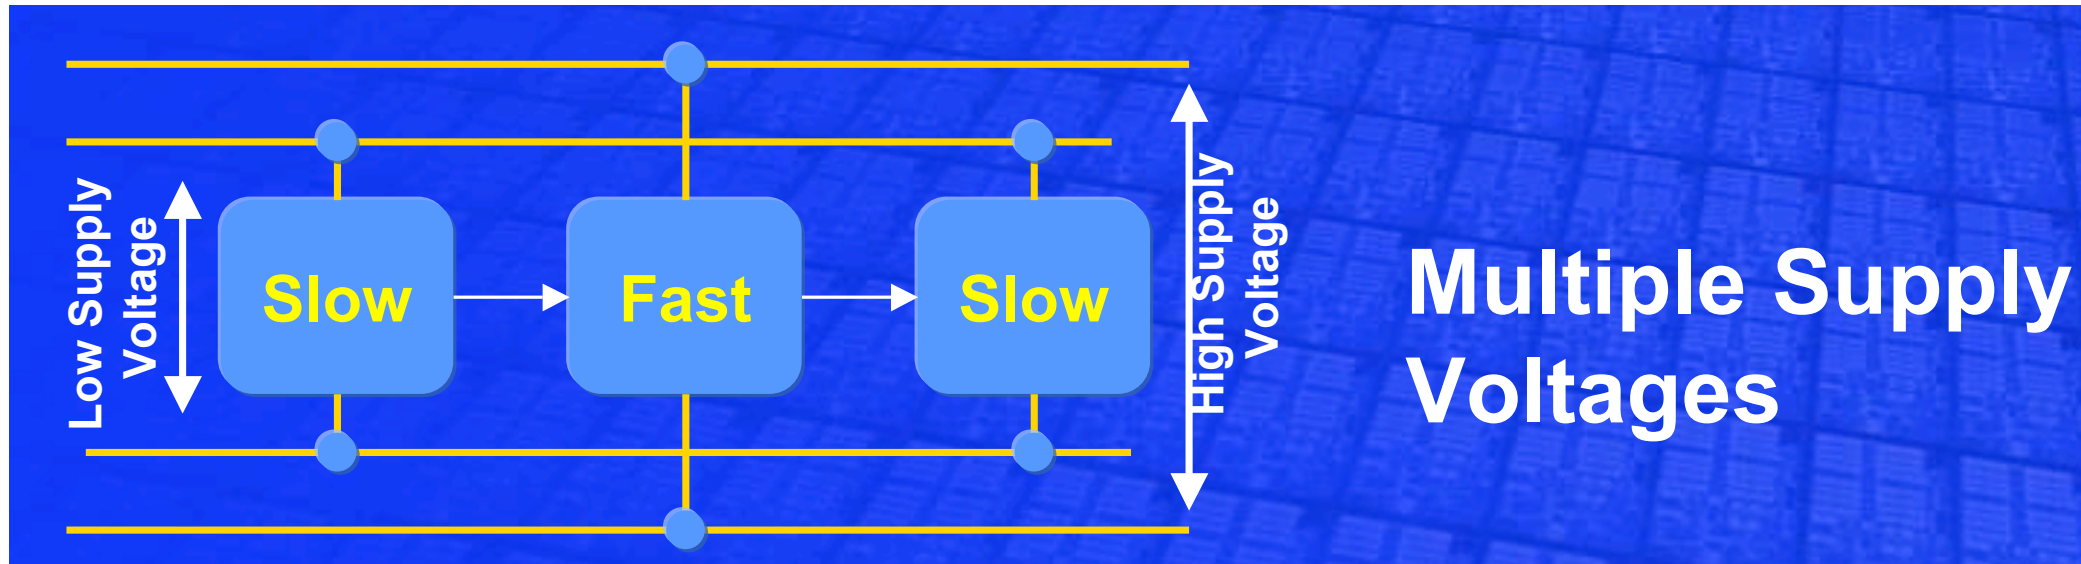
Most wires have hundreds of picoseconds to spare.



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<sub>dd</sub>?** We can reduce **dynamic** power by using low-power V<sub>dd</sub> for logic off the critical path.

**What if we can't do a multi-V<sub>dd</sub> design?**

In a multi-V<sub>t</sub> process, we can reduce **leakage** power on the slow logic by using high-V<sub>th</sub> transistors.

# 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%  
by using high  $V_T$ .



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?

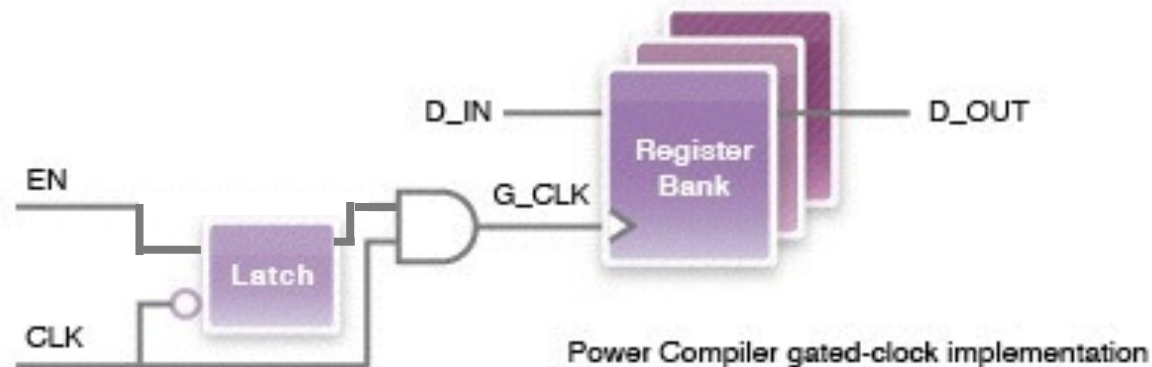
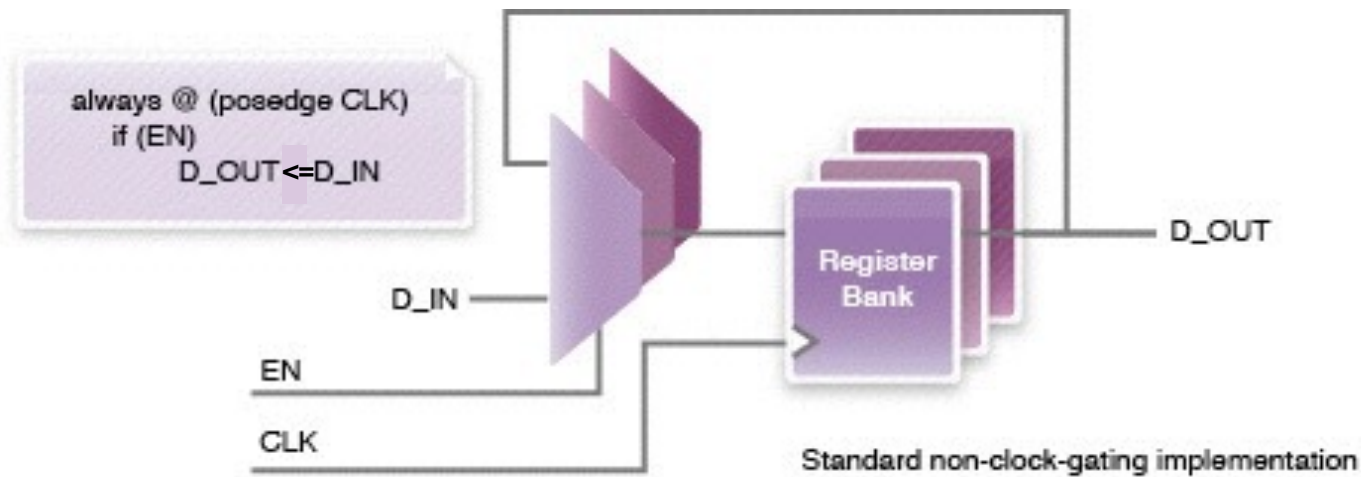


Half of the power goes to latches (Flip-Flops).

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 Design Compiler can do this ...



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

Synopsis Data  
Sheet

# Data-Dependent Processing

---



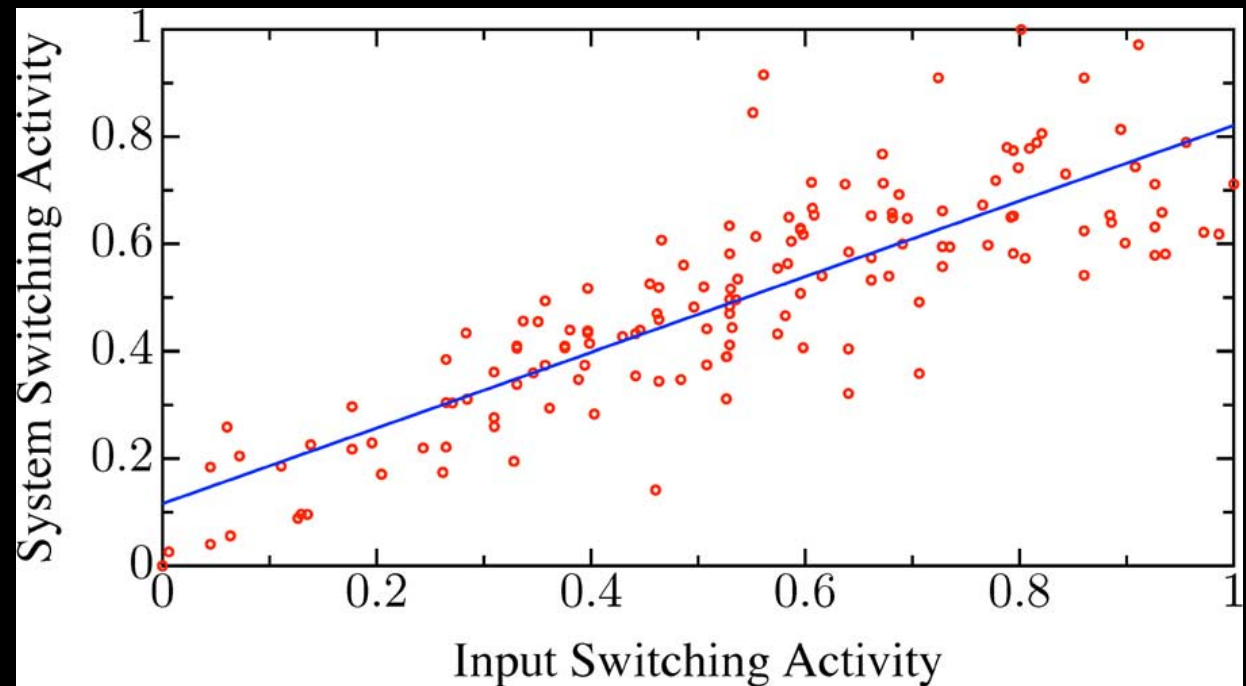

# Example: Video Decode Transform

Most of the time, the inputs flip between small positive and negative integers.

In 2's complement, wastes power:

+1: 0b00001

-1: 0b11110

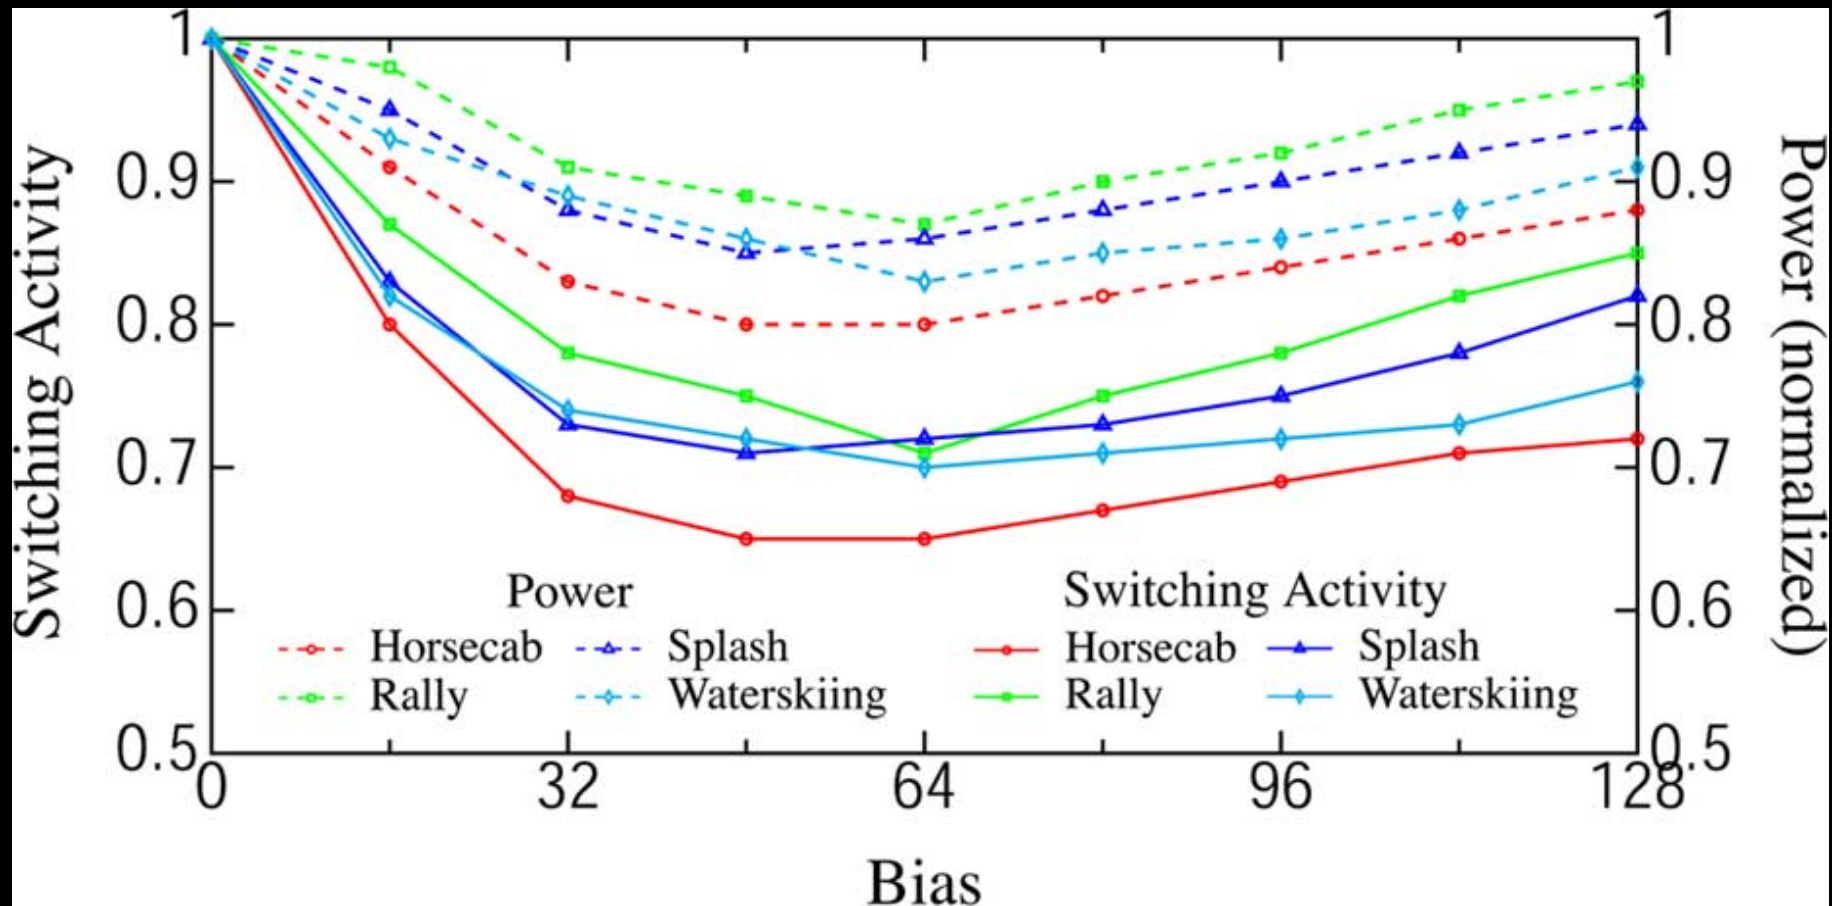


Quad Full-HD Transform Engine for Dual-Standard  
Low-Power Video Coding

Rahul Rithe, *Student Member, IEEE*, Chih-Chi Cheng, *Member, IEEE*, and Anantha P. Chandrakasan, *Fellow, IEEE*

# Solution: Add bias value to all inputs

30+% power reduction for a bias of 64. For this linear transform, correcting the output for the bias is trivial.



Quad Full-HD Transform Engine for Dual-Standard  
Low-Power Video Coding

Rahul Rithe, *Student Member, IEEE*, Chih-Chi Cheng, *Member, IEEE*, and Anantha P. Chandrakasan, *Fellow, IEEE*

# Thermal Management

---



# Keep chip cool to minimize leakage power

Figure 3






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 F

# Idea: Monitor temperature, servo clock speed



# Intel realtime temp monitoring

RealTemp 4.00

Intel Core i7-4700MQ

CPUID 306C3	3591.683 MHz	0:03:31
Threads 8	99.769 x 36.0	VID 1.0463
GPU 36°C	Load 12.9%	15.2 W

Temperature (°C)				Package
45	45	49	62	62


Distance to TJ Max				
55	55	51	38	38

Minimum				
39°C	38°C	39°C	39°C	41°C
14:57:03	14:57:49	14:57:03	14:57:03	14:56:49

Maximum				
54°C	61°C	61°C	62°C	62°C
14:55:45	14:55:52	14:59:06	14:55:51	14:55:51

Thermal Status				
OK	OK	OK	OK	OK

Sensor Test   Testing...   C States   Cancel   Options



# Six low-power design techniques

---

- \* **Parallelism and pipelining**
- \* **Power-down idle transistors**
- \* **Slow down non-critical paths**
- \* **Clock gating**
- \* **Data-dependent processing**
- \* **Thermal management**