CS 150
Digital Design

Lecture 28 – Power and Energy

2011-4-28

John Wawrzynek

today's lecturer: John Lazzaro

TAs: Michael Eastham and Austin Doupnik

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.

This is how electric tea pots work ...

Heats 1 gram of water
0.24 degree C

0.24 Calories per Second

1 Joule of Heat Energy
per Second

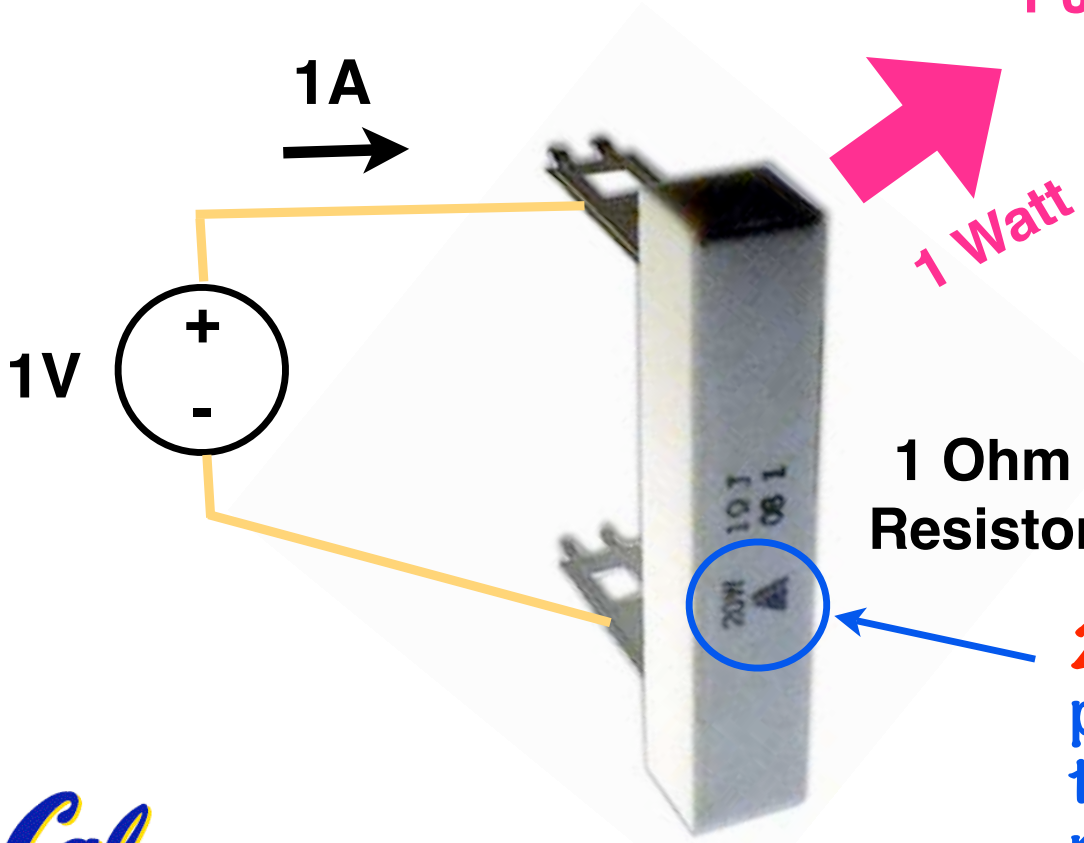

Cooling an iPod nano ...

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)

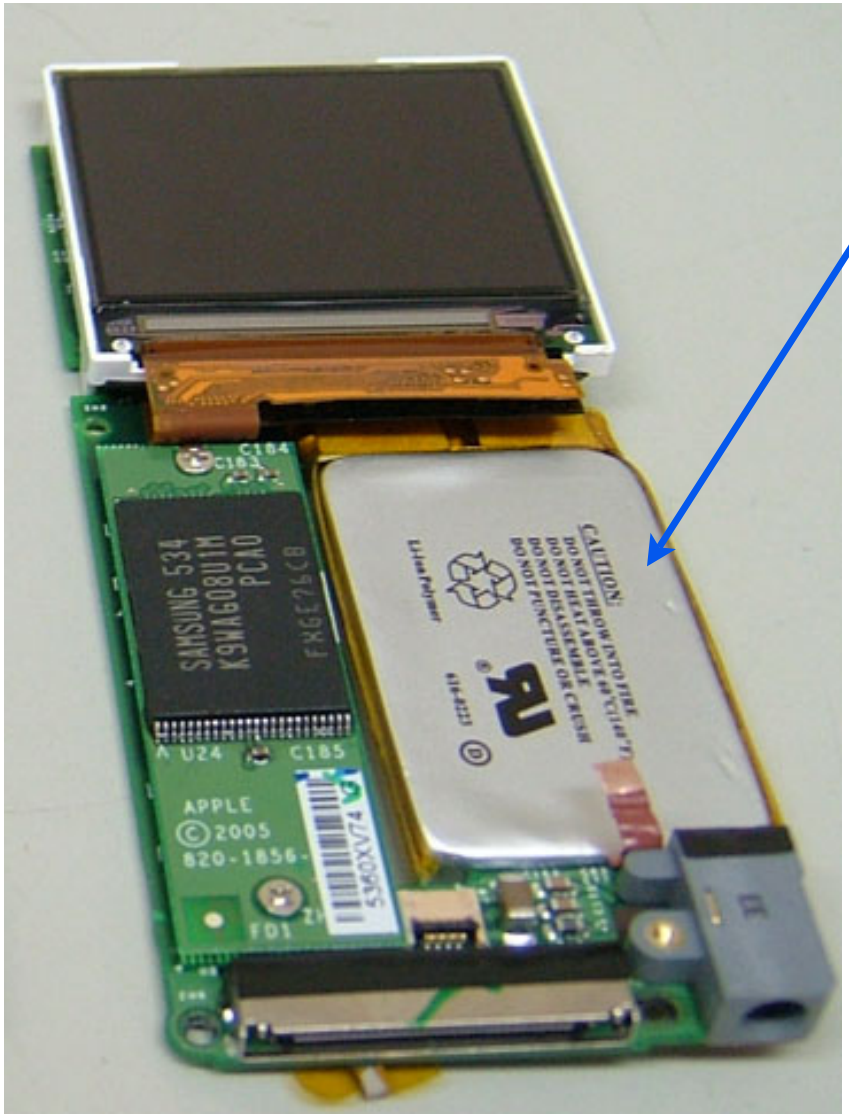

Battery has 1.2 W-hour rating: Can supply 1.2 W of power for 1 hour.

1.2 W / 5 W = 15 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 ...

portalplayer™

A close relative ...

PP5020

digital media management system-on-chip

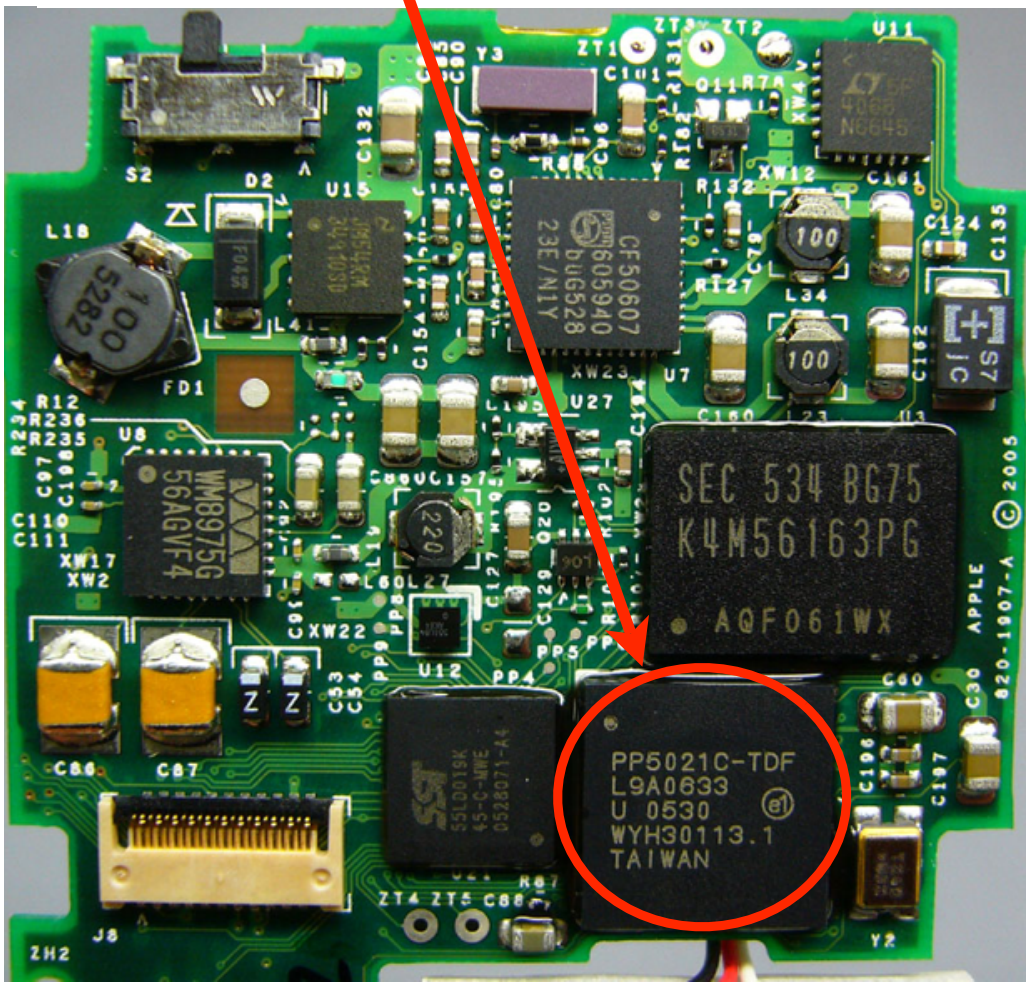

Two 80 MHz CPUs. One CPU used for audio, one for slides.

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 ?

Source: ifixit.com

Source: ifixit.com

Source: ifixit.com

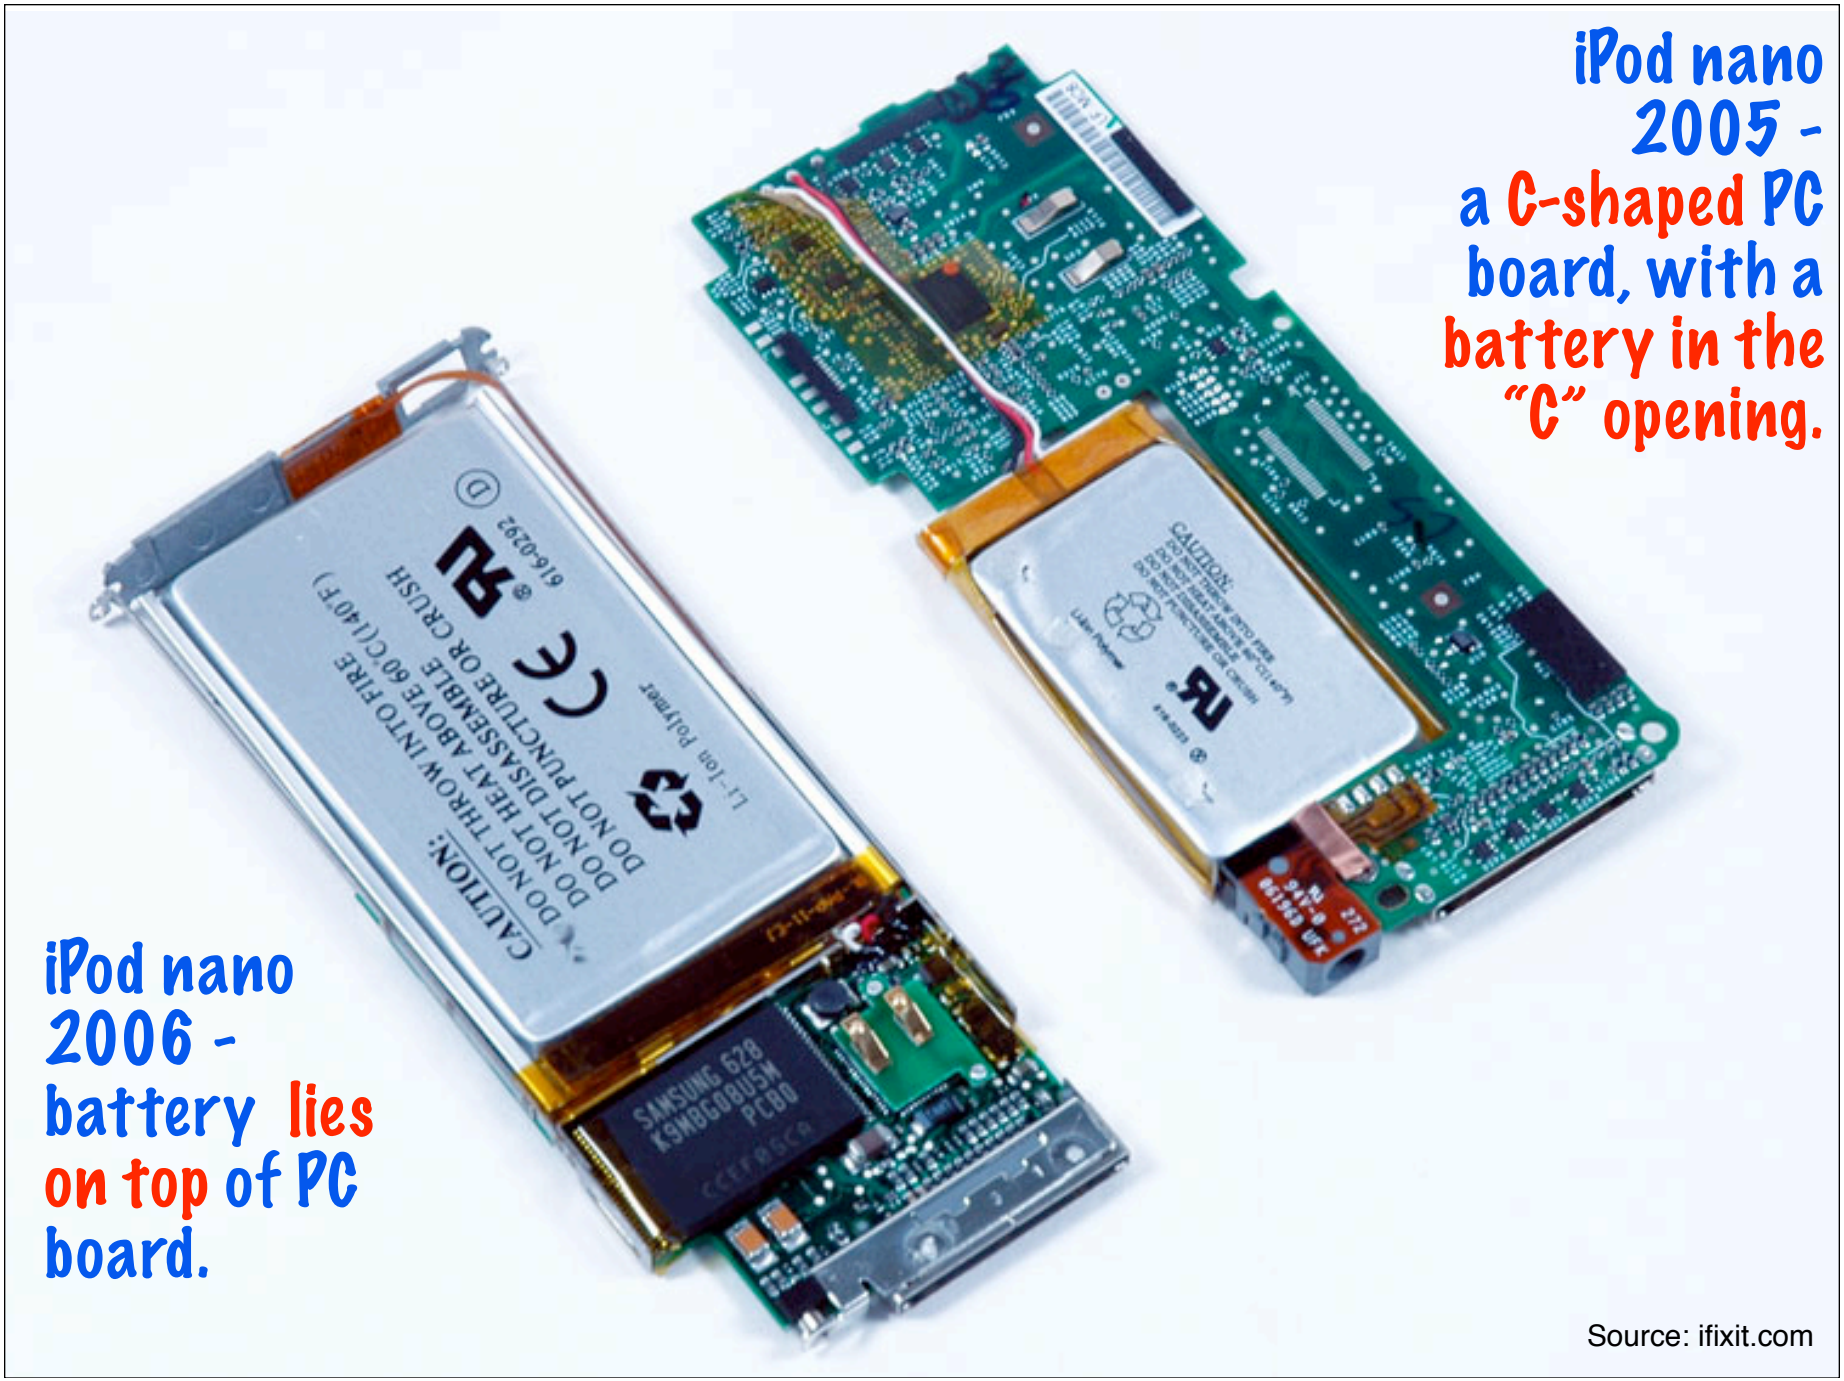
iPod nano 2005 - a C-shaped PC board, with a battery in the "C" opening.
iPod nano 2006 - battery lies on top of PC board.
CAUTION:
DO NOT THROW INTO FIRE
DO NOT HEAT ABOVE 60°C(140°F)
DO NOT DISASSEMBLE OR CRUSH
DO NOT PUNCTURE
SAMSUNG 628
K9MBG08U5M
PCB0
Source: ifixit.com

How? Small IC packages, fewer parts

iPod nano 2006 ⟶

iPod nano 2005 ⟶

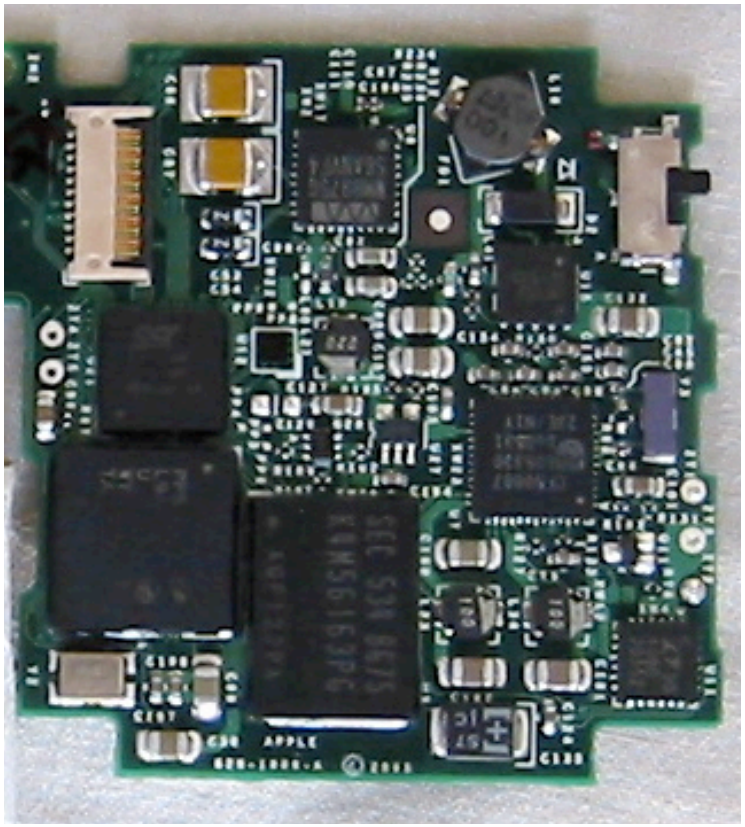

Source: arstechnica.com

Aluminum permits thinner case ...

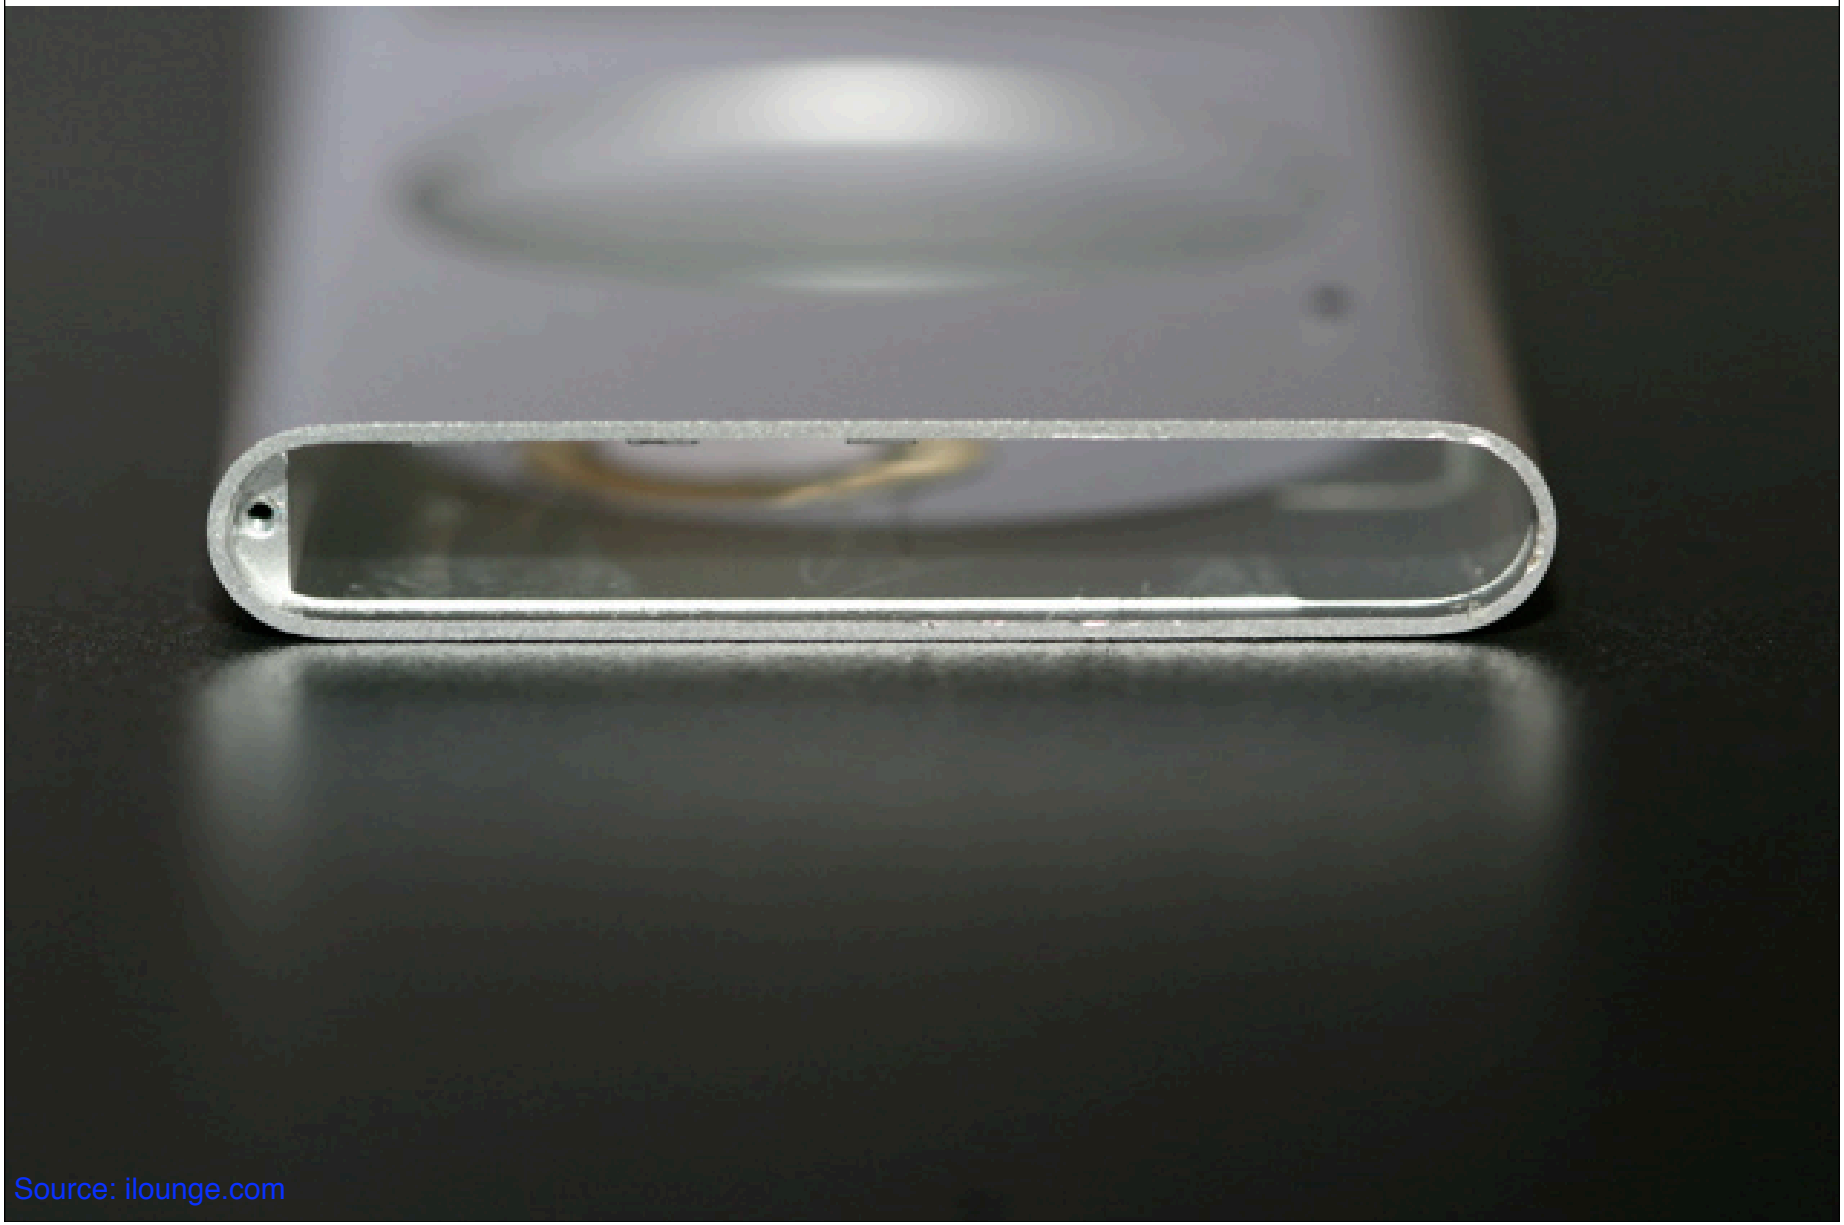

Source: ilounge.com

Year-to-year: continuous improvements

iPod nano
2005
14 hours
battery
life
(audio
playback)

1.2 W-hour
battery.

iPod nano
2006
24 hours
battery
life
(audio
playback)

What's
happened
since ?

Source: ifixit.com

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

0.19 W Hr

Sources: iFixit, Apple

Notebooks ... now most of the PC market.

2006 Apple MacBook -- 5.2 lbs

- **Performance:** Must be "close enough" to desktop performance ... most people no longer use a desktop (including me!)
- **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!

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!

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 ...

MacBook Air ... fits in a manila envelope!

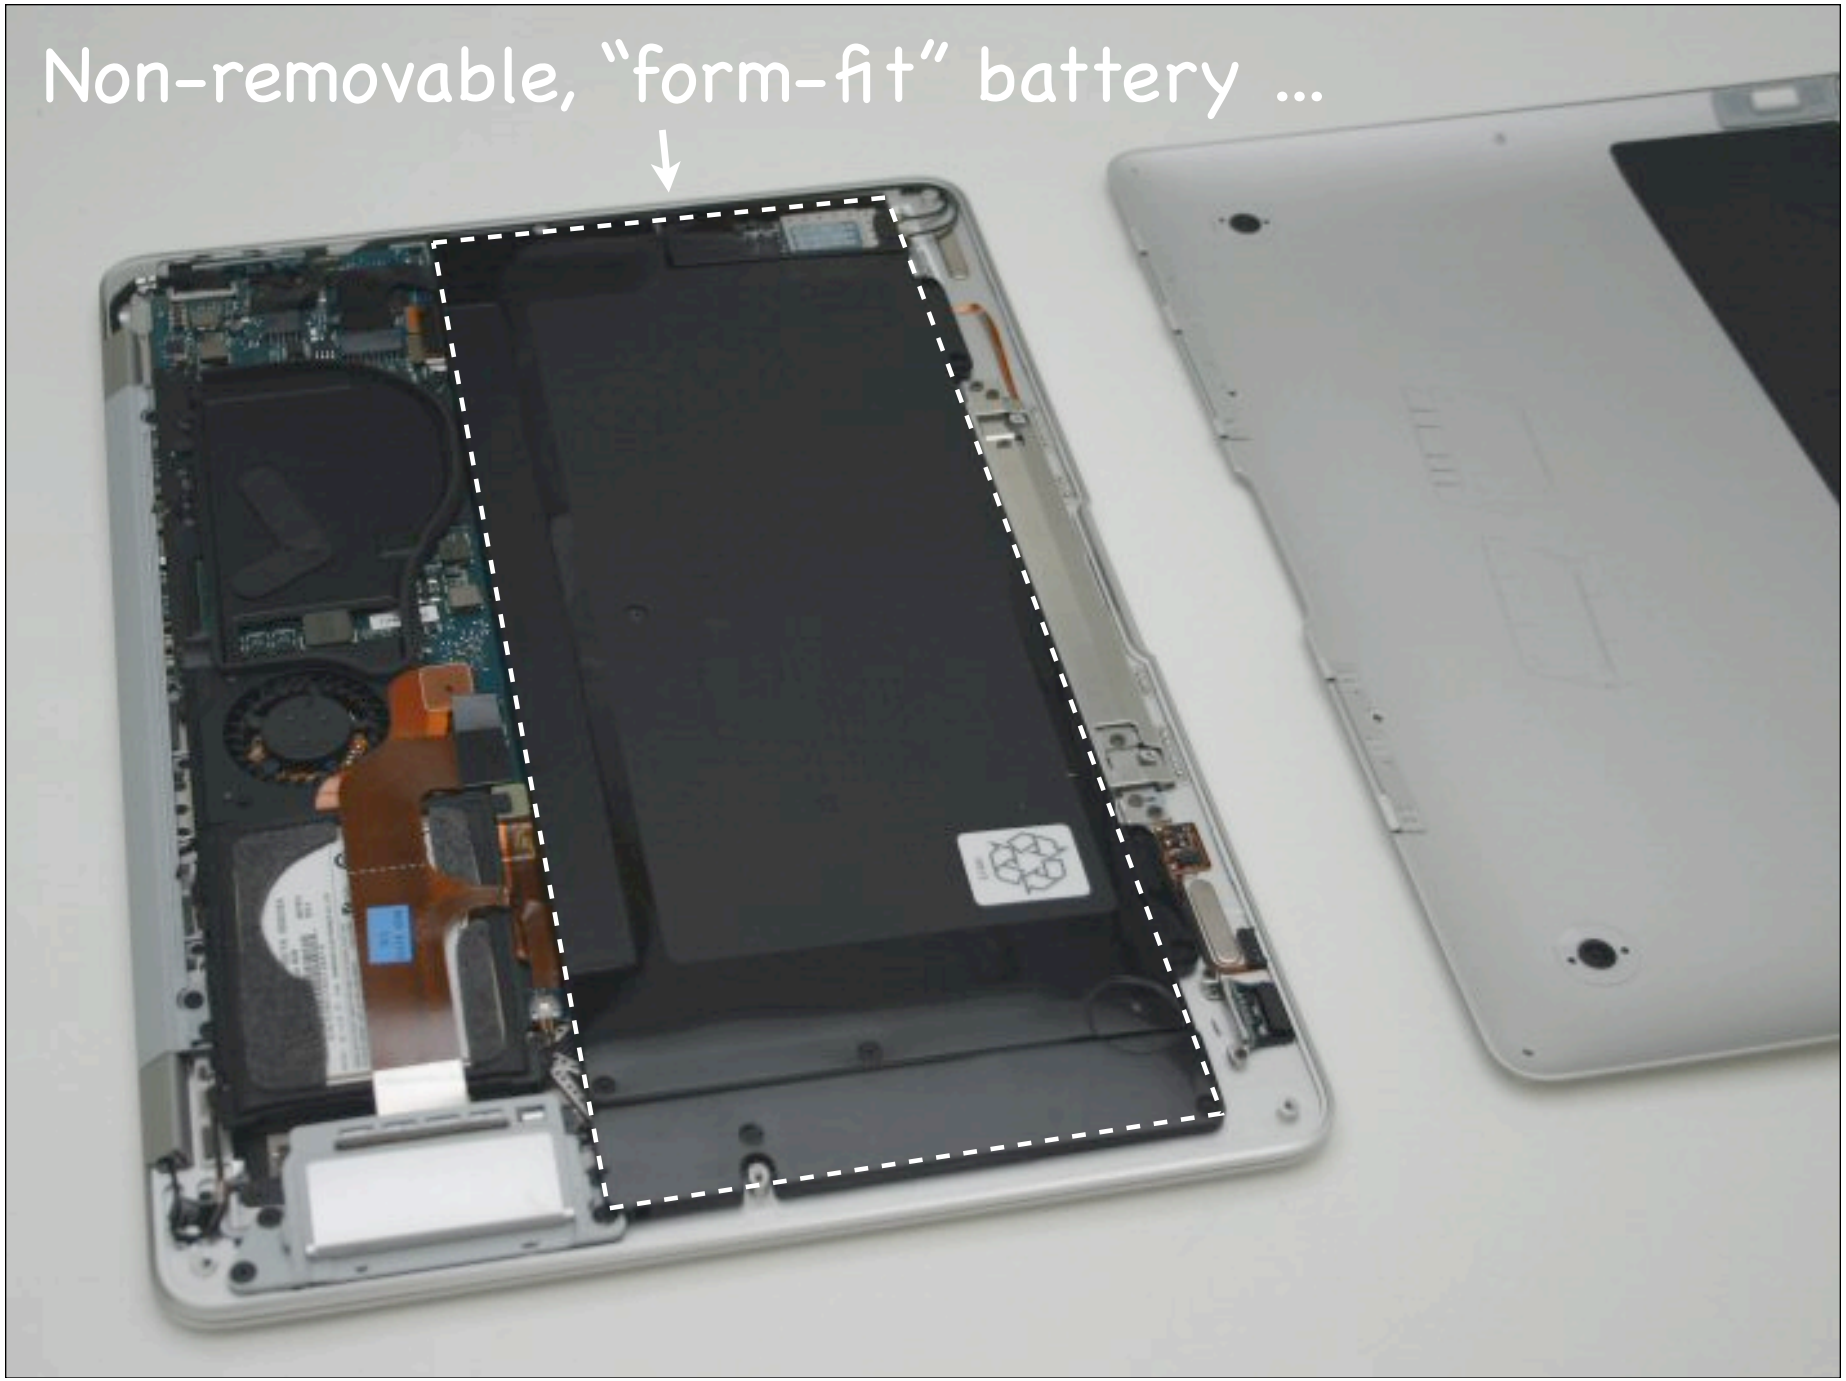
Non-removable, "form-fit" battery …

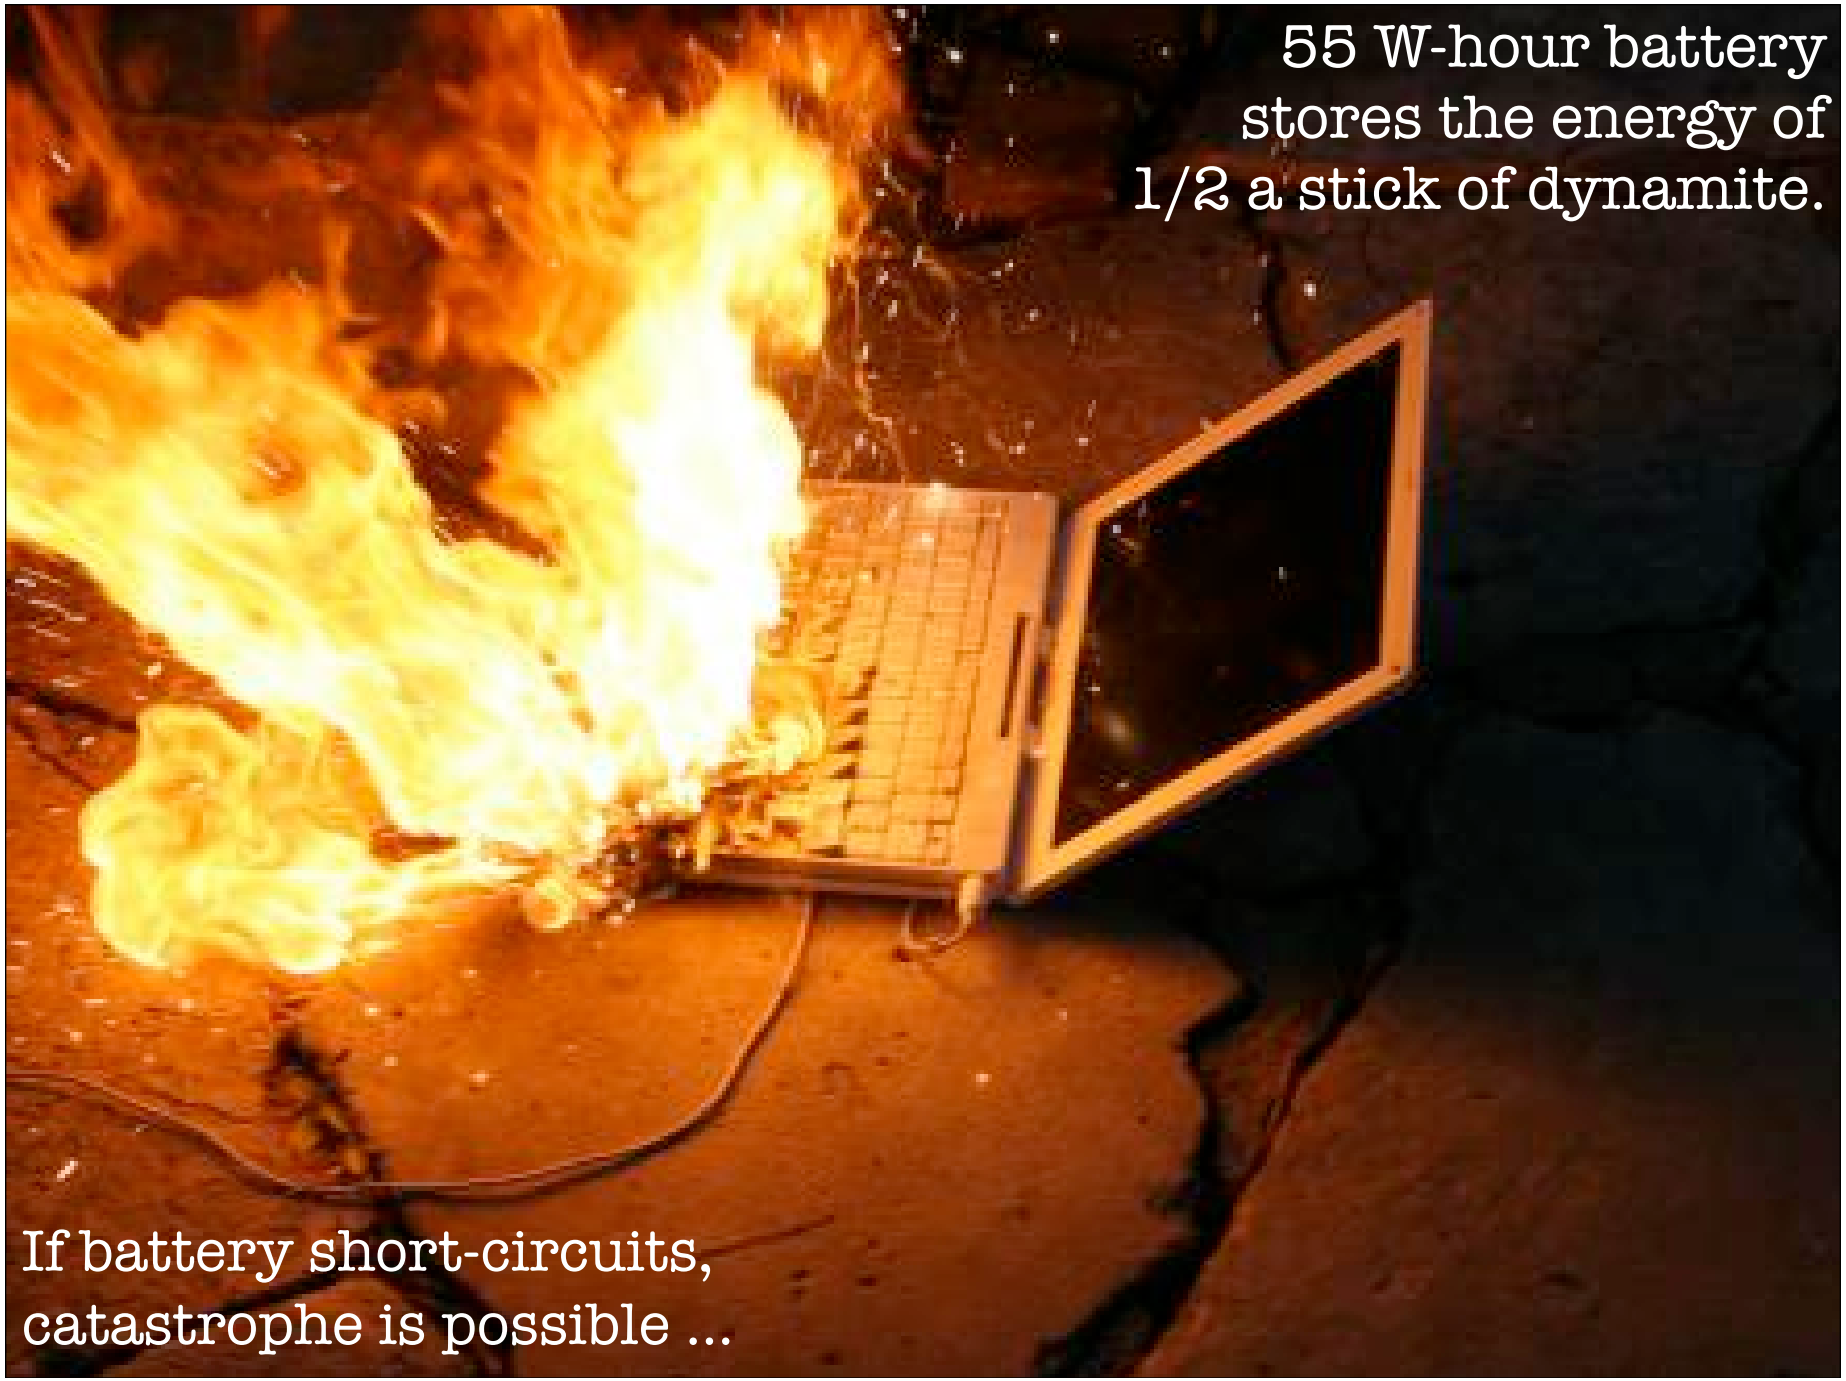
55 W-hour battery
stores the energy of
1/2 a stick of dynamite.
If battery short-circuits,
catastrophe is possible ...

Lithium battery density and mass ...

Energy densities table

Storage type	Specific energy MJ/kg	Energy density MJ./Liter
Uranium-235 used in nuclear weapons	144,000,000	1,500,000,000
Natural gas	53.6	0.0364
Gasoline (petrol)	46.4	34.2
Diesel fuel/residential heating oil	46.2	37.3
Anthracite coal	32.5	72.4
Lithium-ion battery	0.46 to 0.72	0.83 to 3.6

Thus the interest in fuel cells for portable electronics ...

Source: Machine Design magazine

The CPU is only part of power budget!

2004-era notebook
running a full workload.

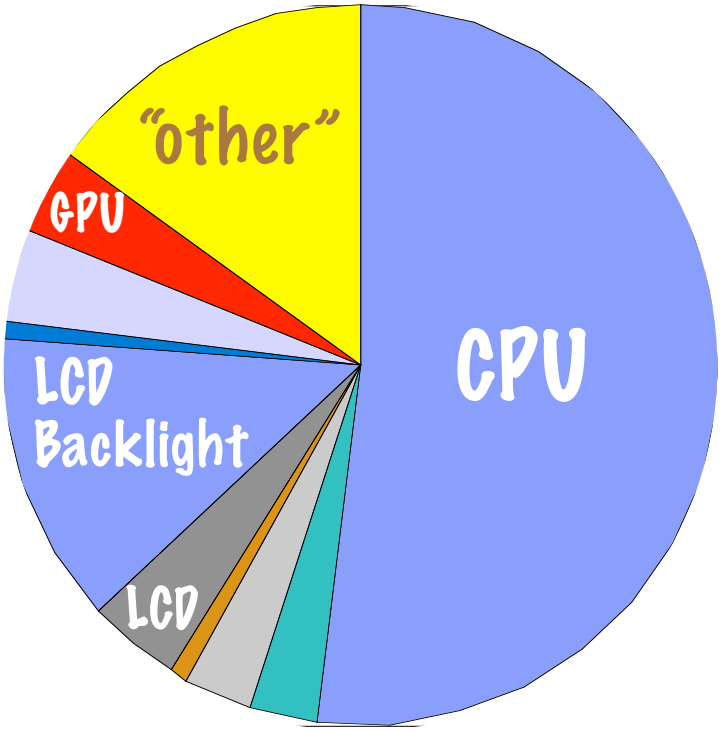

"Amdahl's Law for Power"

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

Data courtesy Mahesri et al., U of Illinois, 2004


Facebook
Prineville | data center

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.

Processors and Energy

Switching Energy: Fundamental Physics

Every logic transition dissipates energy.

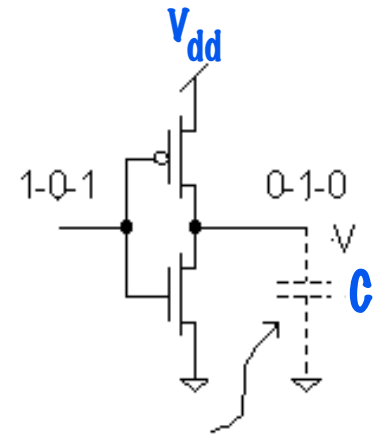

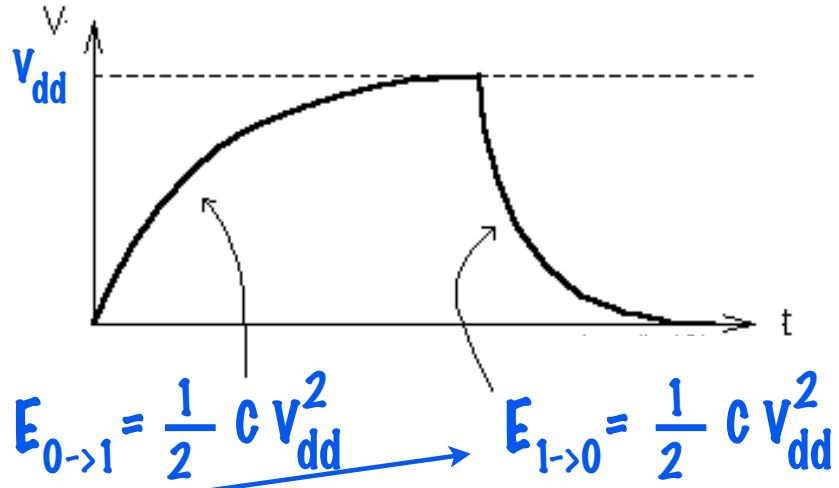

Strong result: Independent of technology.

How can we limit switching energy?

(1) Slow down clock (fewer transitions). But we like speed ...

(2) Reduce Vdd. But lowering Vdd 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 ...

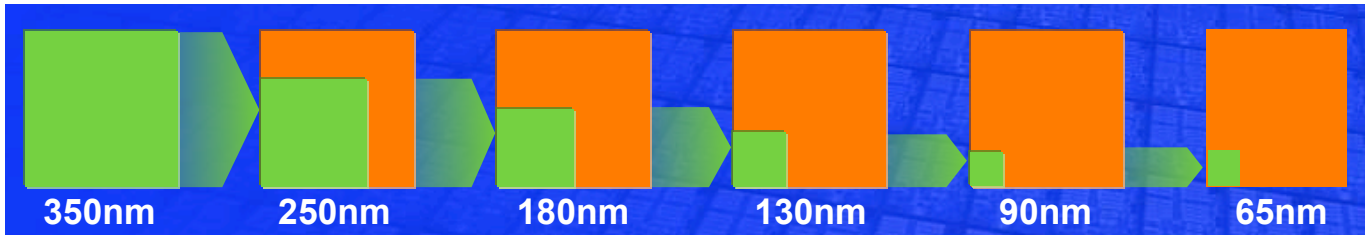

IC process scaling ("Moore's Law")

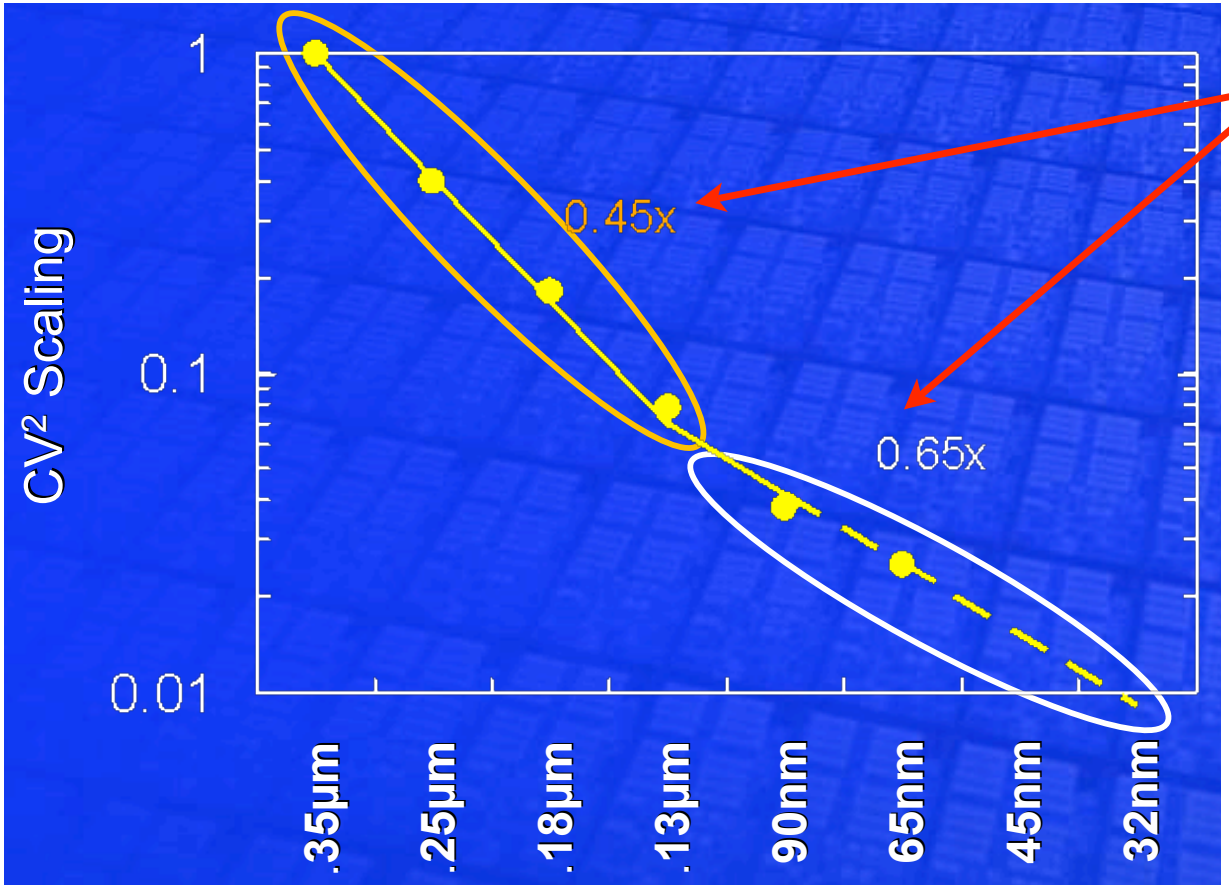

Due to reducing V and C (length and width of Cs decrease, but plate distance gets smaller).

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.

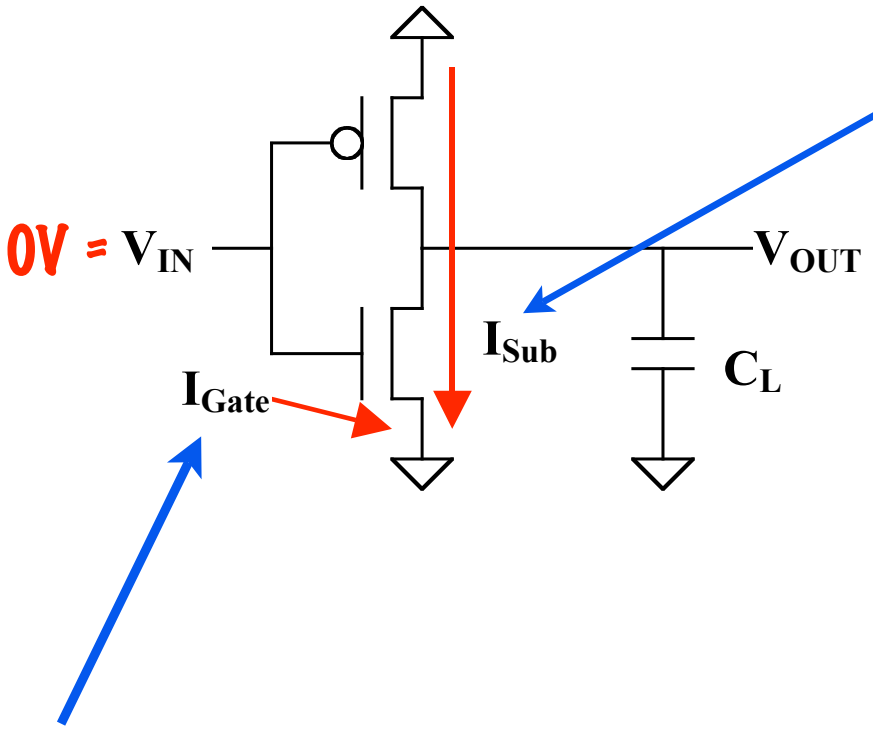

Isub: Even when this nFet is off, it passes an Ioff leakage current.

We can engineer any Ioff we like, but a lower Ioff also results in a lower Ion, and thus a lower maximum clock speed.

Igate: 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

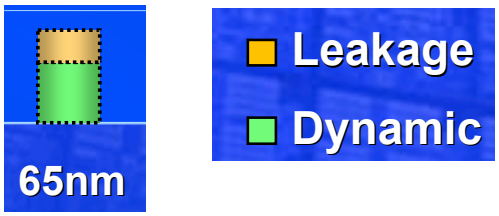

A lot of work was done to get a ratio this good ... 50/50 is common.

Bill Holt, Intel, Hot Chips 17

Engineering "On" Current at 25 nm ...

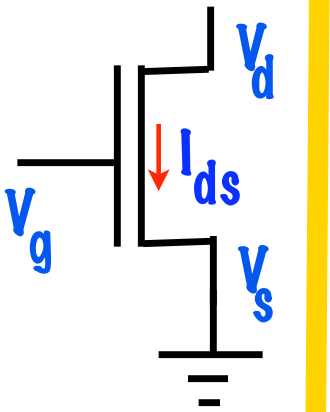

We can increase I_{on} by raising V_{dd} and/or lowering V_t.

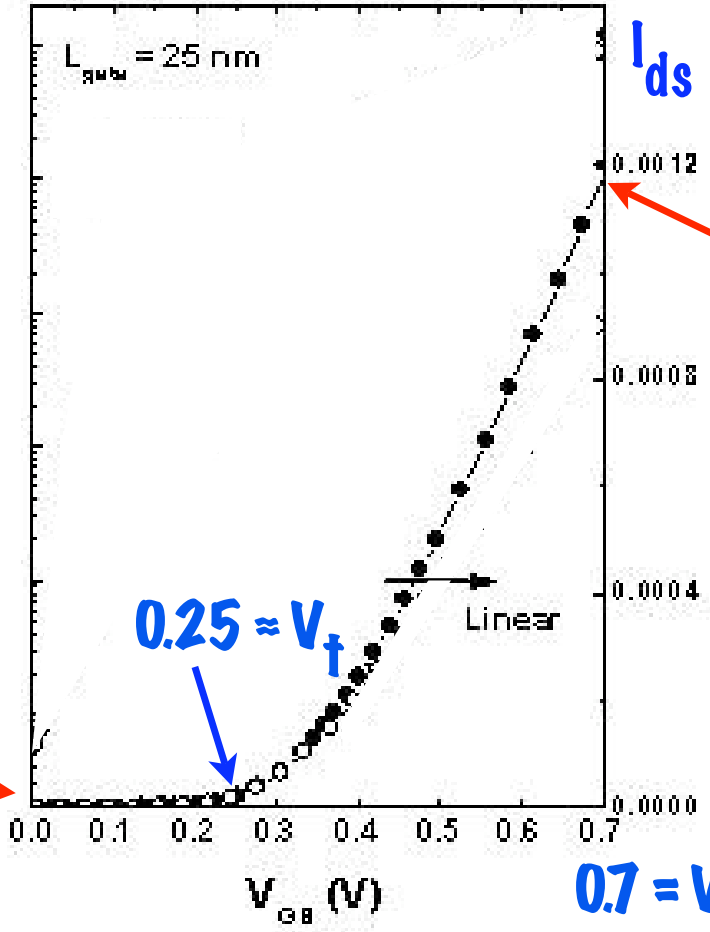

1.2 mA = I_{on}

0.25 ≈ V_t

I_{off} = 0 ???

0.7 = V_{dd}

Plot on a "Log" Scale to See "Off" Current

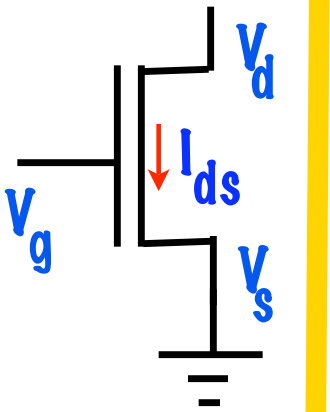

We can decrease I_{off} by raising V_t - but that lowers I_{on}.

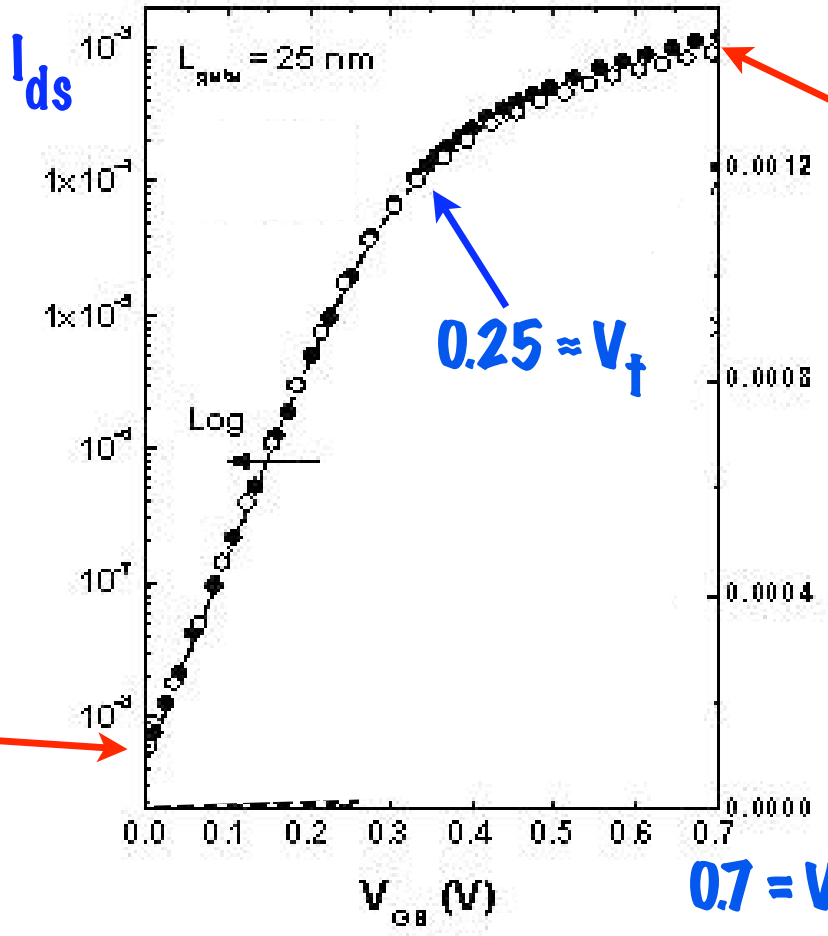

1.2 mA = I_{on}

0.25 ≈ V_t

I_{off} ≈ 10 nA

0.7 = V_{dd}

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.

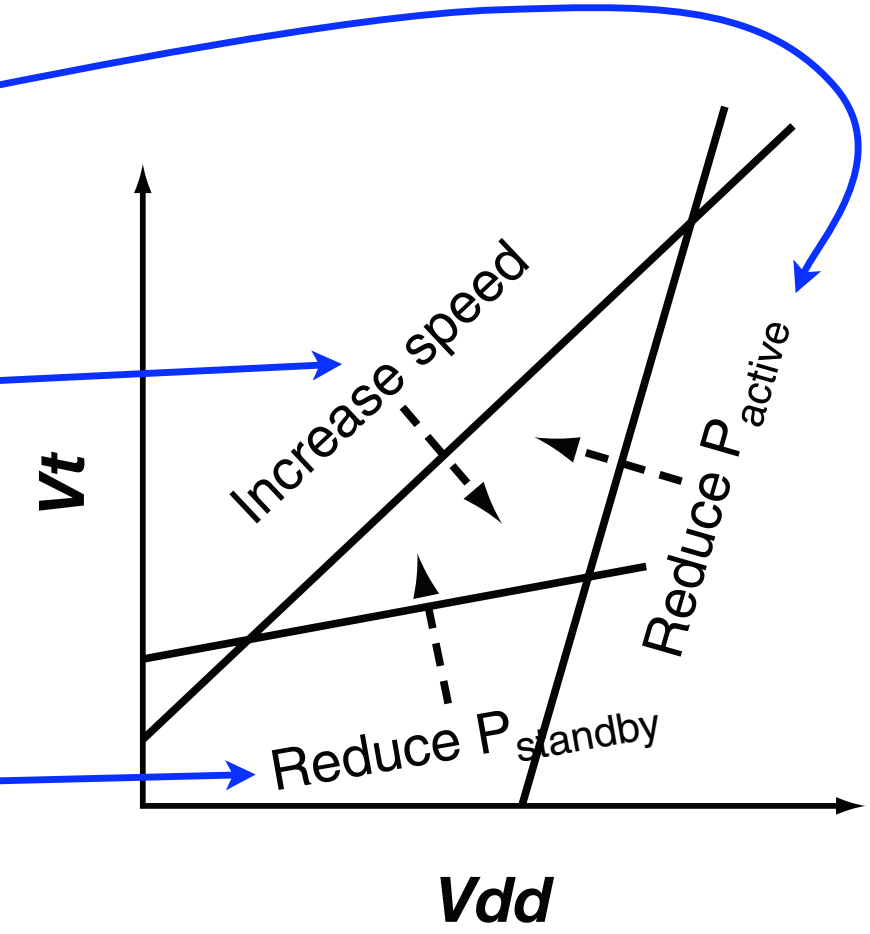

From: Silicon Device Scaling to the Sub-10-nm Regime

Meikei Ieong,[1*] Bruce Doris,[2] Jakub Kedzierski,[1] Ken Rim,[1] Min Yang[1]

Customize processes for product types ...

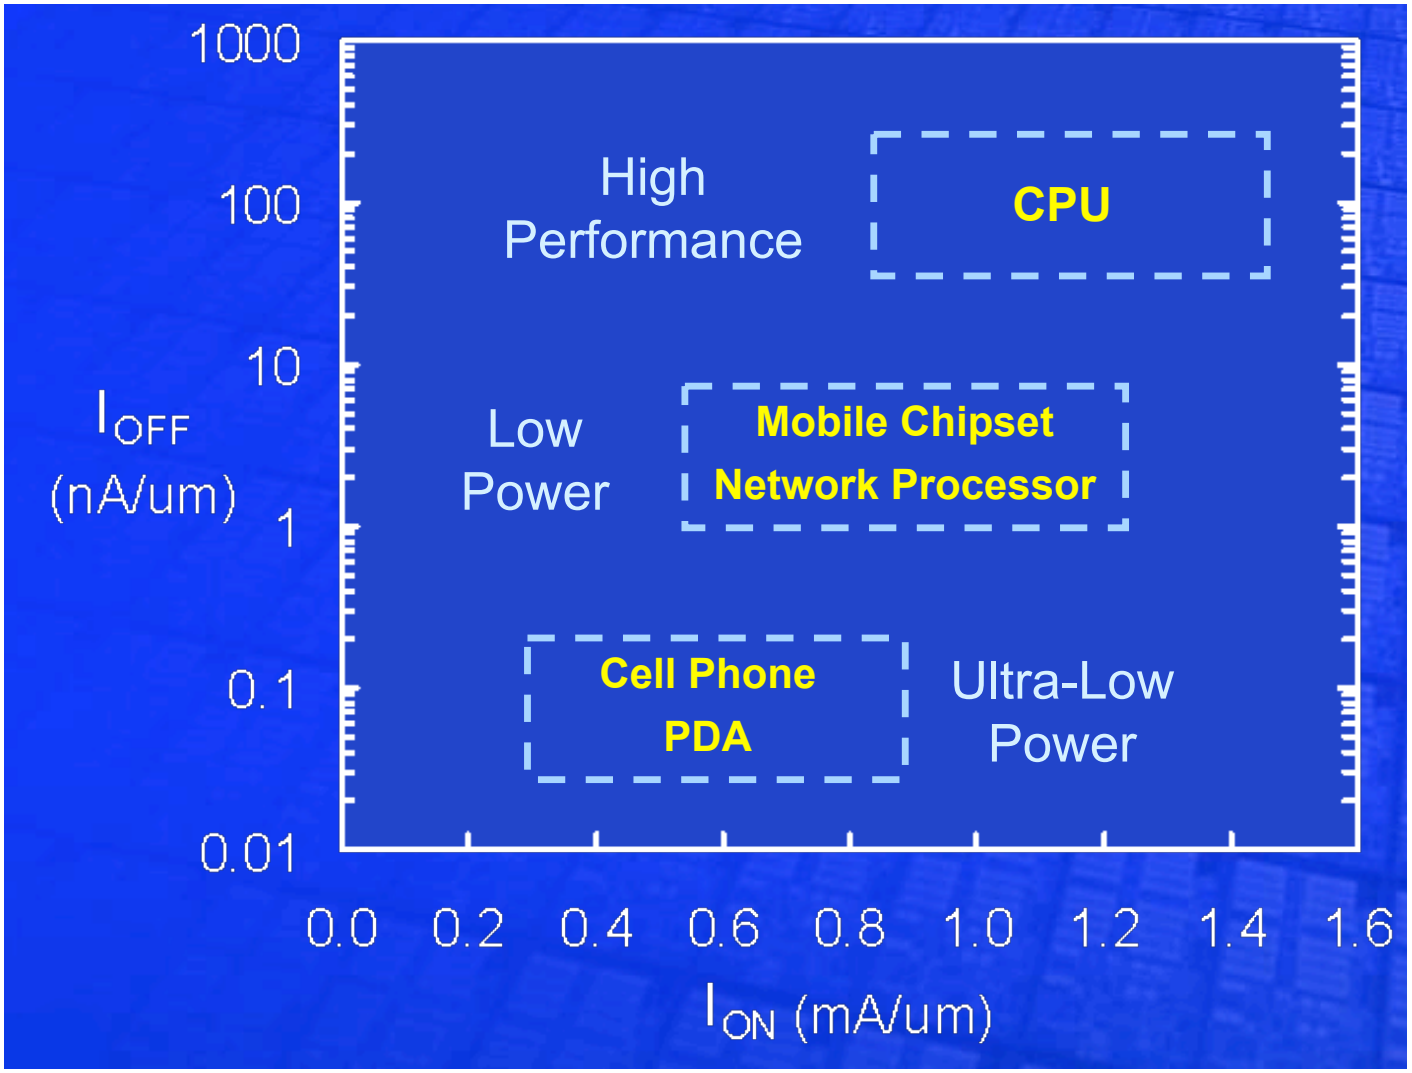

From: "Facing the Hot Chips Challenge Again", Bill Holt, Intel, presented at Hot Chips 17, 2005.

Intel: Comparing 2 CPU generations ...

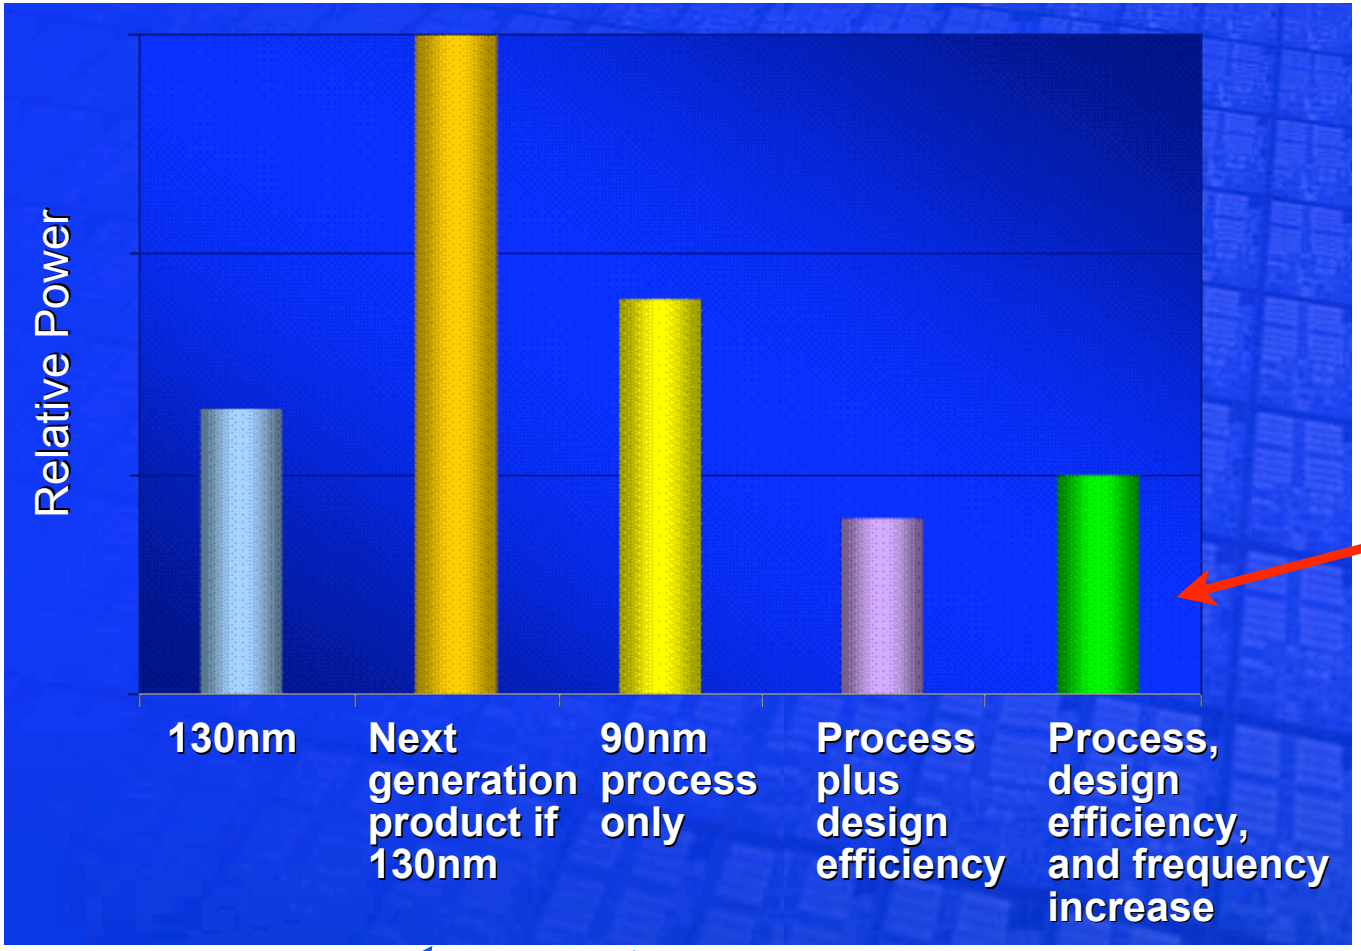

Find enough tricks, and you can afford to raise Vdd a little so that you can raise the clock speed!

Clock speed unchanged ...

Lower Vdd, lower C, but more leakage.

Design tricks: architecture & circuits.

Long-term possibility: New devices

Electrostatic mechanical relays at the nanoscale

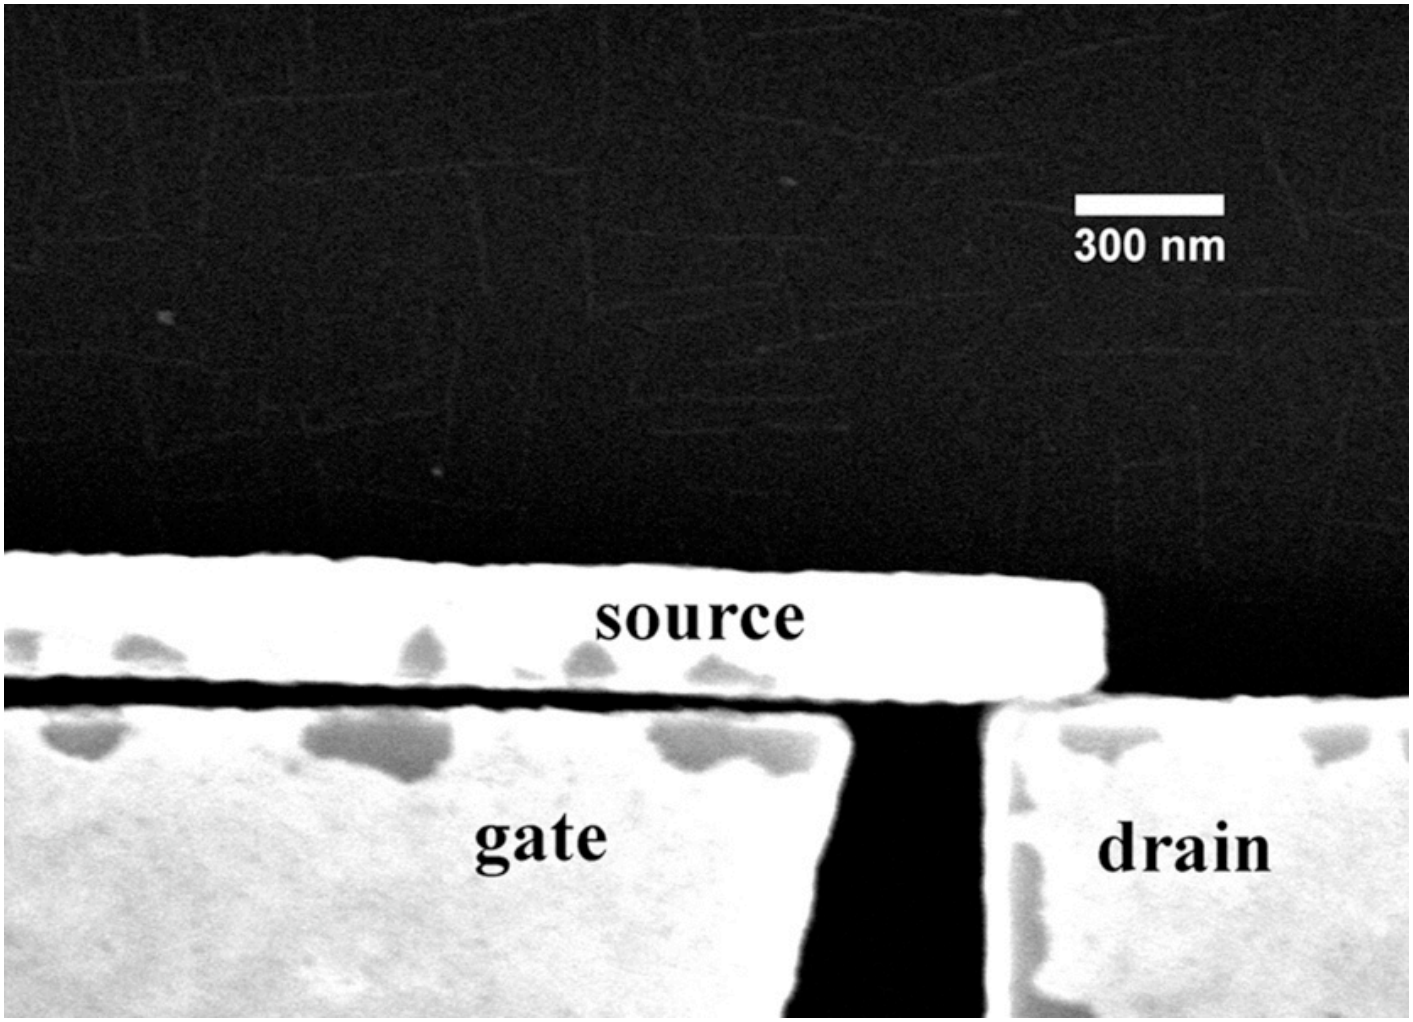

Electromechanical Computing at 500°C with Silicon Carbide. Te-Hao Lee, Swarup Bhunia, Mehran Mehregany

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

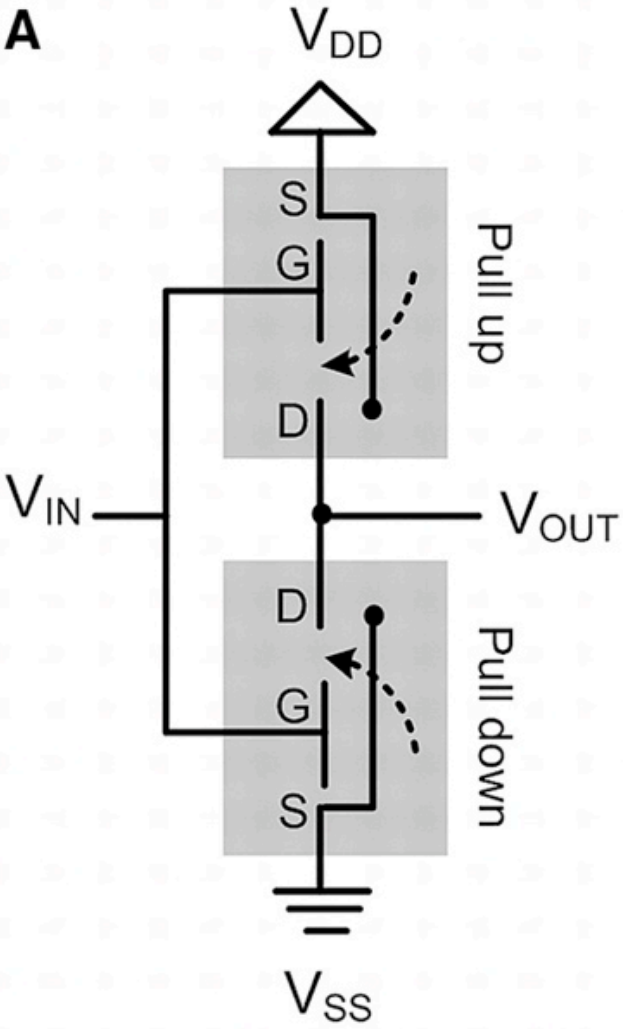

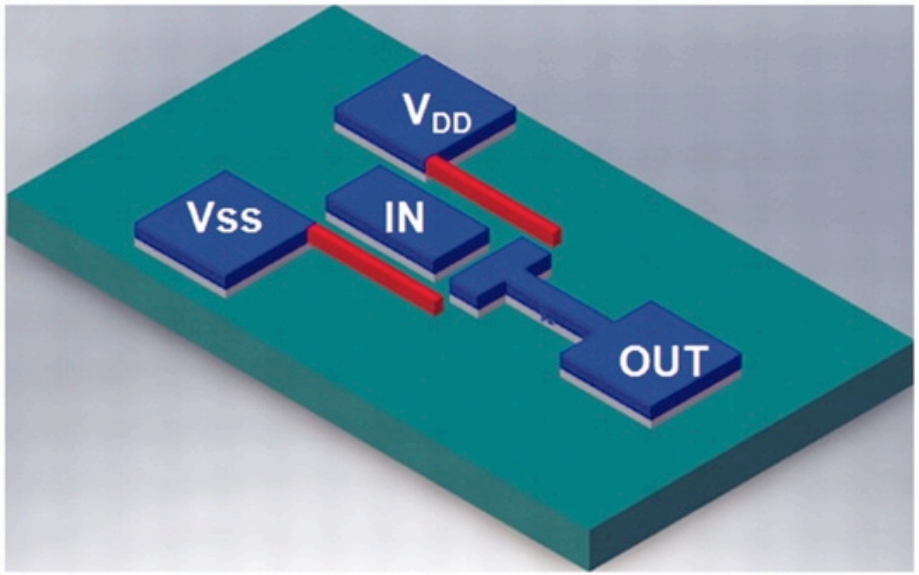

+ 10 fA leakage current
+ Works at 500 degrees C

- Fails after 1-10 days of 500 kHz toggles.
- Switching requires 6V V_{dd}

Electromechanical Computing at 500°C with Silicon Carbide. Te-Hao Lee, Swarup Bhunia, Mehran Mehregany

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 ...

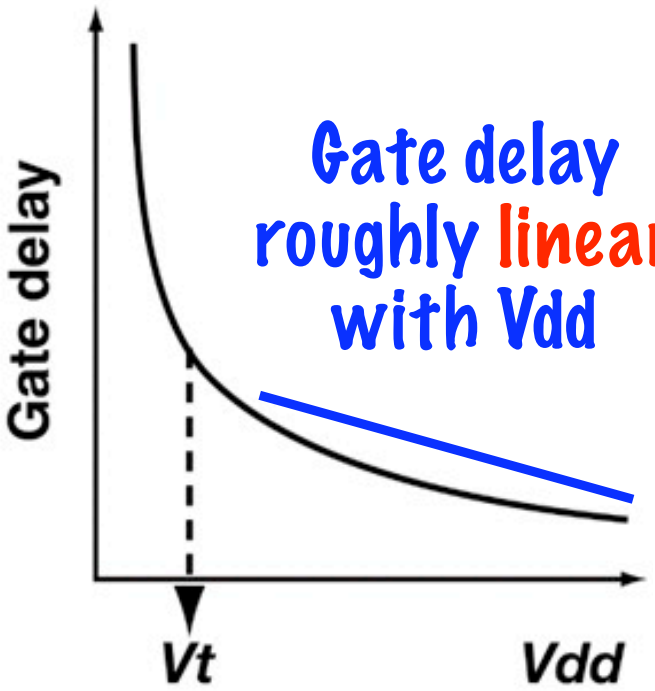

Gate delay roughly linear with Vdd

And so, we can transform this:

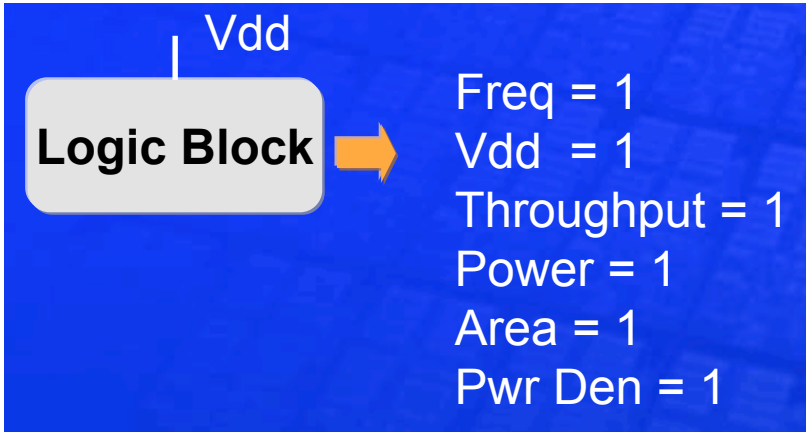

Block processes stereo audio. 1/2 of clocks for "left", 1/2 for "right".

Into this:

CV^2 power only

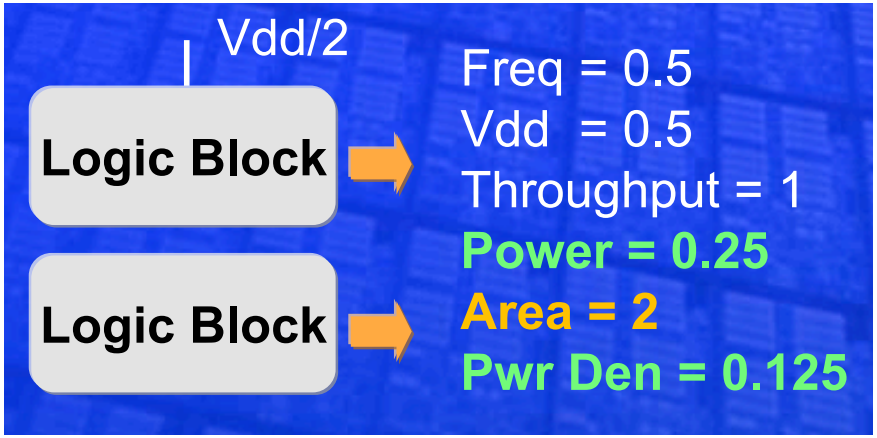

Ex: Top block processes audio channel 1, bottom block processes audio channel 2.

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

Chandrakasan & Brodersen (UCB EECS)

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

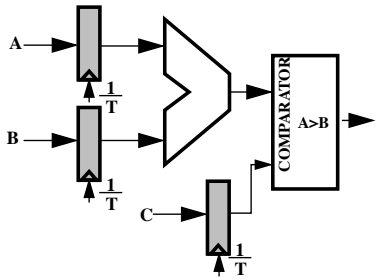

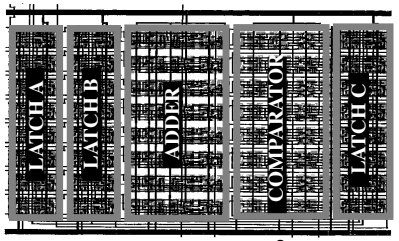

Area = 636 x 833 μ^2

Simple

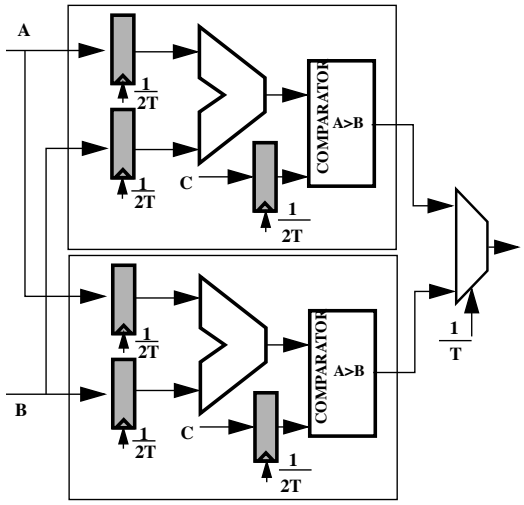

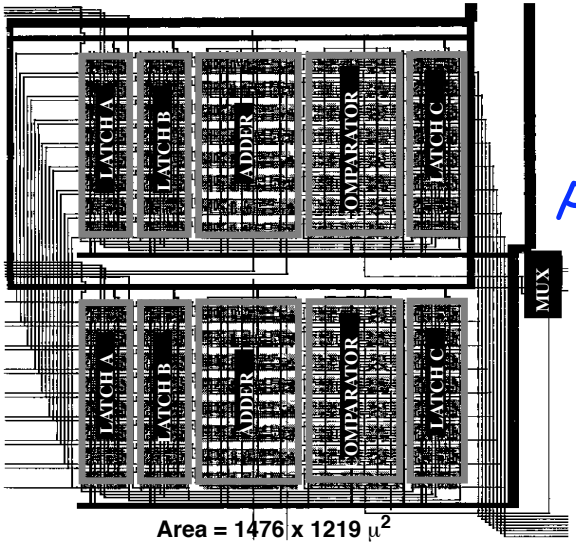

Area = 1476 x 1219 μ^2

Parallel

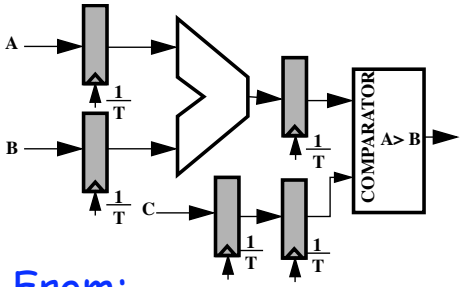

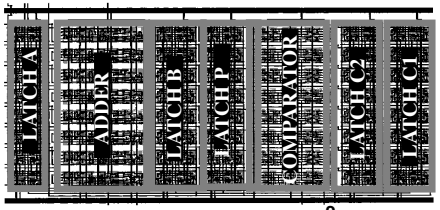

Area = 640 x 1081 μ^2

Pipelined

From:

Minimizing Power Consumption in CMOS Circuits Anantha P. Chandrakasan, Robert W. Brodersen

Multiple Cores for Low Power

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

Cell:
The PS3 chip

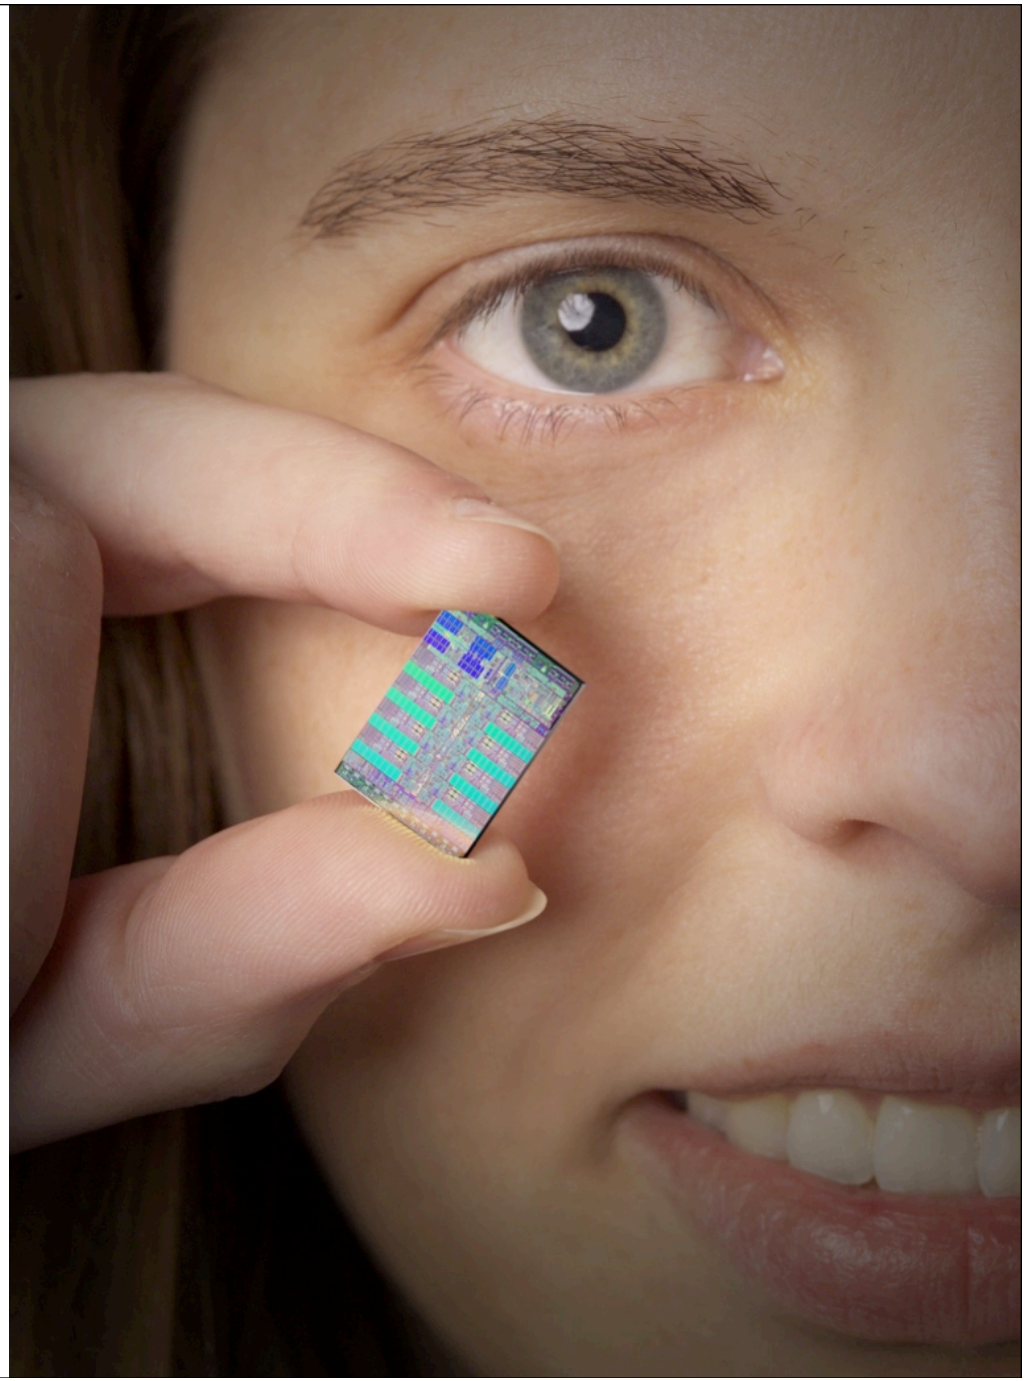

Cell (PS3 Chip): 1 CPU + 8 "SPUs"

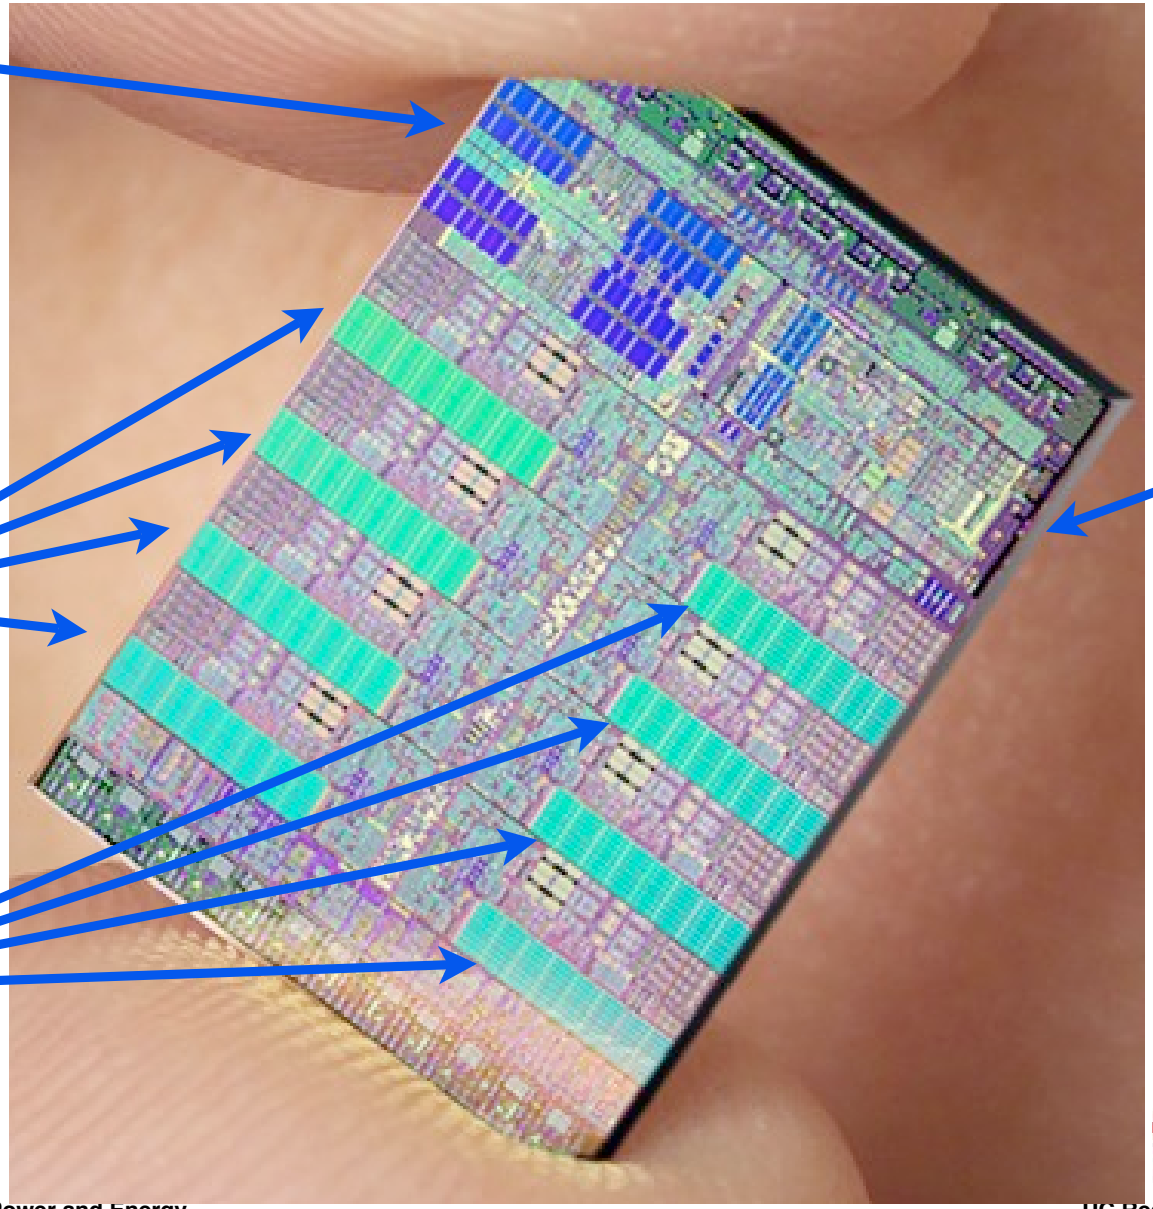

IBM

SONY
COMPUTER
ENTERTAINMENT

TOSHIBA

One Synergistic Processing Unit (SPU)

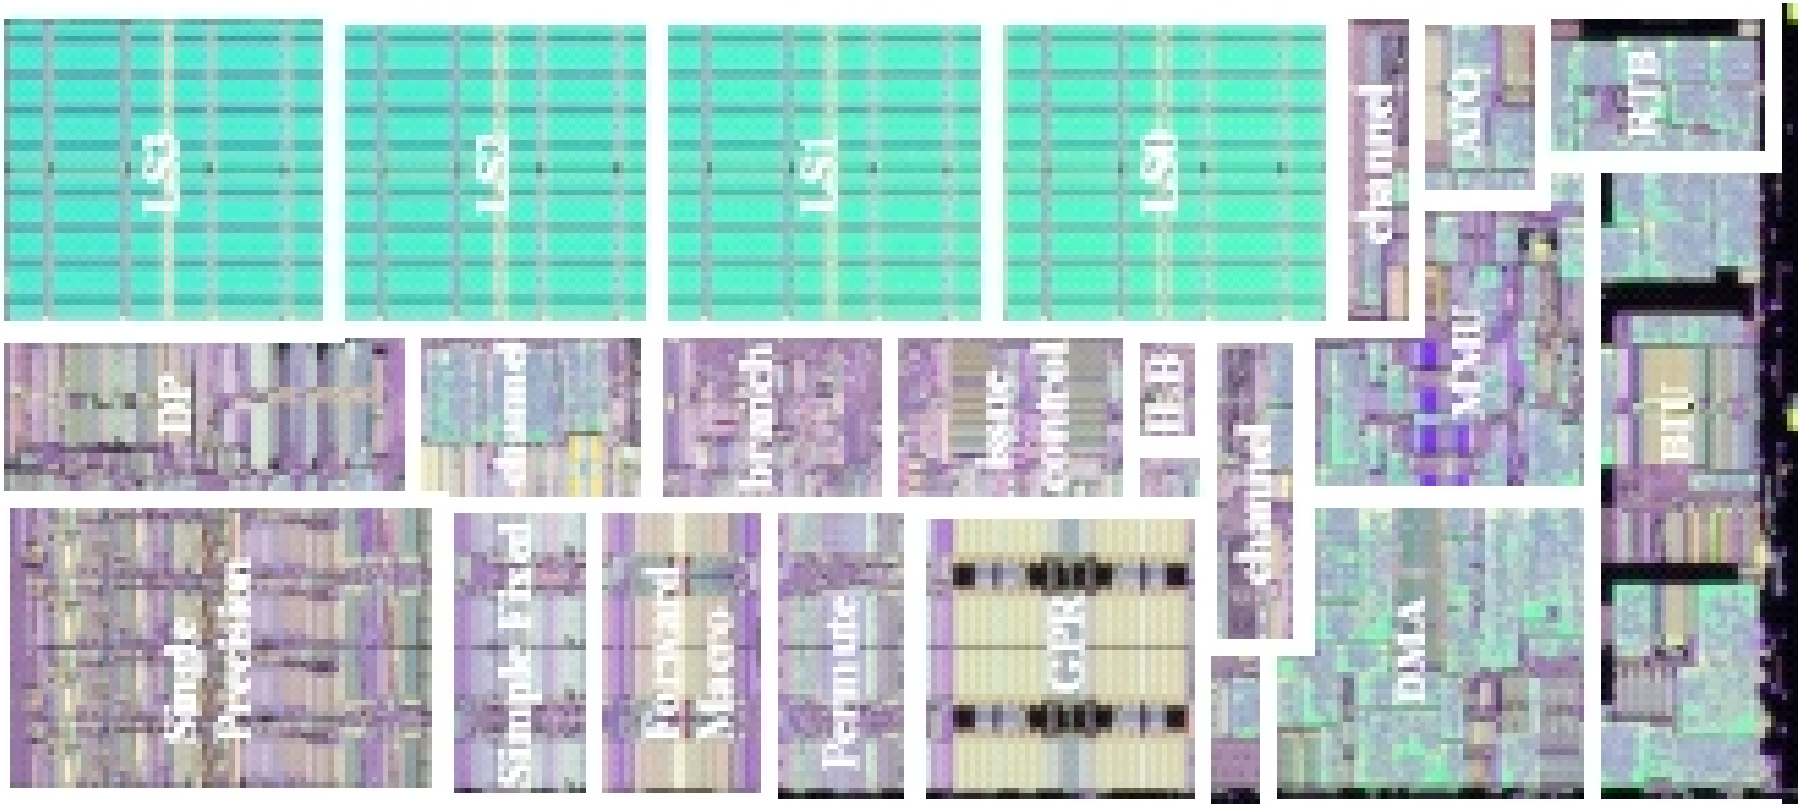

SPU issues 2 inst/cycle (in order) to 7 execution units

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 Vdd, the less dynamic energy consumption.

$$E_{0\to1} = \frac{1}{2} C V_{dd}^2 \qquad E_{1\to0} = \frac{1}{2} C V_{dd}^2$$

The lower Vdd, the longer the maximum clock period, the slower the clock frequency.

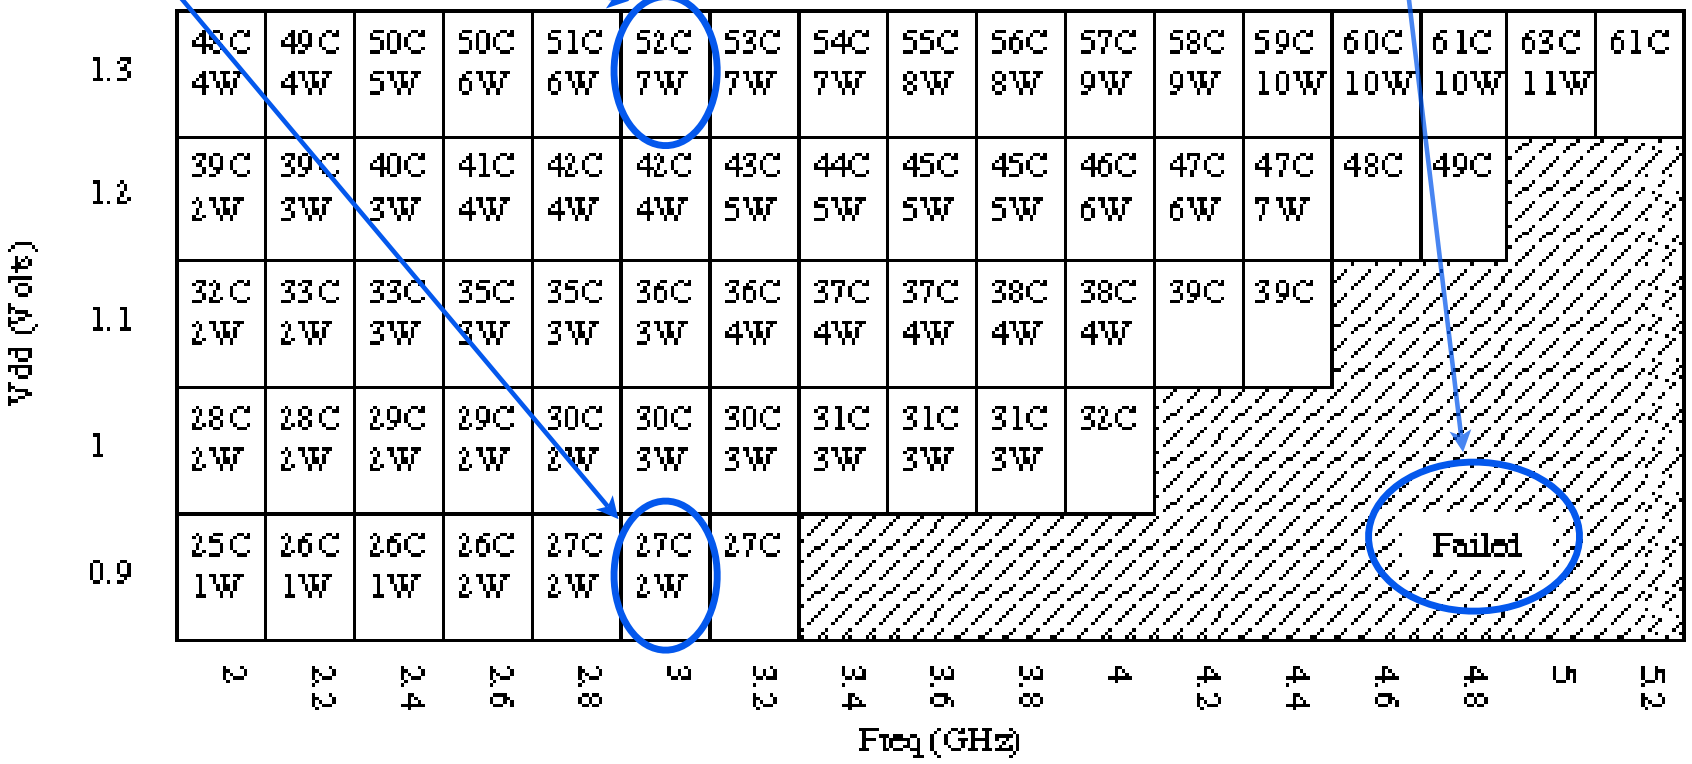

Clock speed alone doesn't help E/op ...

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

$$E_{0->1} = \frac{1}{2} C V_{dd}^2 \qquad E_{1->0} = \frac{1}{2} C V_{dd}^2$$

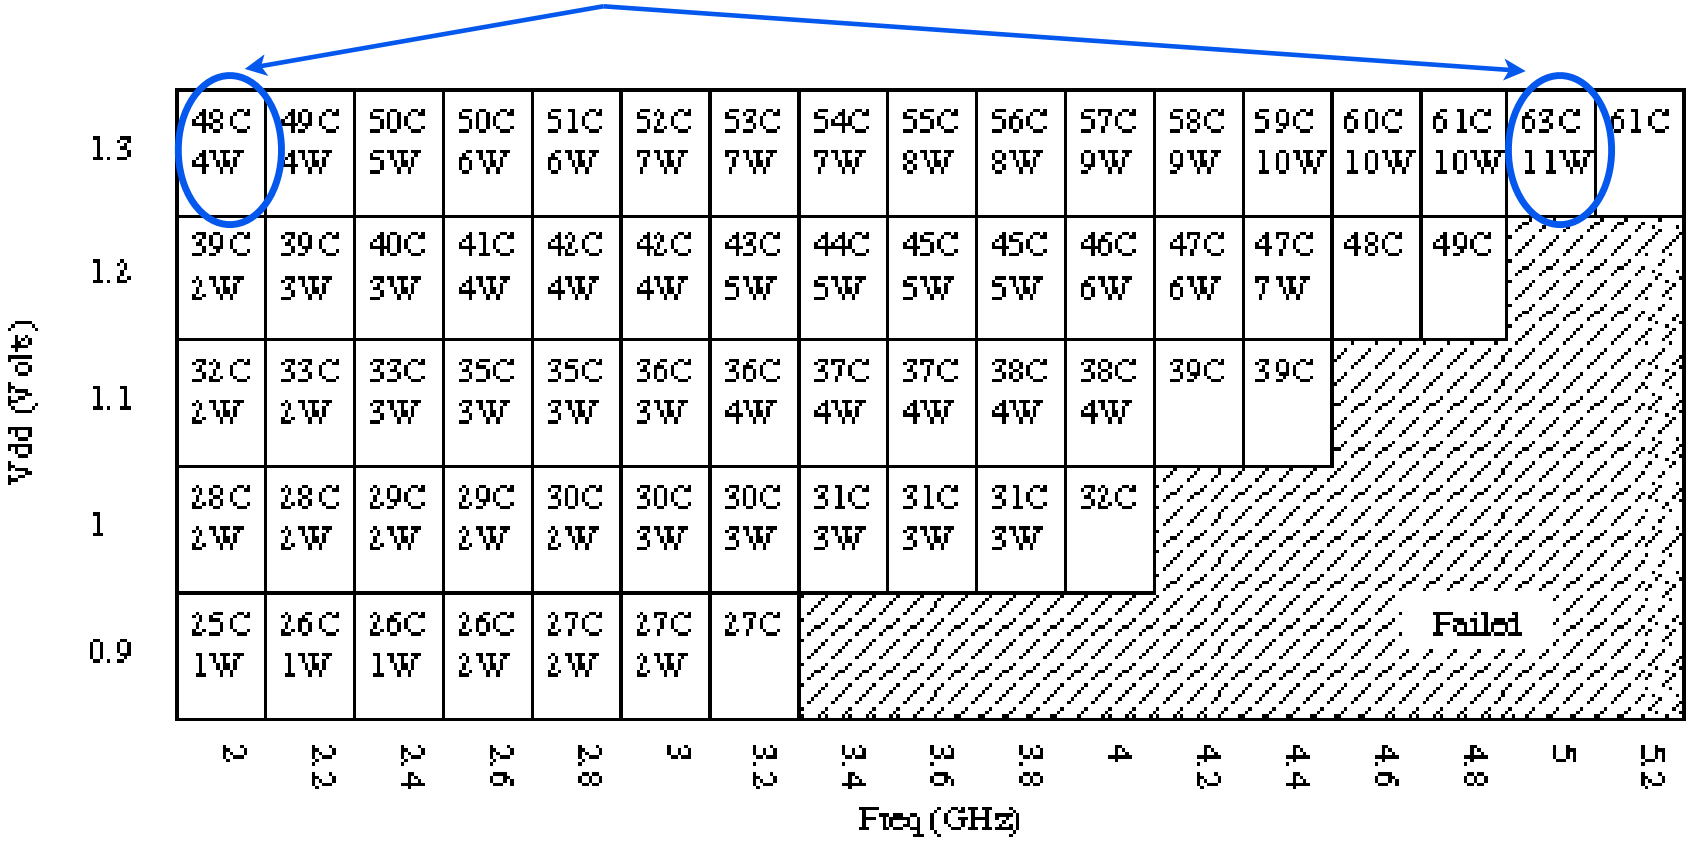

Scaling V and f *does* lower energy/op

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.

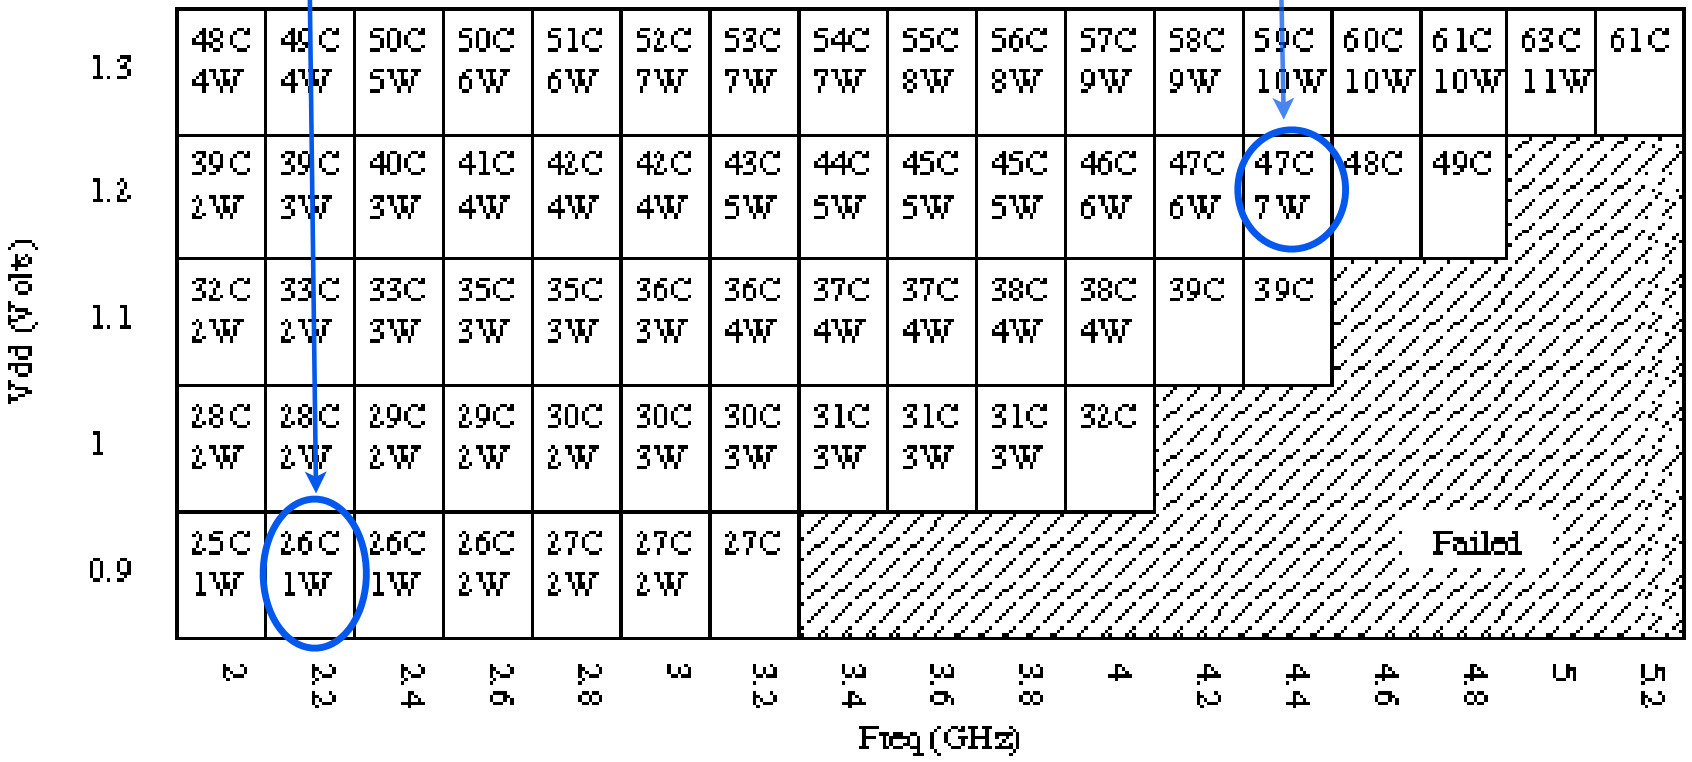

Intel's dual-core analysis ...

But only if your app(s) can put 2 cores to use!

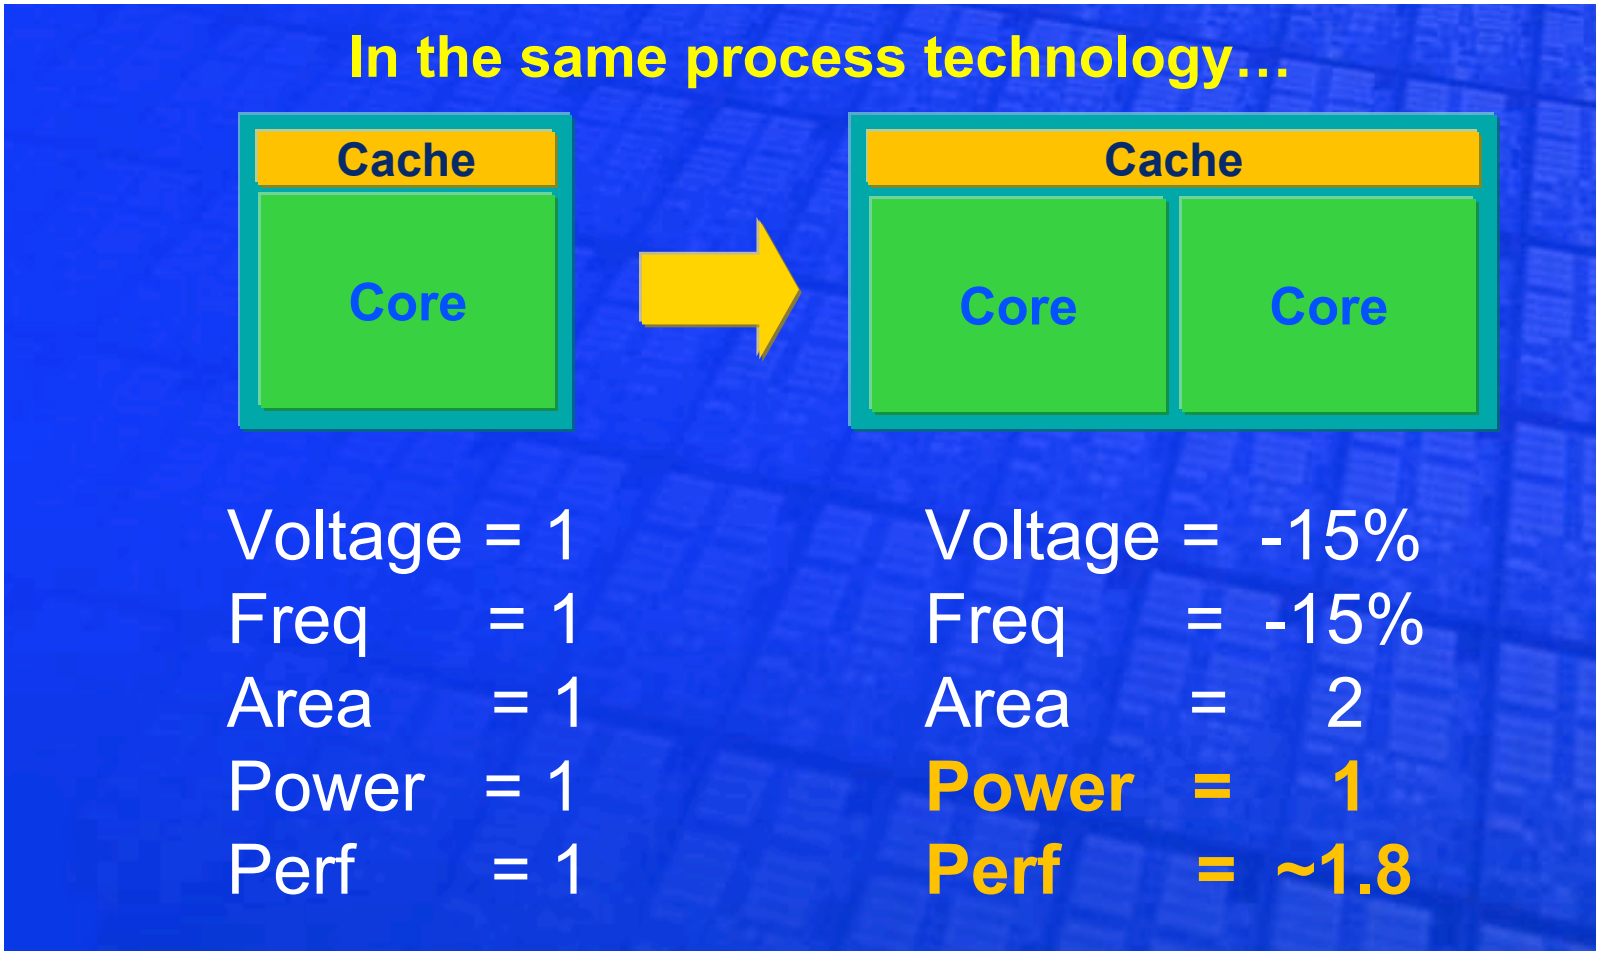

From: "Facing the Hot Chips Challenge Again", Bill Holt, Intel, presented at Hot Chips 17, 2005.

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

portalplayer™

PP5020

digital media management system-on-chip

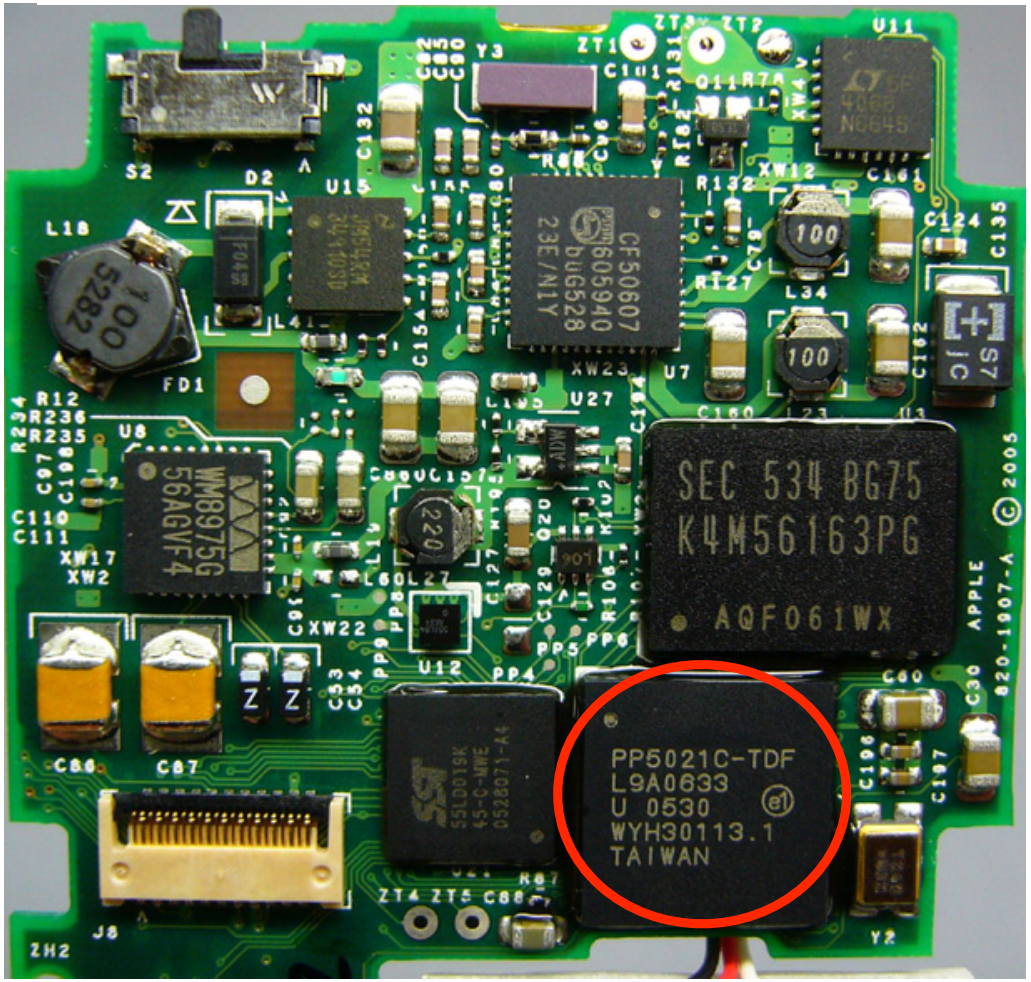

Dual ARM Processors

- 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 ...

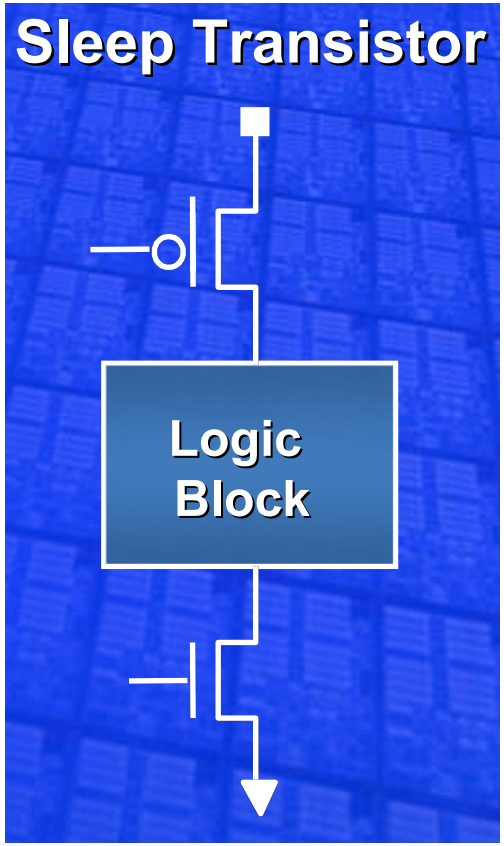

Example: Floating point unit logic.

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

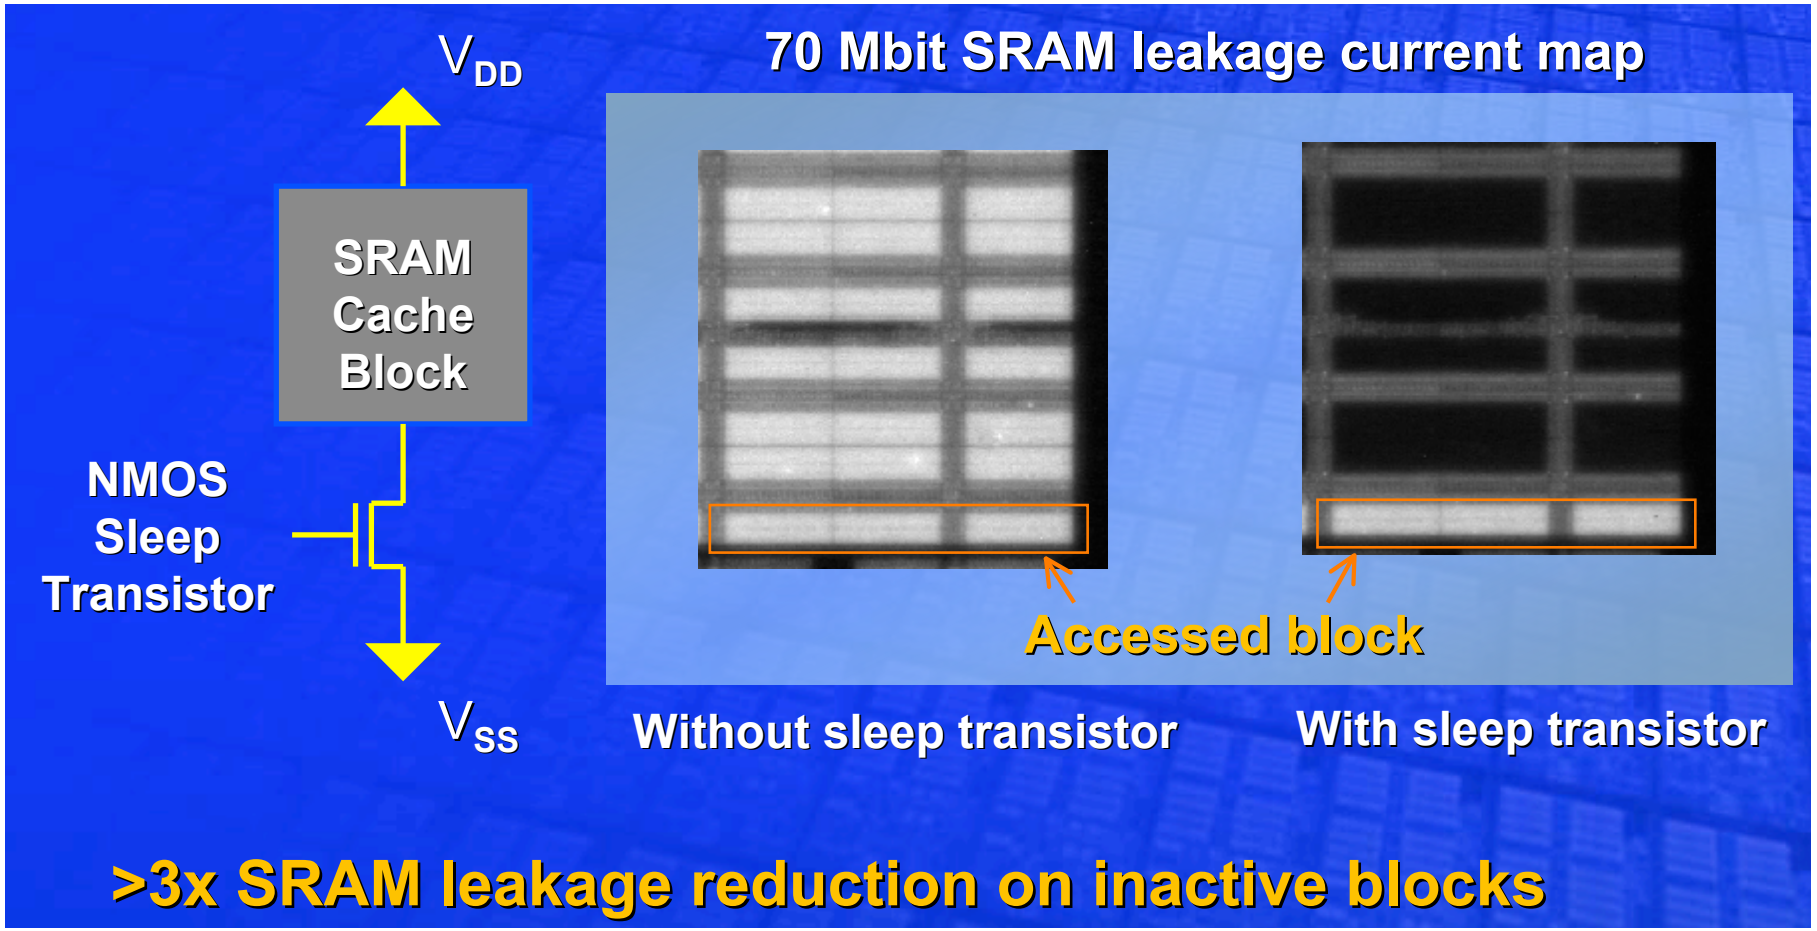

From: "Facing the Hot Chips Challenge Again", Bill Holt, Intel, presented at Hot Chips 17, 2005.

Fall 2009: Intel Mainstream Desktop

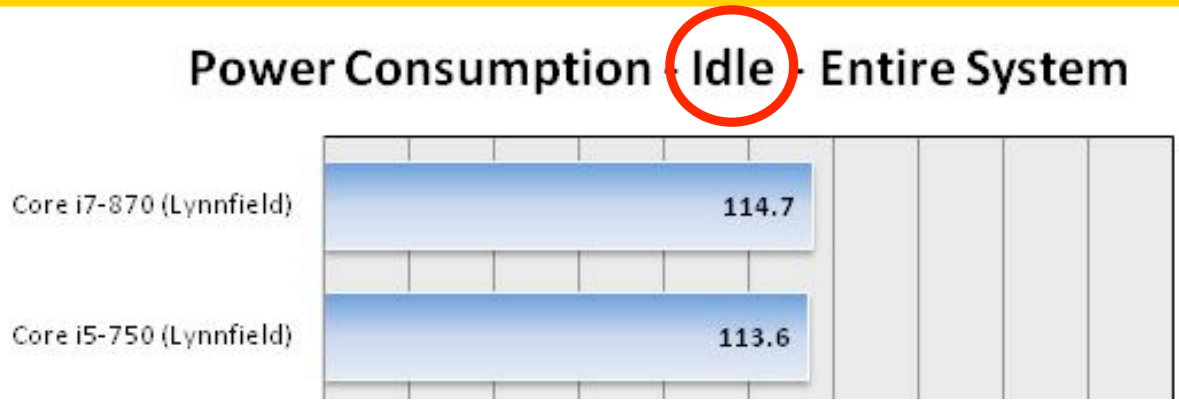

From: PC Perspective website

Slow down "slack paths"

Fact: Most logic on a chip is "too fast"

The critical path

Most wires have hundreds of picoseconds to spare.

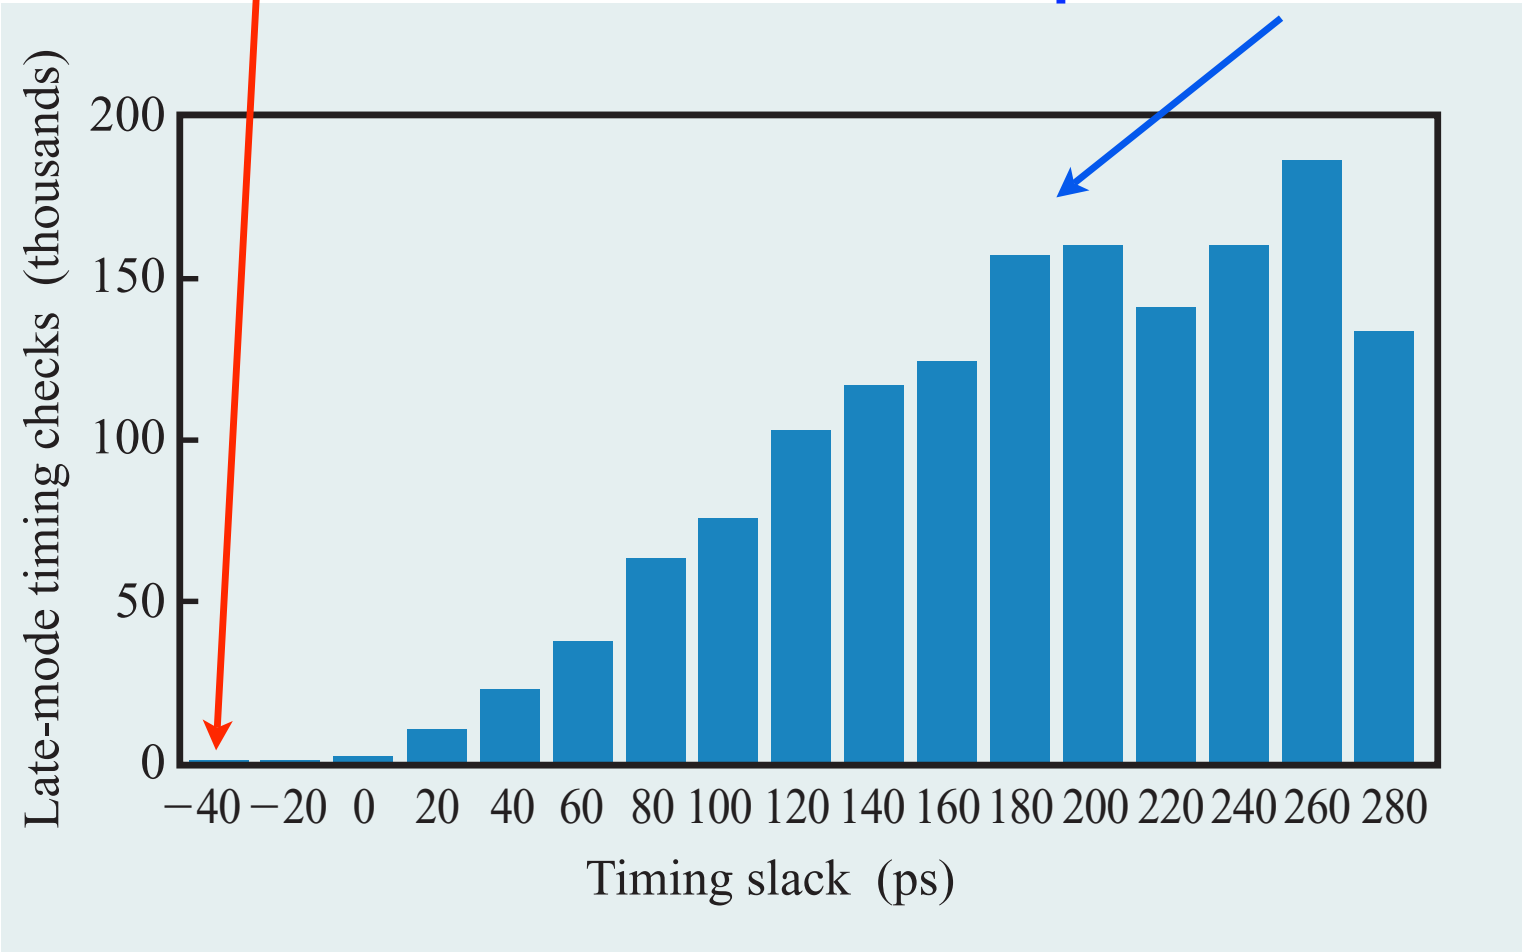

From "The circuit and physical design of the POWER4 microprocessor", IBM J Res and Dev, 46:1, Jan 2002, J.D. Warnock et al.

Use several supply voltages on a chip ...

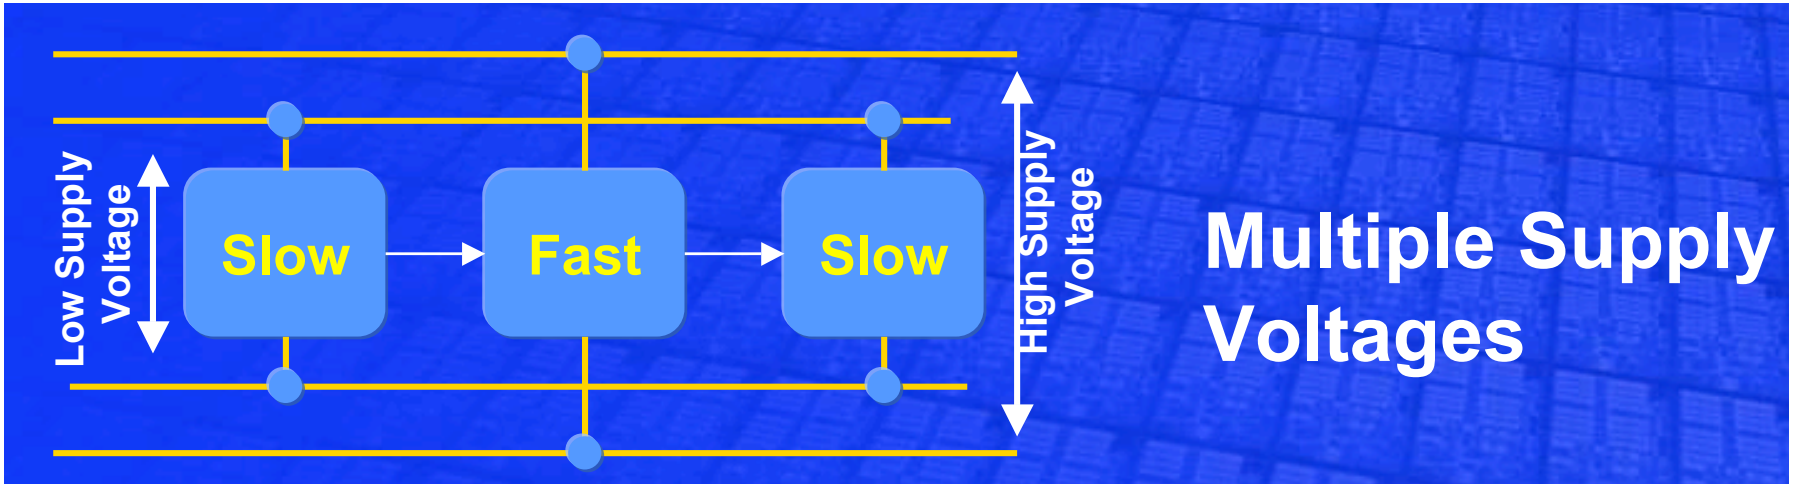

Why use multi-Vdd? We can reduce dynamic power by using low-power Vdd for logic off the critical path.

What if we can't do a multi-Vdd design?
In a multi-Vt process, we can reduce leakage power on the slow logic by using high-Vth 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.

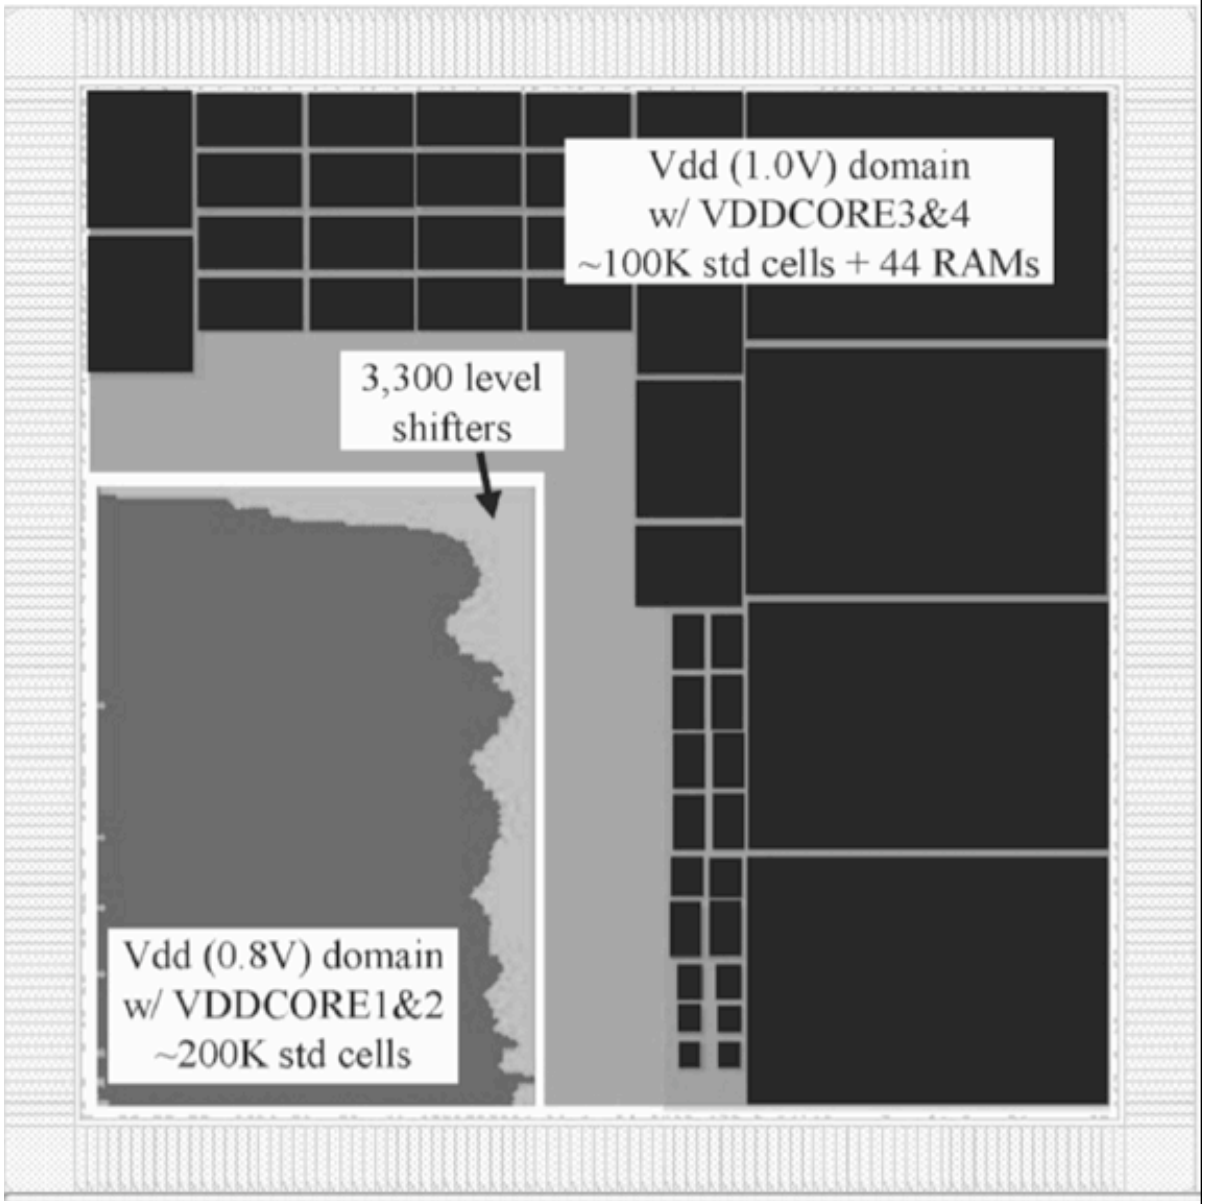

From a chapter from new book on ASIC design by Chinnery and Keutzer (UCB).

Gating clocks to save power

On a CPU, where does the power go?

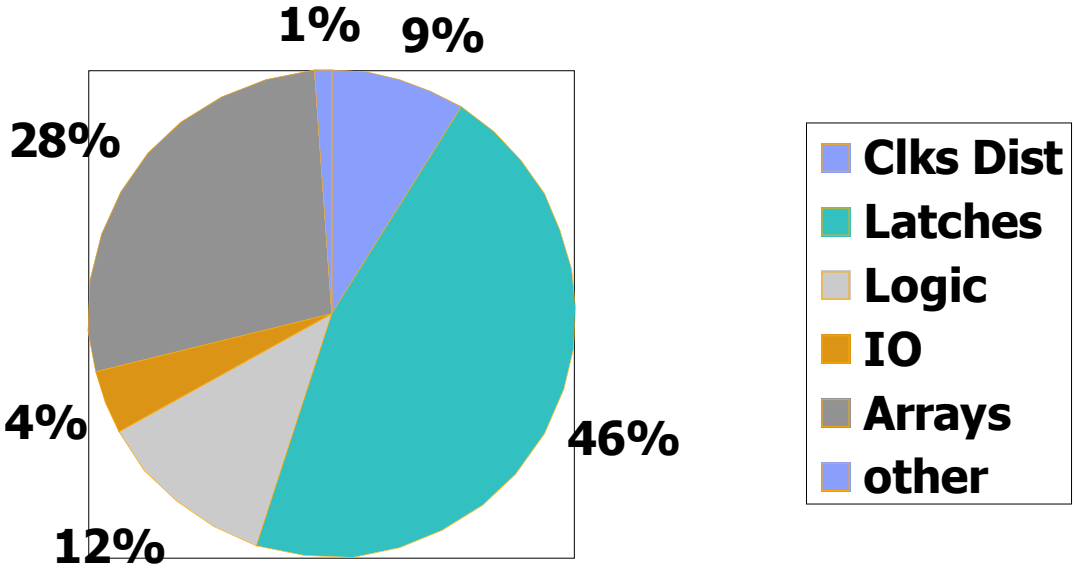

Half of the power goes to latches (Flip-Flops).

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.

From: Bose, Martonosi, Brooks: Sigmetrics-2001 Tutorial

Synopsis Power Compiler can do this ...

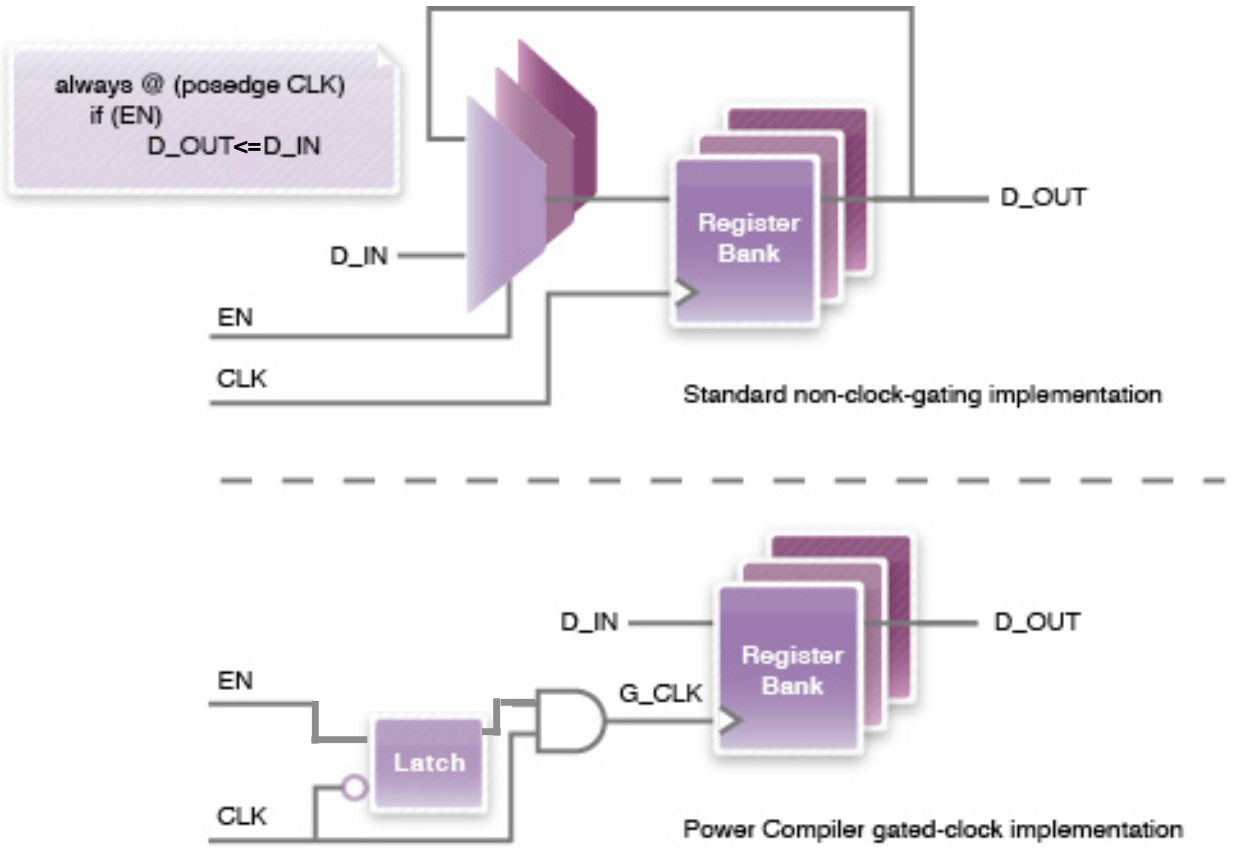

"Up to 70% power savings at the block level, for applicable circuits"
Synopsis Data Sheet

Power Compiler also can do this ...

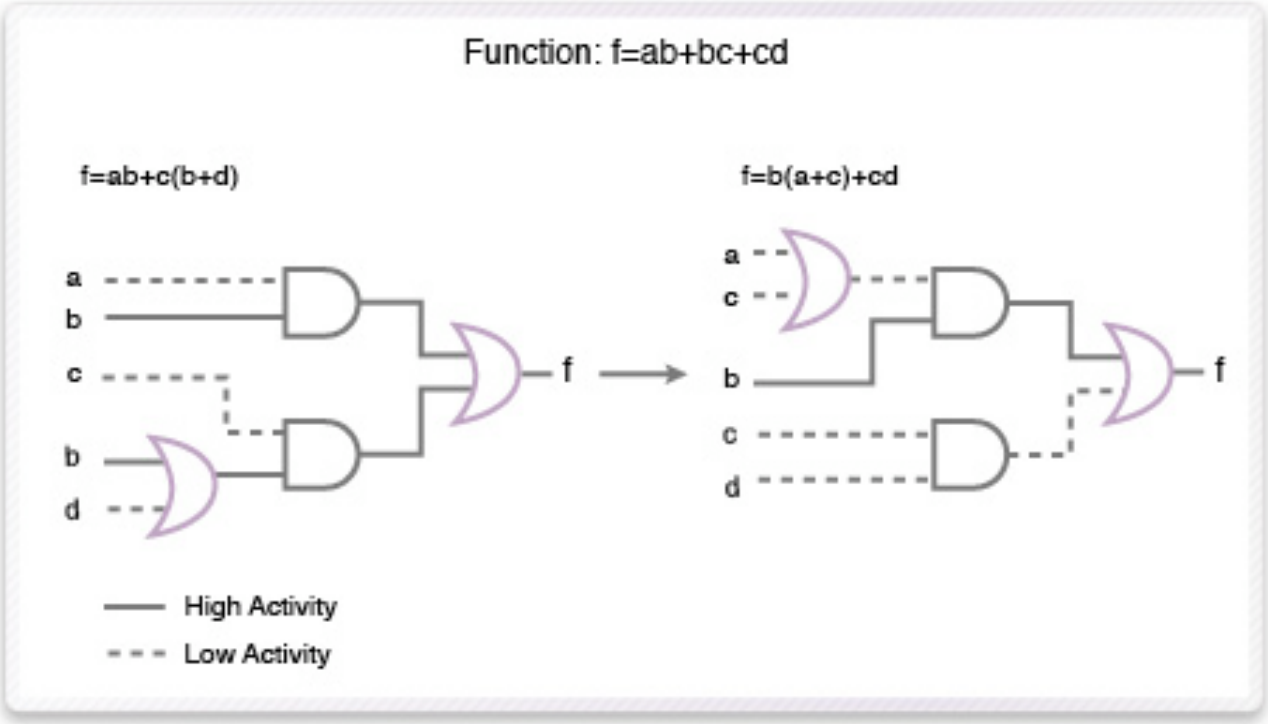

10-20% push-button power savings, using techniques like this one.

Thermal Management

Keep chip cool to minimize leakage power

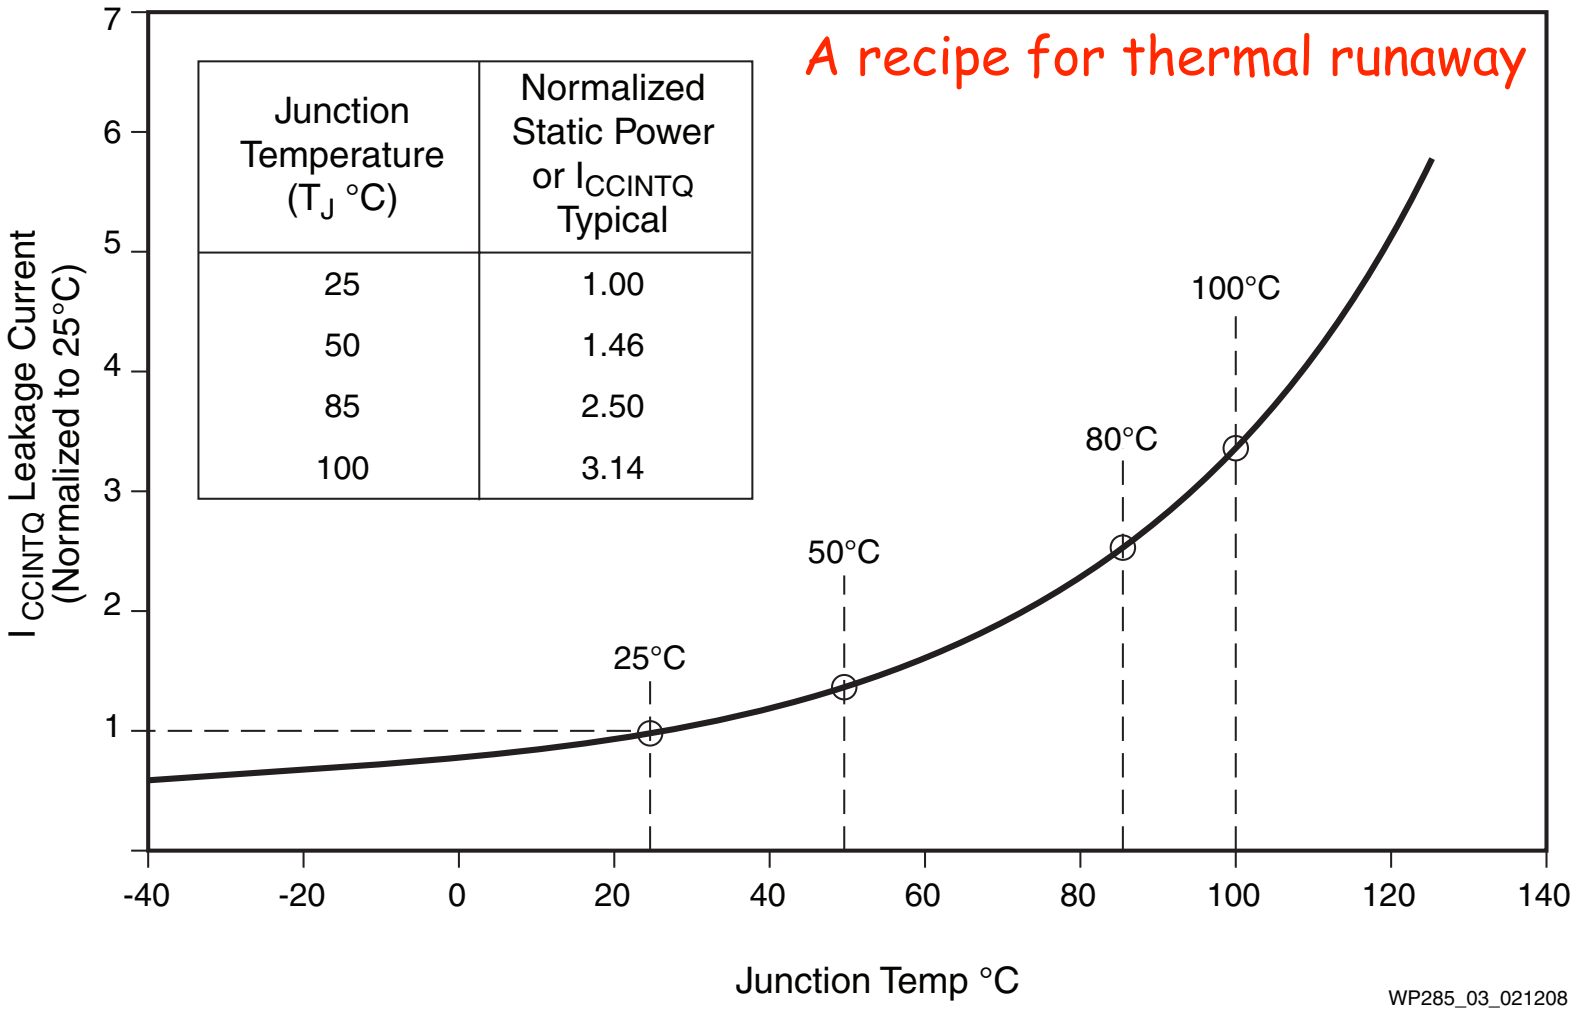

Figure 3: **I_{CCINTQ} vs. Junction Temperature with Increase Relative to 25°C**

Monitor die temperature, servo clock speed

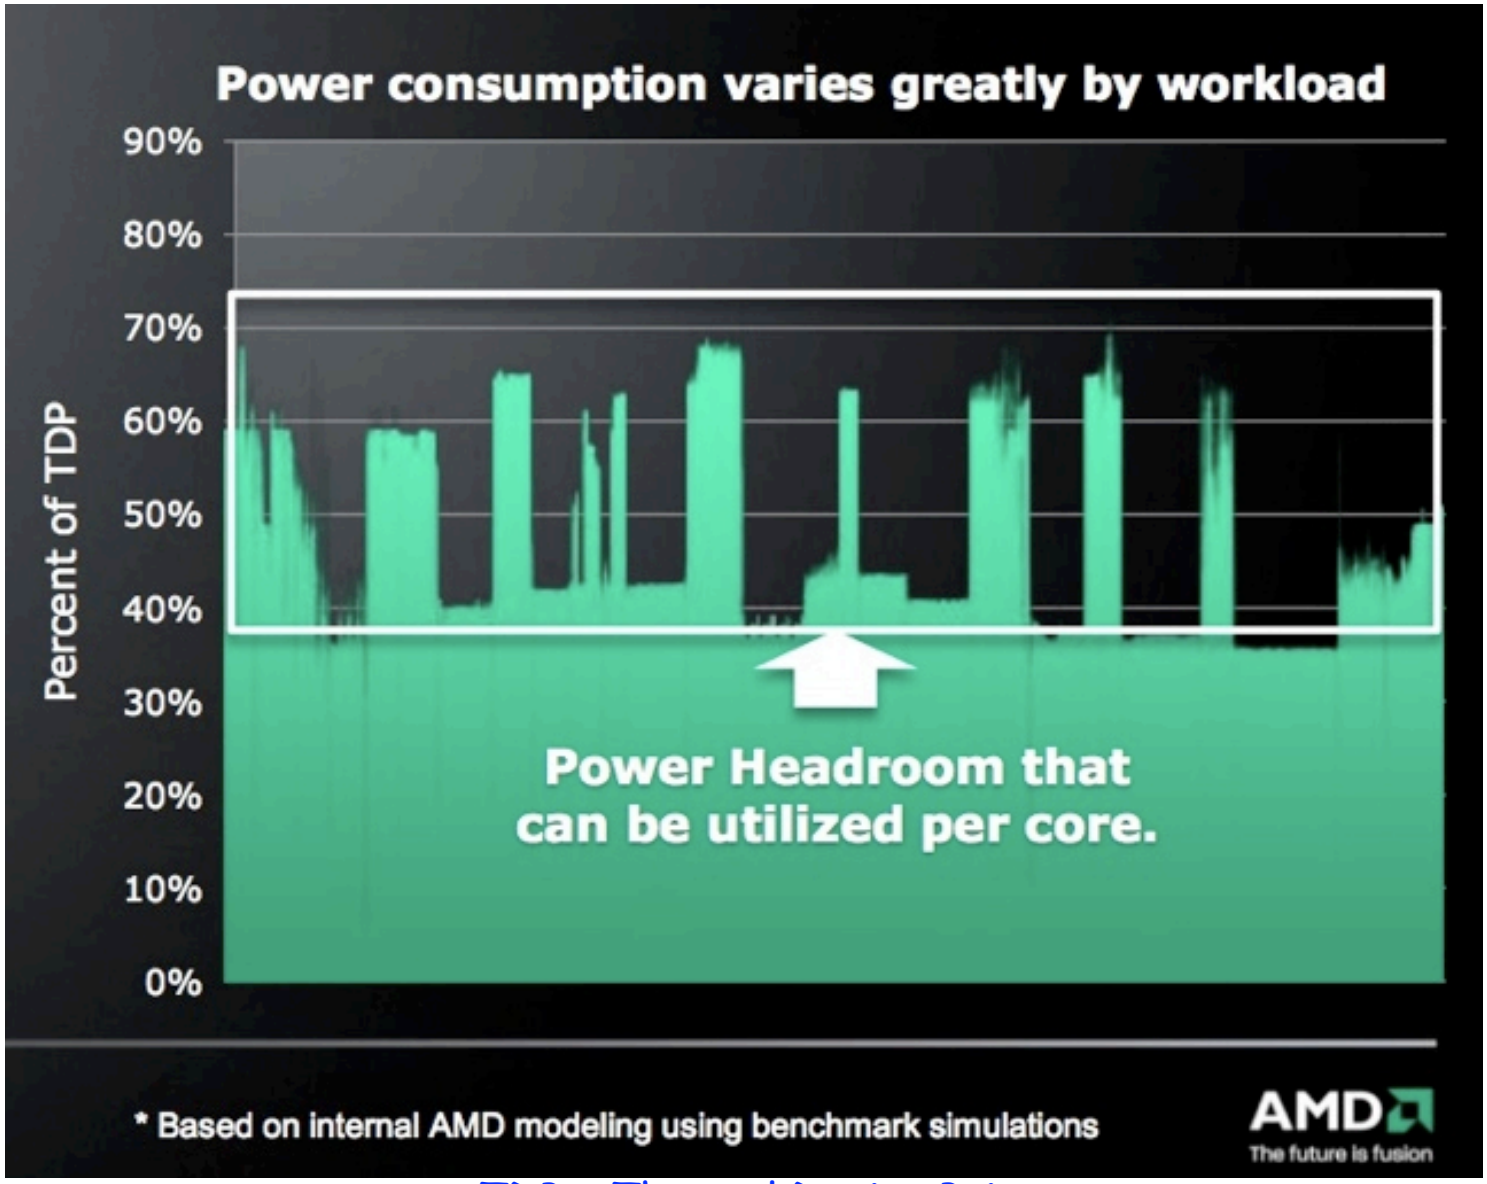

TDP = Thermal Design Point

Preceded by Intel: "Turbo Boost"

Processor	Intel Core i7-870 2.93 GHz				Intel Core i7-860 2.80 GHz				Intel Core i5-750 2.66 GHz			
Processor Cores	4				4				4			
Active Cores	1C	2C	3C	4C	1C	2C	3C	4C	1C	2C	3C	4C
Maximum Intel® Turbo Boost Technology Bin Upside	5	4	2	2	5	4	1	1	4	4	1	1
Maximum Intel® Turbo Boost Technology Frequency	3.6	3.46	3.2	3.2	3.46	3.33	2.93	2.93	3.2	3.2	2.8	2.8

Bin Size = 133 MHz

IBM Power 4: How does die heat up?

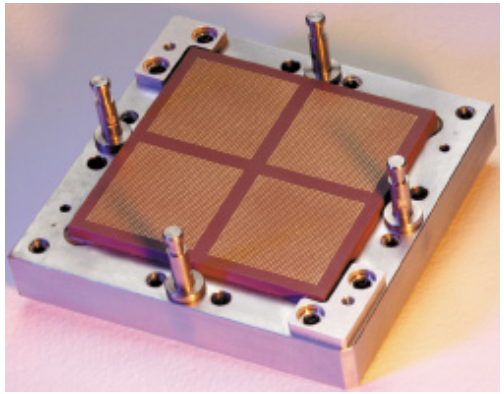

4 dies on a multi-chip module

2 CPUs per die

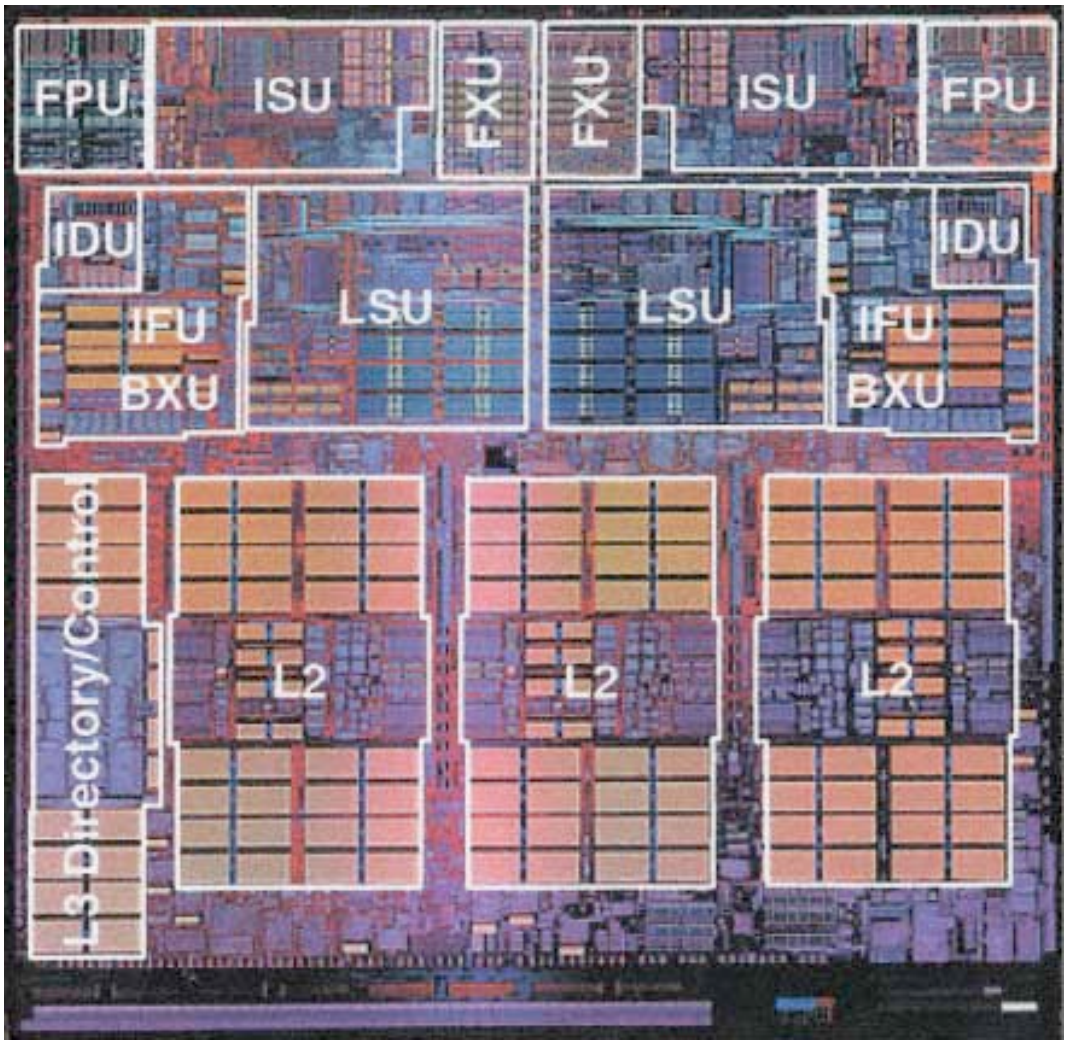

115 Watts: Concentrated in "hot spots"

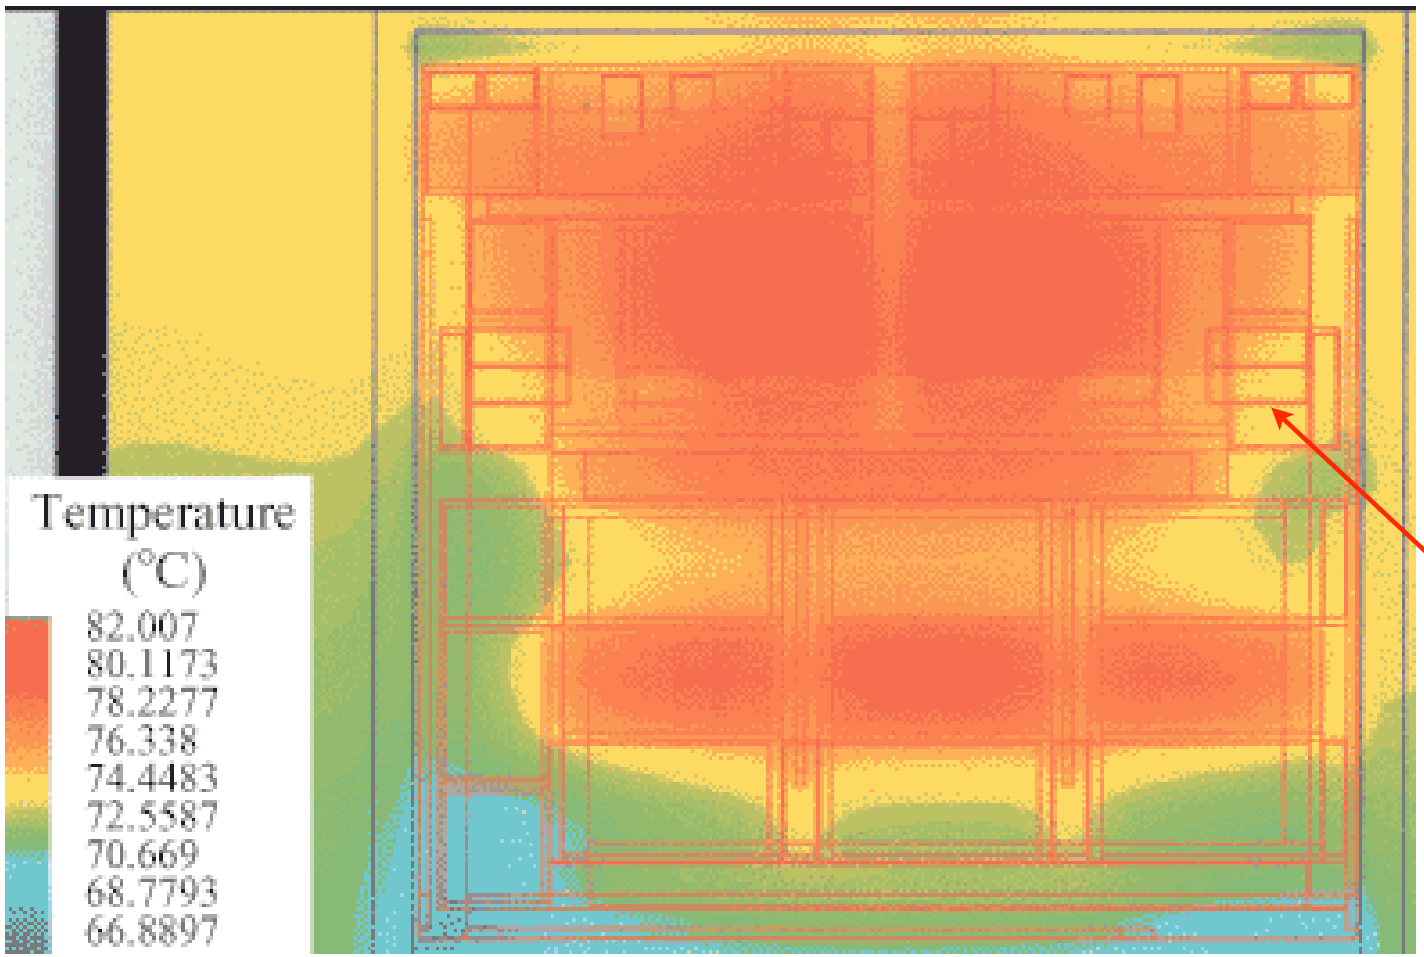

Hot spots

Fixed point units

Cache logic

66.8 C == 152 F 82 C == 179.6

Power and Energy in FPGAs

FPGA: Xilinx Virtex-5 XC5VLX110T

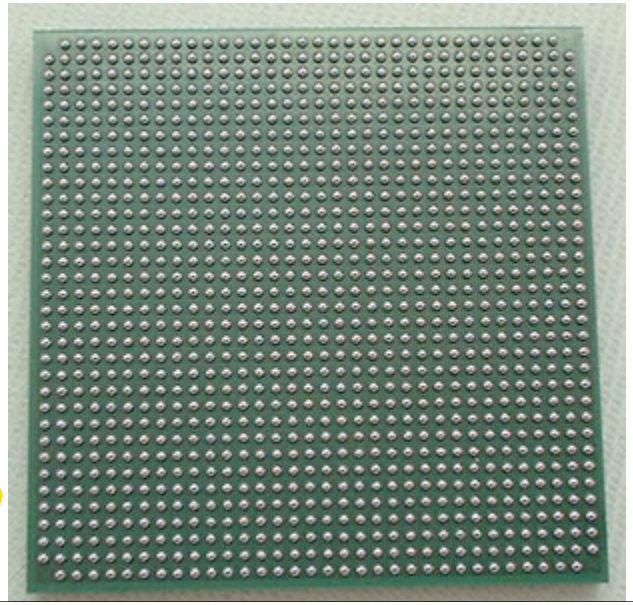

Virtex-5 “die photo”

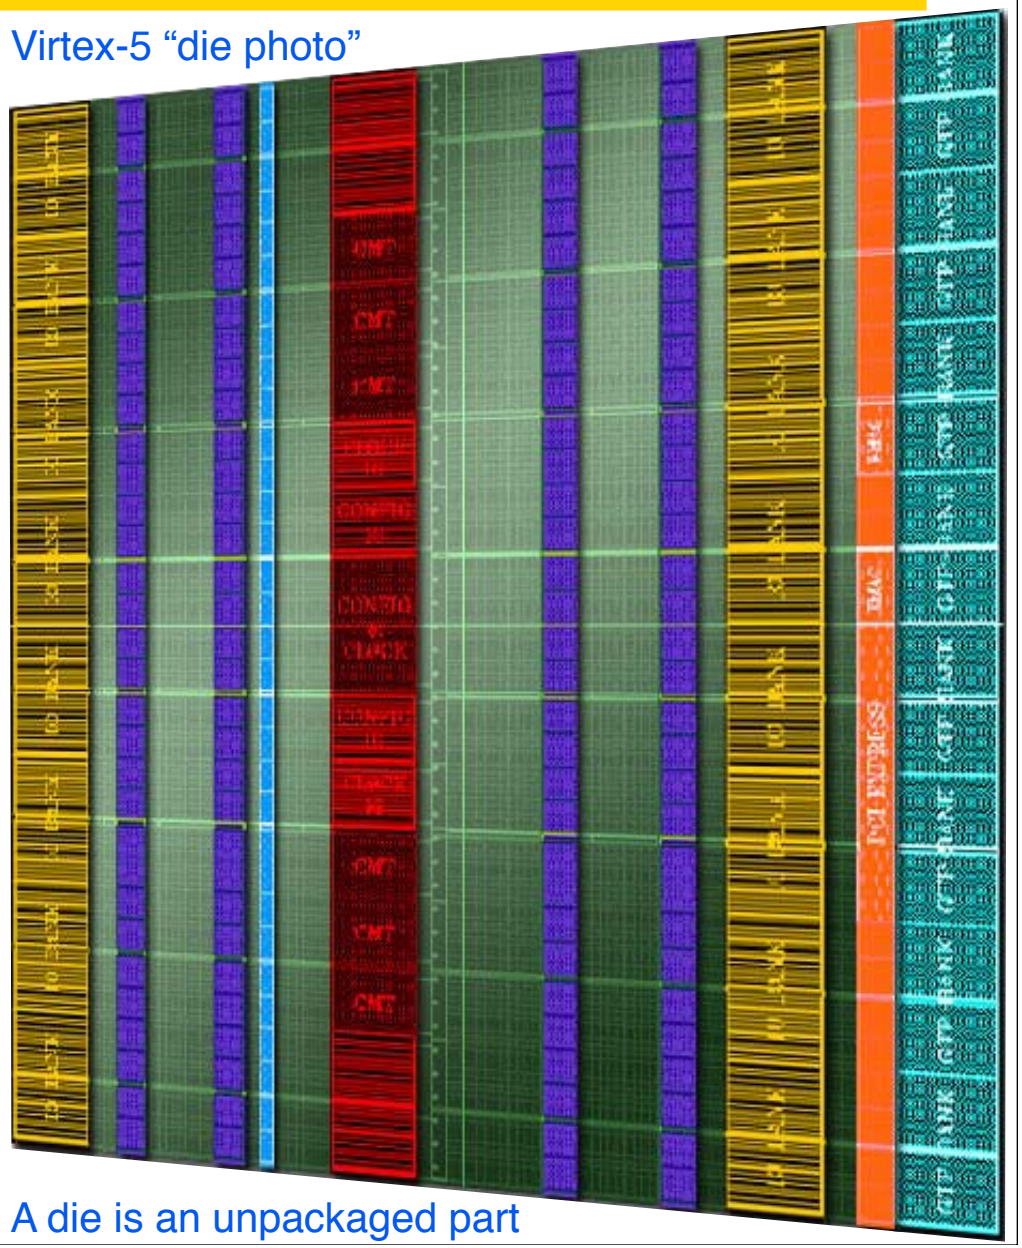

A die is an unpackaged part

Colors represent different types of resources:

Logic
Block RAM
DSP (ALUs)
Clocking
I/O
Serial I/O + PCI

A routing fabric runs throughout the chip to wire everything together.

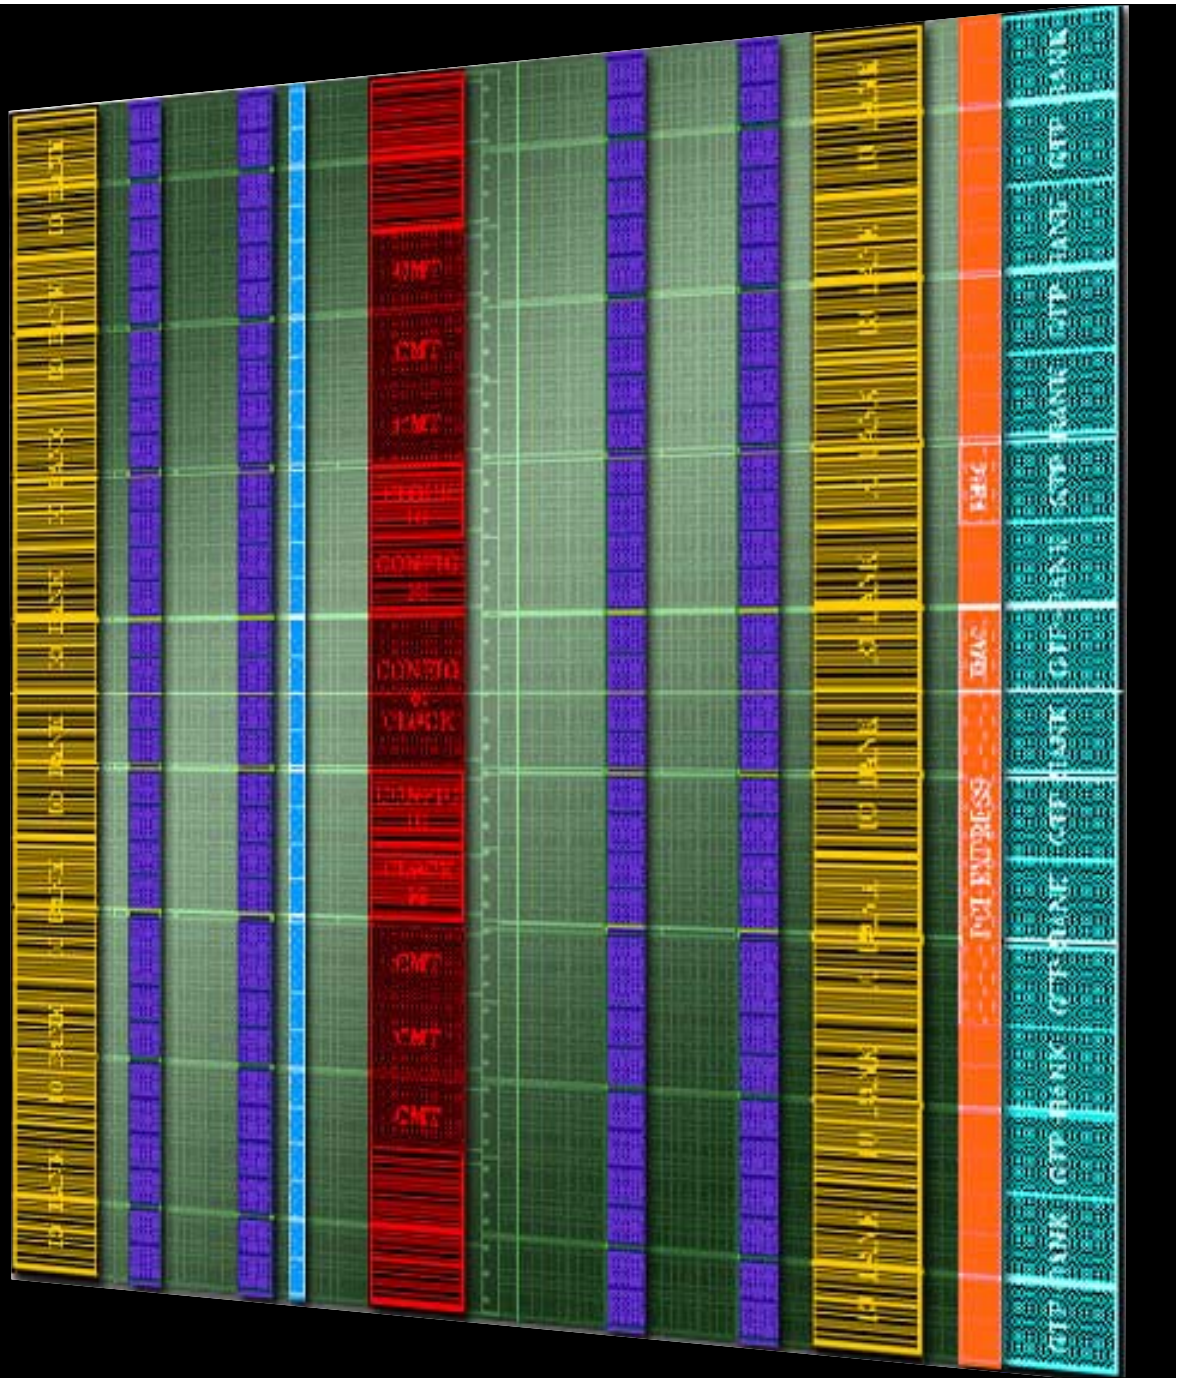

Routing fabric requires many interconnect layers.

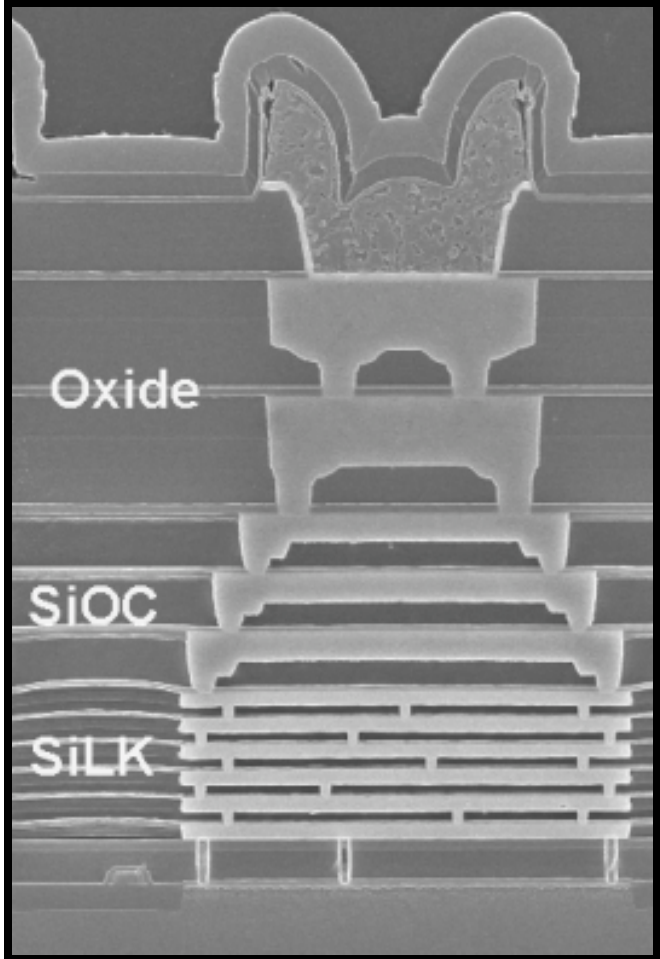

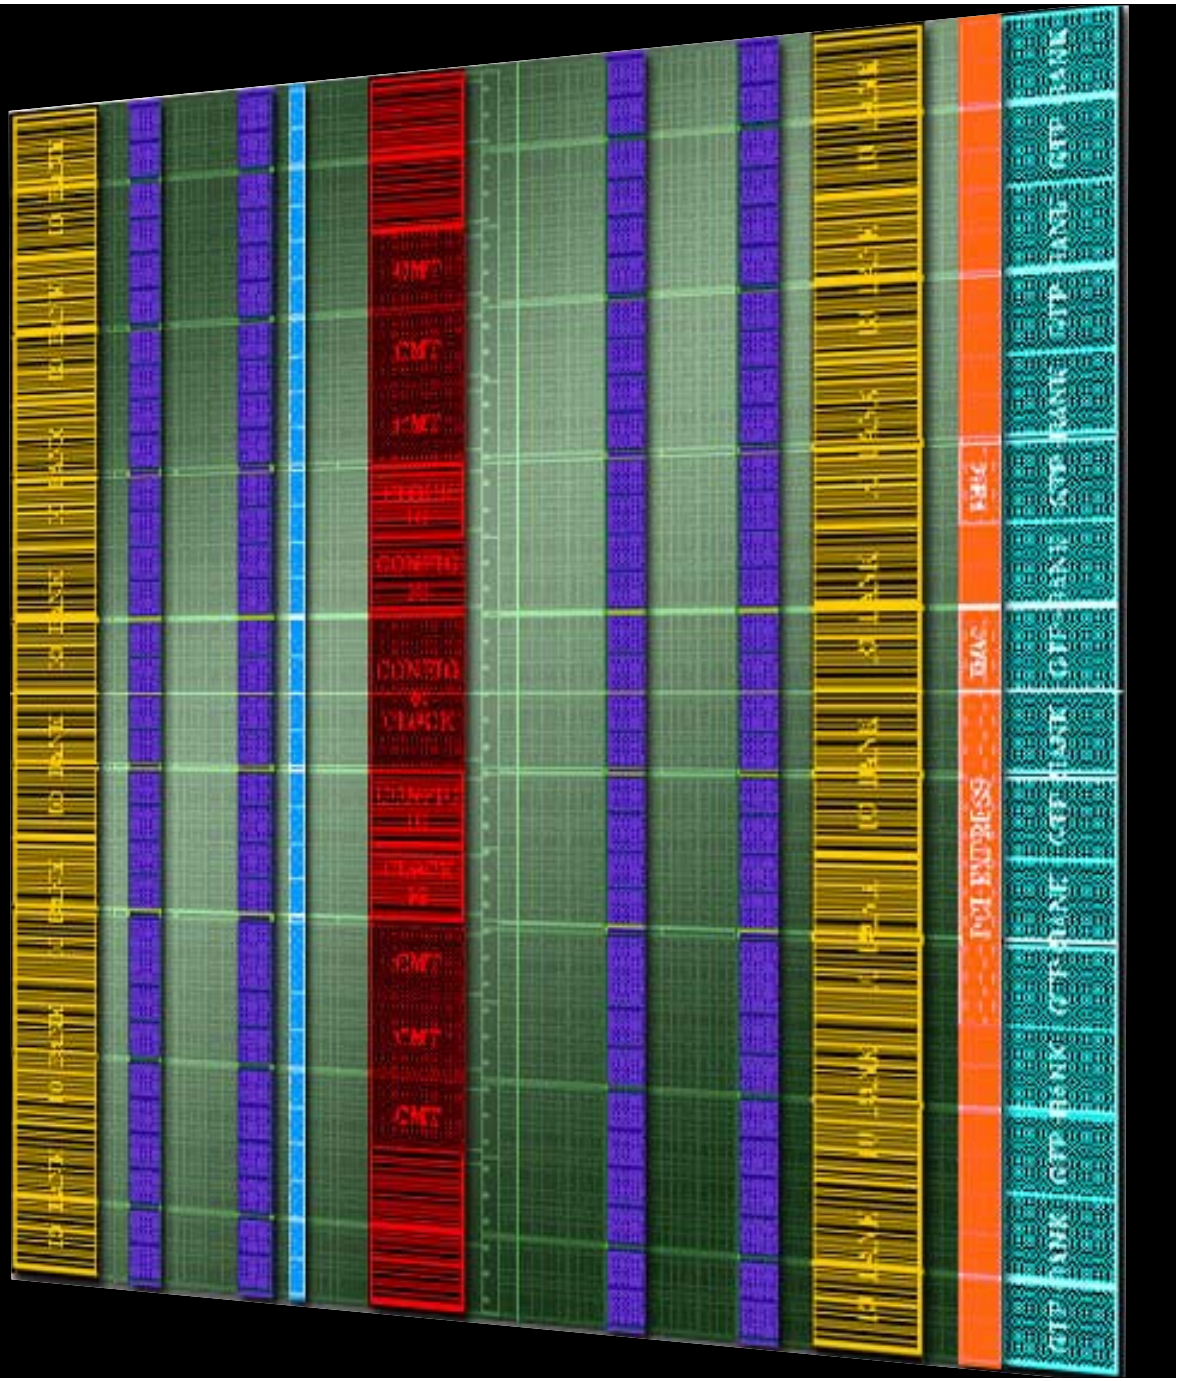

Power Issue #1: Switching Fabric 'C'

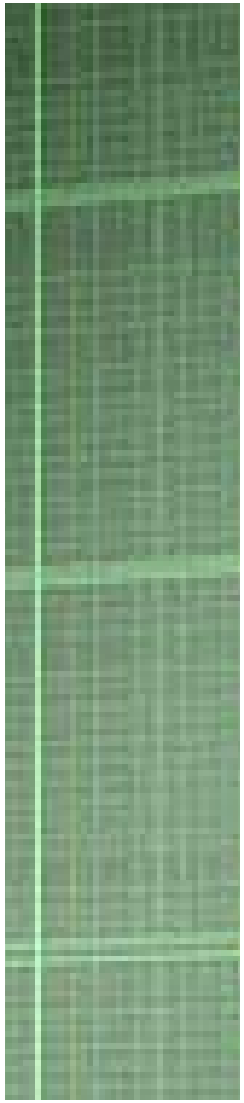

Slices define regular connections to the switching fabric, and to slices in CLBs above and below it on the die.

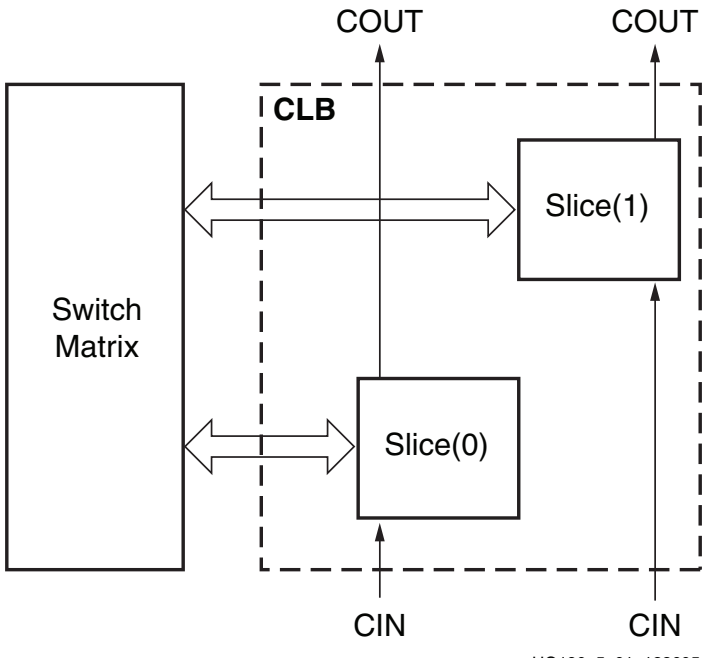

The LX110T has 17,280 slices.

How slices appear on the die ...

X0, X2, ... are lower CLB slices.

X1, X3, ... are upper CLB slices.

Y0, Y1, ... are CLB column positions.

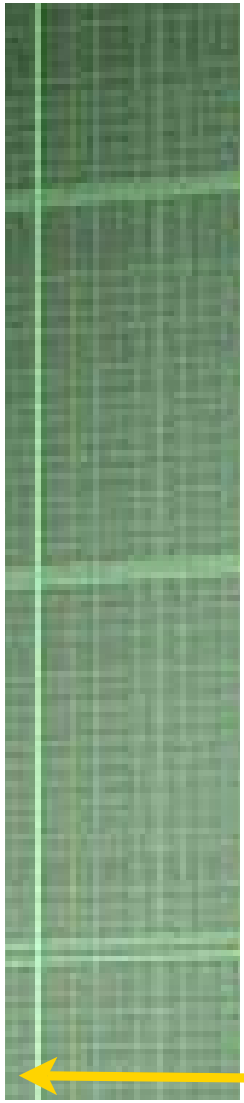

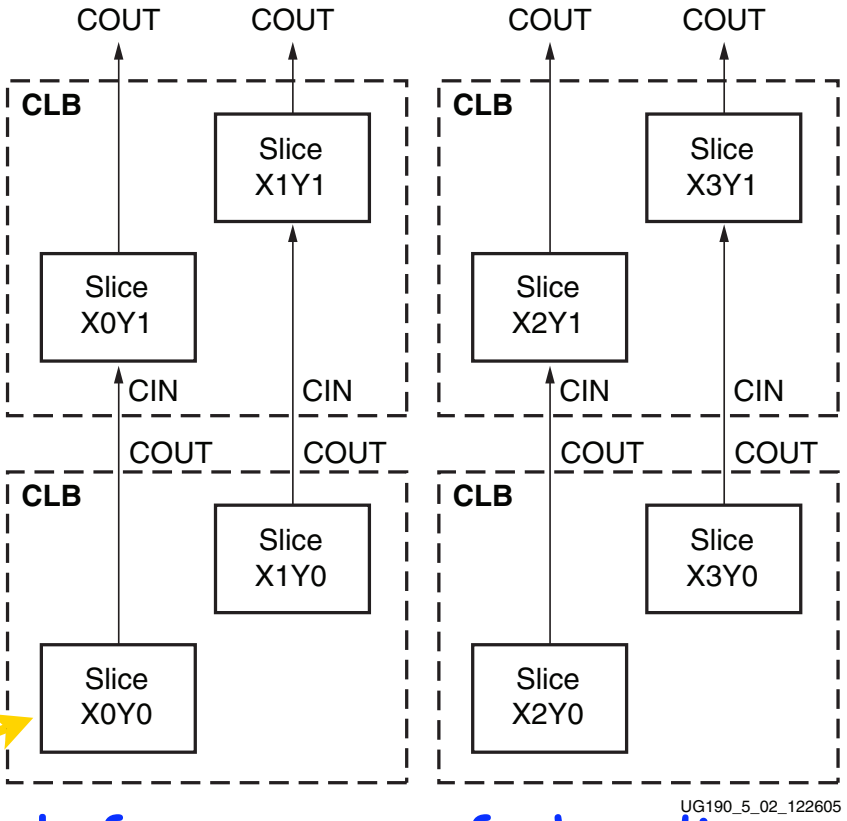

Lower-left corner of the die.

Simple model of FPGA interconnect ...

Why 'C' is so big: (1) each green dot is a transistor switch
(2) path may not be shortest length (3) all wires are too long!

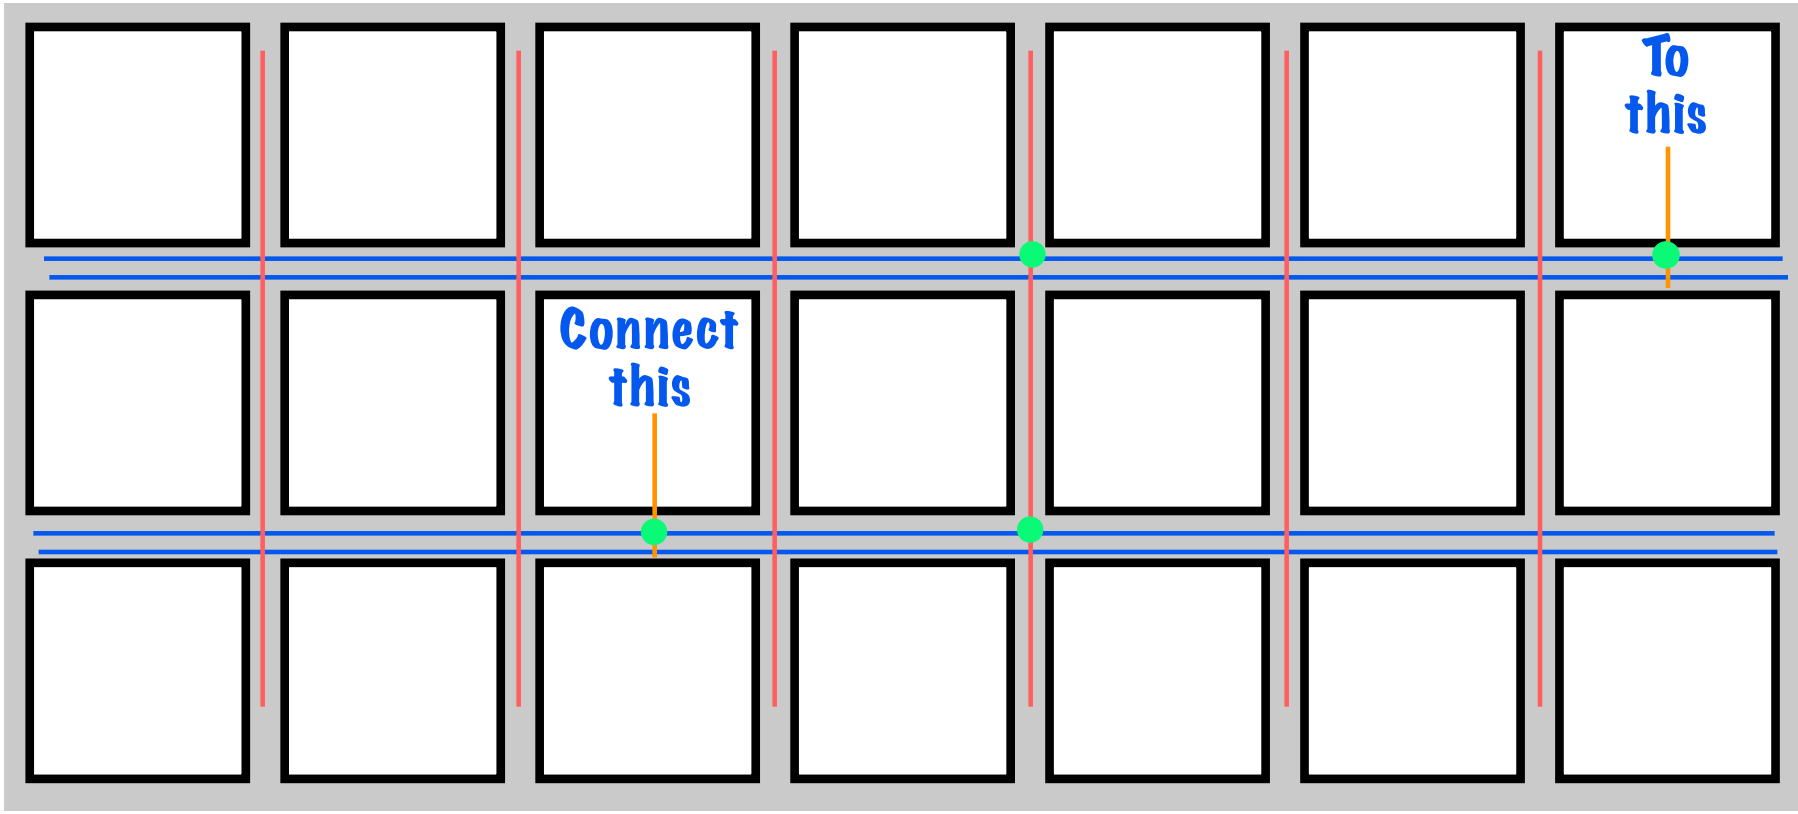

In a non-FPGA (ASIC) chip, this wire may have 10 times less capacitance, and thus use 10 times less power on each flip !

Power Issue #2: Unused slice logic.

SLICEL

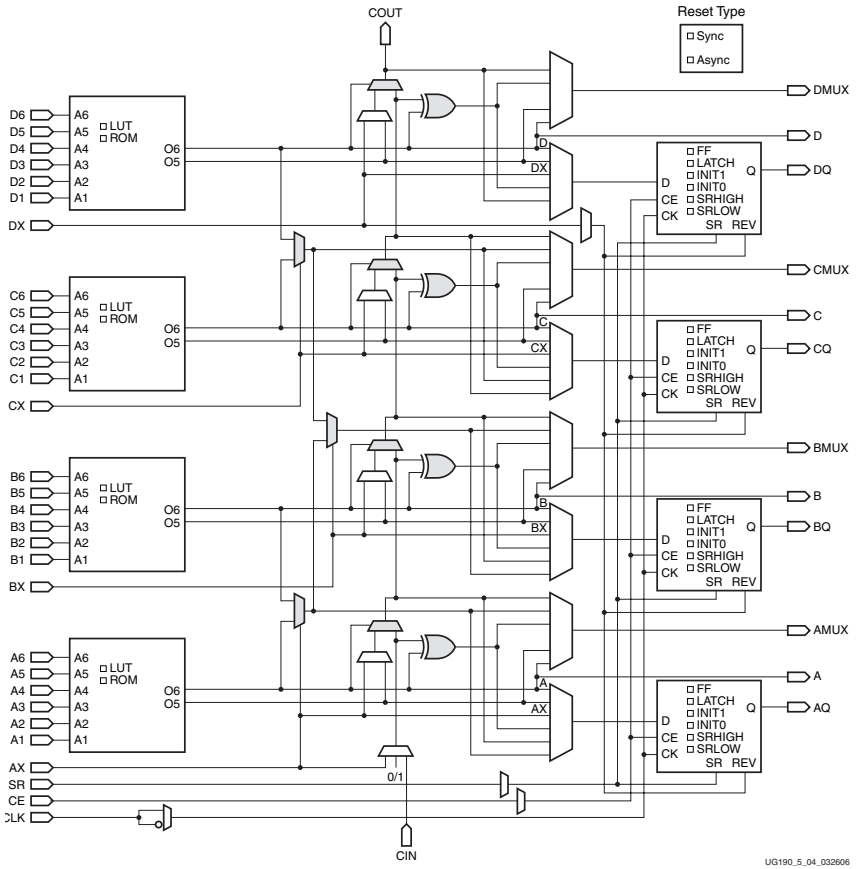

SLICEM

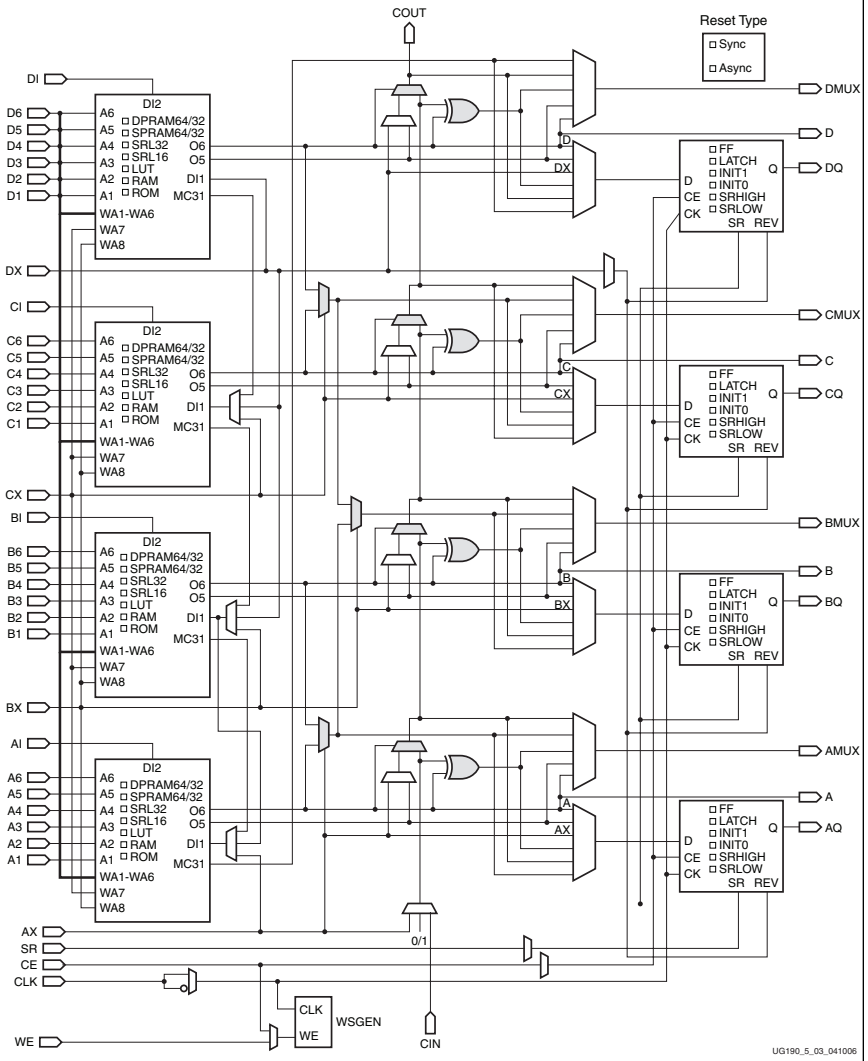

Unused logic that cannot be turned off wastes static power.

Power issue #3:
Clock power.

32 global clock wires go down the red column.

Any 10 may be sent to a clock region.

Much of the clock wire is a clock wire to nowhere.
Big power waste.

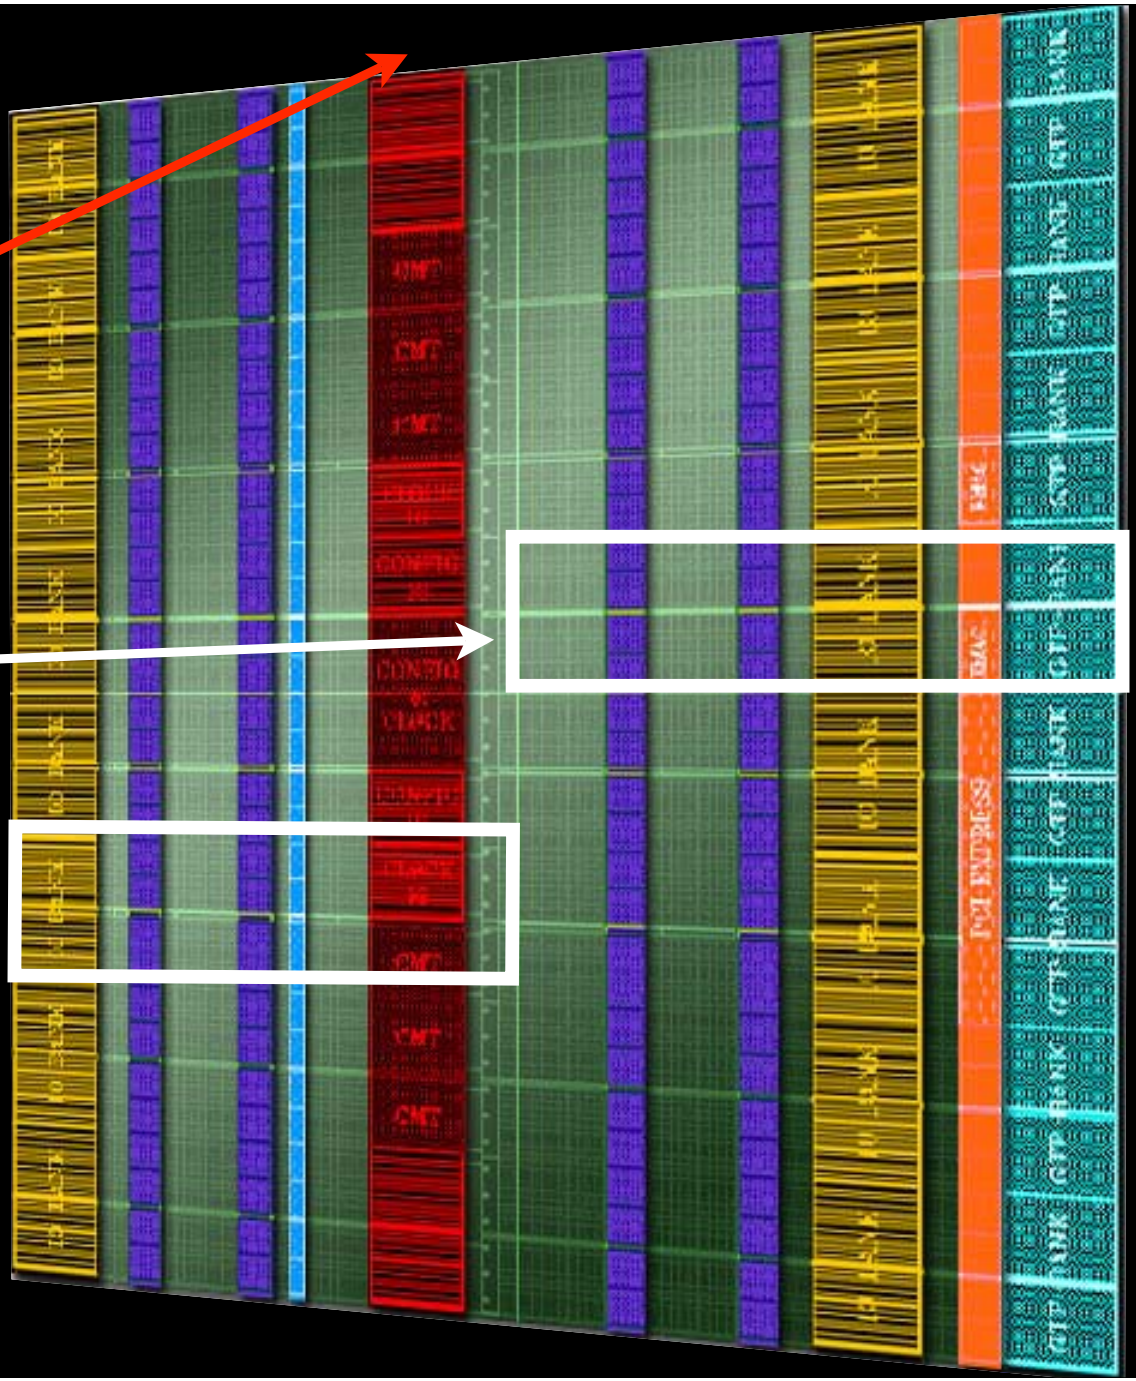

LX110T:

12
Digital Clock
Managers (DCM)

6 Phase Locked
Loops (PLL)

A non-FPGA
(ASIC) chip
can address the
problems DCMs
and PLLs solve
using less power.

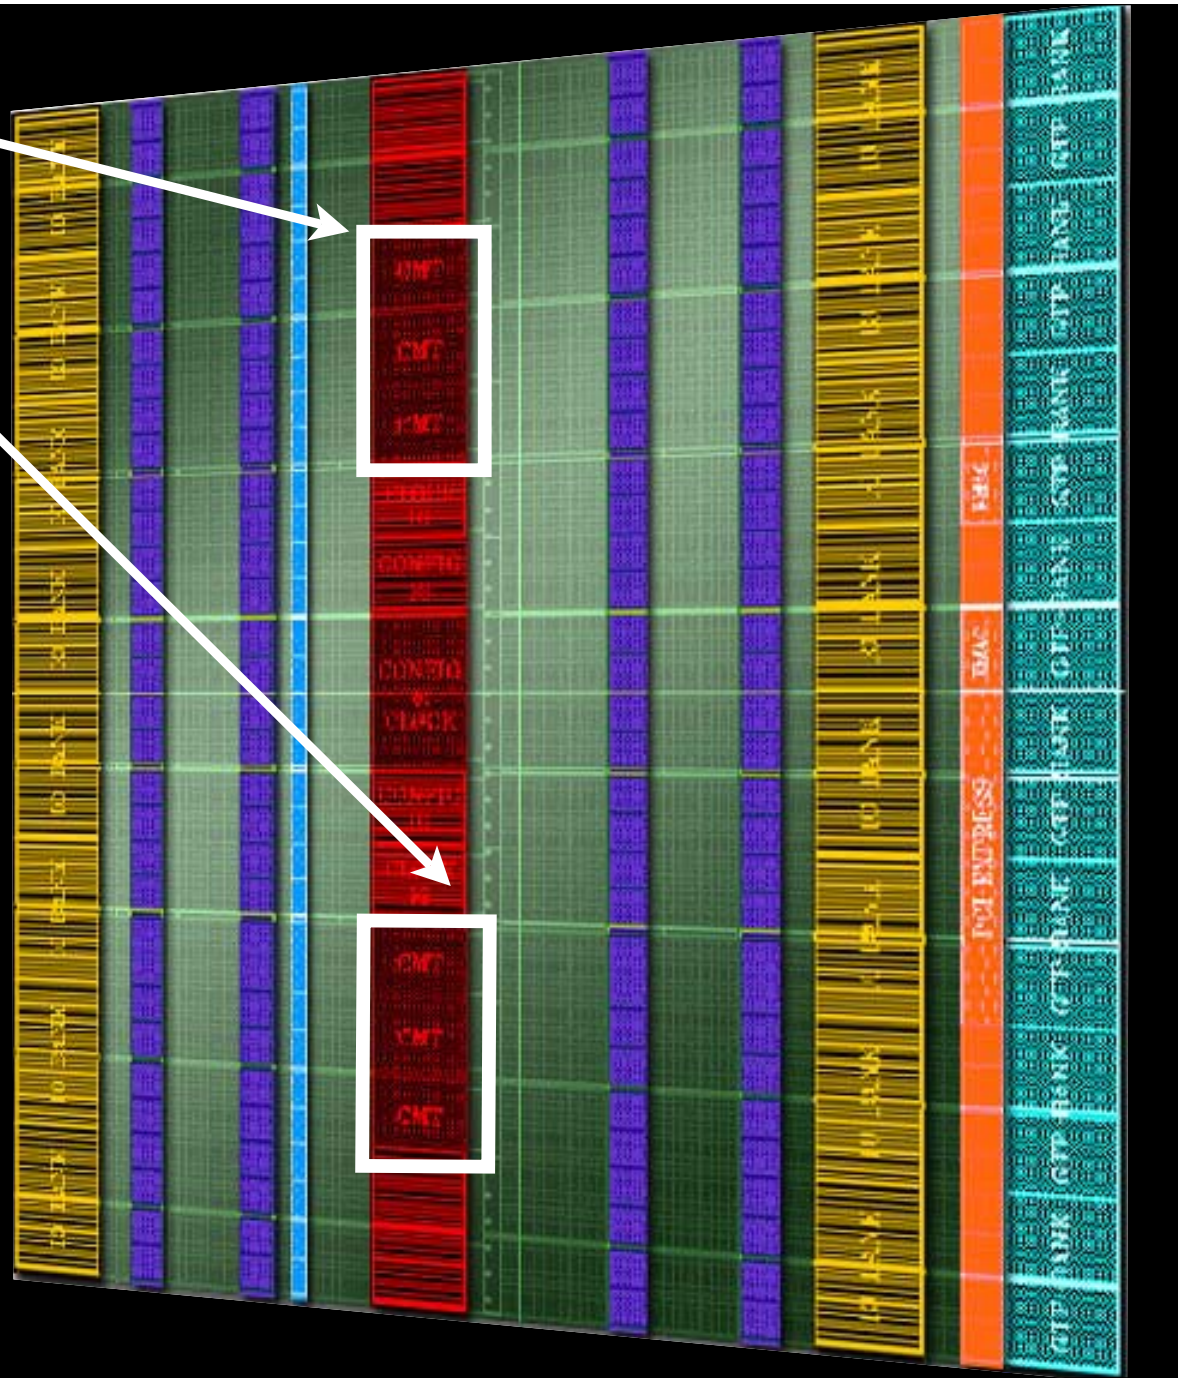

Power Issue #4: CAD speed tradeoffs ...

An example design from a Xilinx "white paper" ...

CAD set for high speed.

1 x 32.5k x 36-bit	Area	Fmax MHz	Total Dynamic Power (mW)	Dynamic RAM Power (mW)
Speed Implementation	70 Slices 73 RAMB16	139	135.6	99.6
XST 9.2i-power	751 Slices 65 RAMB16	123	48	2.4

CAD set for low power.

A small drop in clock speed, but a big power win. But, design size explodes, and the tools run slow.

The final stretch …

Homework due now

	Thu 4/28	Lec #28: Course Wrap-up		now
16	Tue 5/3	**RRR Week**, No Lecture		Cleanup & Optimizations, Early Checkoff
	Thu 5/5	**RRR Week**, No Lecture		
17	Mon 5/9	**Final Exam is Monday 5/9, 11:30-2:30**		**Final Project Checkoff**

Have a great summer!