

CIS 371

Computer Organization and Design

Unit 13: Power & Energy

Slides developed by Milo Martin & Amir Roth at the University of Pennsylvania
with sources that included University of Wisconsin slides
by Mark Hill, Guri Sohi, Jim Smith, and David Wood


Power/Energy Are Increasingly Important

- **Battery life** for mobile devices
 - Laptops, phones, cameras
- **Tolerable temperature** for devices without active cooling
 - Power means temperature, active cooling means **cost**
 - No room for a fan in a cell phone, no market for a hot cell phone
- **Electric bill** for compute/data centers
 - Pay for power twice: once in, once out (to cool)
- **Environmental concerns**
 - “Computers” account for growing fraction of energy consumption

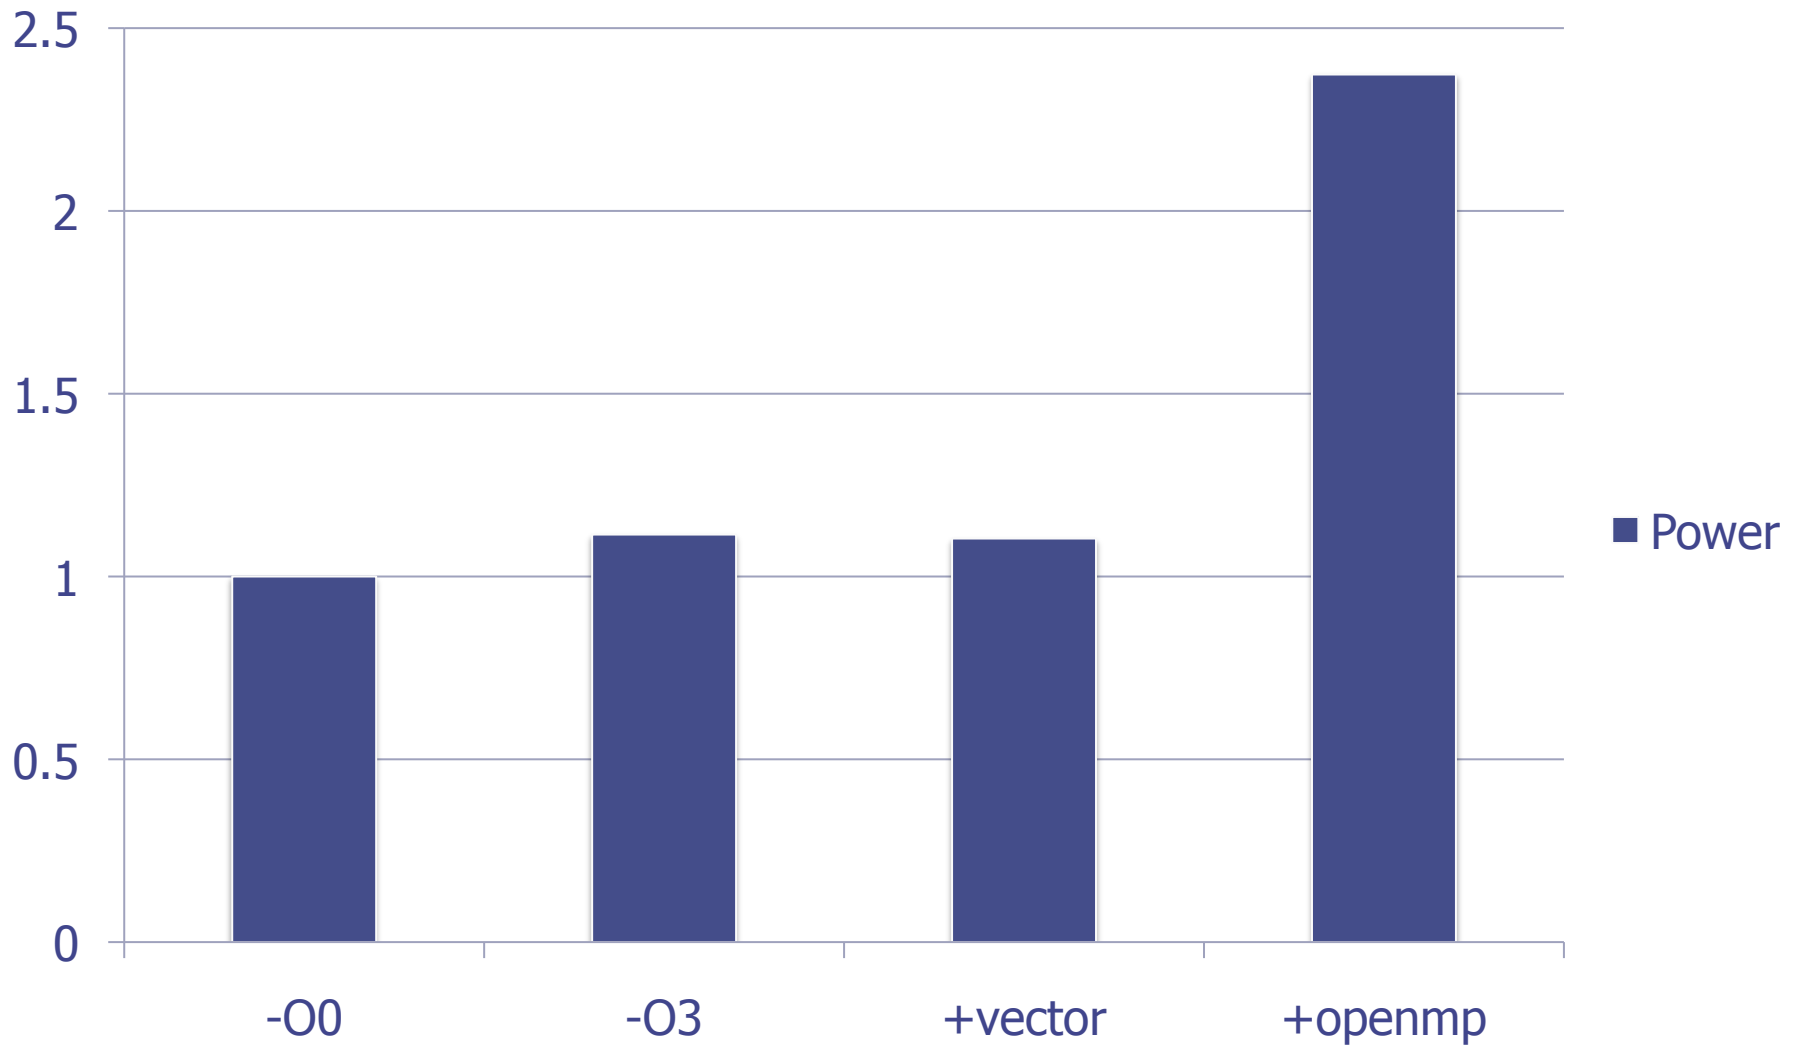
Energy & Power

- **Energy**: measured in Joules or Watt-seconds
 - Total amount of energy stored/used
 - Battery life, electric bill, environmental impact
 - Instructions per Joule (car analogy: miles per gallon)
- **Power**: energy per unit time (measured in Watts)
 - Related to “performance” (which is also a “per unit time” metric)
 - Power impacts power supply and cooling requirements (cost)
 - Power-density (Watt/mm²): important related metric
 - Peak power vs average power
 - E.g., camera, power “spikes” when you actually take a picture
 - Joules per second (car analogy: gallons per hour)
- Two sources:
 - **Dynamic power**: active switching of transistors
 - **Static power**: leakage of transistors even while inactive



Energy Data from Homework 1 (SAXPY)



Power Data from Homework 1 (SAXPY)

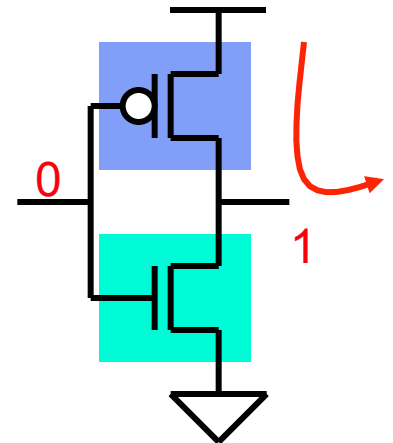


Technology Basis of Transistor Speed

- Physics 101: delay through an electrical component $\sim RC$
 - **Resistance (R)**  \sim length / cross-section area
 - Slows rate of charge flow
 - **Capacitance (C)**  \sim length * area / distance-to-other-plate
 - Stores charge
 - **Voltage (V)**
 - Electrical pressure
 - **Threshold Voltage (V_t)**
 - Voltage at which a transistor turns "on"
 - Property of transistor based on fabrication technology
 - **Switching time \sim to $(R * C) / (V - V_t)$**
- Components contribute to capacitance & resistance
 - Transistors
 - Wires (longer the wire, the more the capacitance & resistance)

Dynamic Power

- **Dynamic power (P_{dynamic}):** aka switching or active power
 - Energy to switch a gate (0 to 1, 1 to 0)
 - Each gate has capacitance (C)
 - Charge stored is $\propto C * V$
 - Energy to charge/discharge a capacitor is $\propto C * V^2$
 - Time to charge/discharge a capacitor is $\propto V$
 - Result: frequency \sim to V
 - **$P_{\text{dynamic}} \approx N * C * V^2 * f * A$**
 - N: number of transistors
 - C: capacitance per transistor (size of transistors)
 - V: voltage (supply voltage for gate)
 - f: frequency (transistor switching freq. is \propto to clock freq.)
 - A: activity factor (not all transistors may switch this cycle)




Reducing Dynamic Power

- Target each component: $P_{\text{dynamic}} \approx N * C * V^2 * f * A$
- **Reduce number of transistors (N)**
 - Use fewer transistors/gates
- **Reduce capacitance (C)**
 - Smaller transistors (Moore's law)
- **Reduce voltage (V)**
 - Quadratic reduction in energy consumption!
 - But also slows transistors (transistor speed is \sim to V)
- **Reduce frequency (f)**
 - Slower clock frequency (reduces power but not energy) Why?
- **Reduce activity (A)**
 - "Clock gating" disable clocks to unused parts of chip
 - Don't switch gates unnecessarily

Static Power

- **Static power (P_{static}):** aka idle or leakage power
 - Transistors don't turn off all the way
 - Transistors "leak"
 - Analogy: leaky valve
 - $P_{\text{static}} \approx N * V * e^{-V_t}$
 - N: number of transistors
 - V: voltage
 - **V_t (threshold voltage):** voltage at which transistor conducts (begins to switch)
- Switching speed vs leakage trade-off
- The lower the **V_t** :
 - Faster transistors (linear)
 - Transistor speed \propto to $V - V_t$
 - Leakier transistors (exponential)



Reducing Static Power

- Target each component: $P_{\text{static}} \approx N * V * e^{-Vt}$
- **Reduce number of transistors** (N)
 - Use fewer transistors/gates
- **Disable transistors** (also targets N)
 - “Power gating” disable power to unused parts (long latency to power up)
 - Power down units (or entire cores) not being used
- **Reduce voltage** (V)
 - Linear reduction in static energy consumption
 - But also slows transistors (transistor speed is \sim to V)
- **Dual V_t** – use a mixture of high and low V_t transistors
 - Use slow, low-leak transistors in SRAM arrays
 - Requires extra fabrication steps (cost)
- **Low-leakage transistors**
 - High-K/Metal-Gates in Intel’s 45nm process, “tri-gate” in Intel’s 22nm
- Reducing frequency can hurt energy efficiency due to leakage power

Continuation of Moore's Law

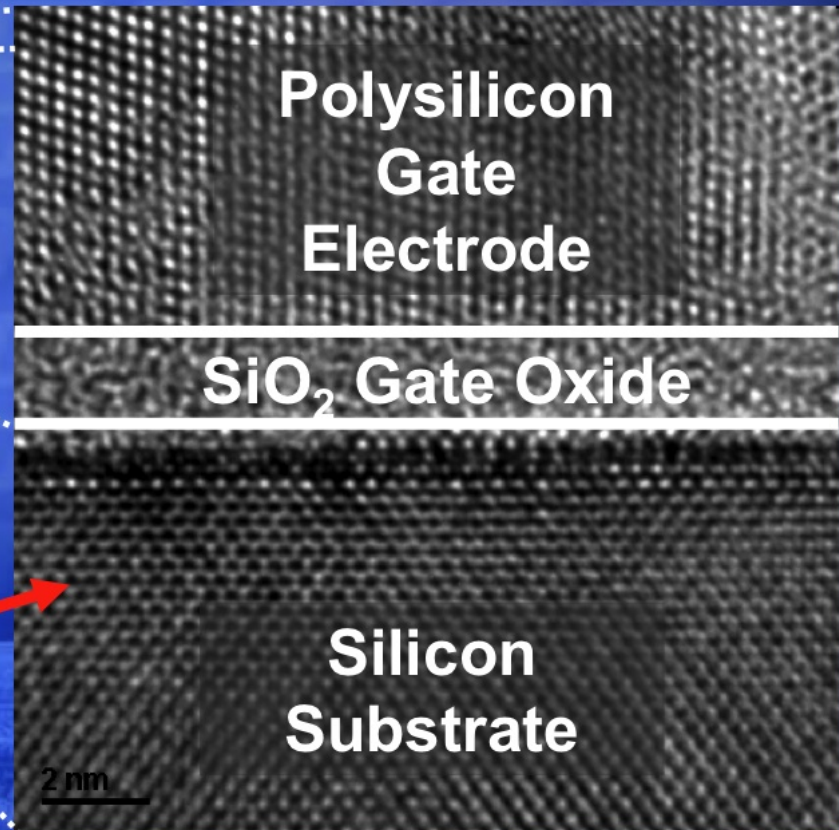
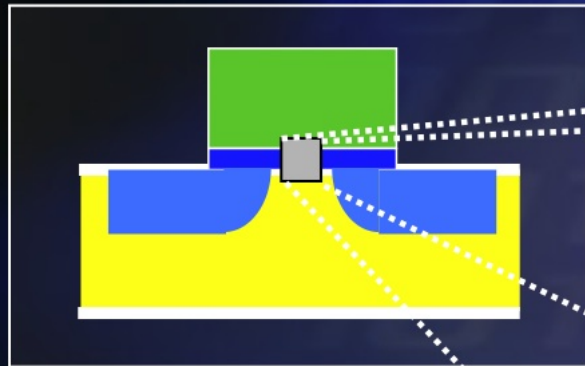
Process Name	P856	P858	Px60	P1262	P1264	P1266	P1268	P1270
1st Production	1997	1999	2001	2003	2005	2007	2009	2011
Process Generation	0.25 μ m	0.18 μ m	0.13 μ m	90 nm	65 nm	45 nm	32 nm	22 nm
Wafer Size (mm)	200	200	200/300	300	300	300	300	300
Inter-connect	Al	Al	Cu	Cu	Cu	Cu	Cu	?
Channel	Si	Si	Si	Strained Si	Strained Si	Strained Si	Strained Si	Strained Si
Gate dielectric	SiO ₂	SiO ₂	SiO ₂	SiO ₂	SiO ₂	High-k	High-k	High-k
Gate electrode	Poly-silicon	Poly-silicon	Poly-silicon	Poly-silicon	Poly-silicon	Metal	Metal	Metal

Introduction targeted at this time

Subject to change


Intel found a solution for High-k and metal gate

Gate dielectric today is only a few molecular layers thick



Individual Atoms

High-k Dielectric reduces leakage substantially



Benefits compared to current process technologies

	High-k vs. SiO ₂	Benefit
Capacitance	60% greater	<i>Much faster transistors</i>
Gate dielectric leakage	> 100x reduction	<i>Far cooler</i>

Dynamic Voltage/Frequency Scaling

- **Dynamically trade-off power for performance**
 - Change the voltage and frequency at runtime
 - Under control of operating system
- Recall: $P_{\text{dynamic}} \approx N * C * V^2 * f * A$
 - Because frequency \propto to $V - V_t$...
 - $P_{\text{dynamic}} \propto$ to $V^2(V - V_t) \approx V^3$
- Reduce both voltage and frequency linearly
 - **Cubic decrease in dynamic power**
 - Linear decrease in performance (actually sub-linear)
 - Thus, only about quadratic in energy
 - Linear decrease in static power
 - Thus, static energy can become dominant
- Newer chips can adjust frequency on a per-core basis

Dynamic Voltage/Frequency Scaling

	Mobile PentiumIII "SpeedStep"	Transmeta 5400 "LongRun"	Intel X-Scale (StrongARM2)
f (MHz)	300–1000 (step=50)	200–700 (step=33)	50–800 (step=50)
V (V)	0.9–1.7 (step=0.1)	1.1–1.6V (cont)	0.7–1.65 (cont)
High-speed	3400MIPS @ 34W	1600MIPS @ 2W	800MIPS @ 0.9W
Low-power	1100MIPS @ 4.5W	300MIPS @ 0.25W	62MIPS @ 0.01W

- Dynamic voltage/frequency scaling
 - **Favors parallelism**
- Example: Intel Xscale
 - 1 GHz → 200 MHz reduces energy used by 30x
 - But around 5x slower
 - 5 x 200 MHz in parallel, use **1/6th the energy**
 - Power is driving the trend toward multi-core

Moore's Effect on Power

- + Moore's Law reduces power/transistor...
 - Reduced sizes and surface areas reduce capacitance (C)
- ...but increases power density and total power
 - By increasing transistors/area and total transistors
 - Faster transistors → higher frequency → more power
 - Hotter transistors leak more (thermal runaway)
- What to do? Reduce voltage (V)
 - + Reduces dynamic power quadratically, static power linearly
 - Already happening: Intel 486 (5V) → Core2 (1.3V)
 - Trade-off: reducing V means either...
 - Keeping V_t the same and reducing frequency (f)
 - Lowering V_t and increasing leakage exponentially
 - Use techniques like high-K and dual- V_T
- The end of voltage scaling & “dark silicon”


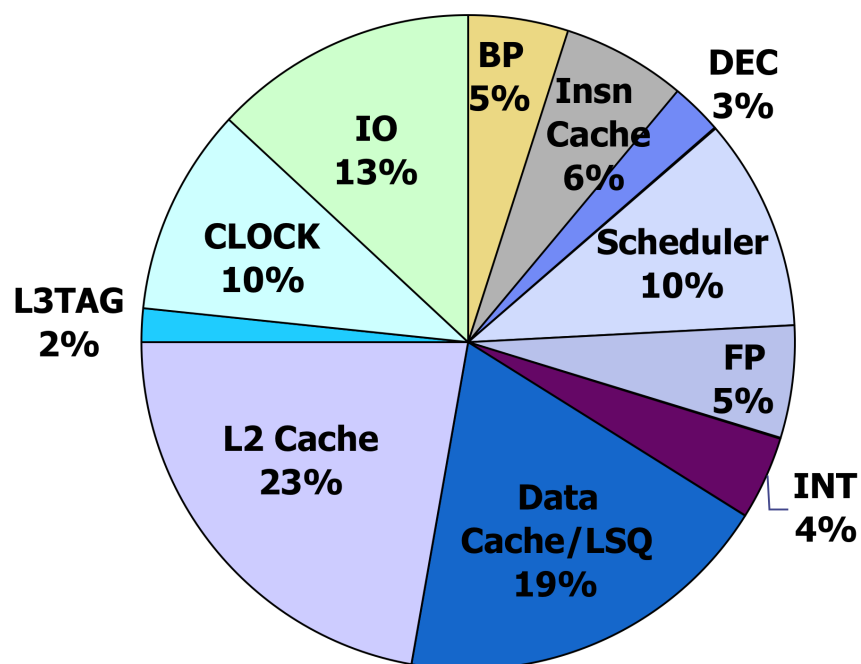
Trends in Power

	386	486	Pentium	Pentium II	Pentium4	Core2	Core i7
Year	1985	1989	1993	1998	2001	2006	2009
Technode (nm)	1500	800	350	180	130	65	45
Transistors (M)	0.3	1.2	3.1	5.5	42	291	731
Voltage (V)	5	5	3.3	2.9	1.7	1.3	1.2
Clock (MHz)	16	25	66	200	1500	3000	3300
Power (W)	1	5	16	35	80	75	130
Peak MIPS	6	25	132	600	4500	24000	52800
MIPS/W	6	5	8	17	56	320	406

- Supply voltage decreasing over time
 - But “voltage scaling” is perhaps reaching its limits
- Emphasis on power starting around 2000
 - Resulting in slower frequency increases
 - Also note number of cores increasing (2 in Core 2, 4 in Core i7)

Processor Power Breakdown

- Power breakdown for IBM POWER4
 - Two 4-way superscalar, 2-way multi-threaded cores, 1.5MB L2
 - Big power components are L2, data cache, scheduler, clock, I/O
 - Implications on “complicated” versus “simple” cores




Implications on Software

- Software-controlled dynamic voltage/frequency scaling
 - Example: video decoding
 - Too high a clock frequency – wasted energy (battery life)
 - Too low a clock frequency – quality of video suffers
 - “Race to sleep” versus “slow and steady” approaches
- Managing low-power modes
 - Don’t want to “wake up” the processor every millisecond
- Tuning software
 - Faster algorithms can be converted to lower-power algorithms
 - Via dynamic voltage/frequency scaling
- Exploiting parallelism & heterogeneous cores
 - NVIDIA Tegra 3: 5 cores (4 “normal” cores & 1 “low power” core)
- Specialized hardware accelerators


Recent Technology Update From Intel

Reduced Channel Doping

Planar

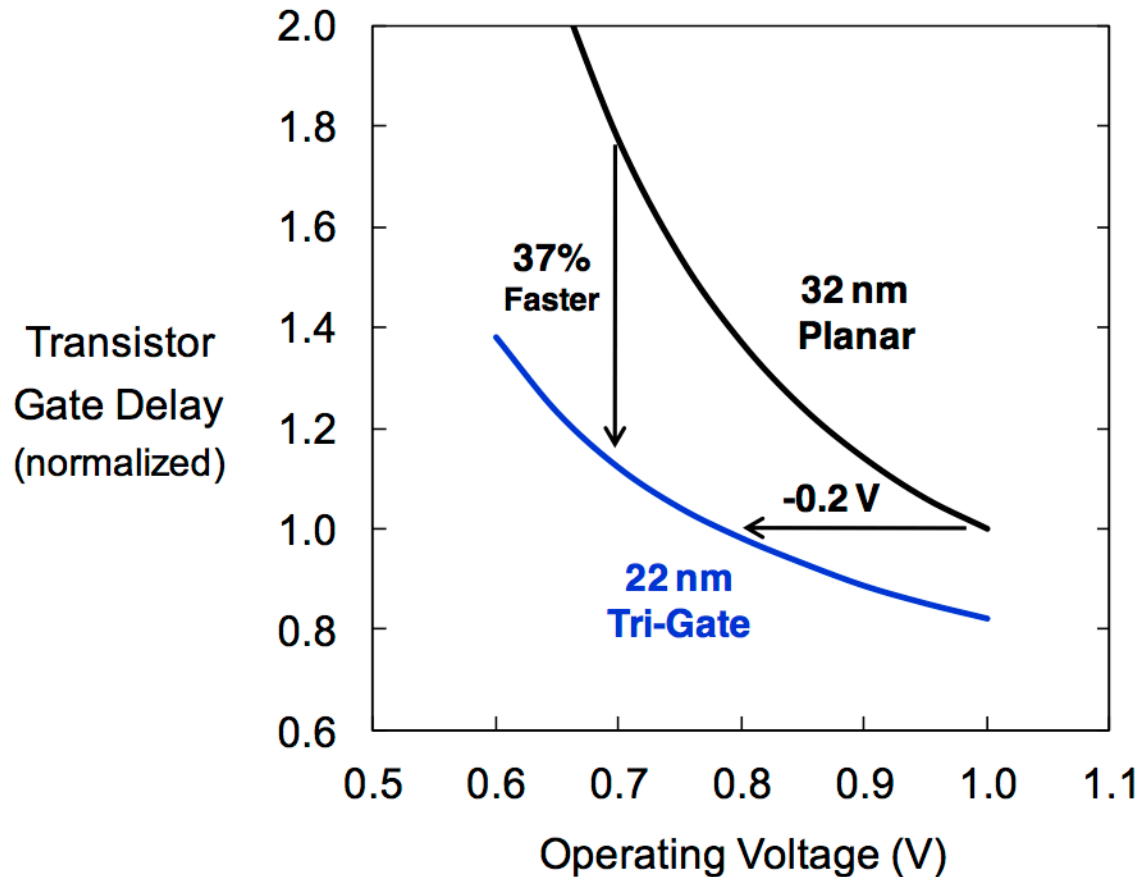


Tri-Gate



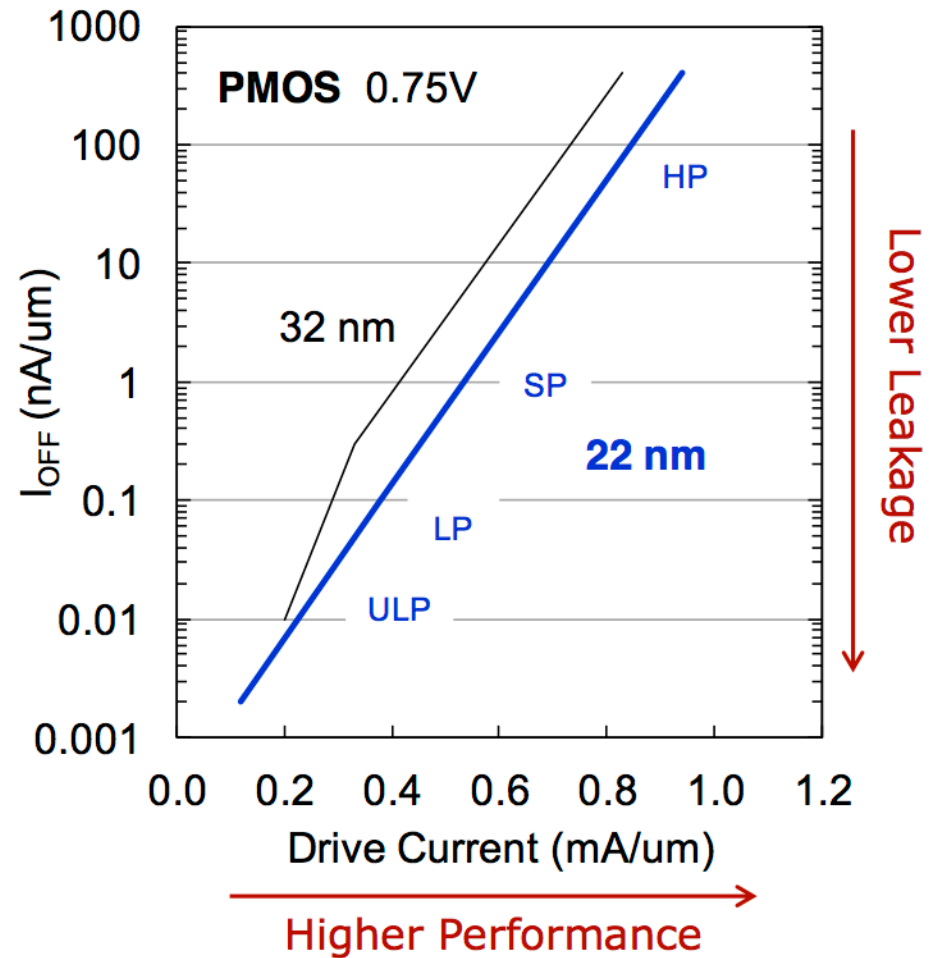
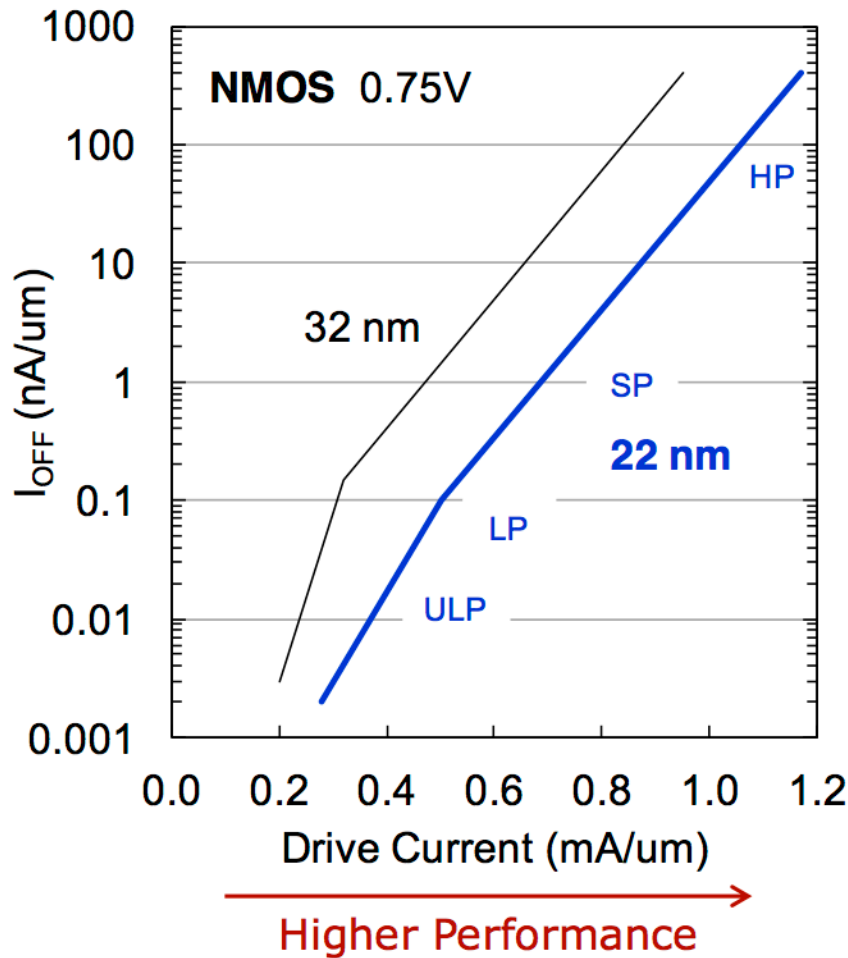
Fully depleted Tri-Gate structure has reduced channel doping, providing improved performance and reduced variability

Performance/Power Benefits



**Tri-Gate provides 37% speed up at low voltage
or 50% active power reduction at same performance**


Transistor Performance vs. Leakage



22 nm SoC technology offers a wide range of transistors

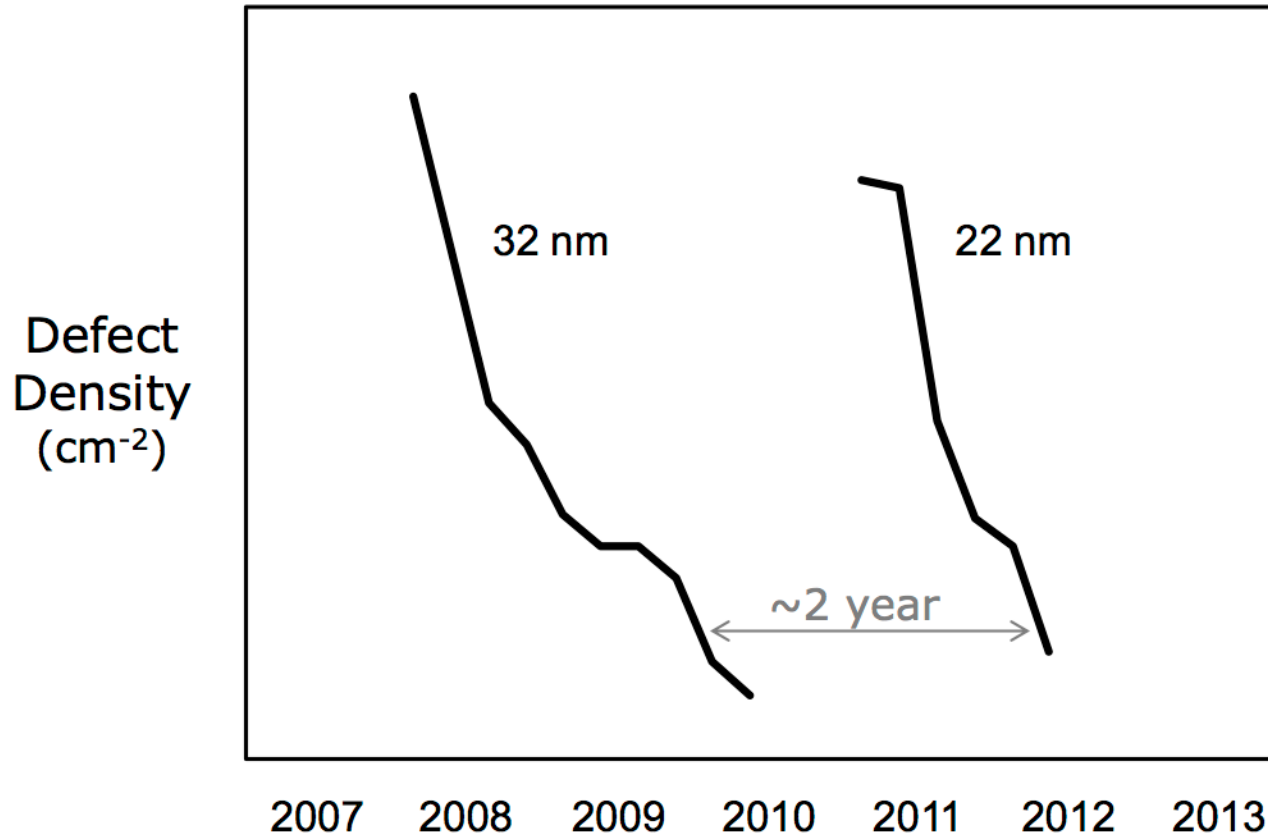
22 nm Interconnects

<u>Layer</u>	<u>Pitch</u>
TM	14 μm
M8	360 nm
M7	320 nm
M6	240 nm
M5	160 nm
M4	112 nm
M3	80 nm
M2	80 nm
M1	90 nm




***Minimum pitch scaled $\sim 0.7x$ from 32 nm
for $\sim 2x$ transistor density improvement***

22 nm Defect Density Trend



22 nm defect density now at low levels needed for volume manufacturing

3RD Generation Intel[®] Core[™] Processor



22 nm Tri-Gate Technology
4 Cores + Integrated Graphics
1.4 Billion Transistors, 160 mm²