Microprocessor Design in the Nanoscale Era

Stefan Rusu

Senior Principal Engineer
Intel Corporation

IEEE Fellow

stefan.rusu@intel.com

Agenda

- Microprocessor Design Trends
- Process Technology Directions
- Active Power Management
- Leakage Reduction Techniques
- Packaging and Thermal Modeling
- Future Directions and Summary

Microprocessor Evolution

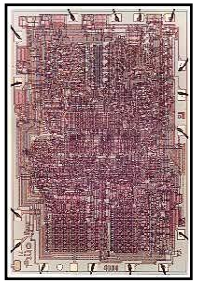

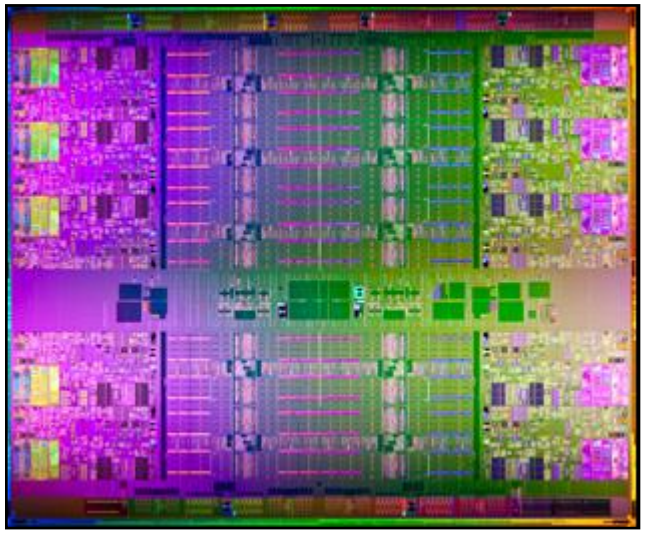

	4004 Processor	Westmere-EX Processor
Year	1971	2011
Transistors	2300	2.6 B
Process	10 µm	32 nm
Die area	12 mm^2	513 mm^2

Die photos not at scale

Scaling Trends

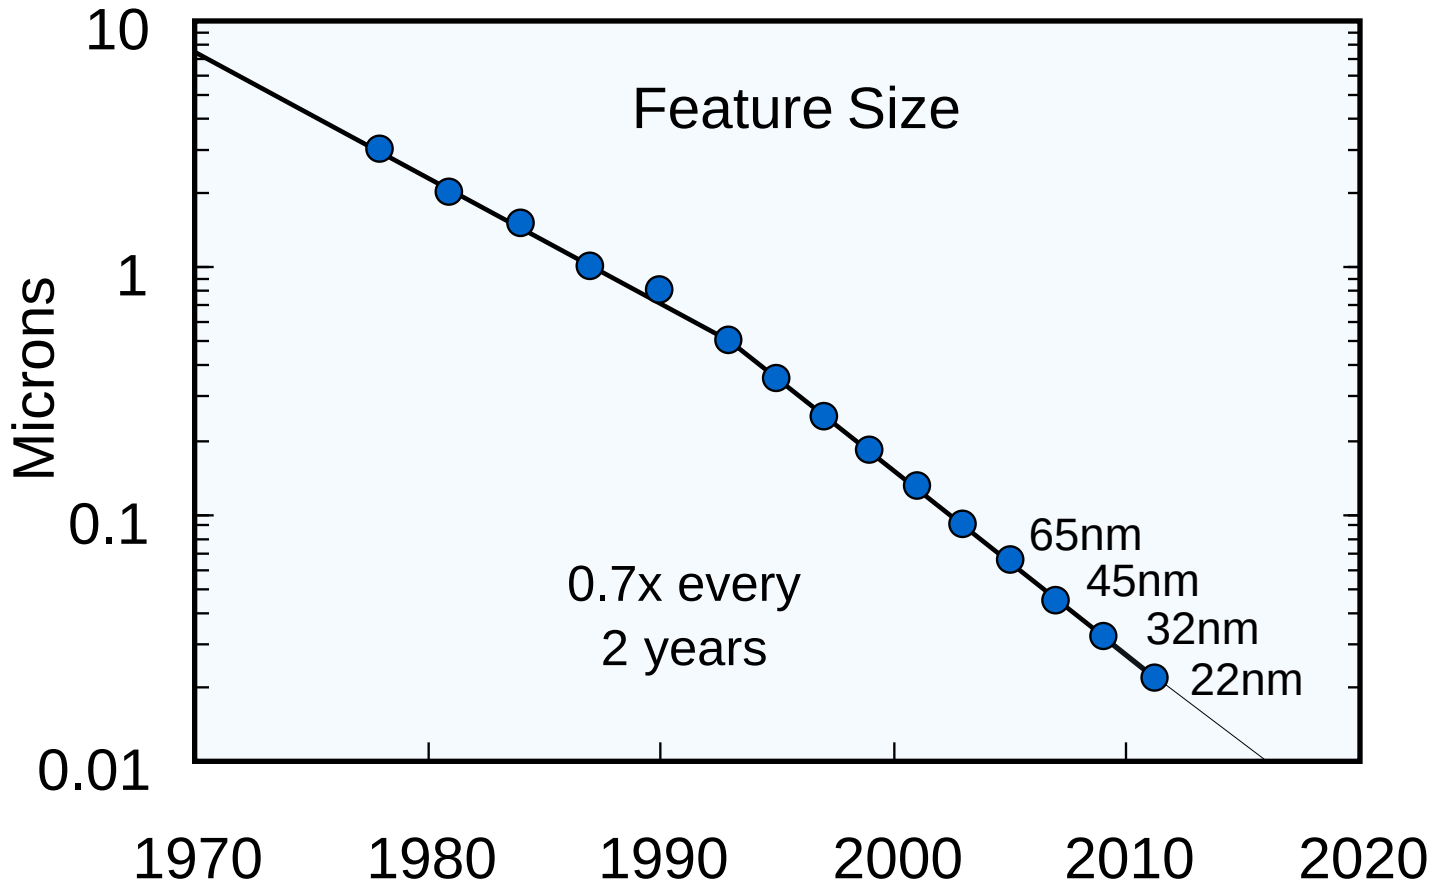

M. Bohr

Client Processor Trend: Integrated Graphics

- Ivy Bridge 22nm client processor with monolithic integrated graphics
- Up to 4 dual-threaded cores and 8MB L3 cache
- Dual channel DDR3 memory controller at 1600MT/s
- Integrated PCIe interface (16 Gen3 + 4 Gen2 + 4 DMI lanes)
  - First Client CPU to support PCIe Gen3
- Three independent displays
- 1.4B transistors in 160mm^2 die

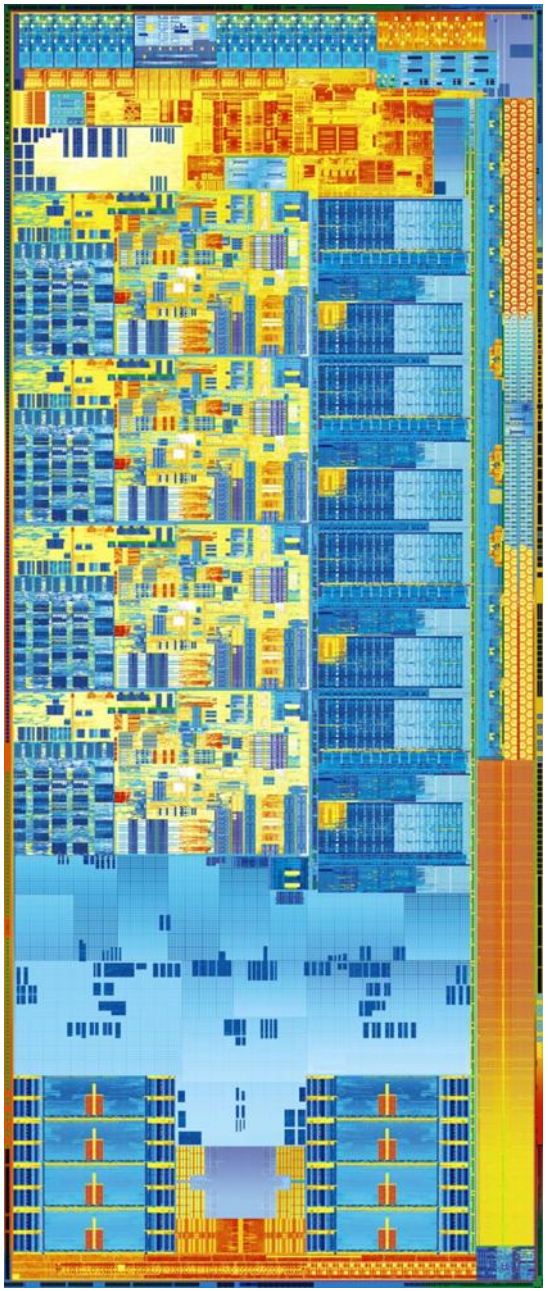

S. Damaraju, ISSCC 2012

Client Processor Trend: Integrated WiFi RF

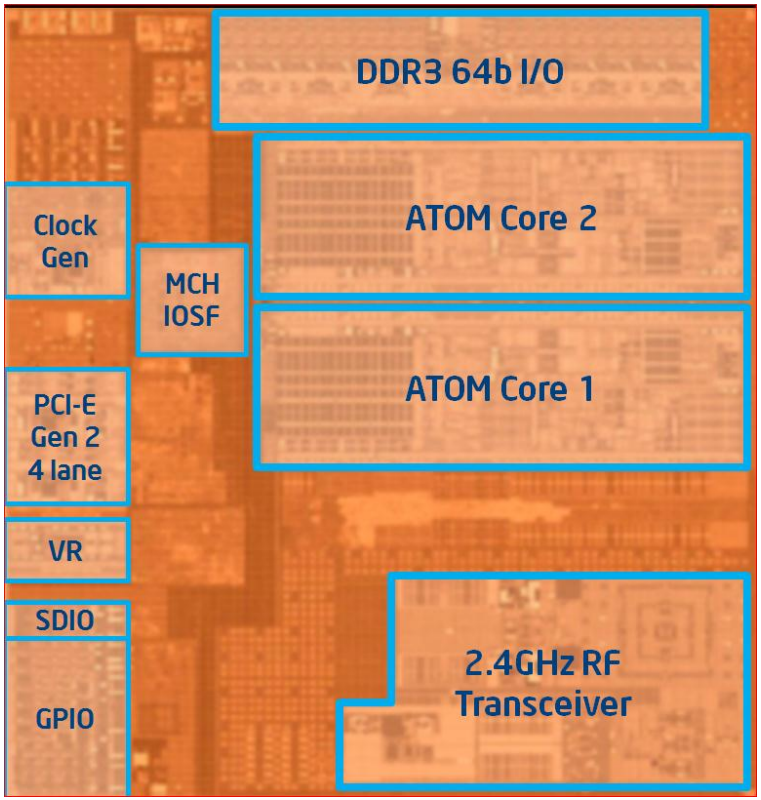

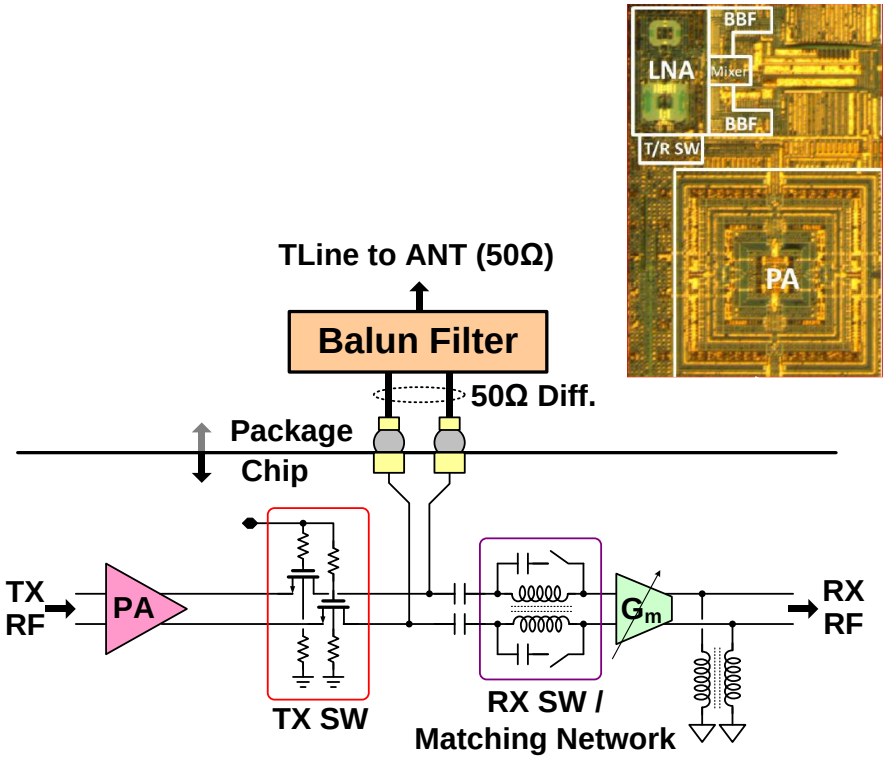

- Sensitive RF circuits integrated with 32nm ATOM and PCH
- Integration of traditional III-V RF components
  - 21dBm Power amp, and 34dBm T/R switch, 3.5dB NF LNA

H. Lakdawala, ISSCC 2012

Server Processor Trends: More Cores

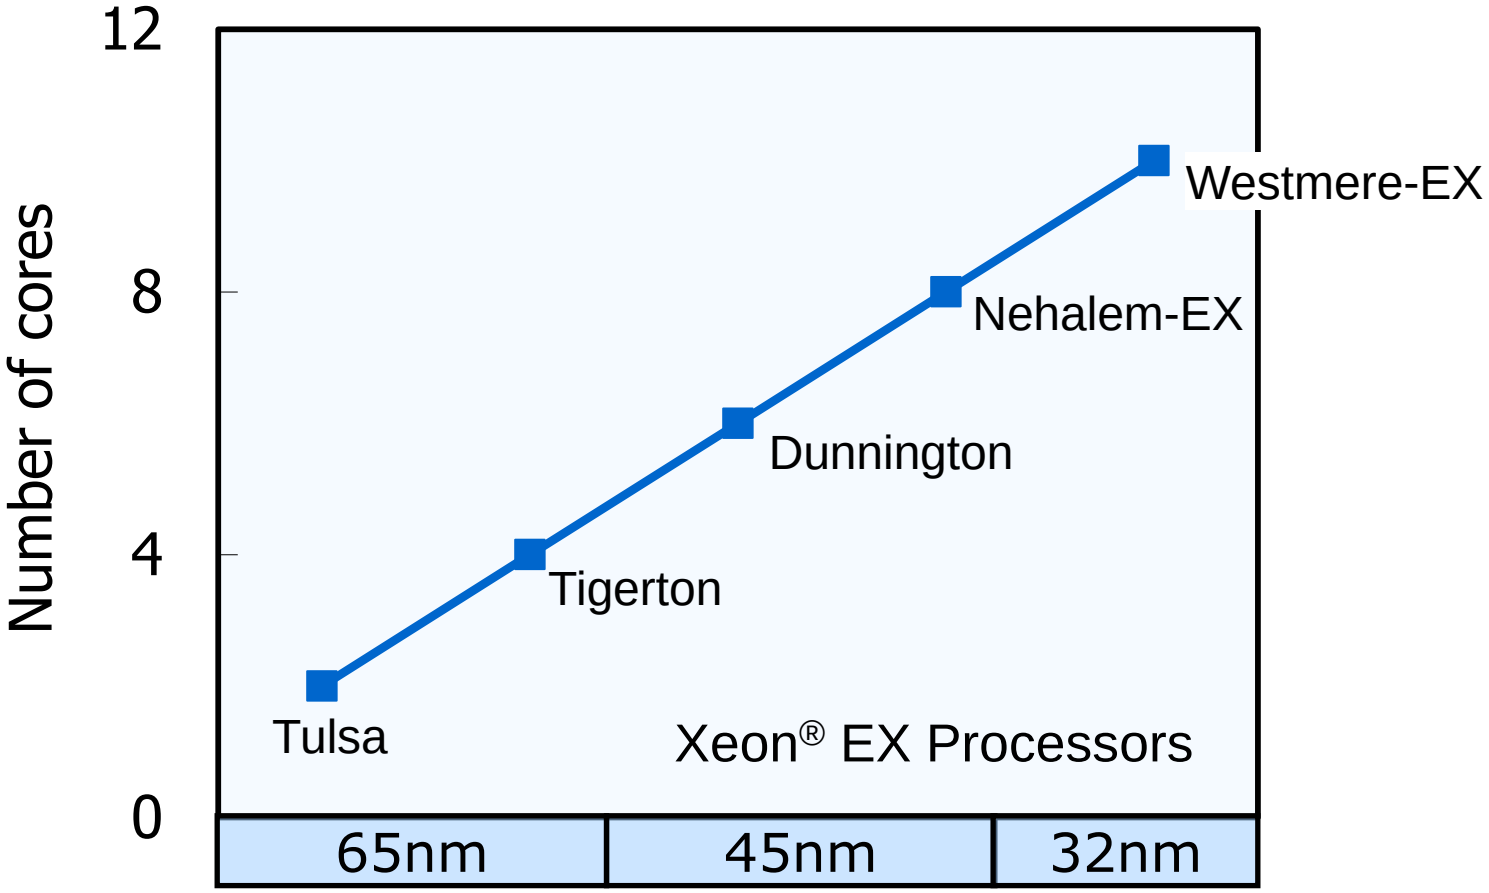

Server core count increases every generation, while keeping within flat power budget

Server Processor Trends: More Cache

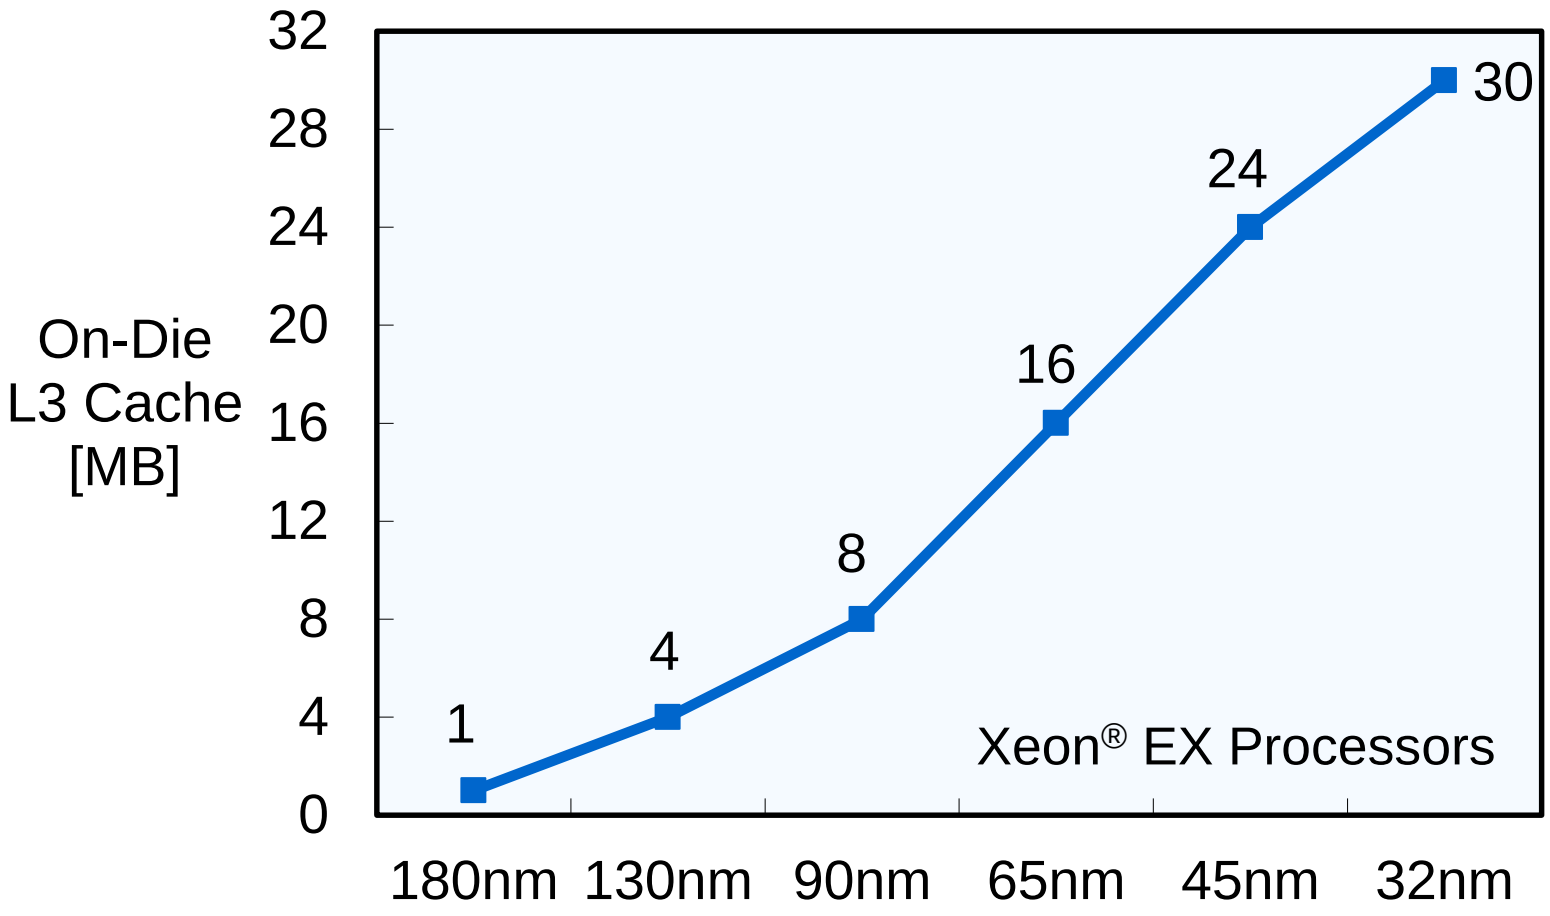

Cache size increases with every process generation

Server Processors Power Trends

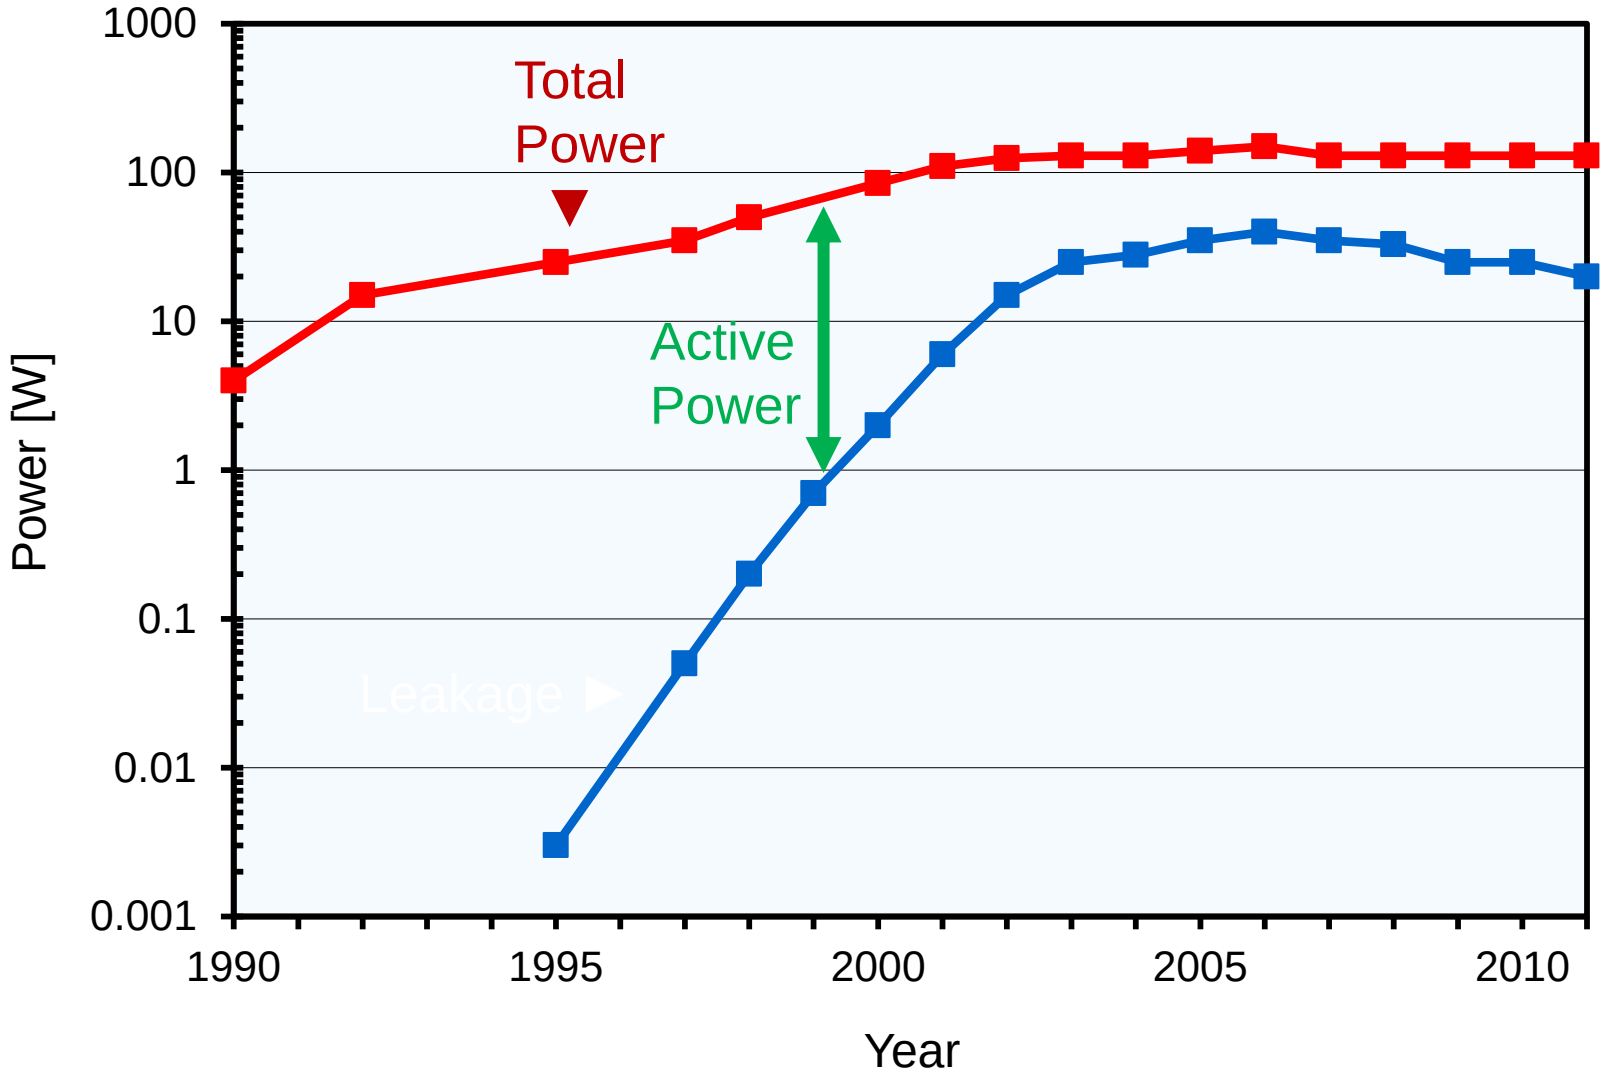

Voltage Scaling Has Slowed Down

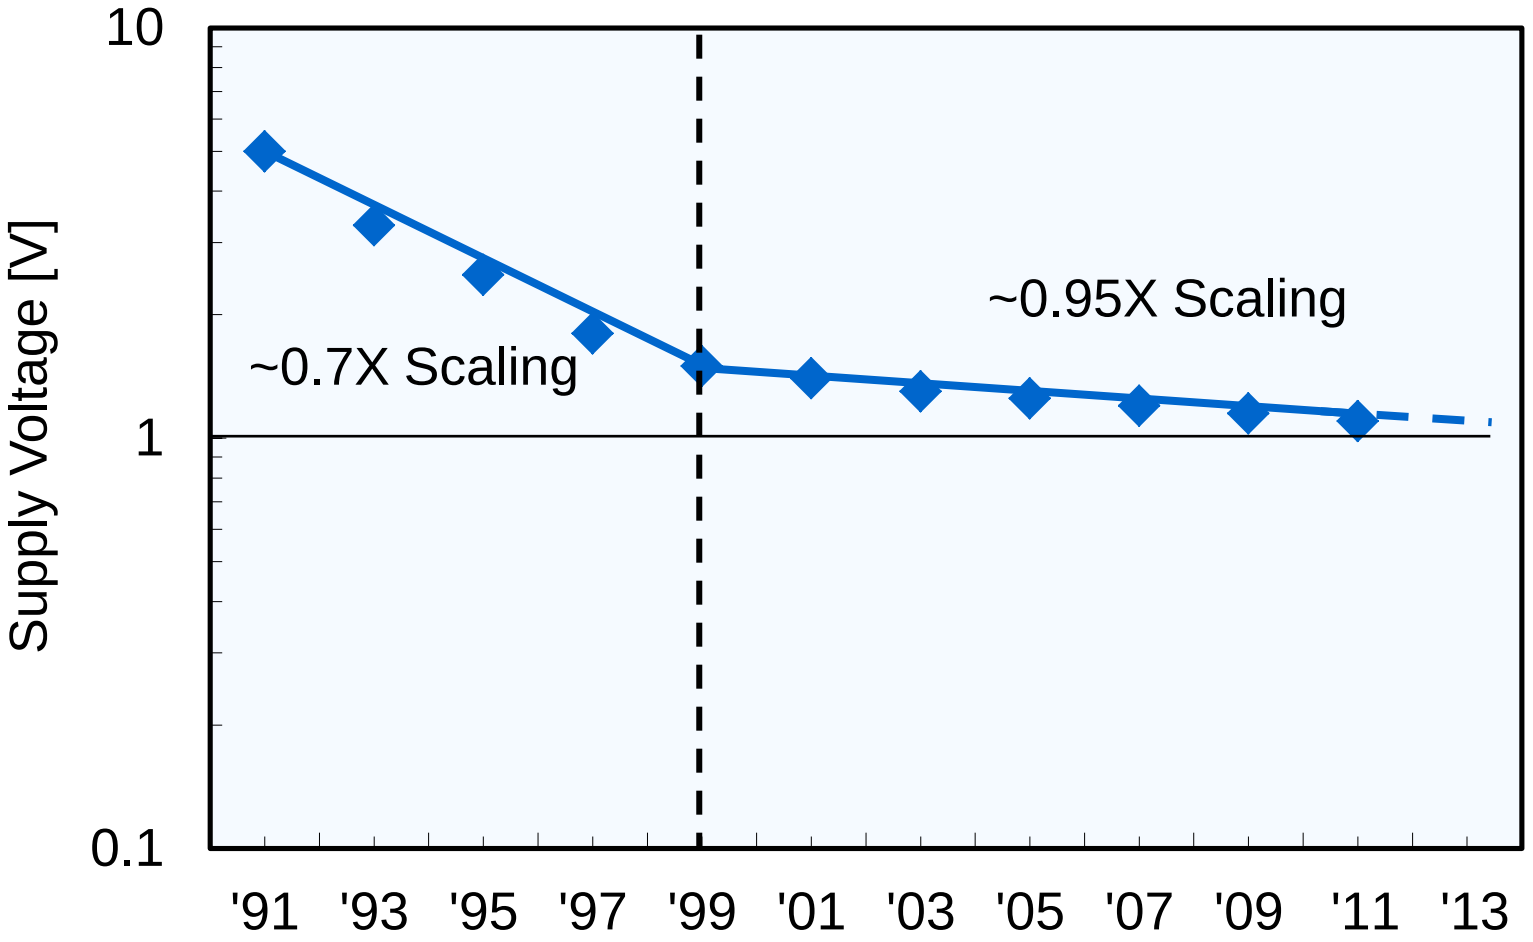

Agenda

- Microprocessor Design Trends
- Process Technology Directions
- Active Power Management
- Leakage Reduction Techniques
- Packaging and Thermal Modeling
- Future Directions and Summary

30 Years of MOSFET Scaling

Dennard 1974

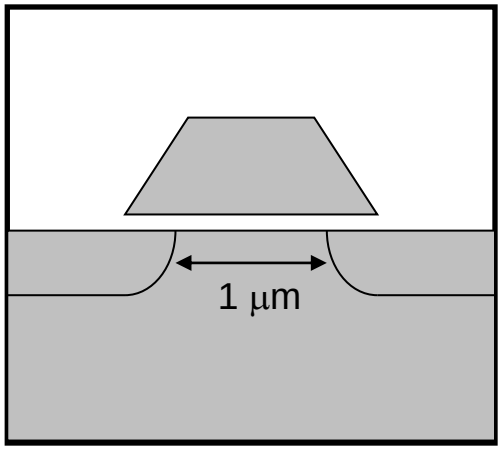

Intel 2005

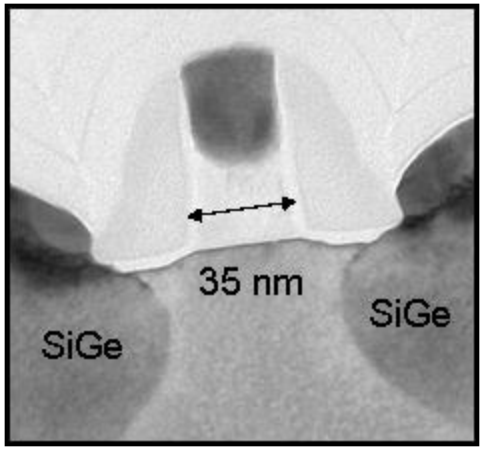

Gate Length:	1.0 μm	35 nm
Gate Oxide Thickness:	35 nm	1.2 nm
Operating Voltage:	4.0 V	1.2 V

M. Bohr, ISSCC 2009

90 nm Strained Silicon Transistors

NMOS

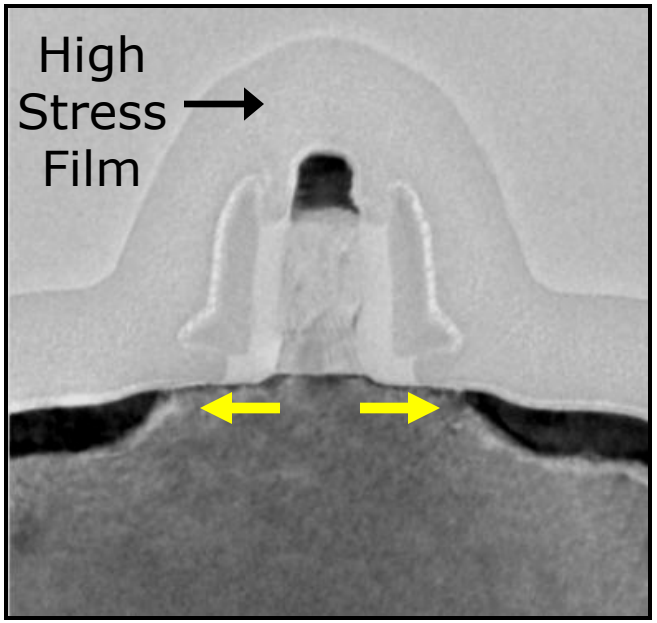

SiN cap layer
Tensile channel strain

PMOS

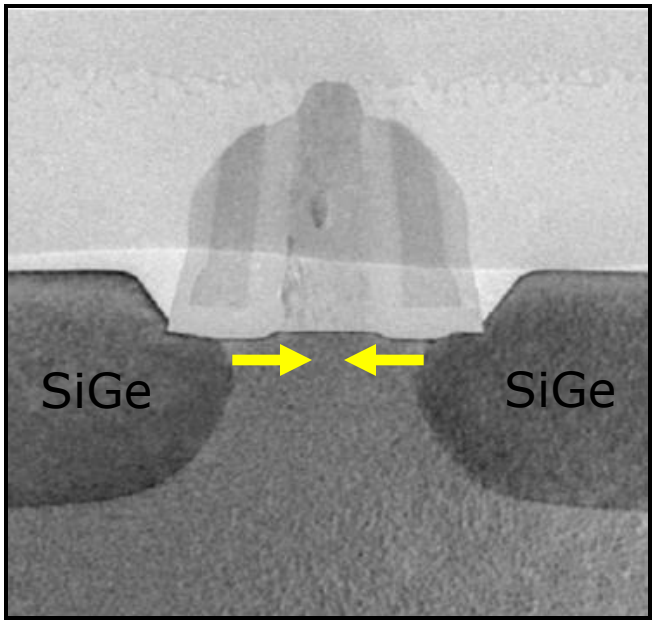

SiGe source-drain
Compressive channel strain

M. Bohr, ISSCC 2009

45 nm High-k + Metal Gate Transistors

65 nm Transistor

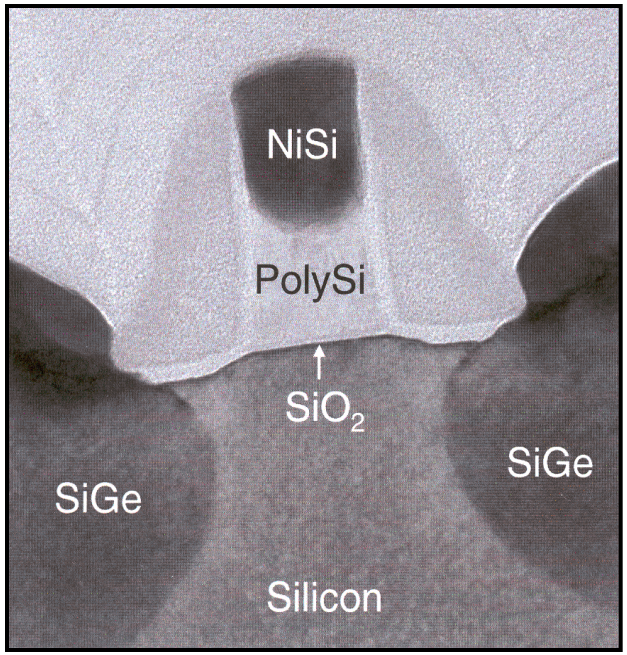

SiO_2 dielectric
Polysilicon gate electrode

45 nm HK+MG

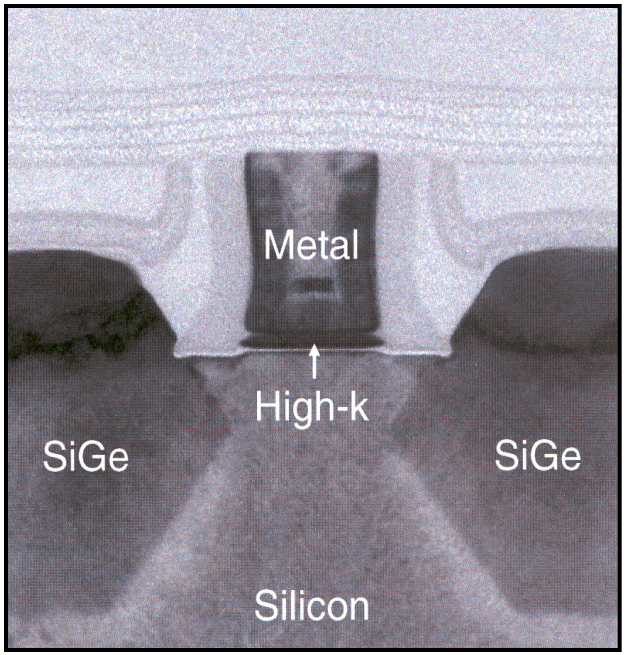

Hafnium-based dielectric
Metal gate electrode

M. Bohr, ISSCC 2009

HK/MG Gate Leakage Reduction

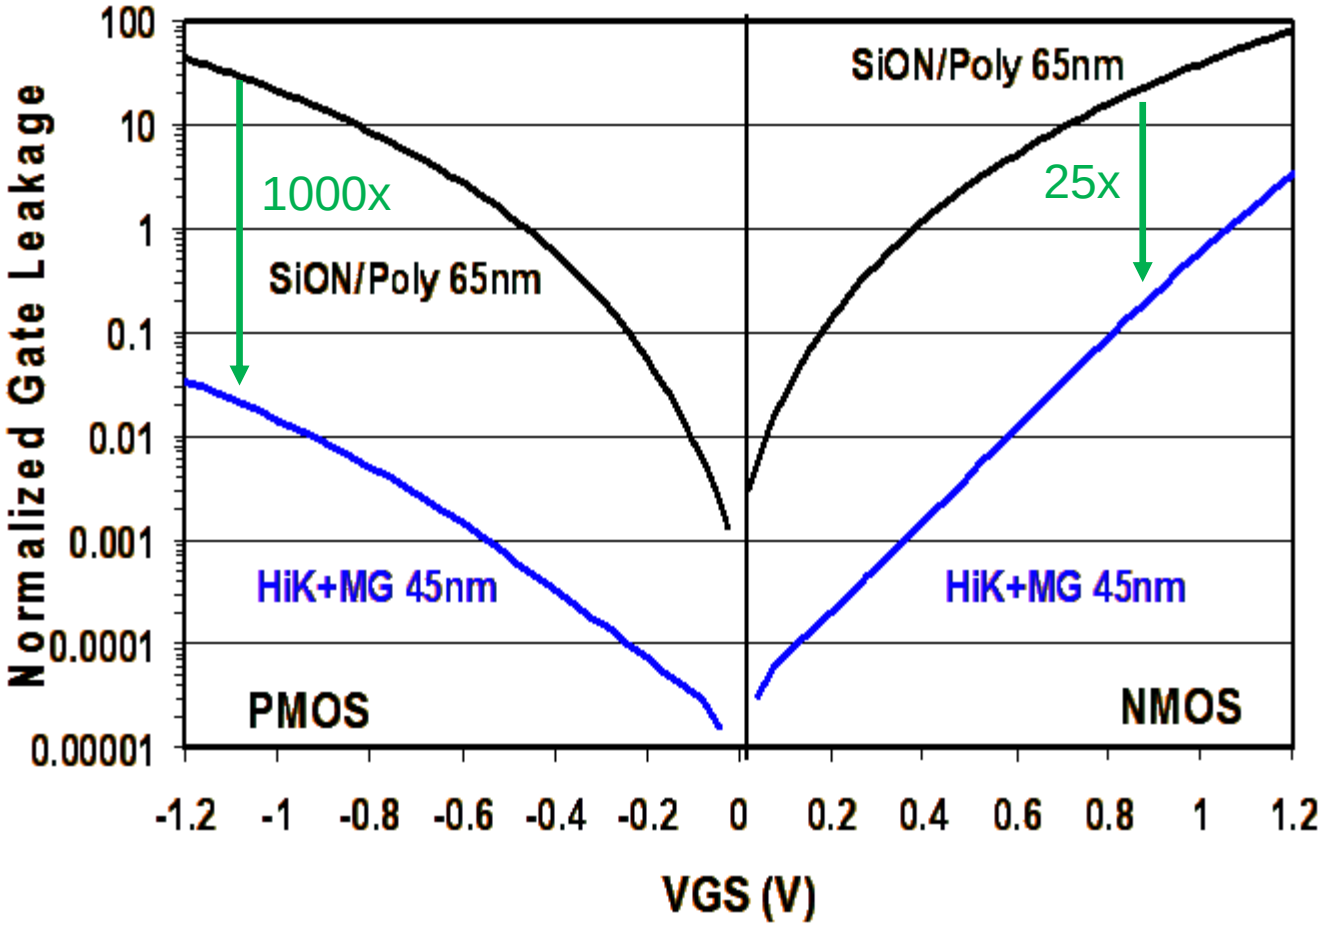

K. Mistry, IEDM 2007

6T SRAM Bit Cell Leakage Reduction

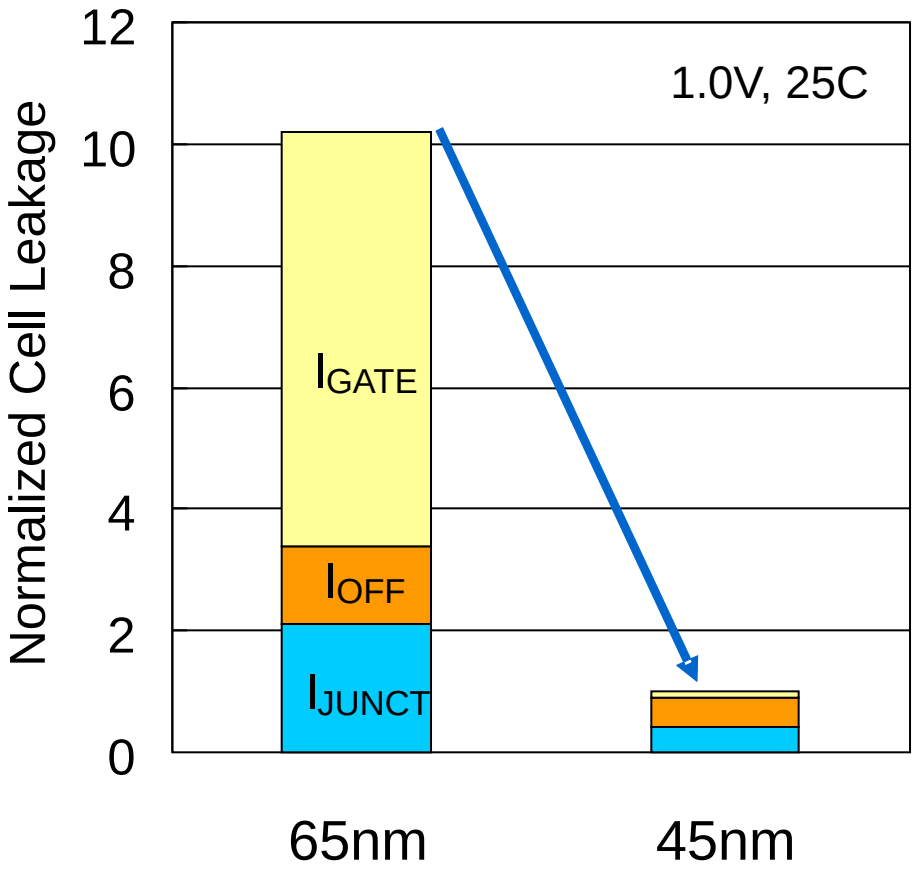

M. Bohr, ISSCC 2009

Traditional Planar Transistor

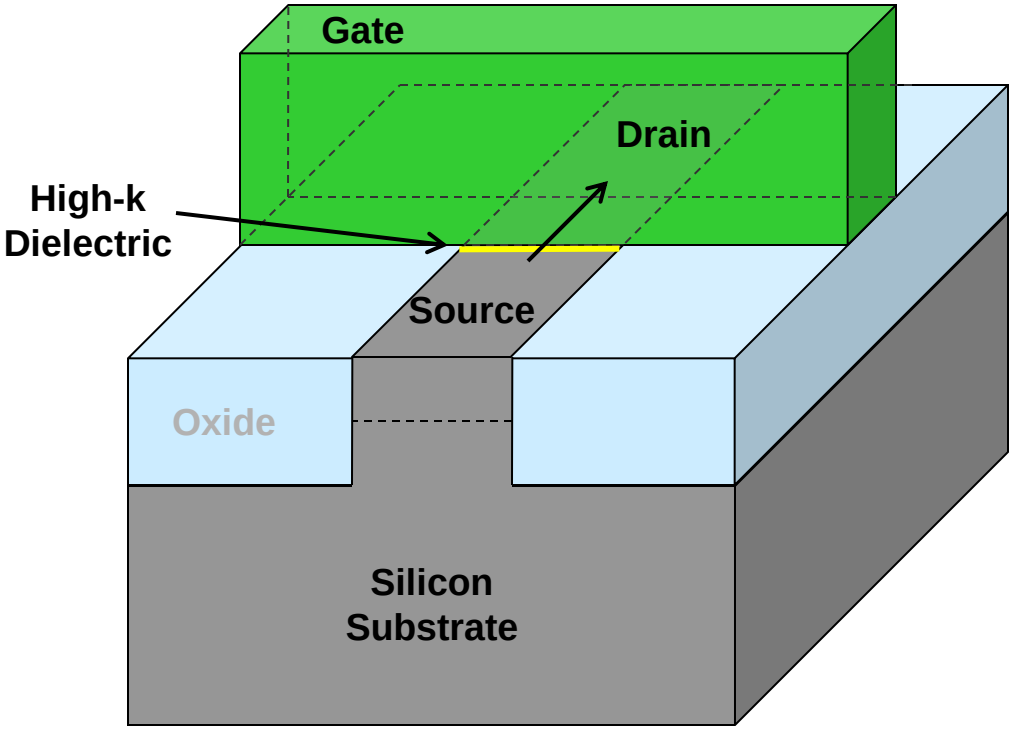

Traditional 2-D planar transistors form a conducting channel in the silicon region under the gate electrode when in the "on" state

M. Bohr, 2011

22 nm Tri-Gate Transistor

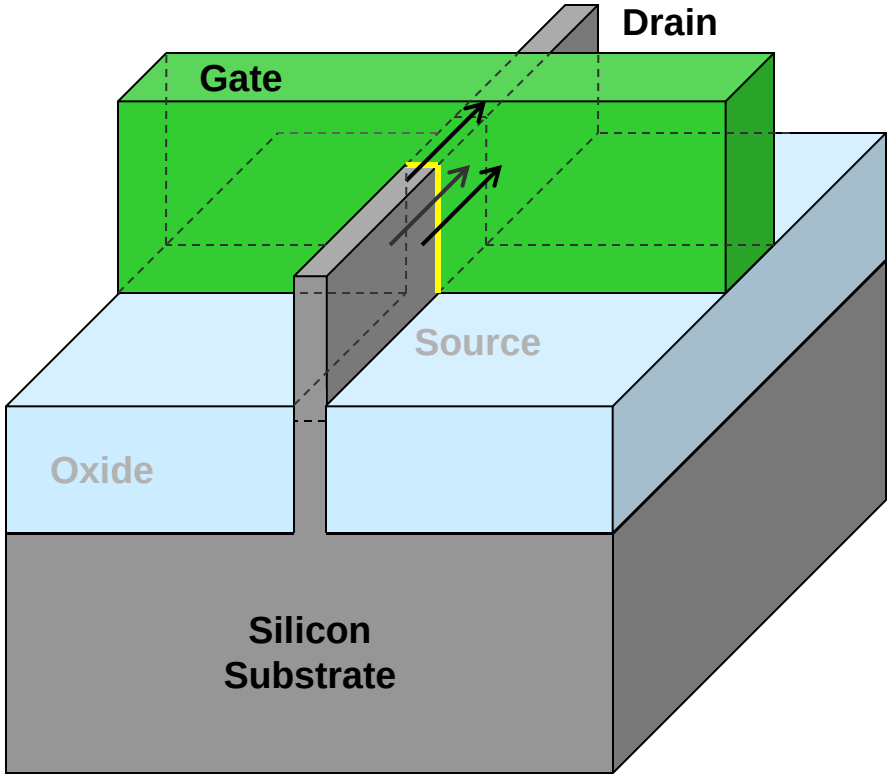

3-D Tri-Gate transistors form conducting channels on three sides of a vertical fin structure, providing “fully depleted” operation

Transistor Scaling Trends

32 nm Planar Transistors

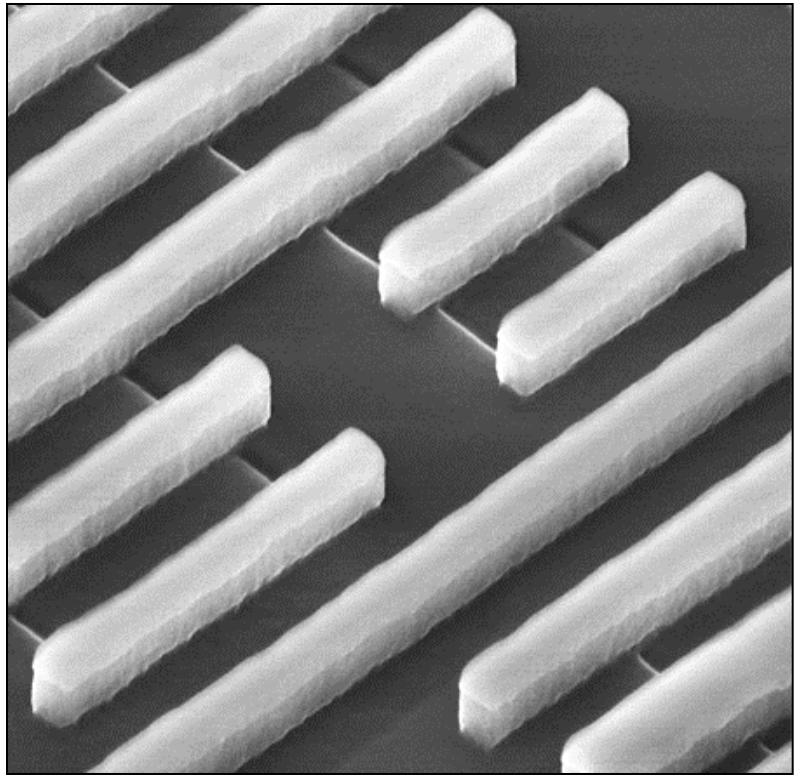

22 nm Tri-Gate Transistors

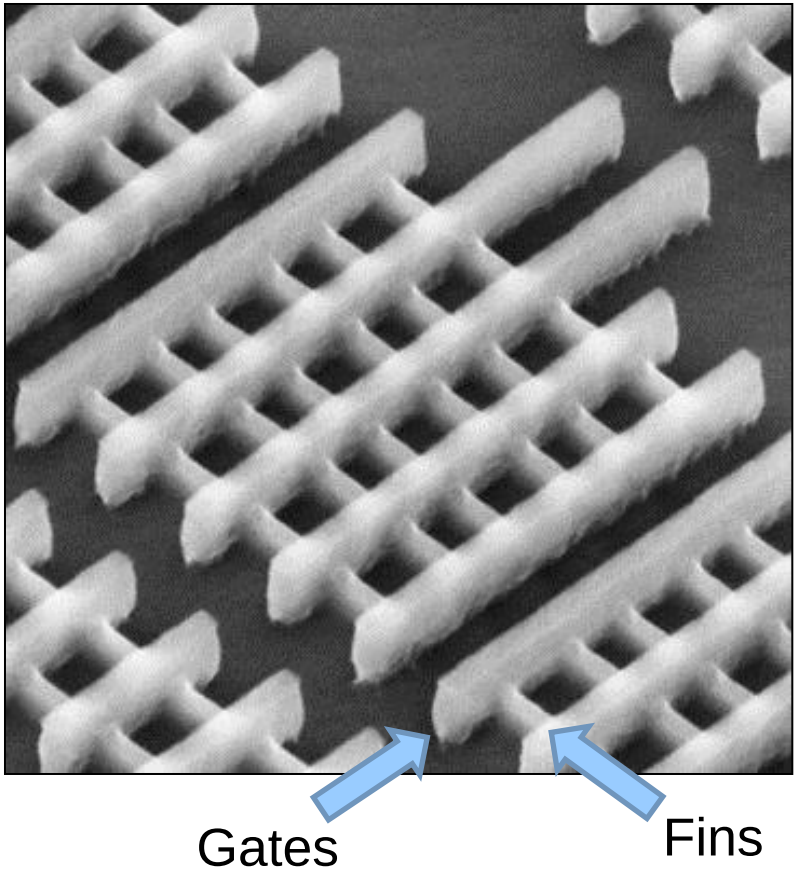

M. Bohr, 2011

Transistor Gate Delay

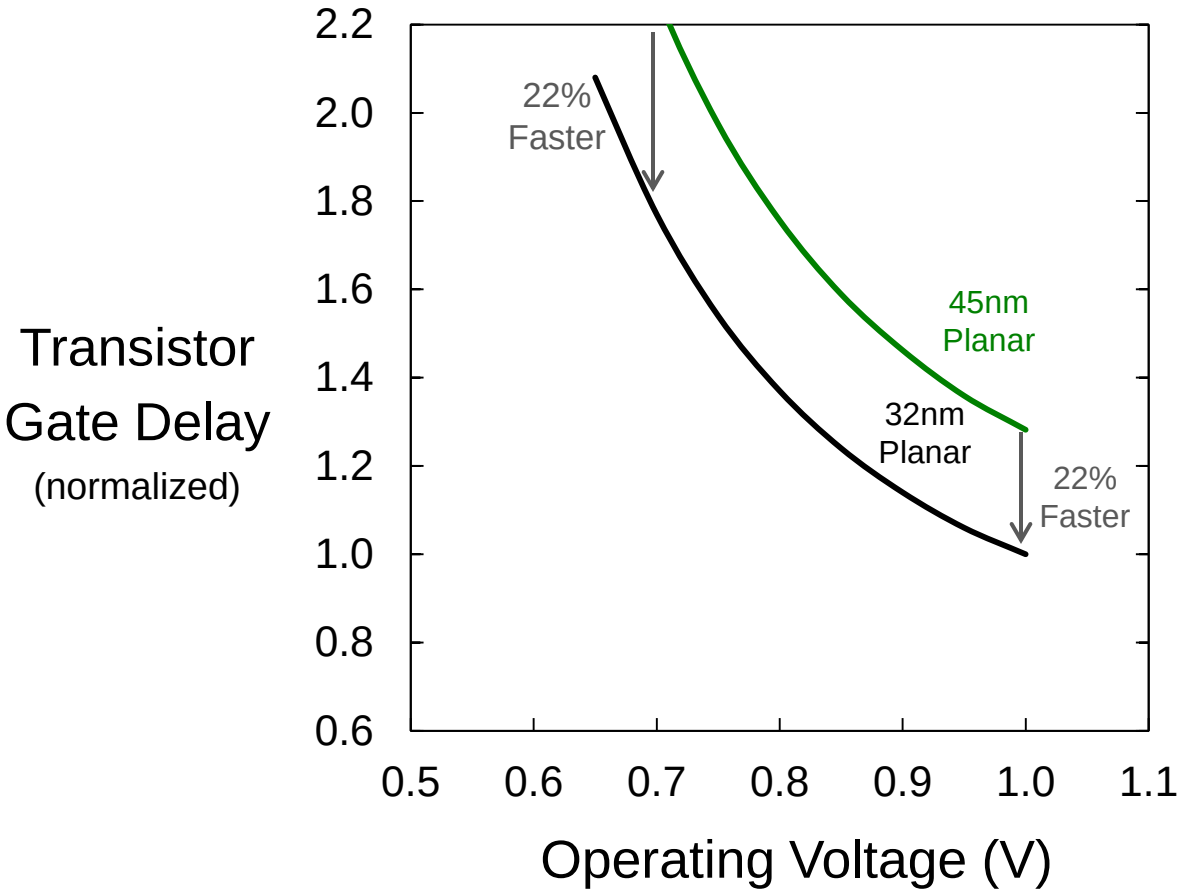

32nm planar transistors 22% faster than 45nm planar

D. Perlmutter, ISSCC 2012

Transistor Gate Delay

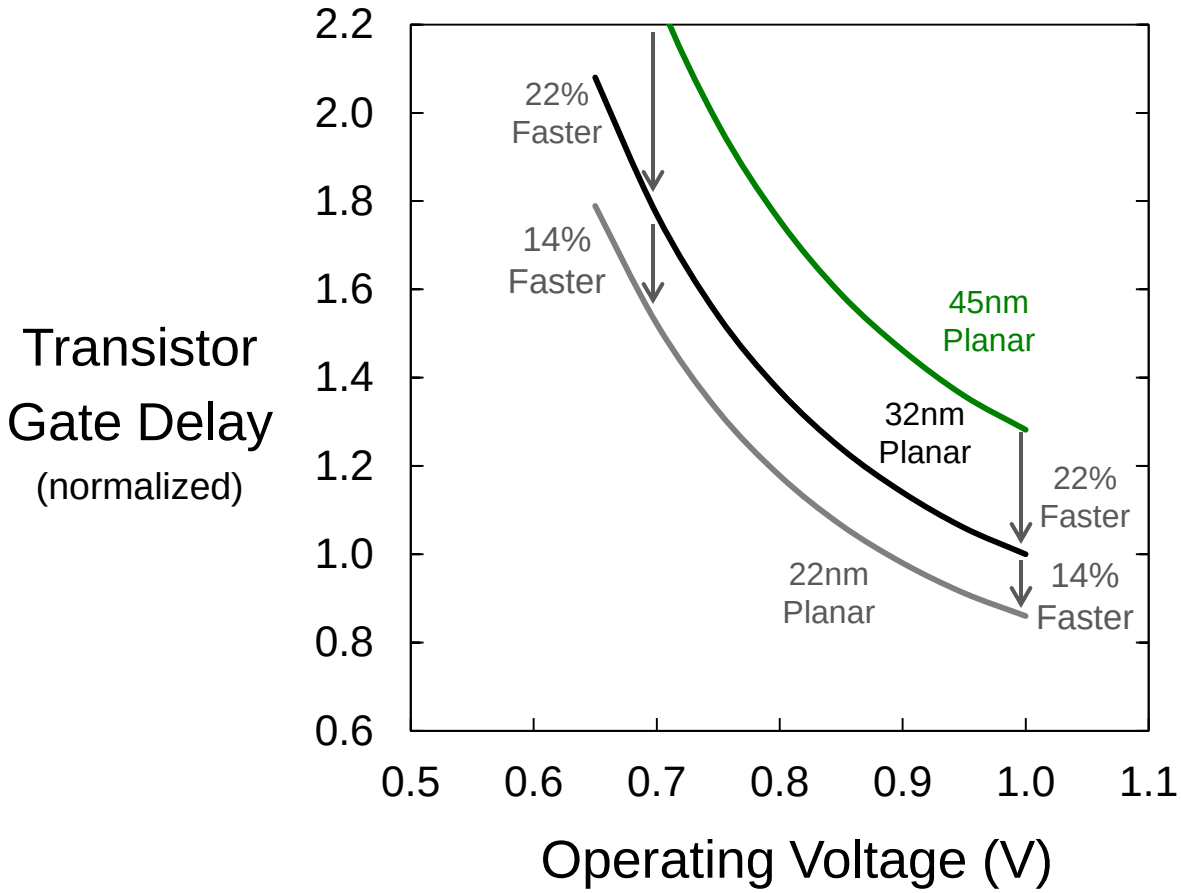

22nm planar transistors would have been only 14% faster

Transistor Gate Delay

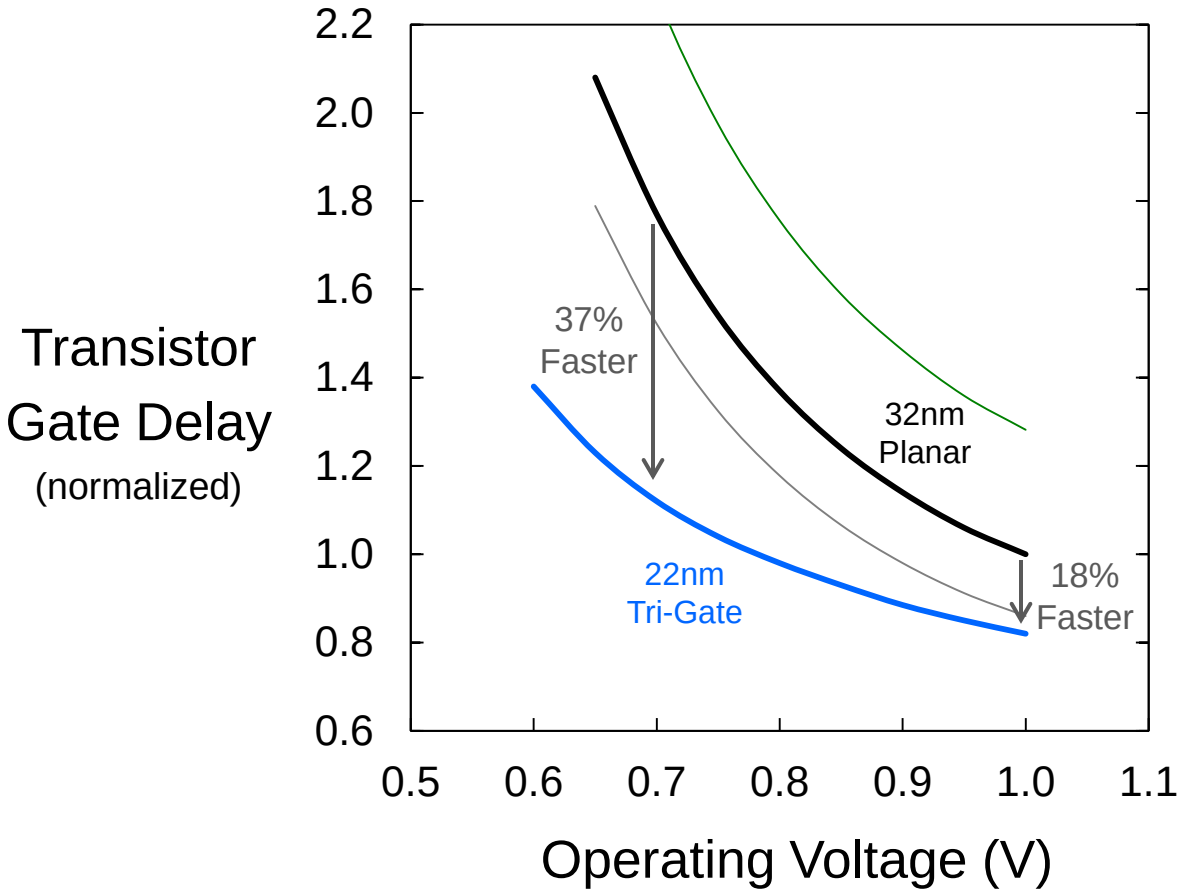

22nm Tri-Gate transistors provide improved performance at high voltage and unprecedented 37% speedup at low voltage

Intel Transistor Leadership

2003	2005	2007	2009	2011
90 nm	65 nm	45 nm	32 nm	22 nm
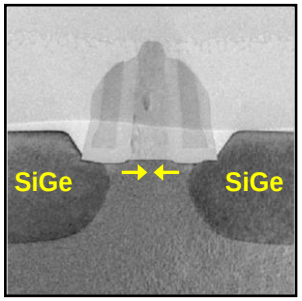	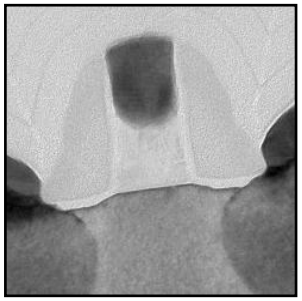	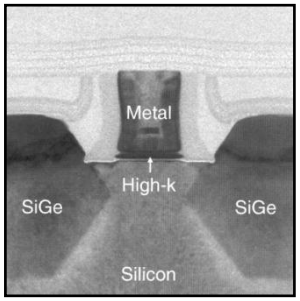	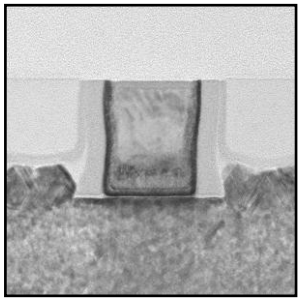	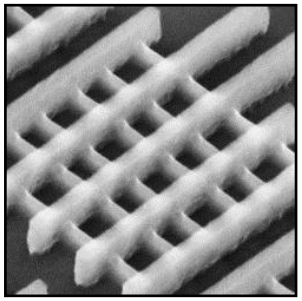
Invented SiGe Strained Silicon	2nd Gen. SiGe Strained Silicon	Invented Gate-Last High-k Metal Gate	2nd Gen. Gate-Last High-k Metal Gate	First to Implement Tri-Gate

Strained Silicon

High-k Metal Gate

Tri-Gate

Lithography Challenges

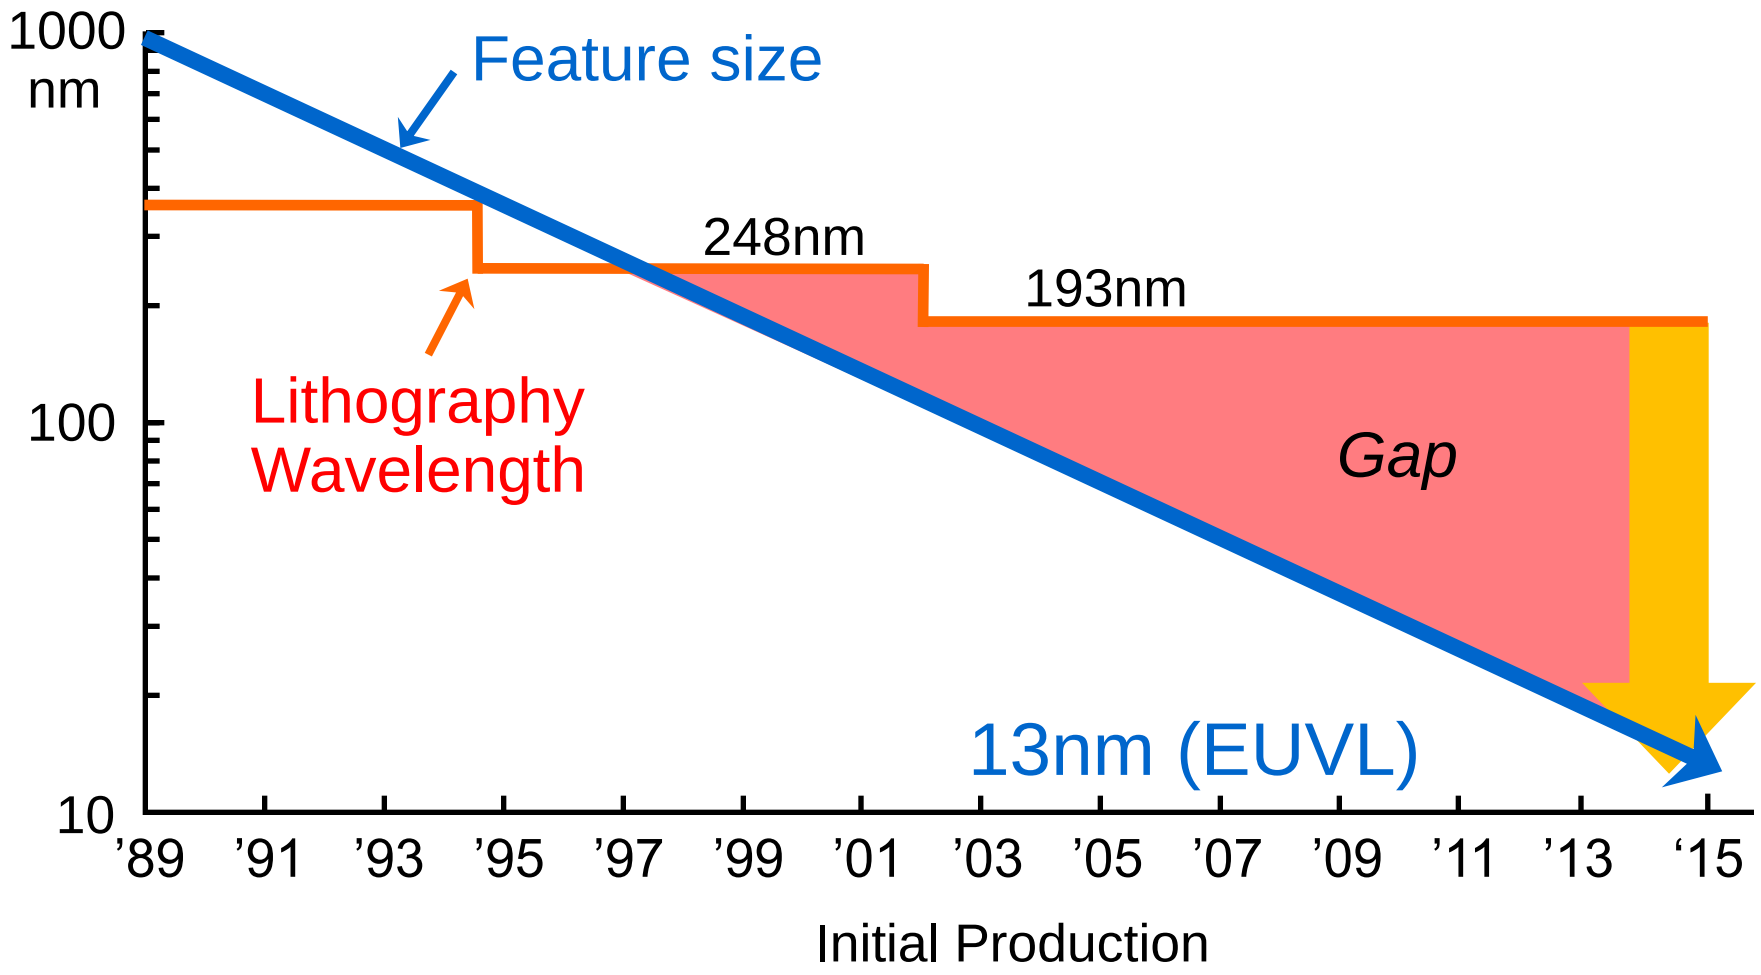

193 nm enhancements enable the 22 nm generation

Extreme Ultraviolet Lithography

- EUV lithography uses extremely short wavelength light
  - *Visible light – 400 to 700 nm*
  - *DUV lithography – 193 and 248 nm*
  - *EUV lithography – 13 nm*

EUV Micro Exposure Tool

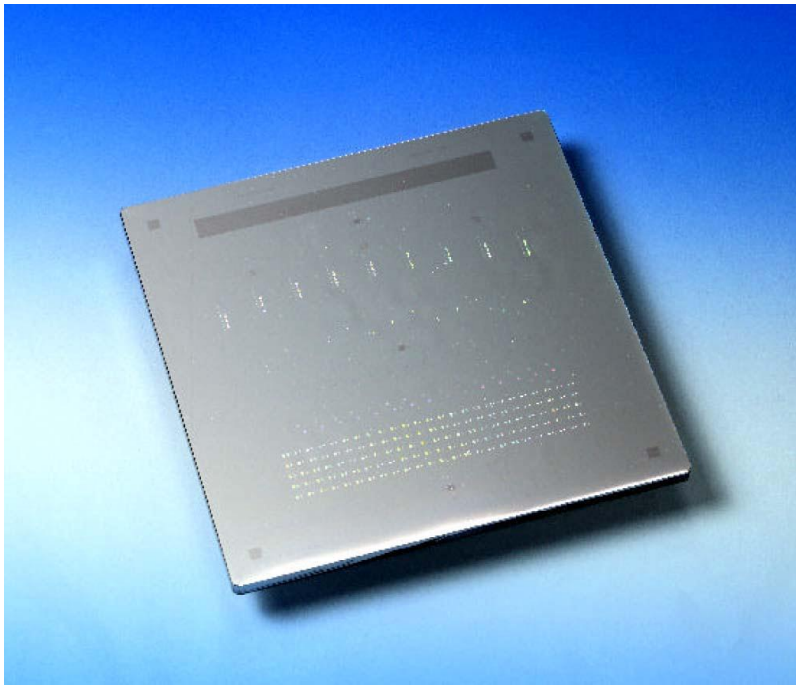

World's First EUV Mask

Layout Restrictions

65 nm Layout Style

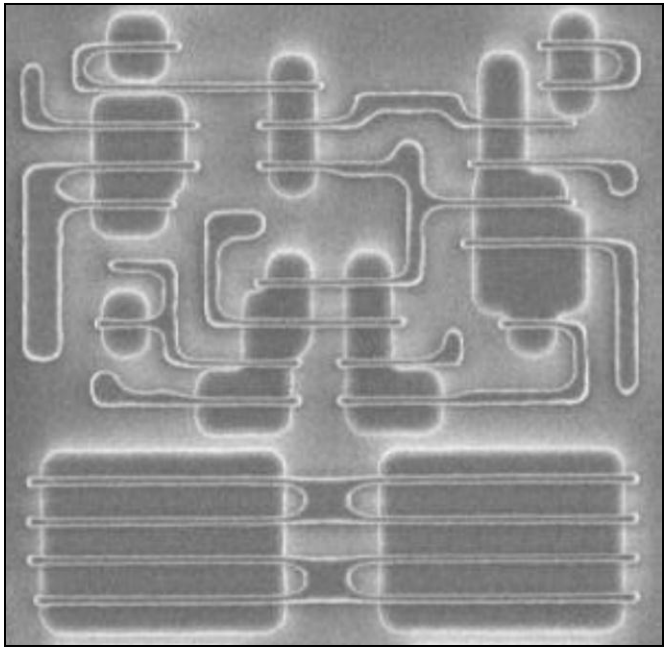

Bi-directional features

Varied gate dimensions

Varied pitches

32 nm Layout Style

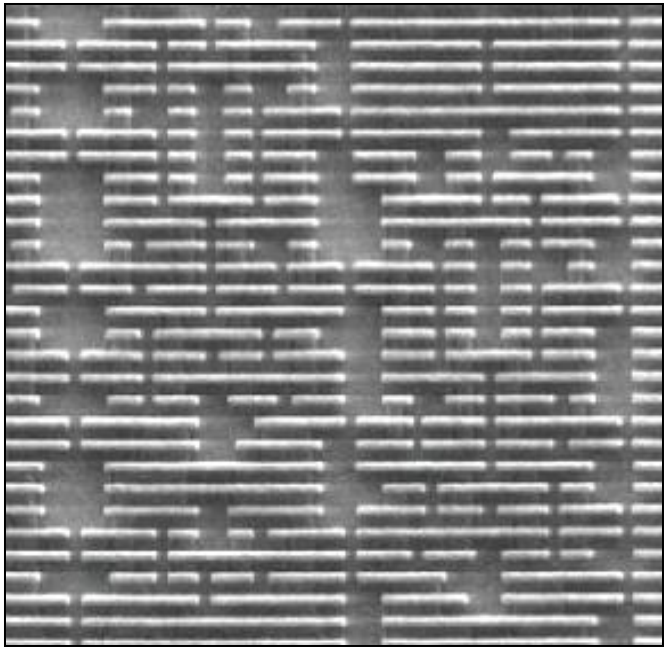

Uni-directional features

Uniform gate dimension

Gridded layout

M. Bohr, ISSCC 2009

450mm in the Era of Complex Scaling:
Must coordinate demand drivers, technical requirements and resources

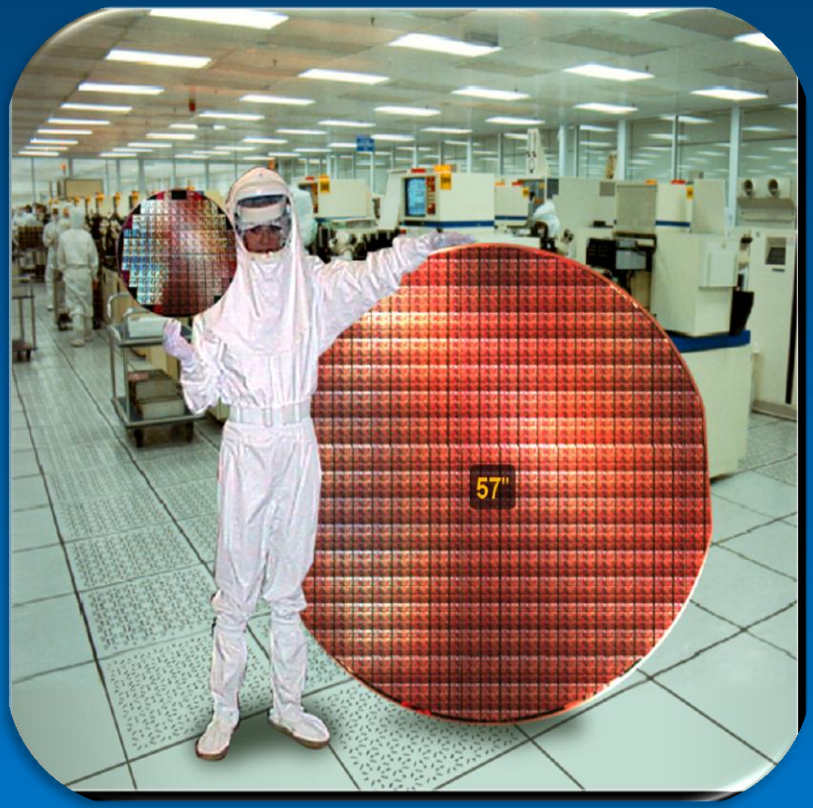

Projected 2000 Wafer, circa 1975 (Gordon Moore, ISSCC '03)

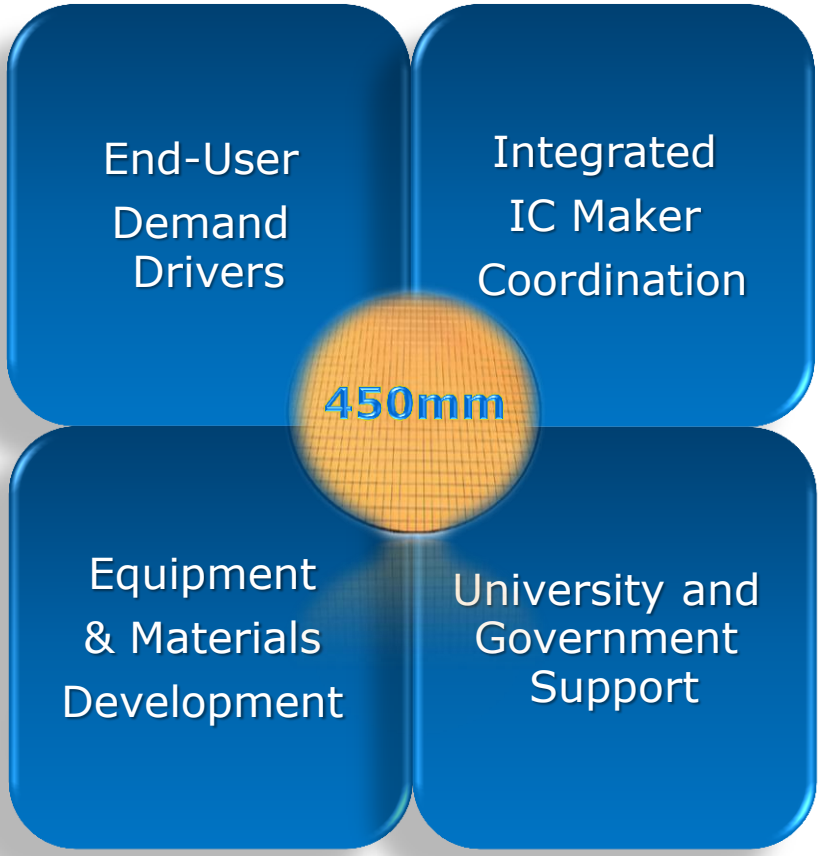

Process Variations

Die-to-Die Variations

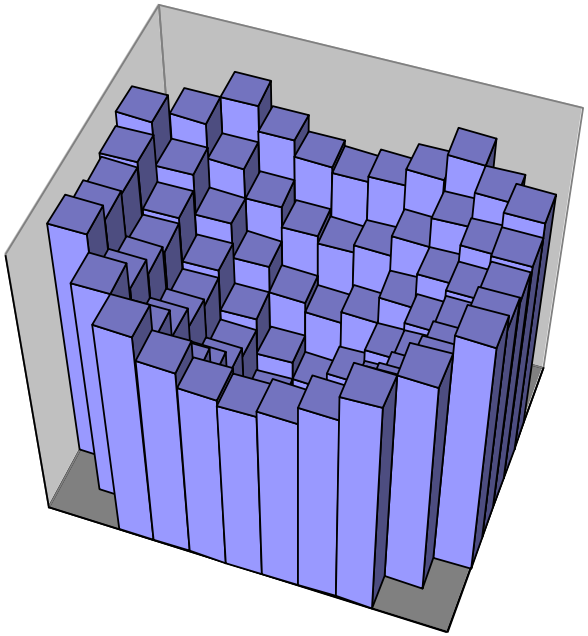

Resist Thickness

Within-Die Variations

Systematic

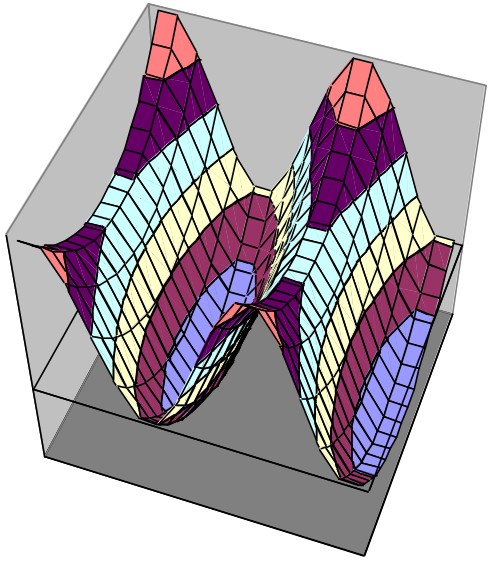

Lens Aberrations

Random

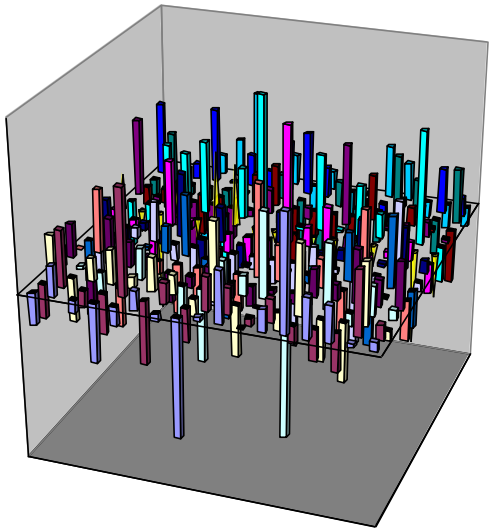

Random Placement of Dopant Atoms

Voltage and Temperature Variations

- Voltage
  - Chip activity change
  - Current delivery—RLC
  - Dynamic: ns to 10-100µs
  - Within-die variation

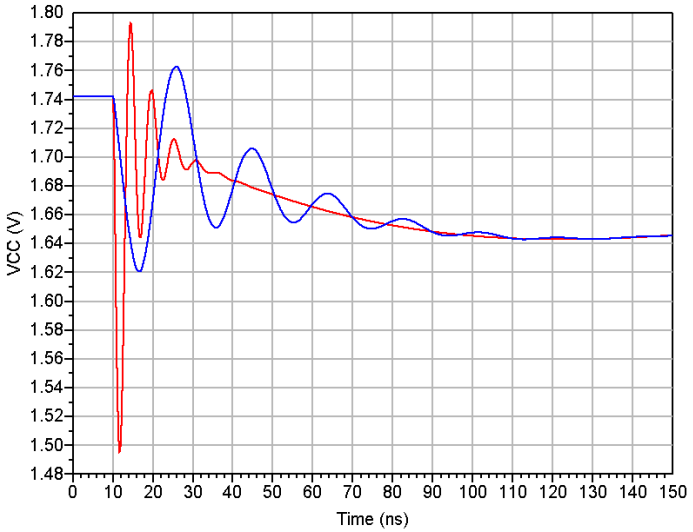

- Temperature
  - Activity & ambient change
  - Dynamic: 100-1000µs
  - Within-die variation

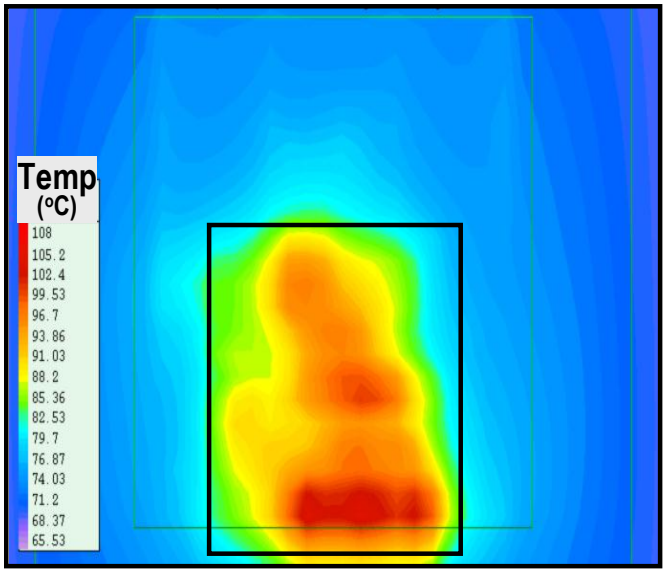

Impact on Design Methodology

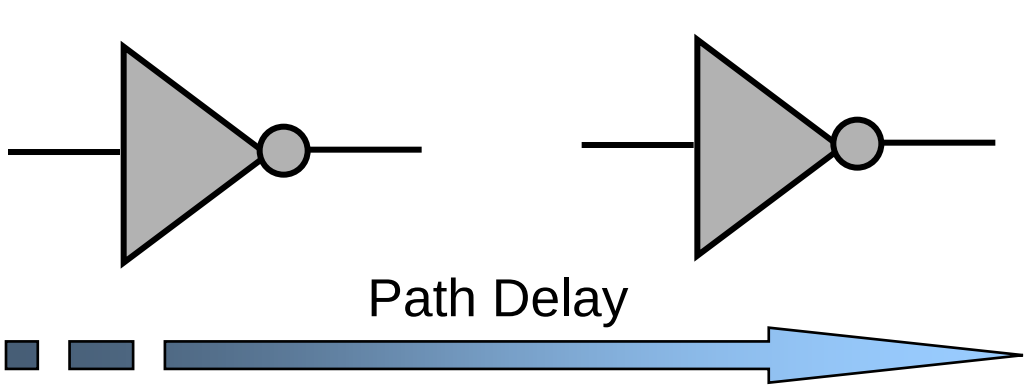

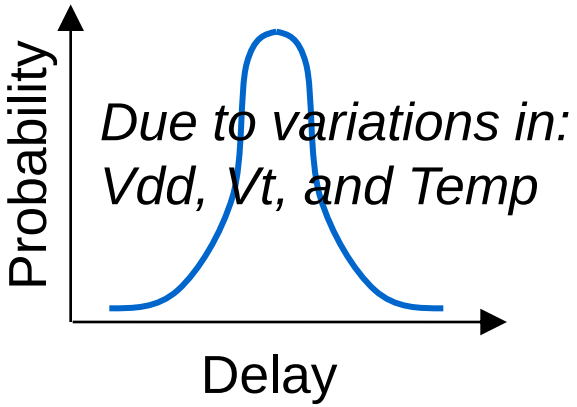

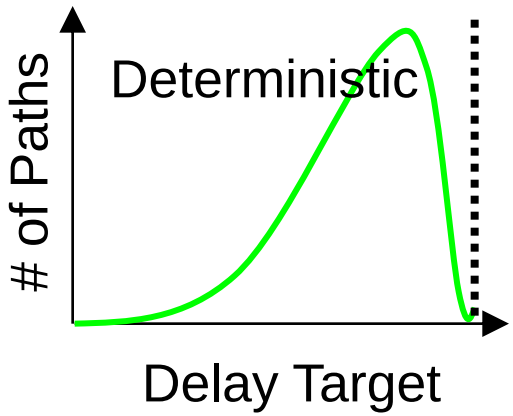

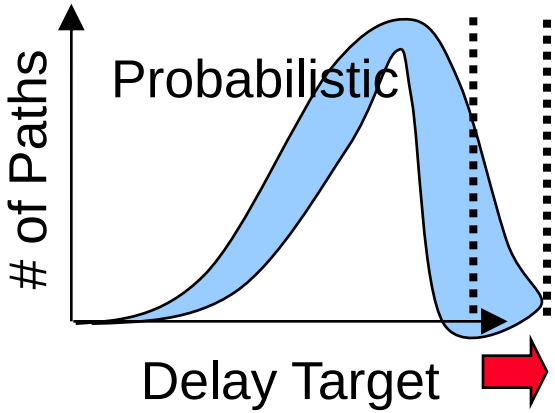

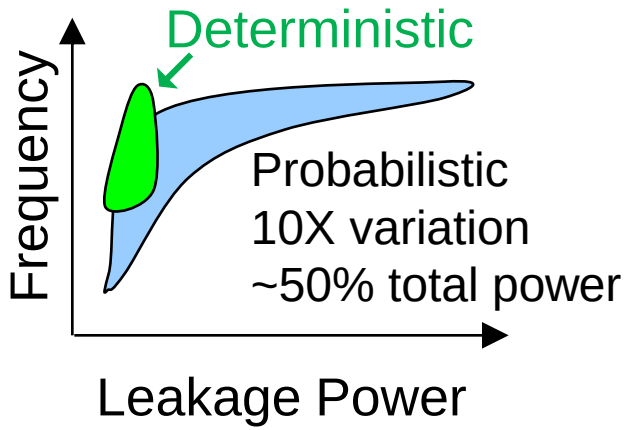

Major paradigm shift from deterministic design to probabilistic / statistical design

SRAM Cell Size Scaling

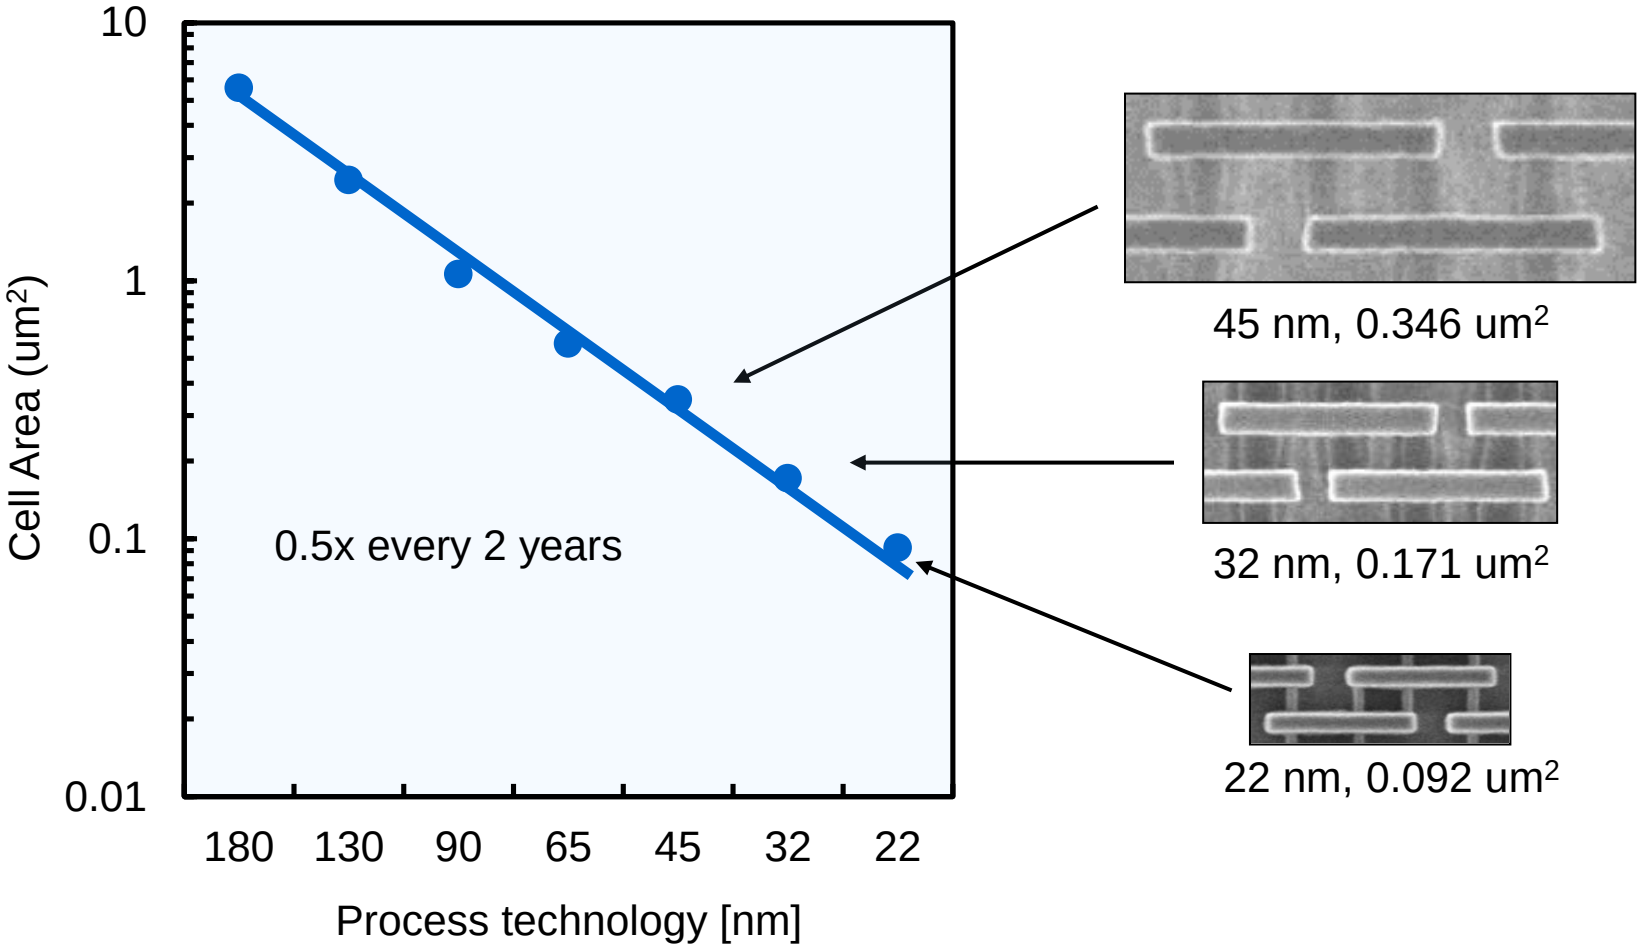

Memory density continues to double every 2 years

Interconnect Trends

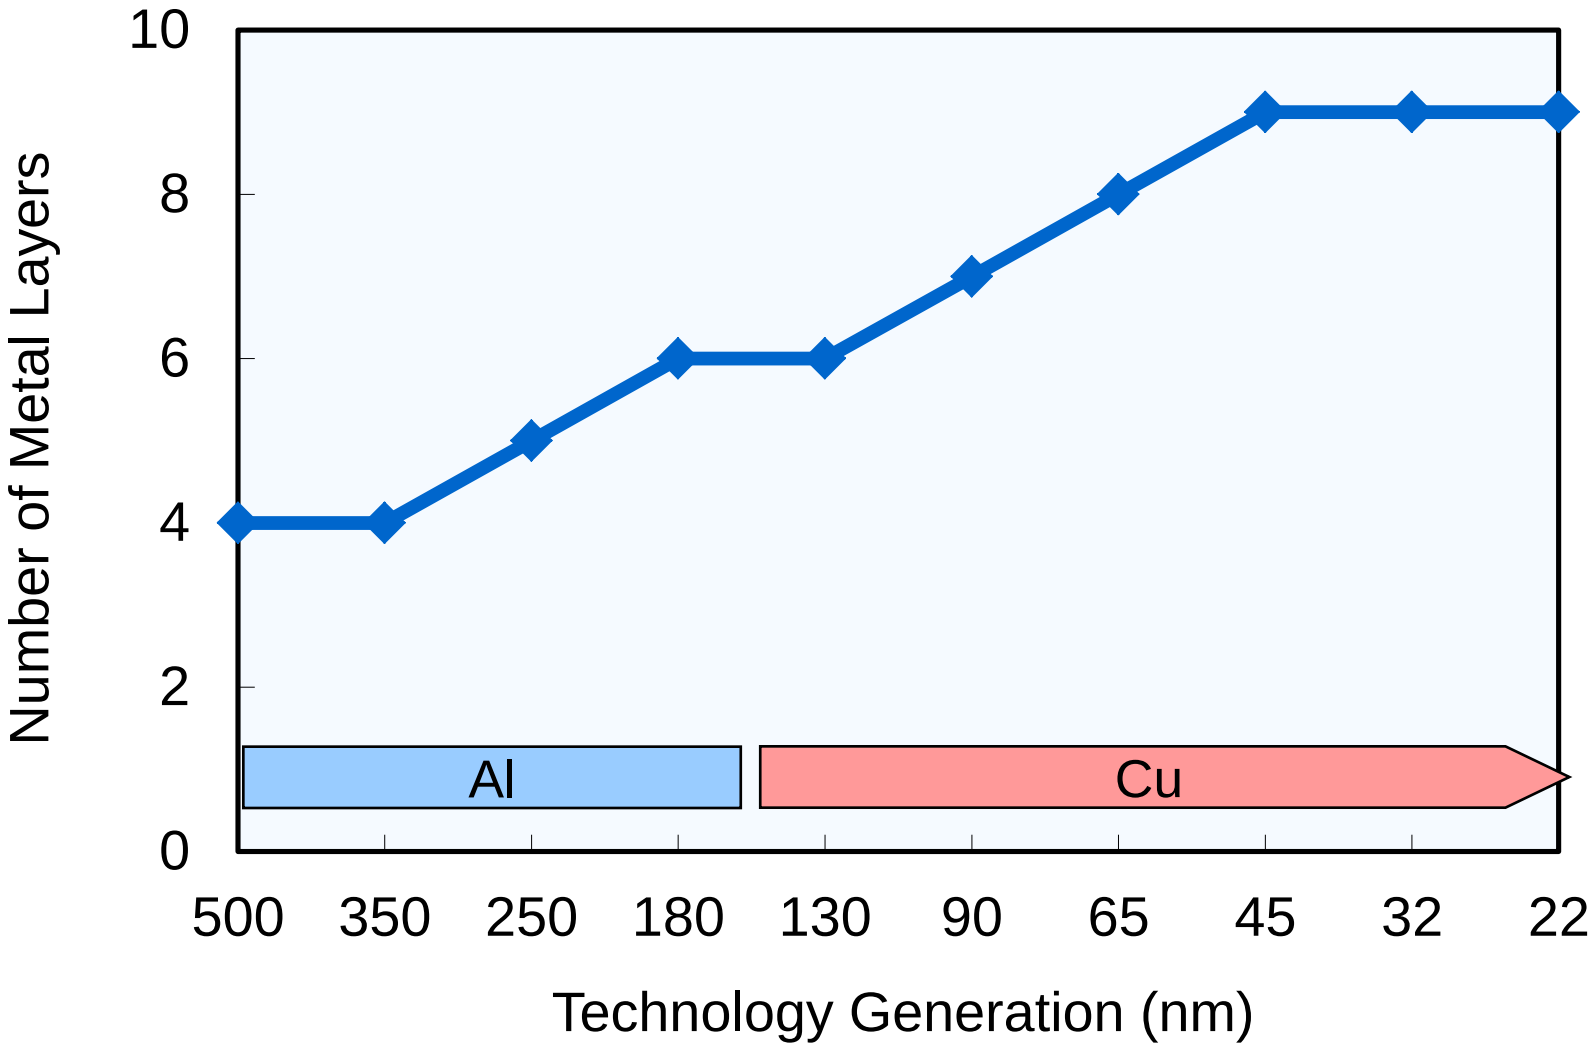

22nm Interconnects

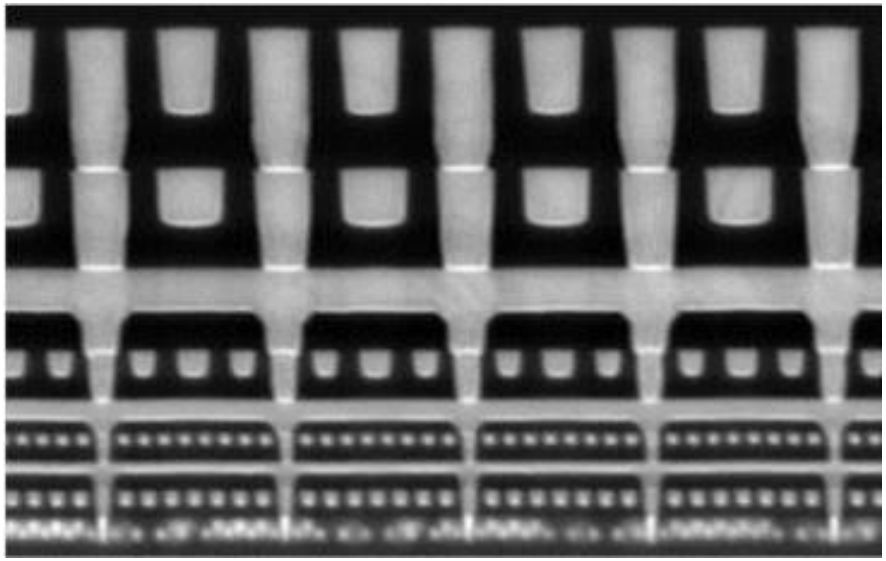

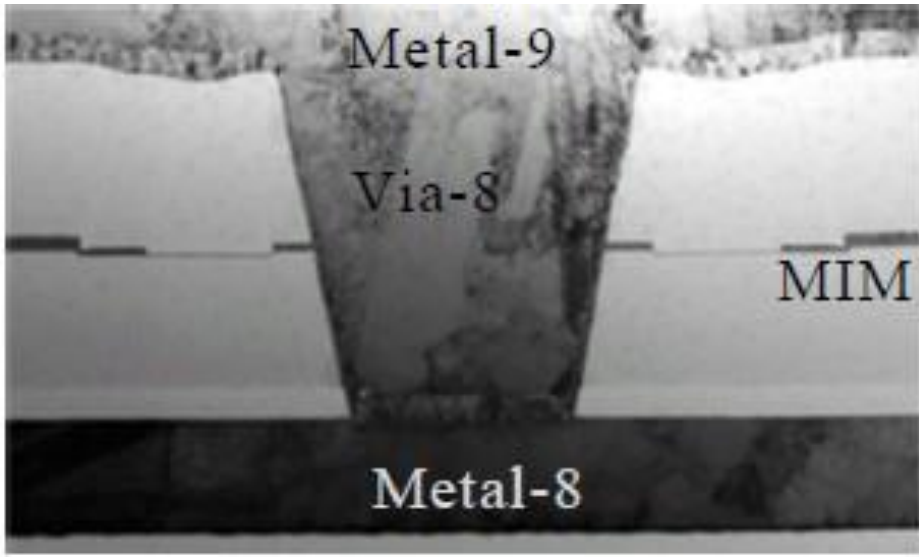

- M1 to M8 cross-section
- M1-M6 use ultra-low-k ILD and self-aligned vias providing 13-18% capacitance reduction

- Cross-section of integrated MIM capacitor
- Enables capacitance density of >20fF/mm^2

C. Auth, VLSI Symposium 2012

On-chip Interconnect Trend

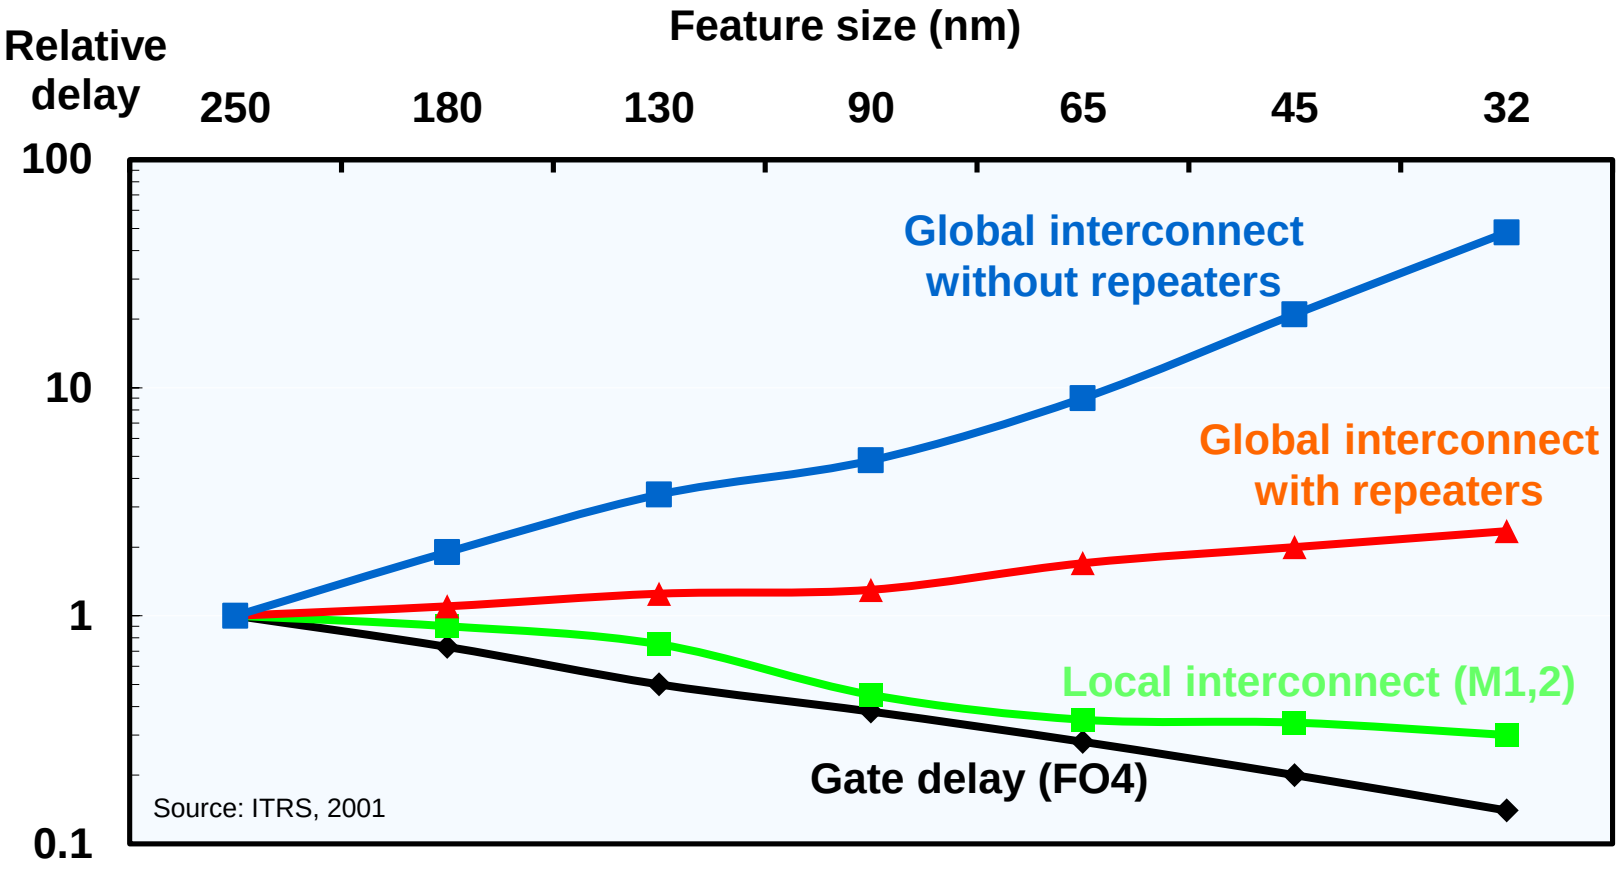

- Local interconnects scale with gate delay
- Global interconnects do not keep up with scaling

Agenda

- Microprocessor Design Trends
- Process Technology Directions
- Active Power Management
- Leakage Reduction Techniques
- Packaging and Thermal Modeling
- Future Directions and Summary

Voltage and Frequency Scaling

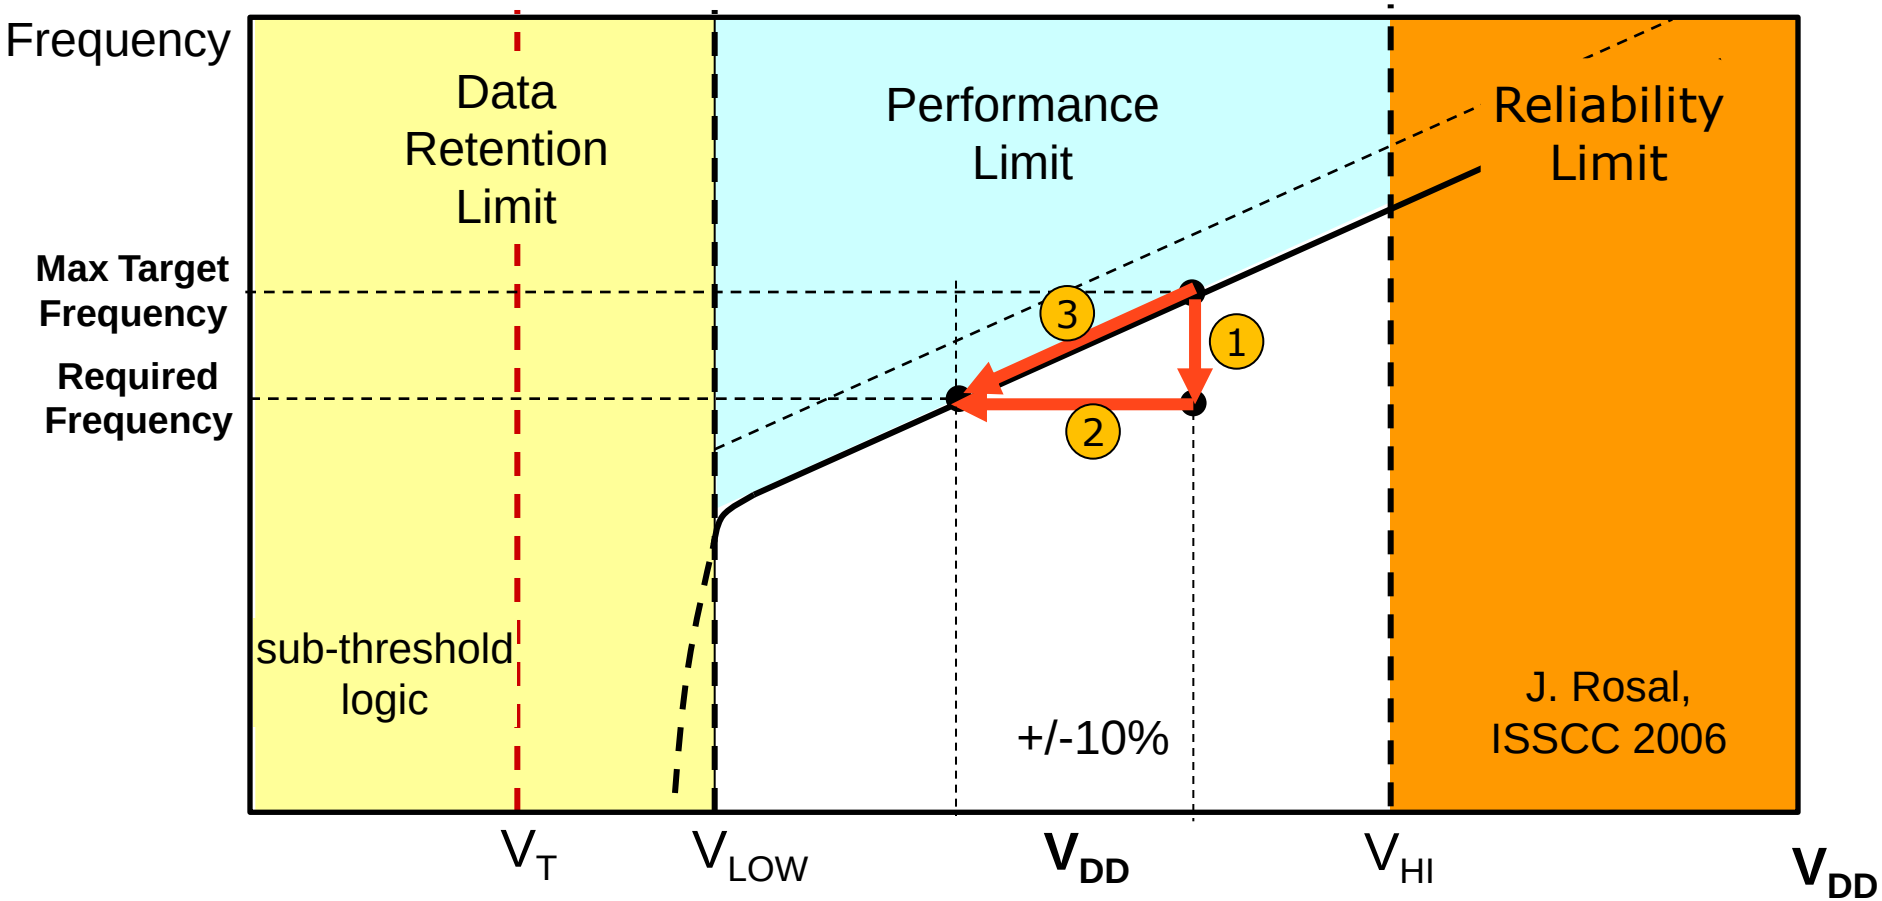

1 - Fixed V_{DD}, Frequency Scaling: Linear Power Reduction

2 - Fixed Frequency, V_{DD} Scaling: Square Power Reduction

3 - Voltage and Frequency scaling: Cubic Power Reduction

Memory and RF Vmin Reduction

- Write Assist circuit temporarily drops the array supply node to make it easier to write into the bit-cell
- Both Cache and Register Files use this technique to improve write Vmin in 22nm Ivy Bridge processor

22nm transistor and circuit improvements enable Vmin reduction of >100mV for Cache and 60mV for RF

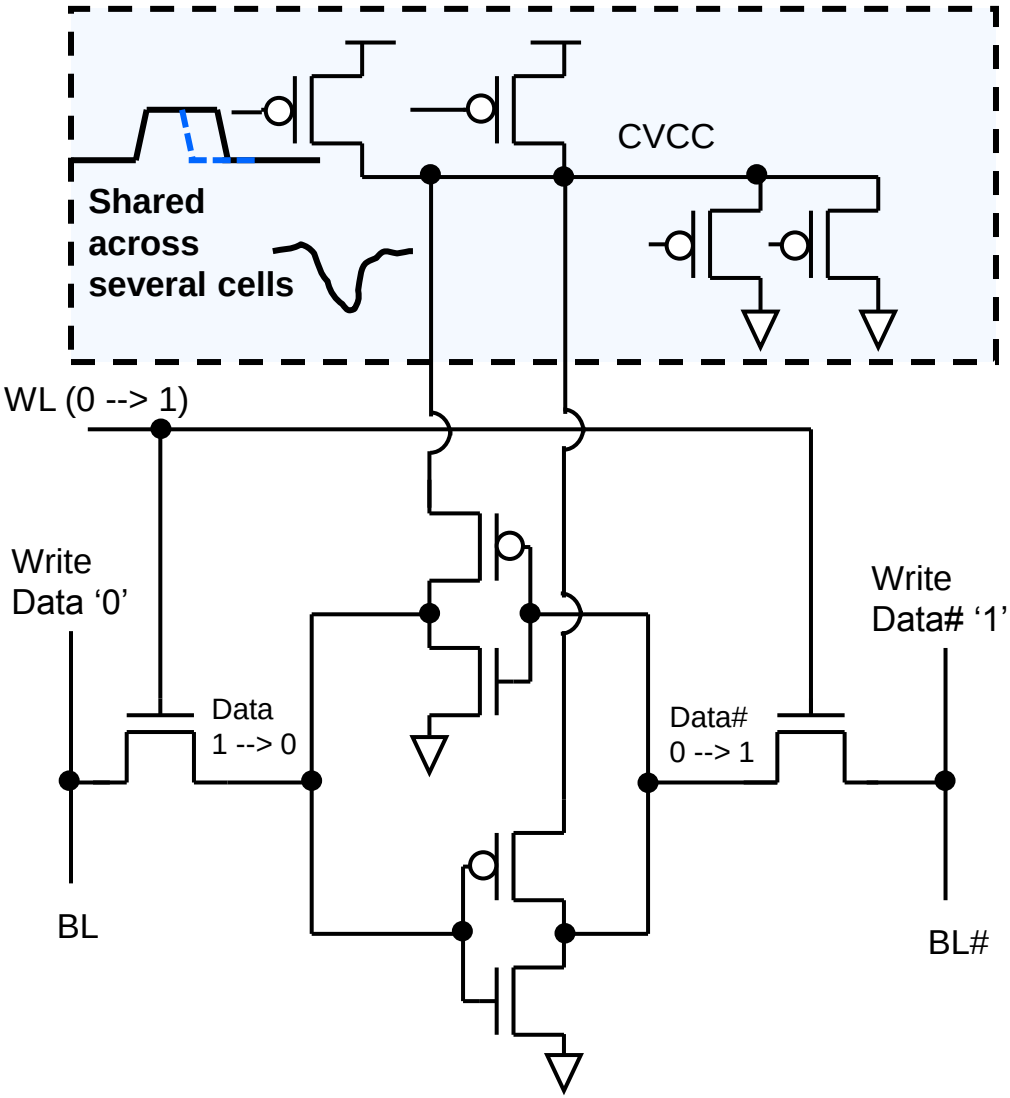

S. Damaraju, ISSCC 2012

NTV Pentium® Processor

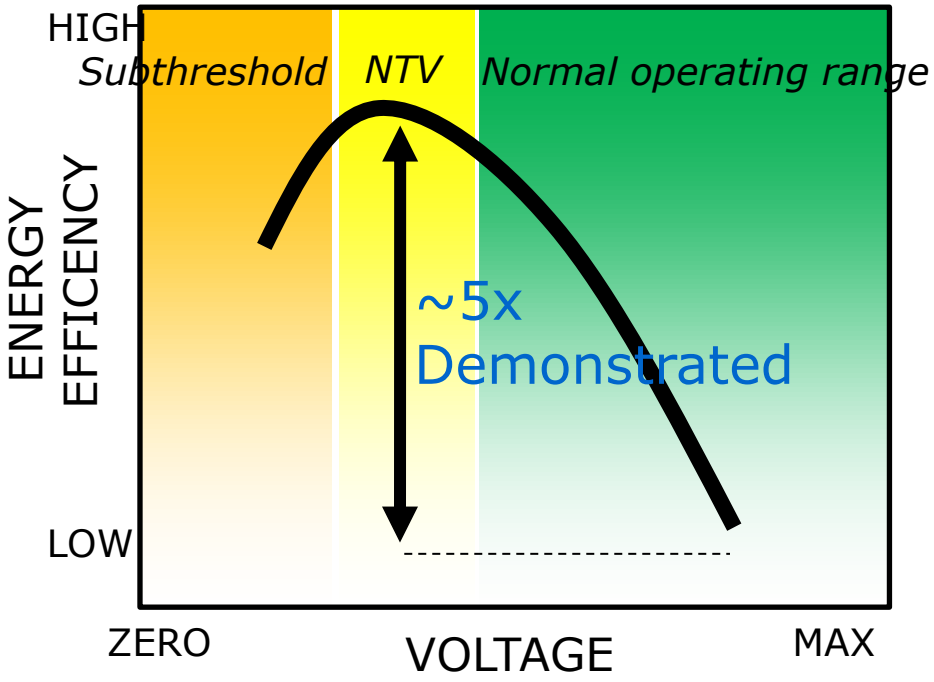

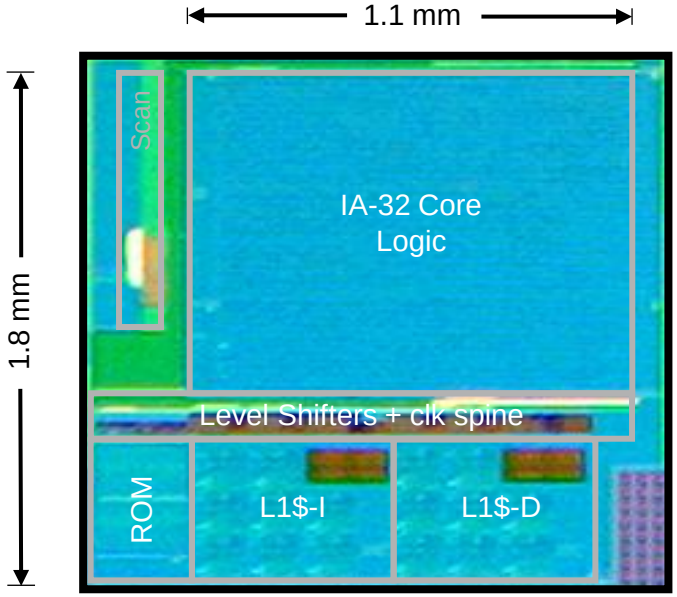

Ultra-low Power	Energy Efficient	High Performance
280 mV	0.45 V	1.2 V
3 MHz	60 MHz	915 MHz
2 mW	10 mW	737 mW
1500 Mips/W	5830 Mips/W	1240 Mips/W

Technology	32nm High-K Metal Gate
Interconnect	1 Poly, 9 Metal (Cu)
Transistors	Core:6M
Core Area	2mm^2
Package	951 Pins FCBGA11

S. Jain, ISSCC 2012

Clock Gating

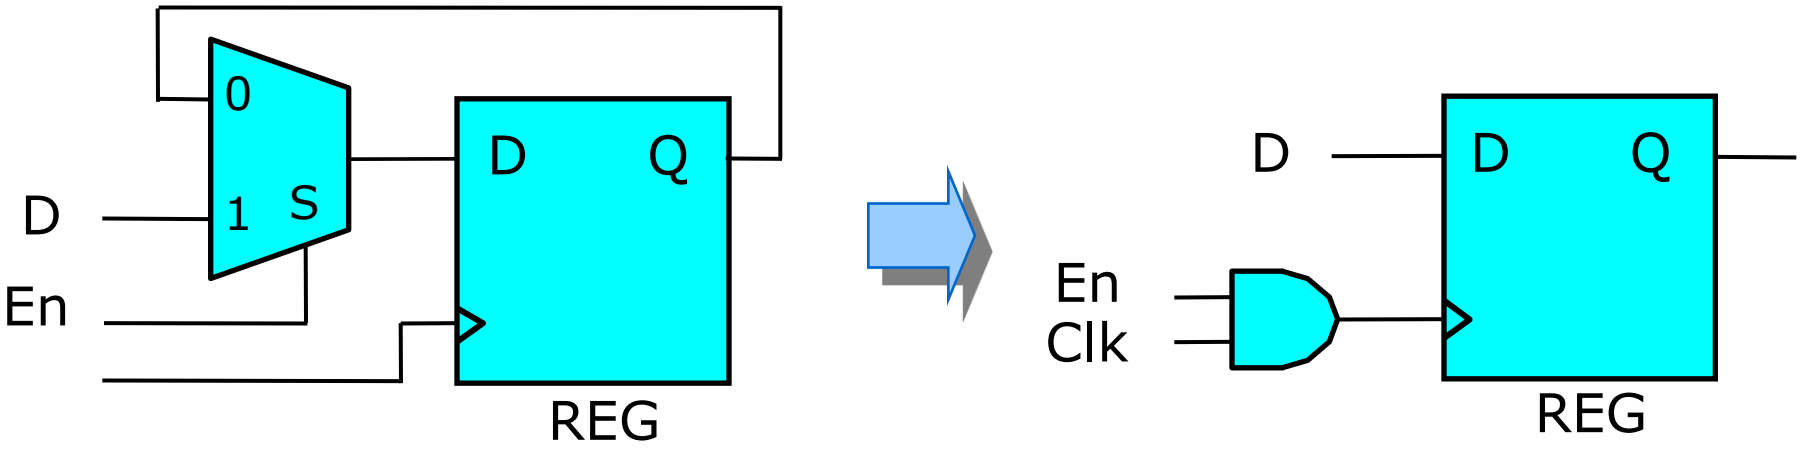

- Save power by gating the clock when data activity is low
- Requires detailed logic validation

Core Power Management

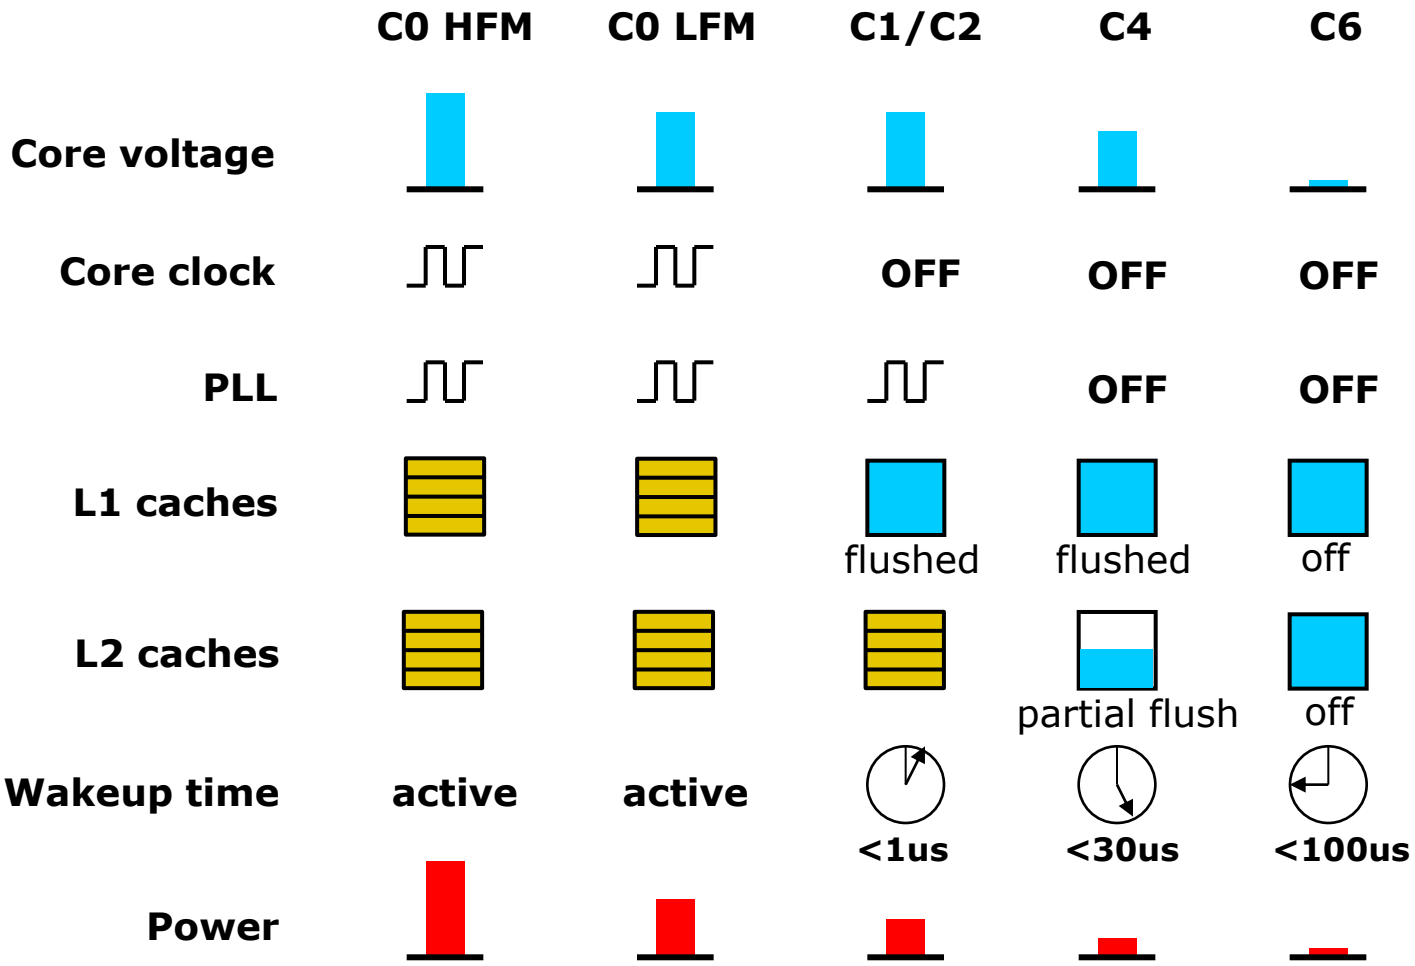

Modulating the processor core voltage and frequency enables lower power states

Gerosa, A-SSCC 2008

Multiple Voltage Domains

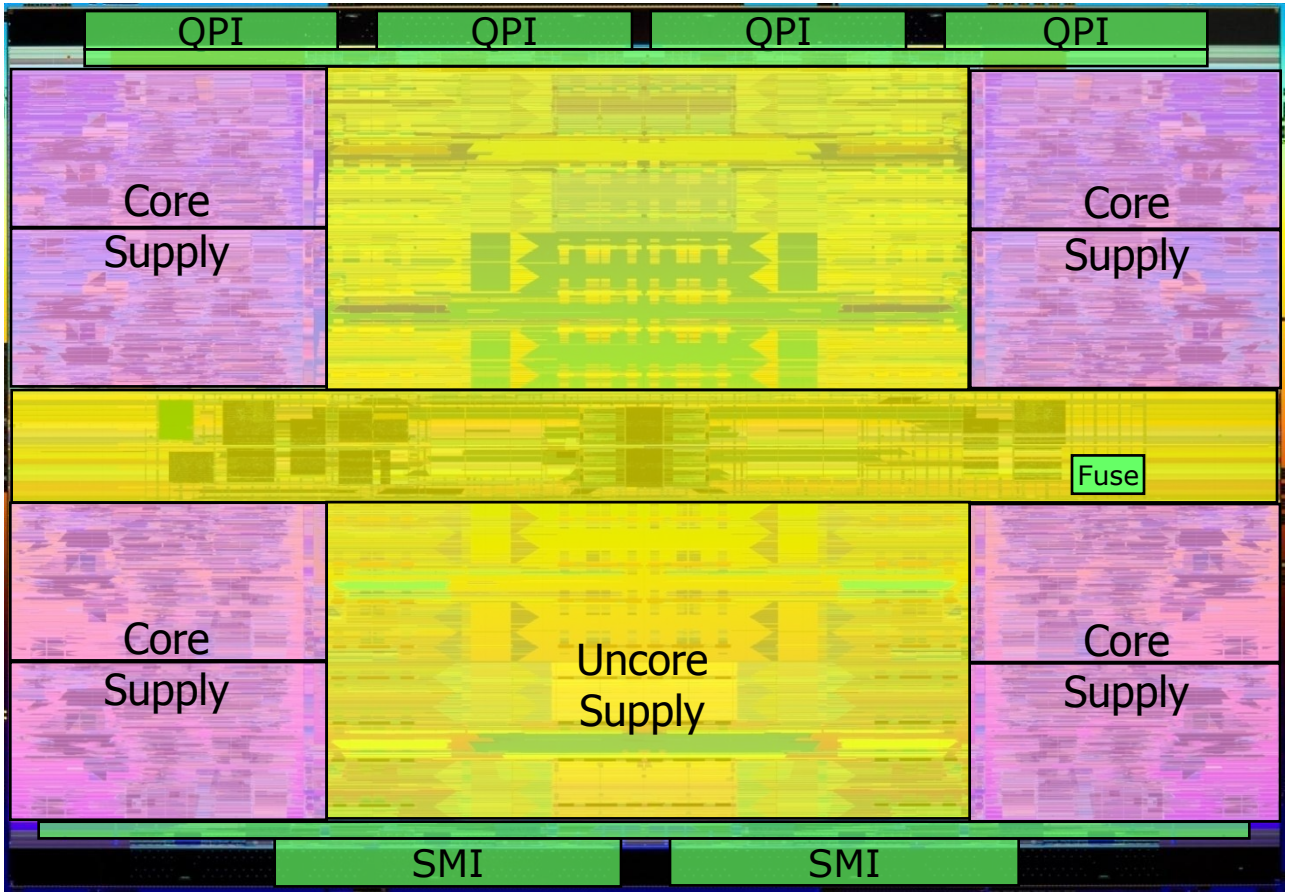

I/O Domain
1.1V
fixed

Un-Core Domain
0.9-1.1V
fixed

Core Domain
0.85-1.1V
variable

Multiple voltage domains minimize power consumption across the core and uncore areas

Rusu, ISSCC 2009

Multiple Clock Domains

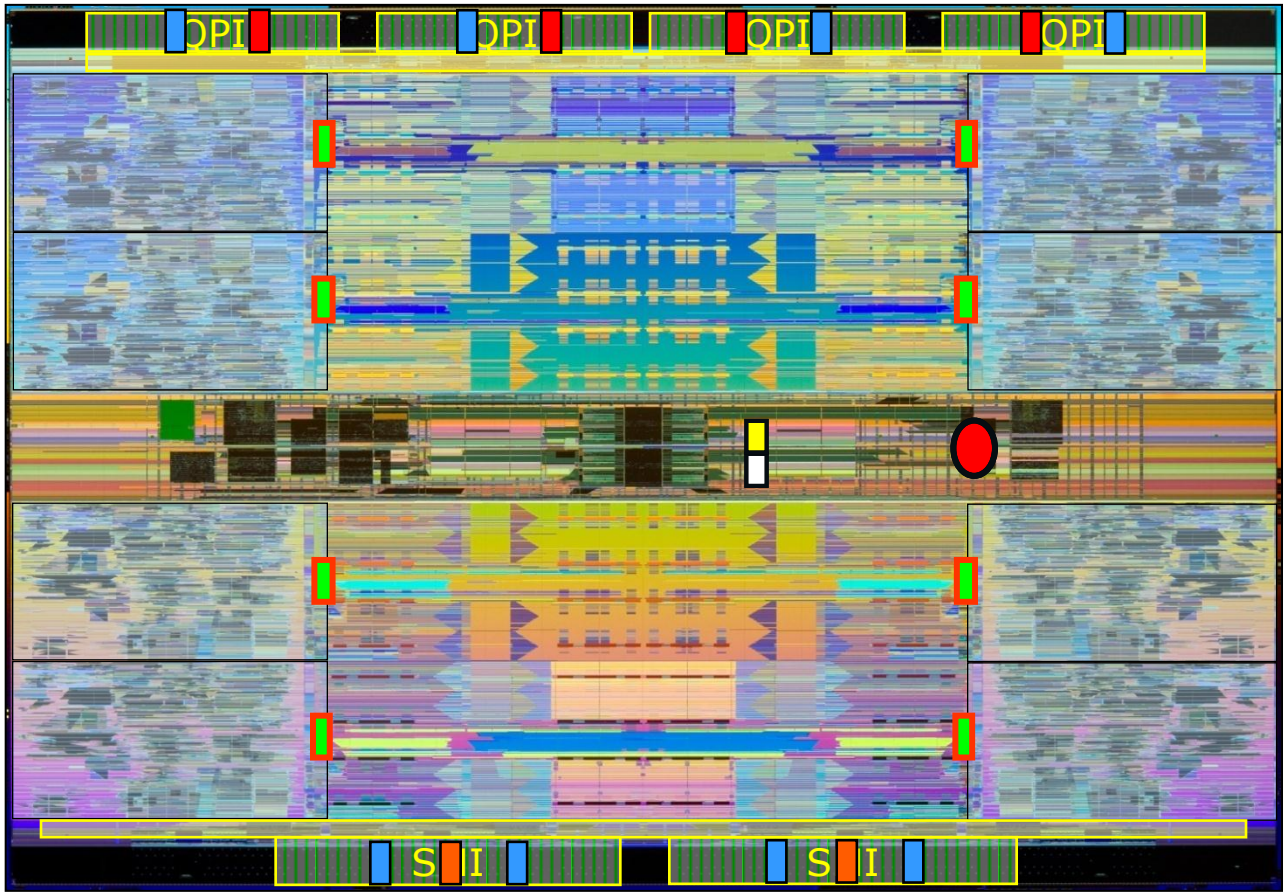

BCLK

IO PLLs

Filter PLL

Un-core PLL

IO DLLs

Core PLLs

Three primary clock domains: core, un-core, I/O

Total of 16 PLLs and 8 DLLs

Rusu, ISSCC 2009

Agenda

- Microprocessor Design Trends
- Process Technology Directions
- Active Power Management
- Leakage Reduction Techniques
- Packaging and Thermal Modeling
- Future Directions and Summary

Subthreshold Leakage Trend

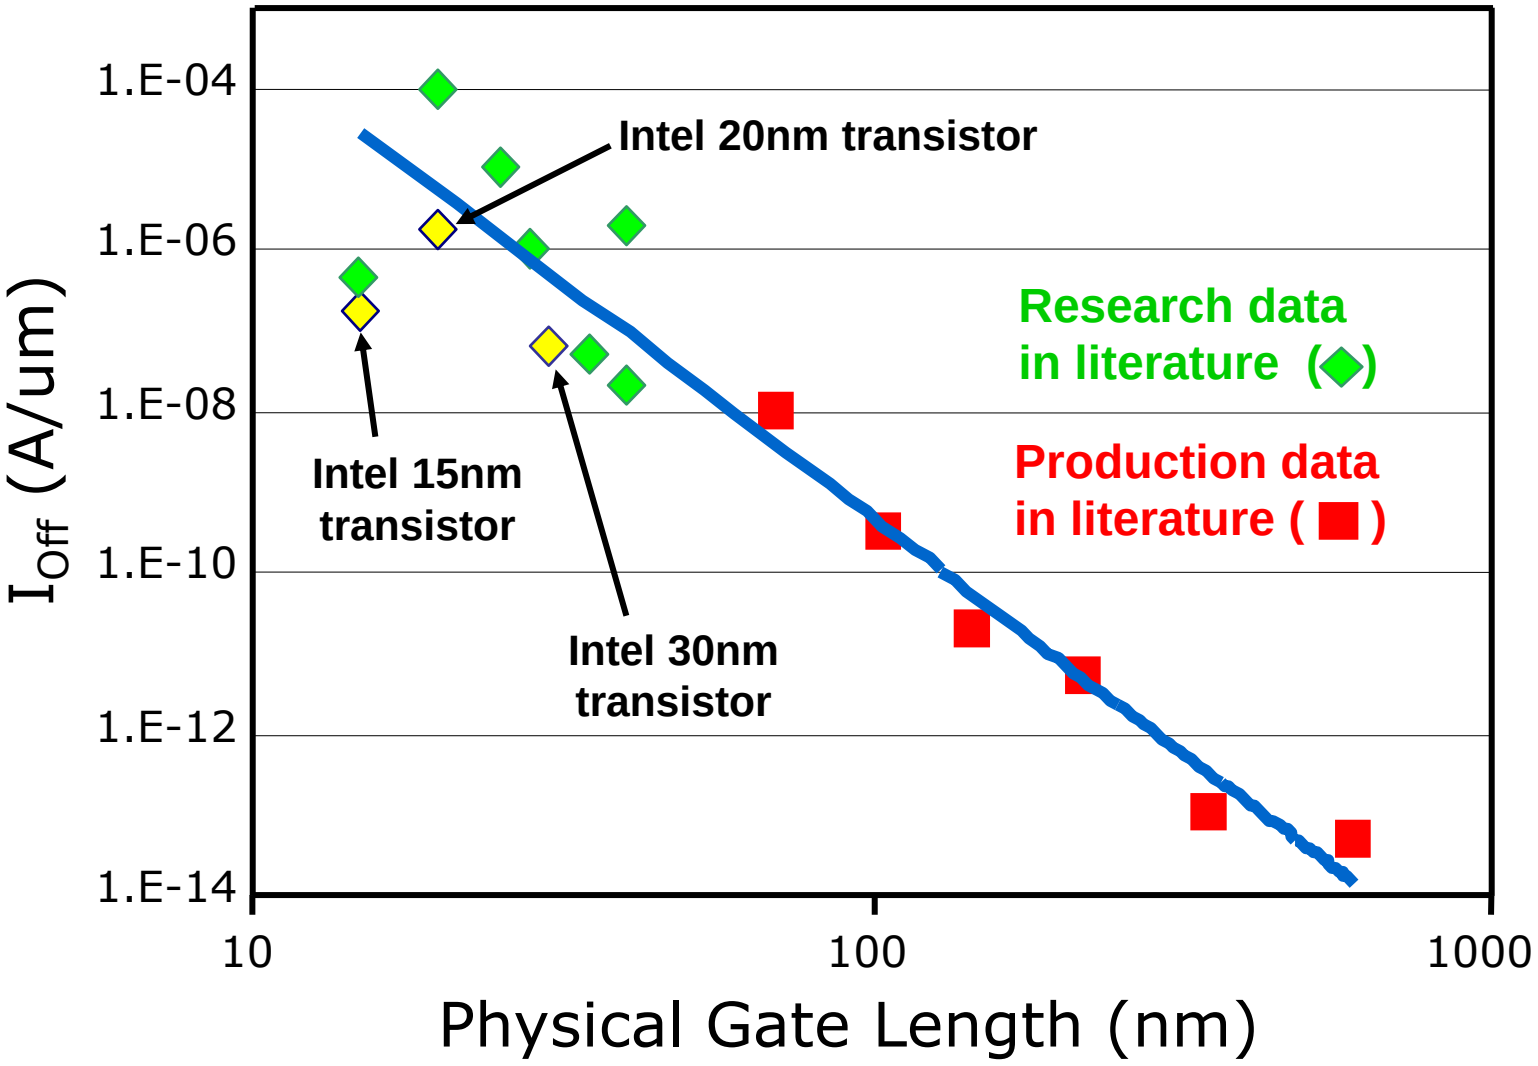

Leakage Reduction Techniques

Body Bias

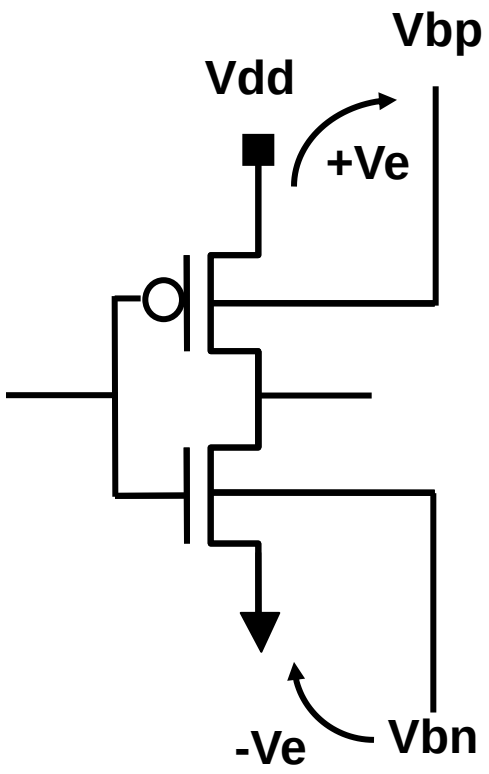

2-3X reduction

Stack Effect

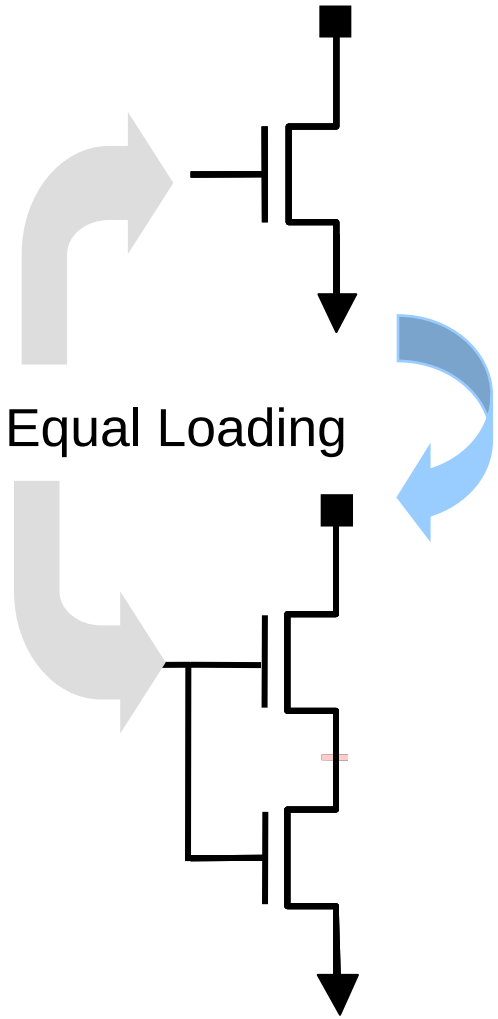

Sleep Transistor

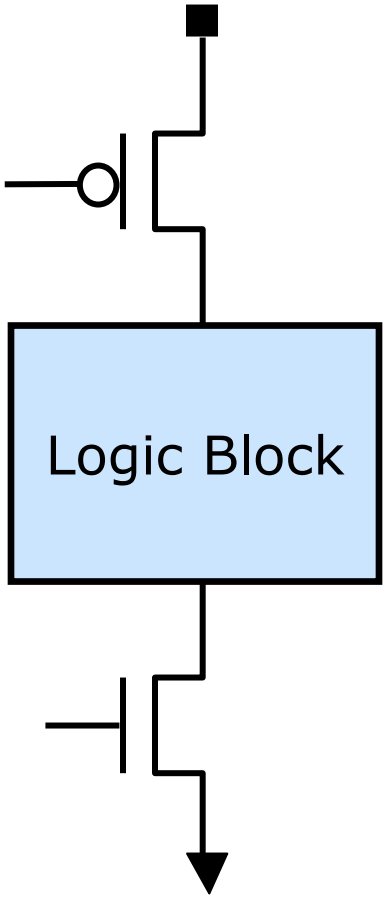

2-1000X reduction

Leakage is a Strong Function of Voltage

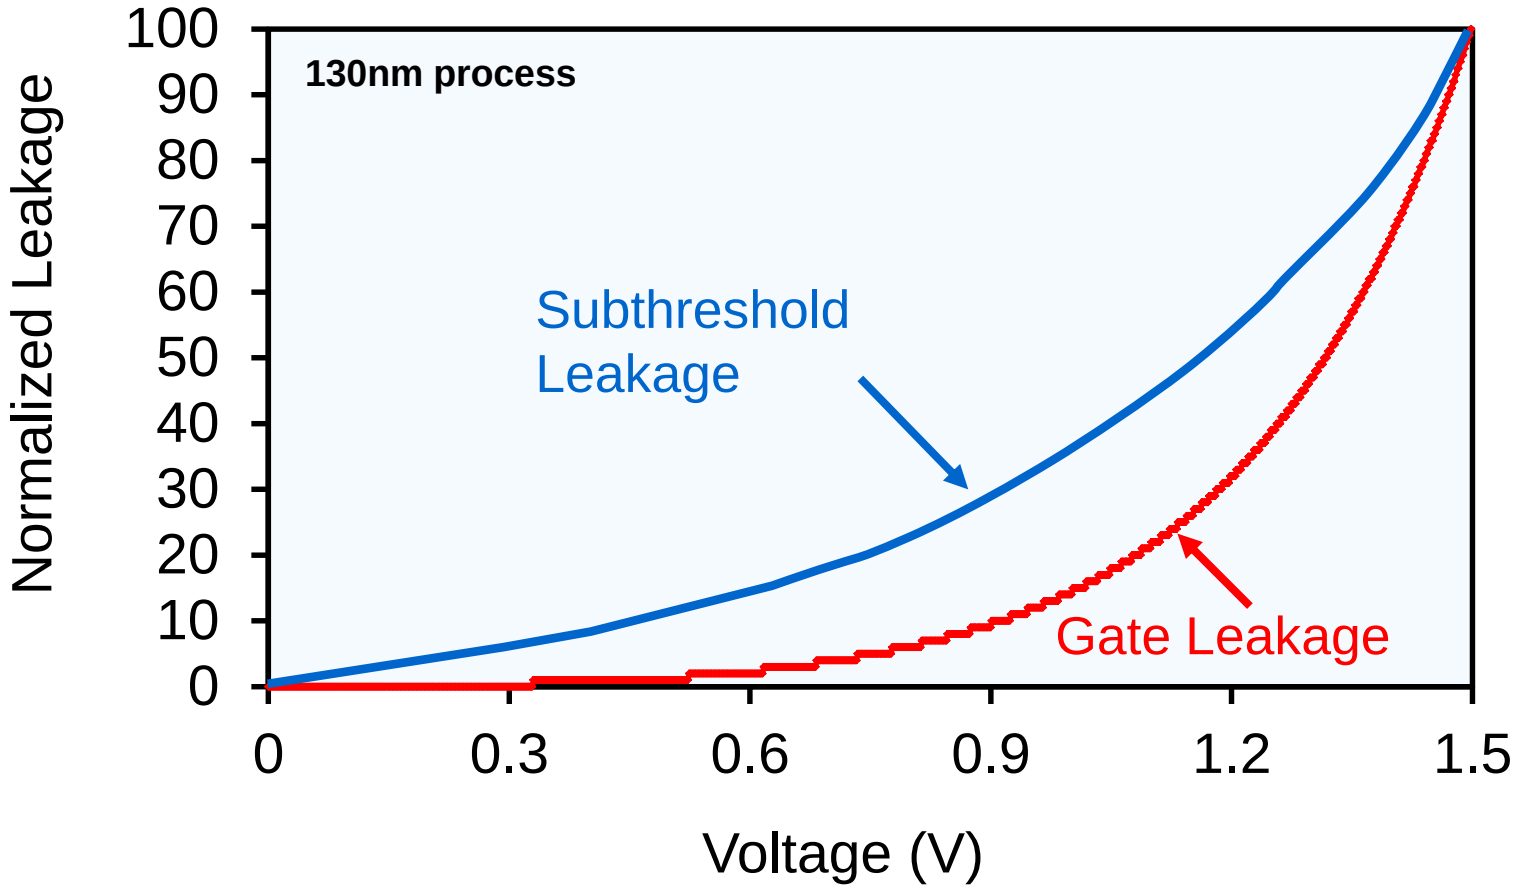

Sub-threshold and gate leakage reduce with lower supply voltage

Cache Sleep and Shut-off Modes

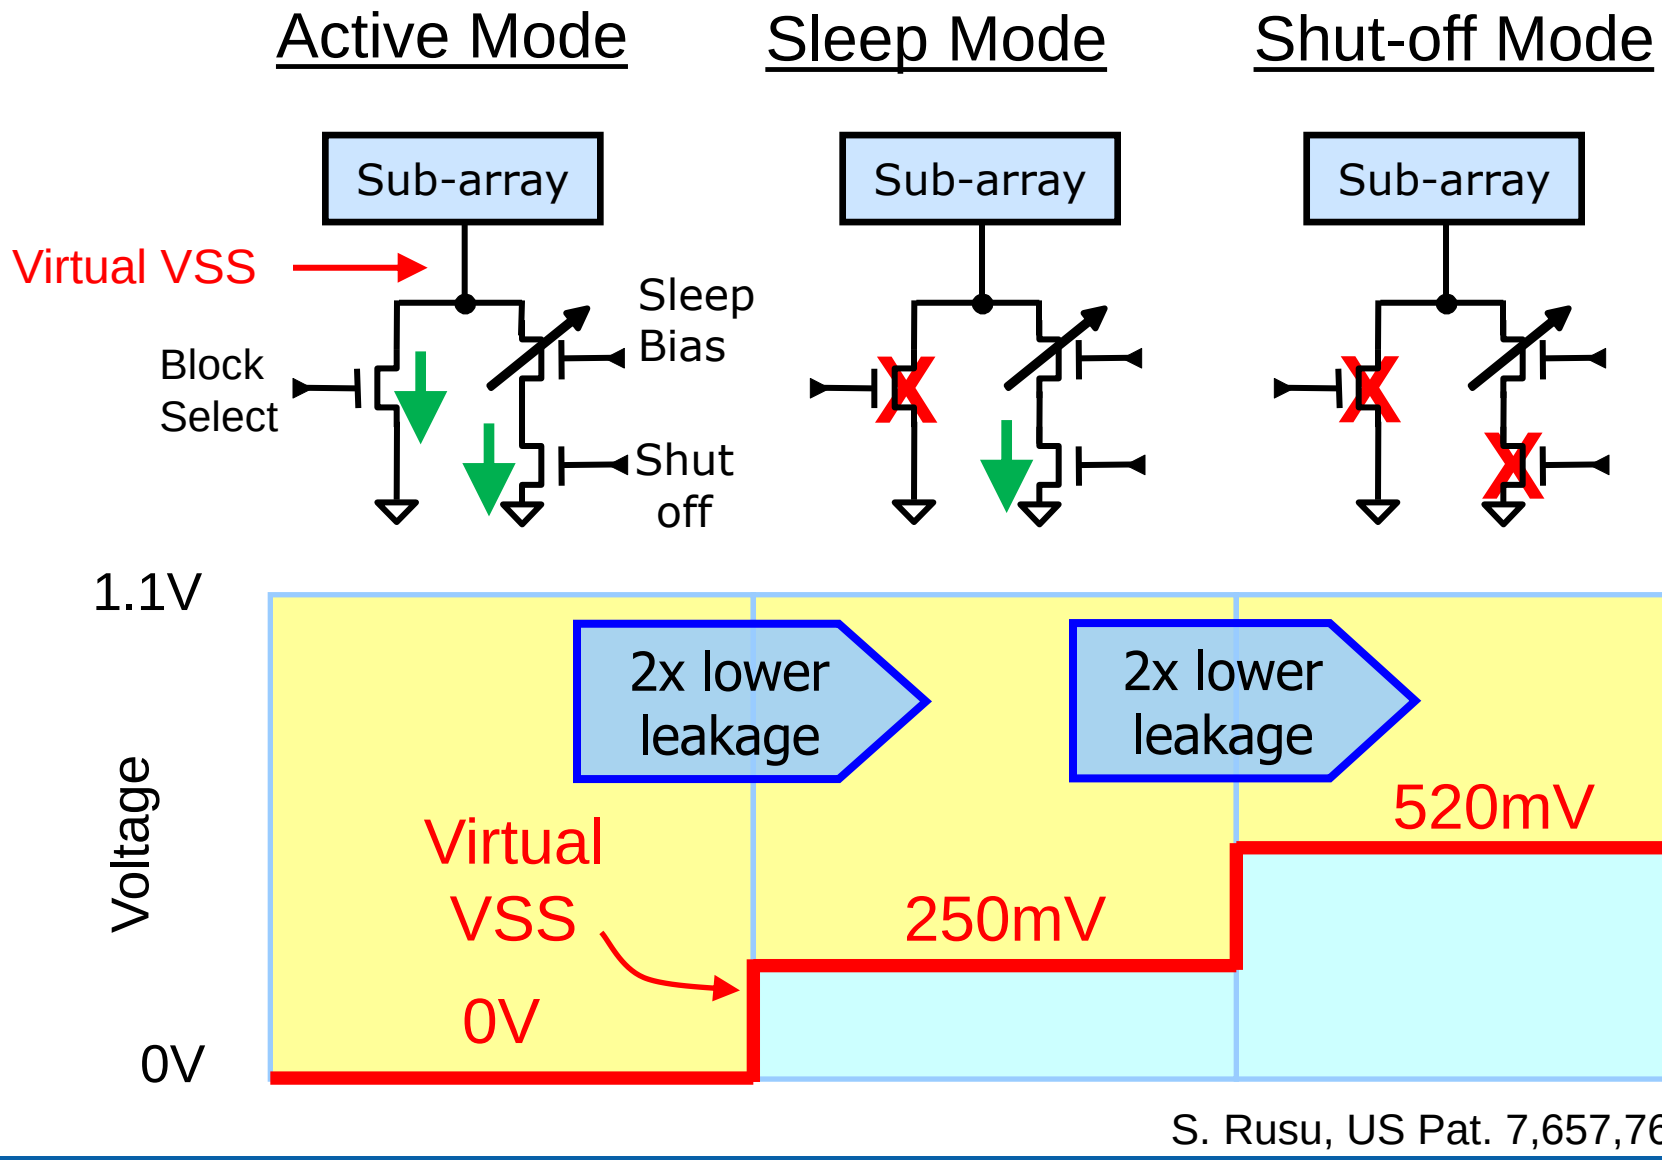

S. Rusu, US Pat. 7,657,767

Leakage Shut-off Infrared Images

16MB part

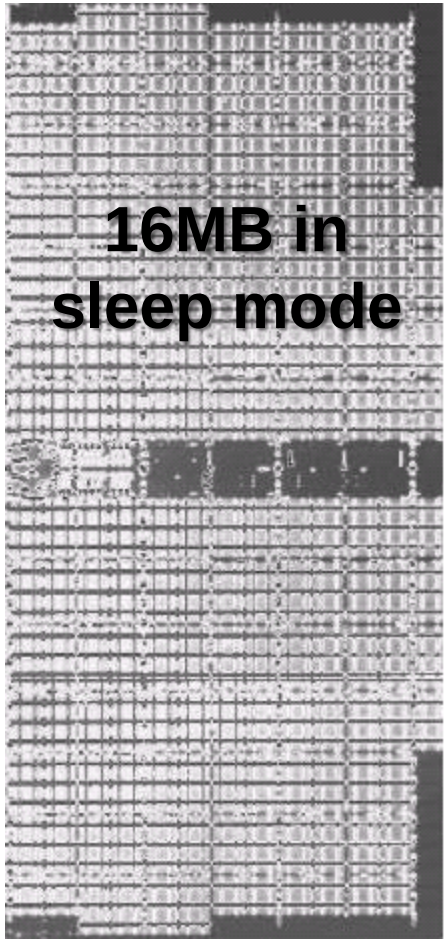

8MB part

4MB part

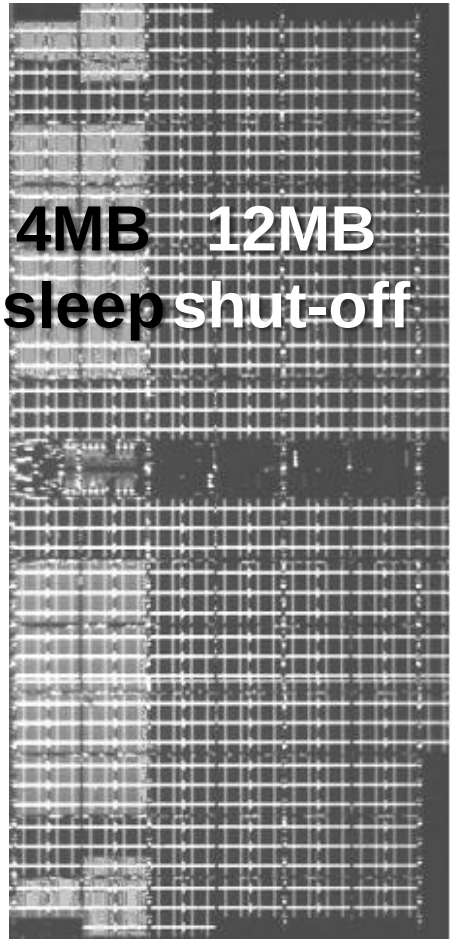

Leakage reduction ► 3W (8MB) 5W (4MB)

Cache Dynamic Shut-off

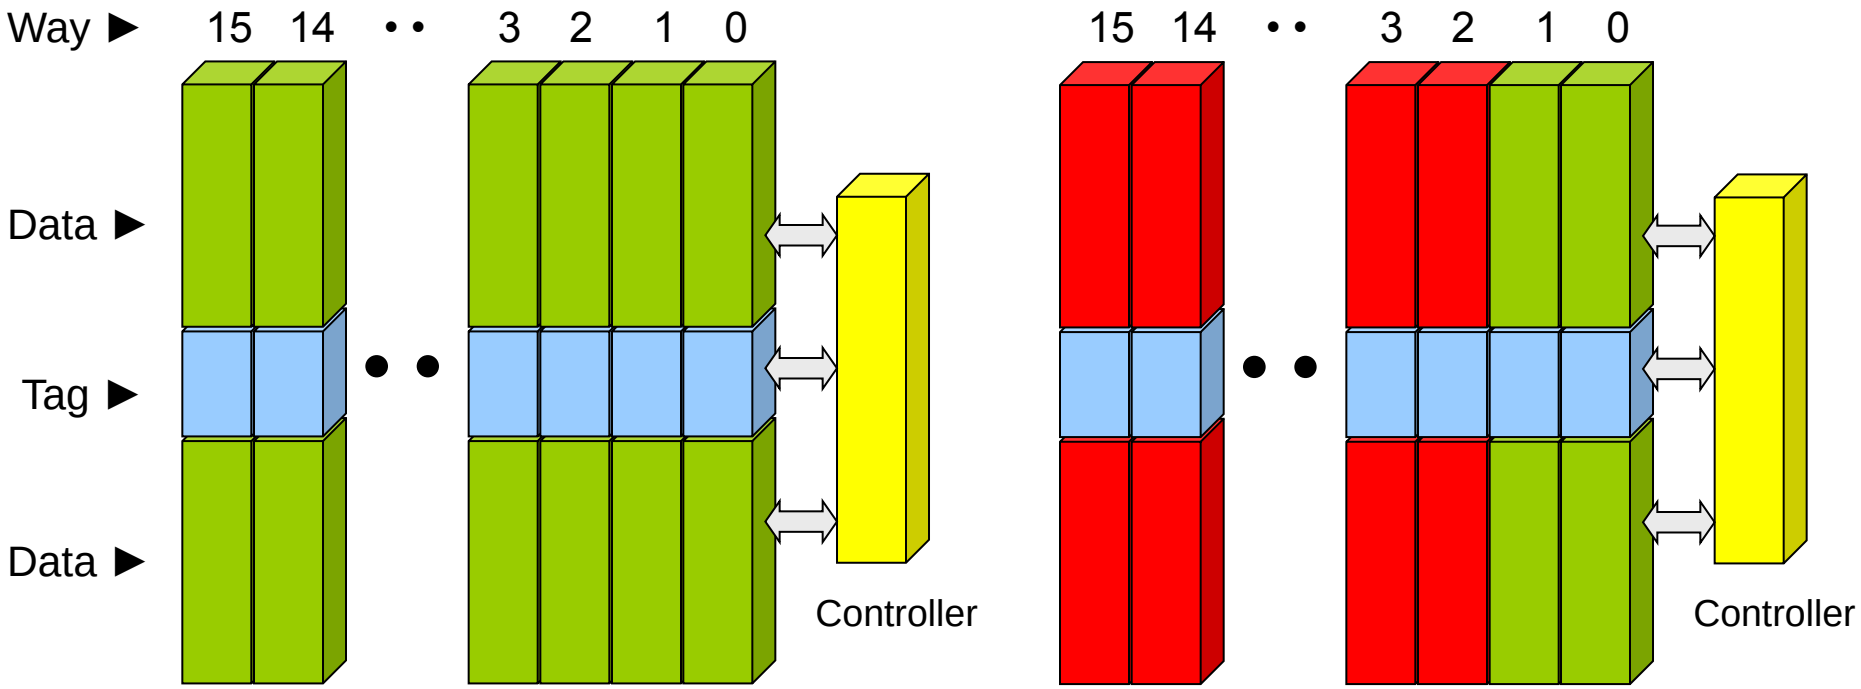

Normal Operation

In the full-load state, all 16 ways are enabled (green)

Cache-by-Demand Operation

Under idle or low-load states, cache ways are dynamically flushed out and put in shut-off mode (red)

Cache Leakage Management

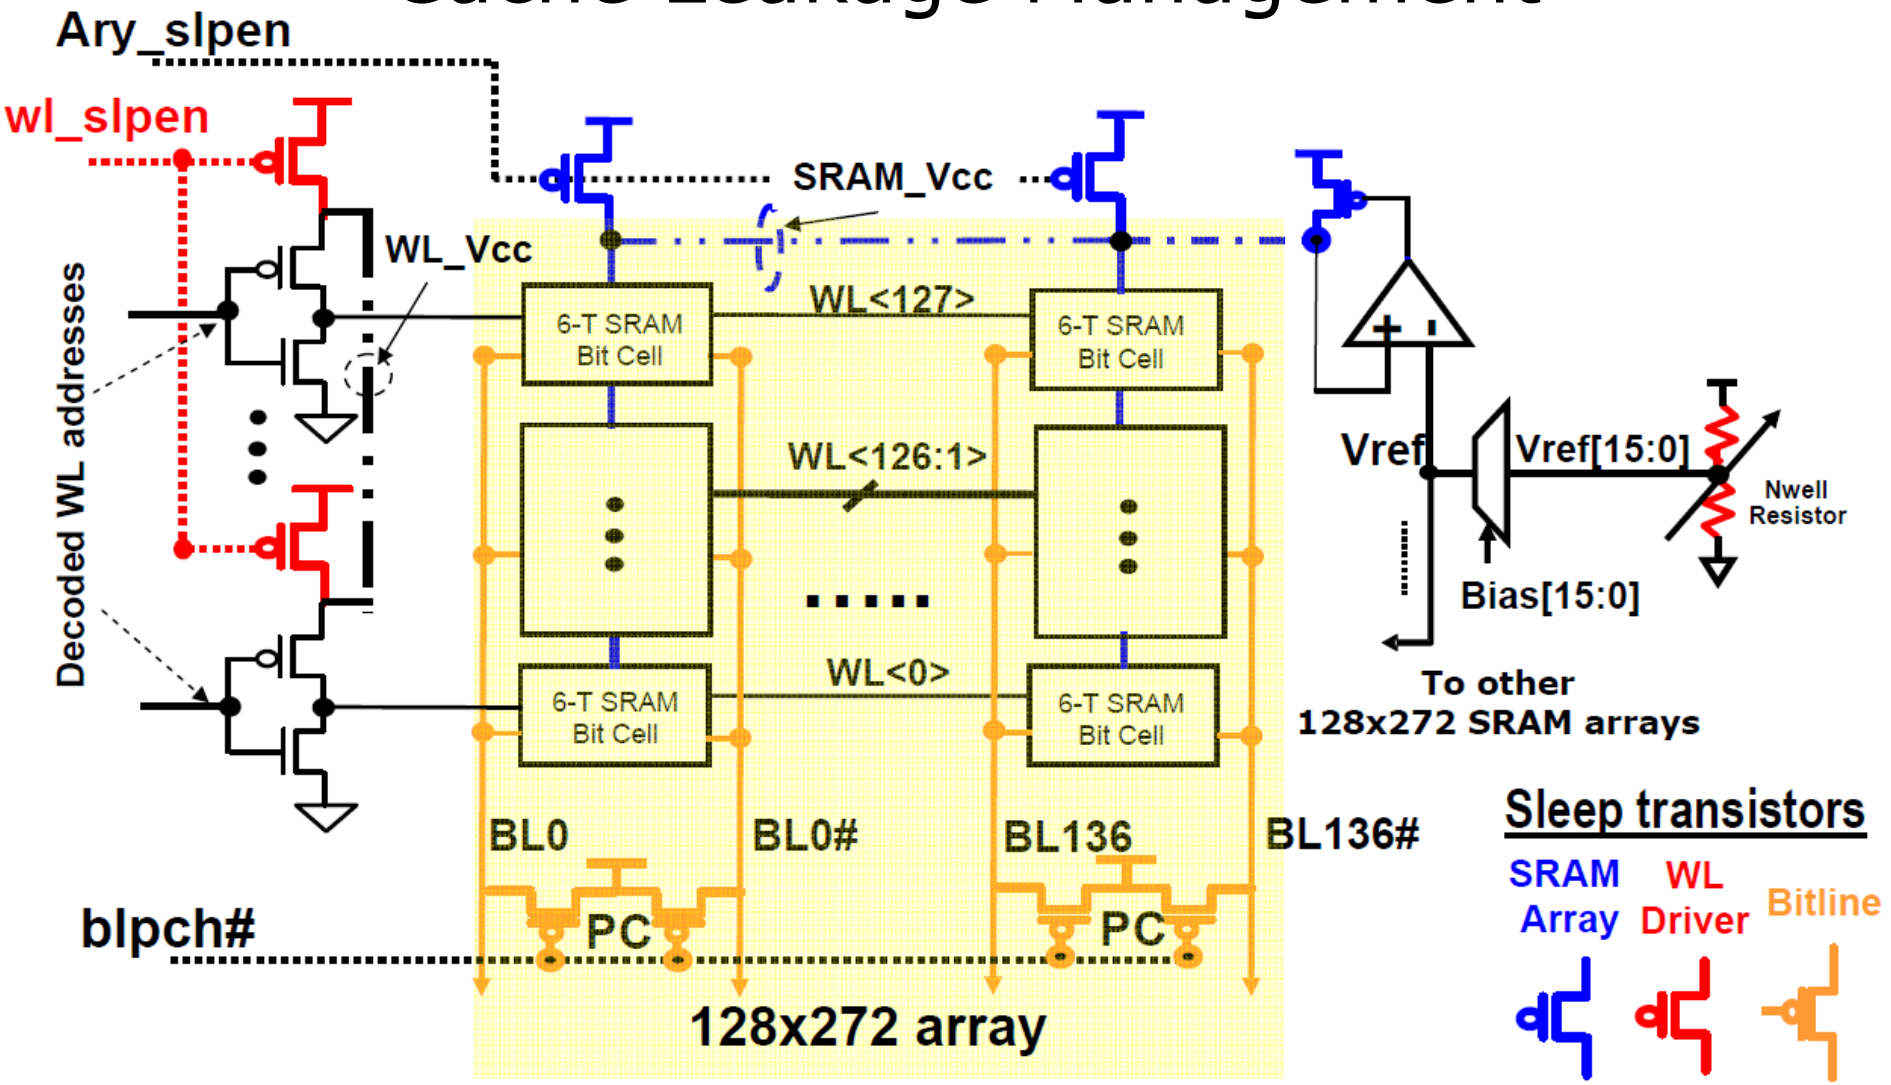

Three PMOS sleep transistor groups for sub-array leakage reduction

Y. Wang, ISSCC 2009

Cache Leakage Reduction Benefit

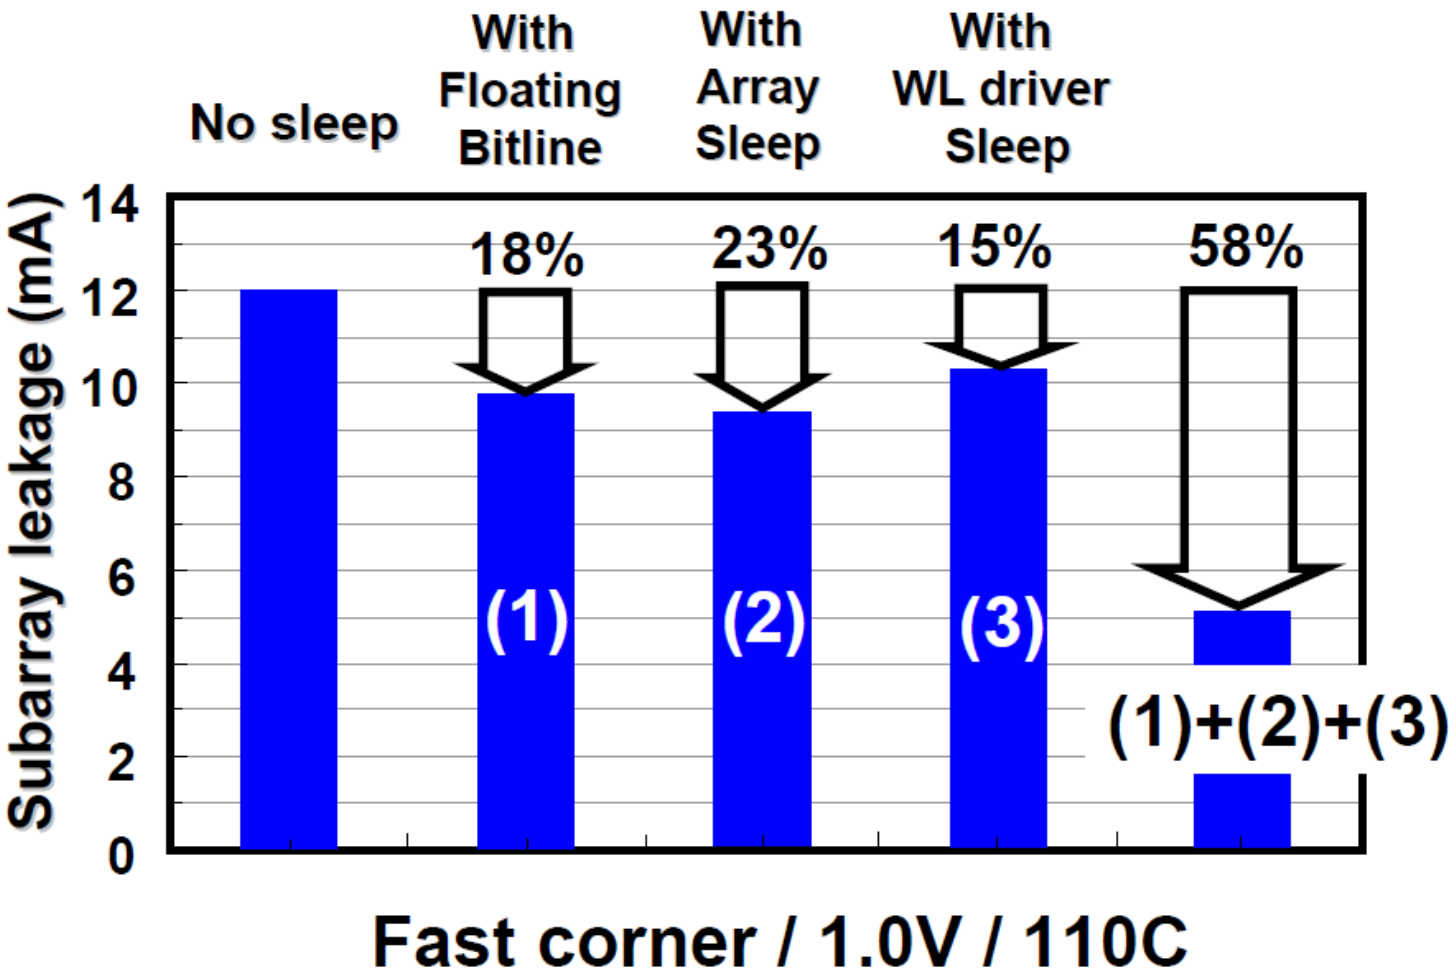

Leakage management circuit reduces sub-array leakage by 58%

Leakage Mitigation: Long-Le Transistors

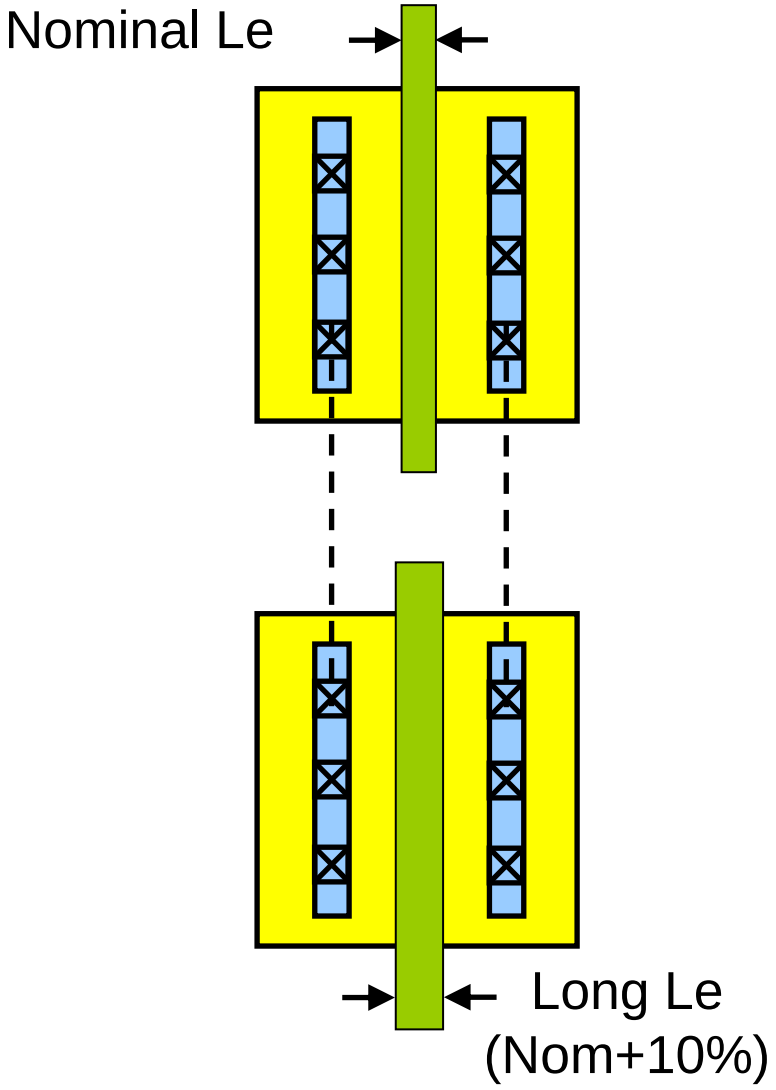

- All transistors can be either nominal or long-Le
- Most library cells are available in both flavors
- Long-Le transistors are ~10% slower, but have 3x lower leakage
- All paths with timing slack use long-Le transistors
- Initial design uses only long channel devices

S. Rusu, ISSCC 2006

Long-Le Transistors Usage Map

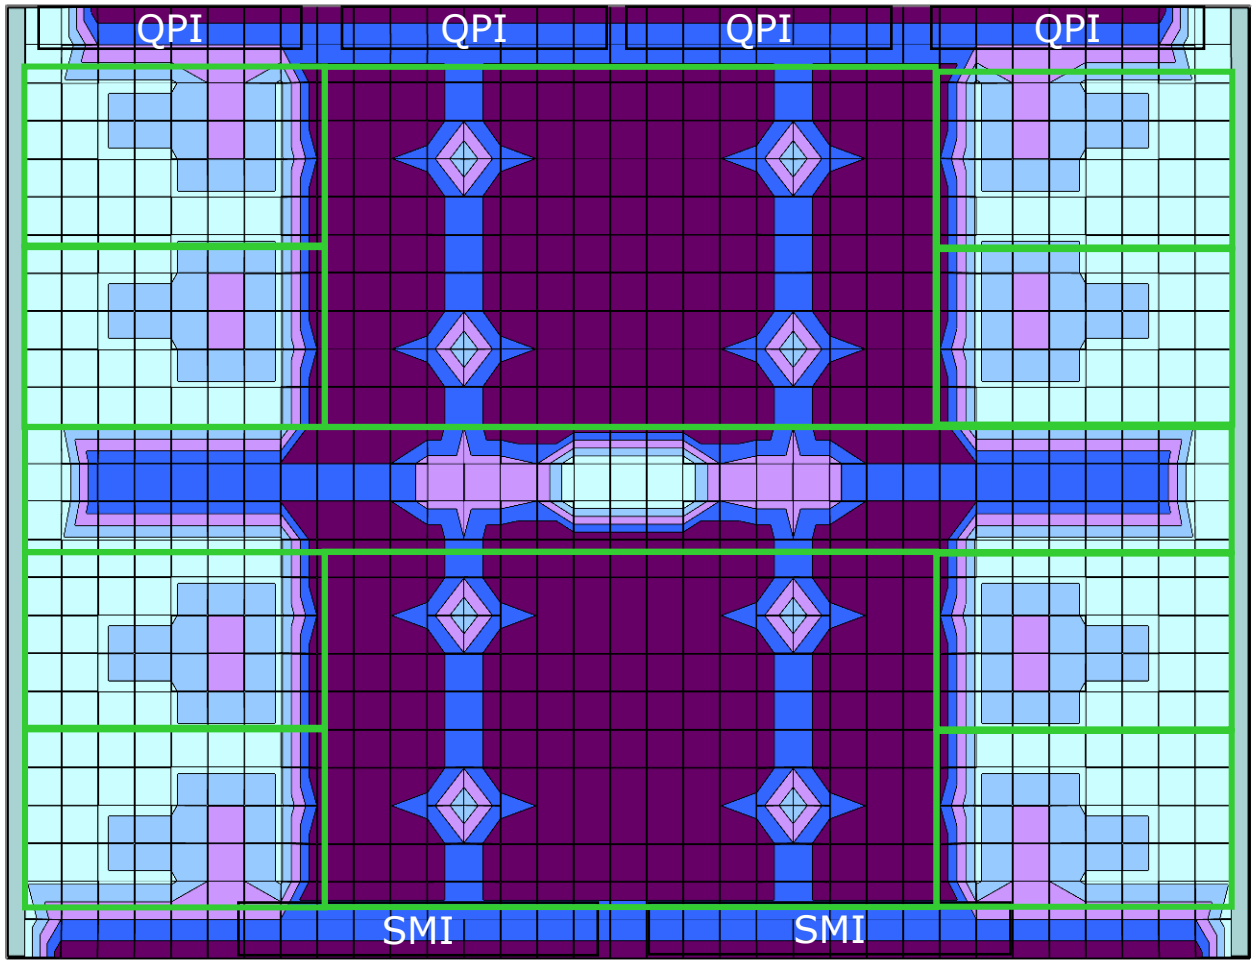

Massive long-channel usage in uncore reduces leakage

Power & Leakage Breakdown

Nehalem-EX 45nm example

Power Breakdown

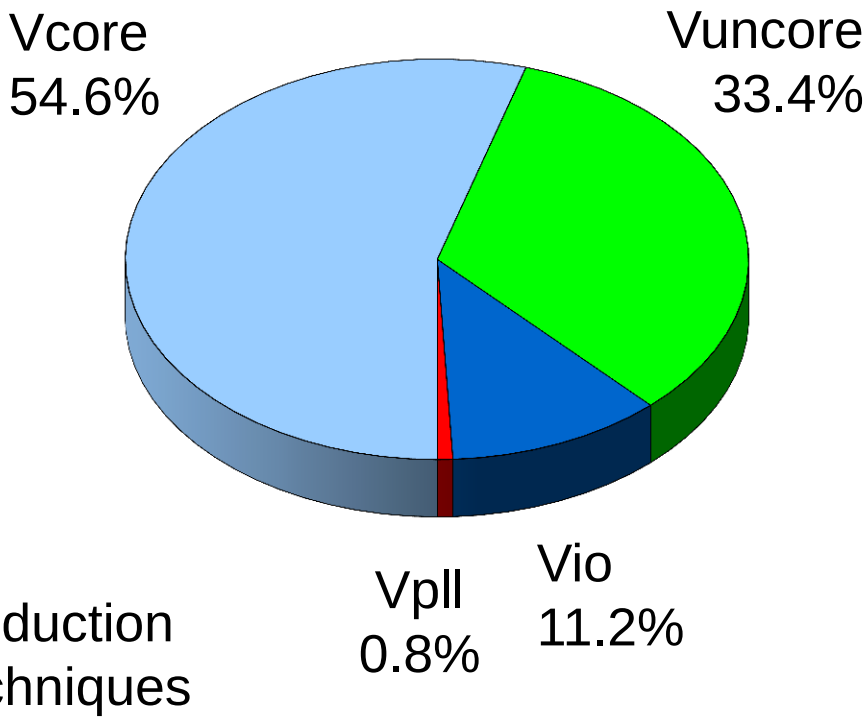

Leakage Breakdown

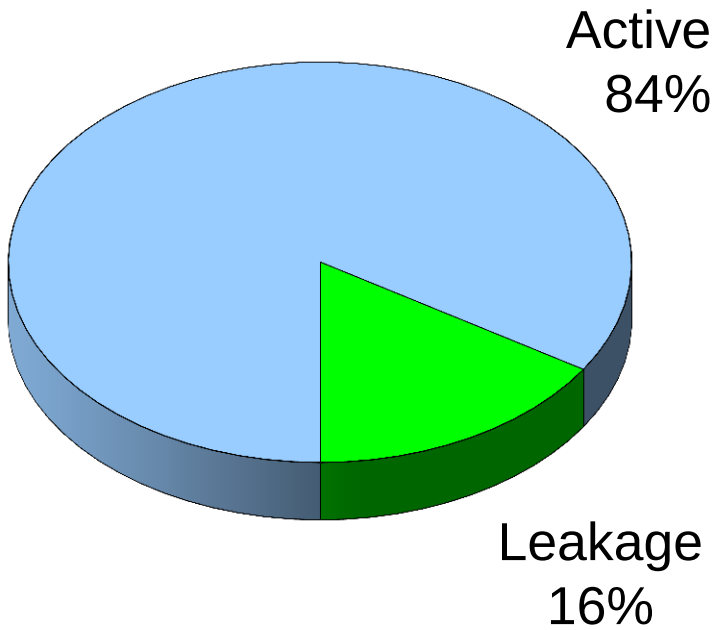

Reduction techniques

Clock gating
Run uncore at 0.9V

Long channel device usage:
58% cores, 85% uncore

S. Rusu, ISSCC 2009

Core and Cache Recovery Example

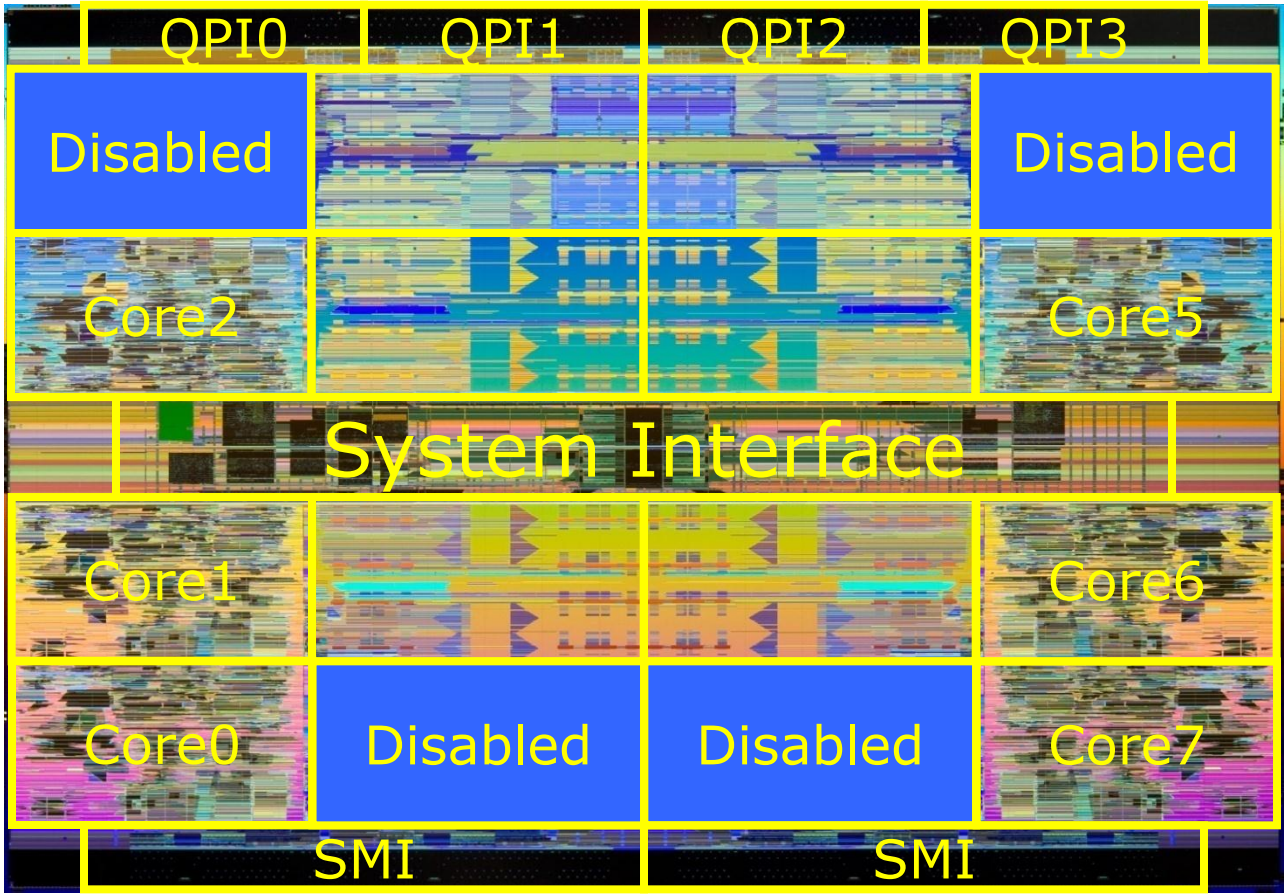

Defective core and cache slices can be disabled in horizontal pairs

S. Rusu, ISSCC 2009

Minimize Leakage in Disabled Blocks

- Disabled cores ► Power gated

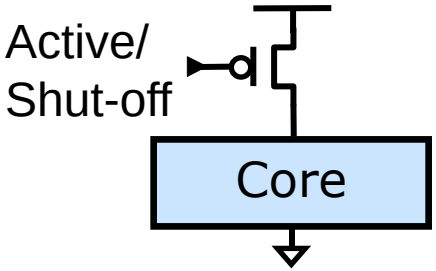

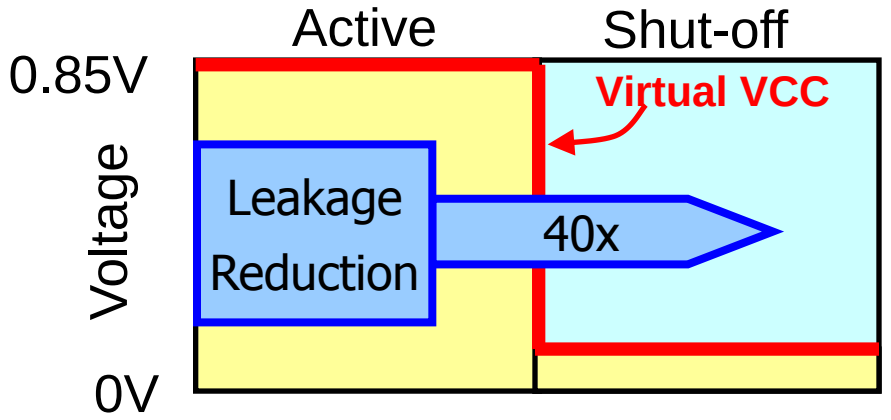

- Disabled cache slices ► All major arrays in shut-off

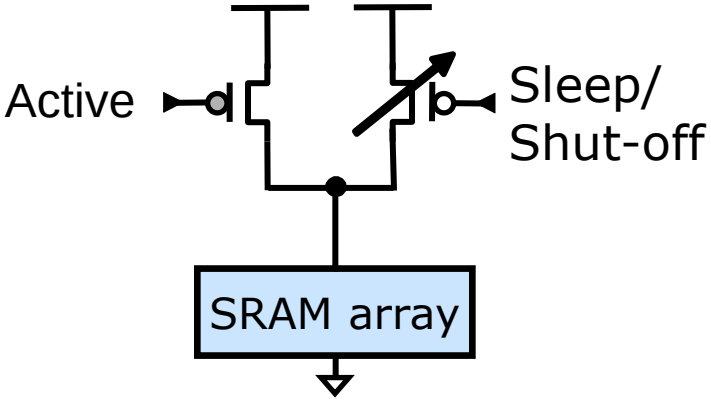

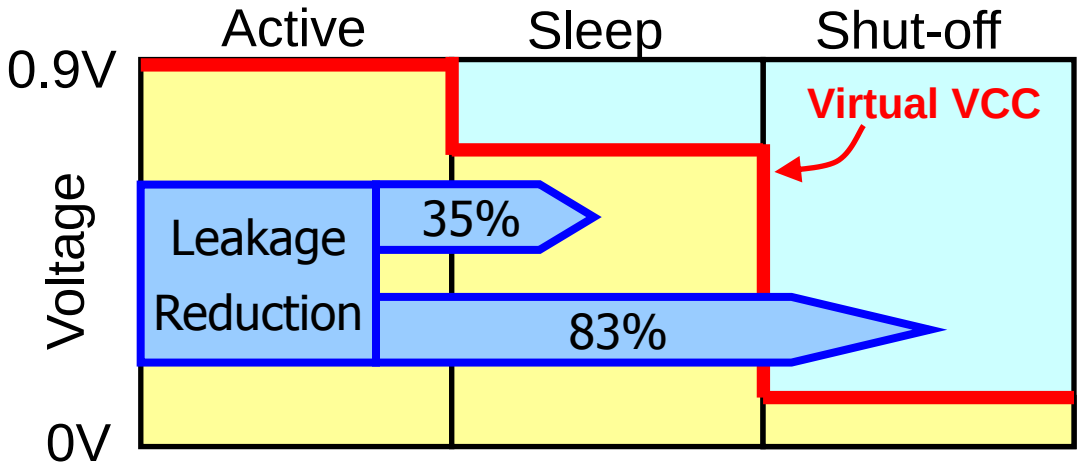

S. Rusu, ISSCC 2009

Core/Cache Recovery – Infrared Image

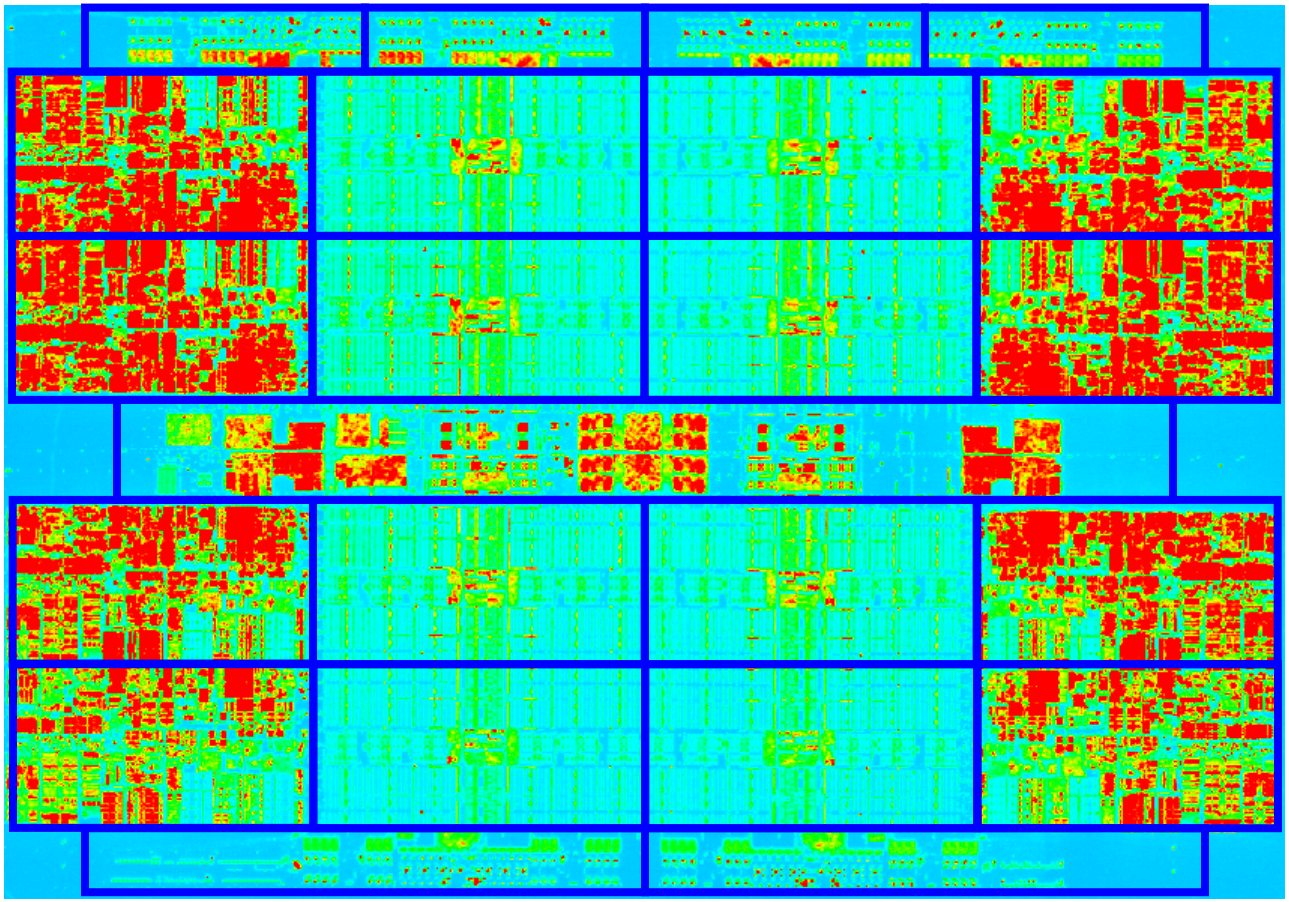

All cores and cache slices are enabled

S. Rusu, ISSCC 2009

Core/Cache Recovery – Infrared Image

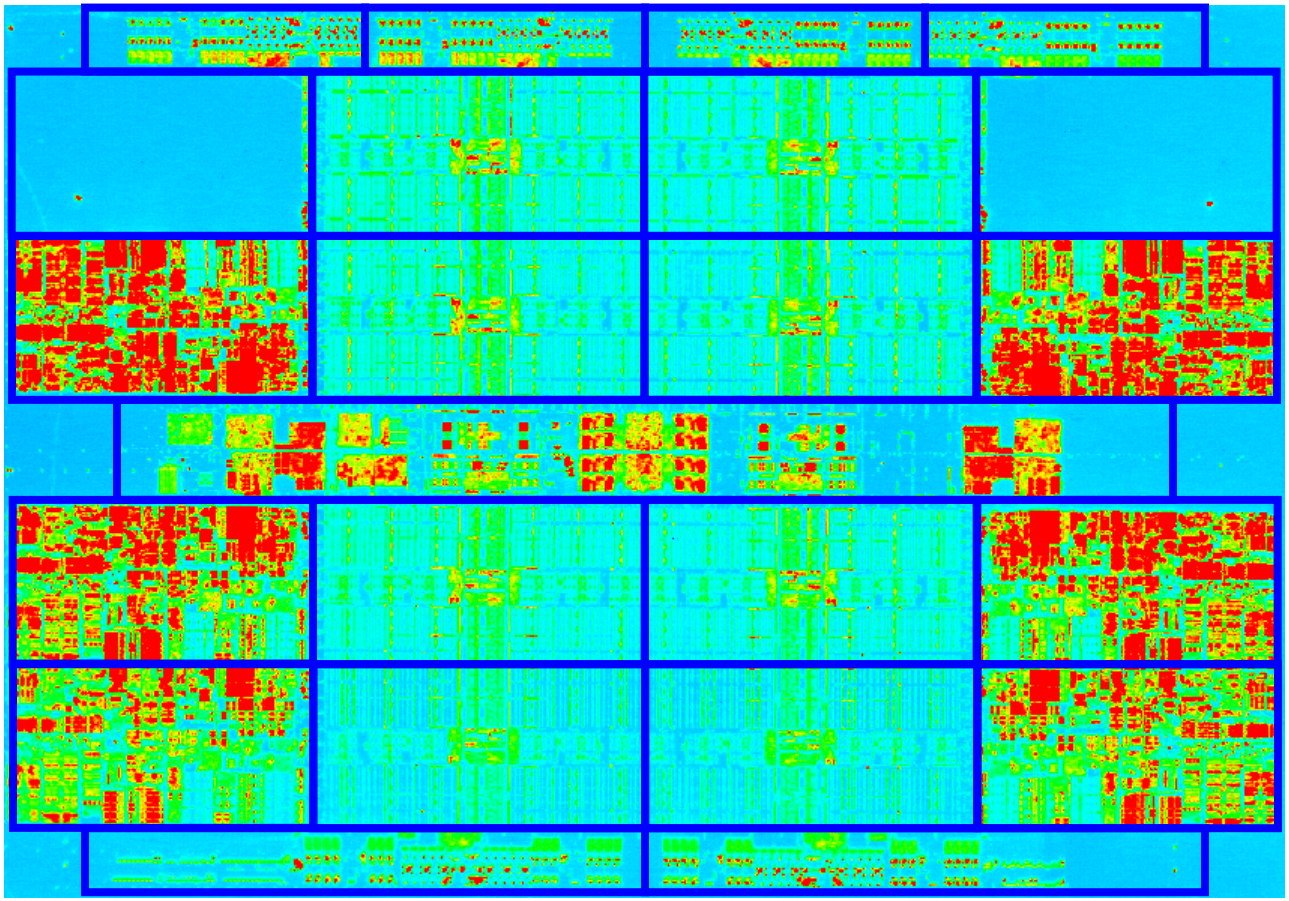

Shut-off 2 cores (top row) and 2 cache slices (bottom row)

Disabled blocks are clock and power gated

S. Rusu, ISSCC 2009

Agenda

- Microprocessor Design Trends
- Process Technology Directions
- Active Power Management
- Leakage Reduction Techniques
- Packaging and Thermal Modeling
- Future Directions and Summary

Microprocessor Package Evolution

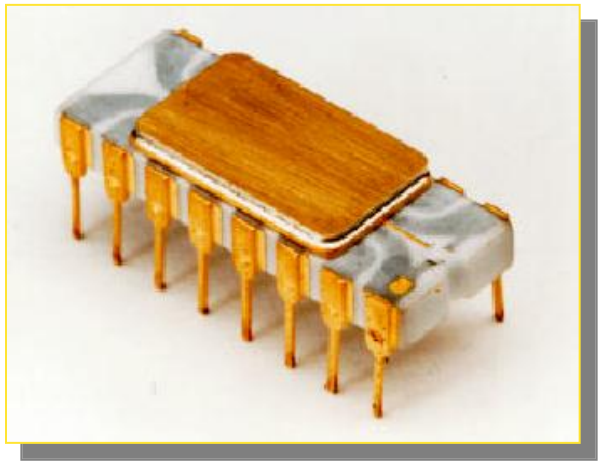

- 1971 – 4004 Processor
  - 16-pin ceramic package
  - Wire bond attach
  - 750 kHz I/O
- 2012 – Xeon® E5 Processor
  - 2011-contact organic package
  - Flip-chip attach
  - 8.0 GHz I/O

Power Density Models

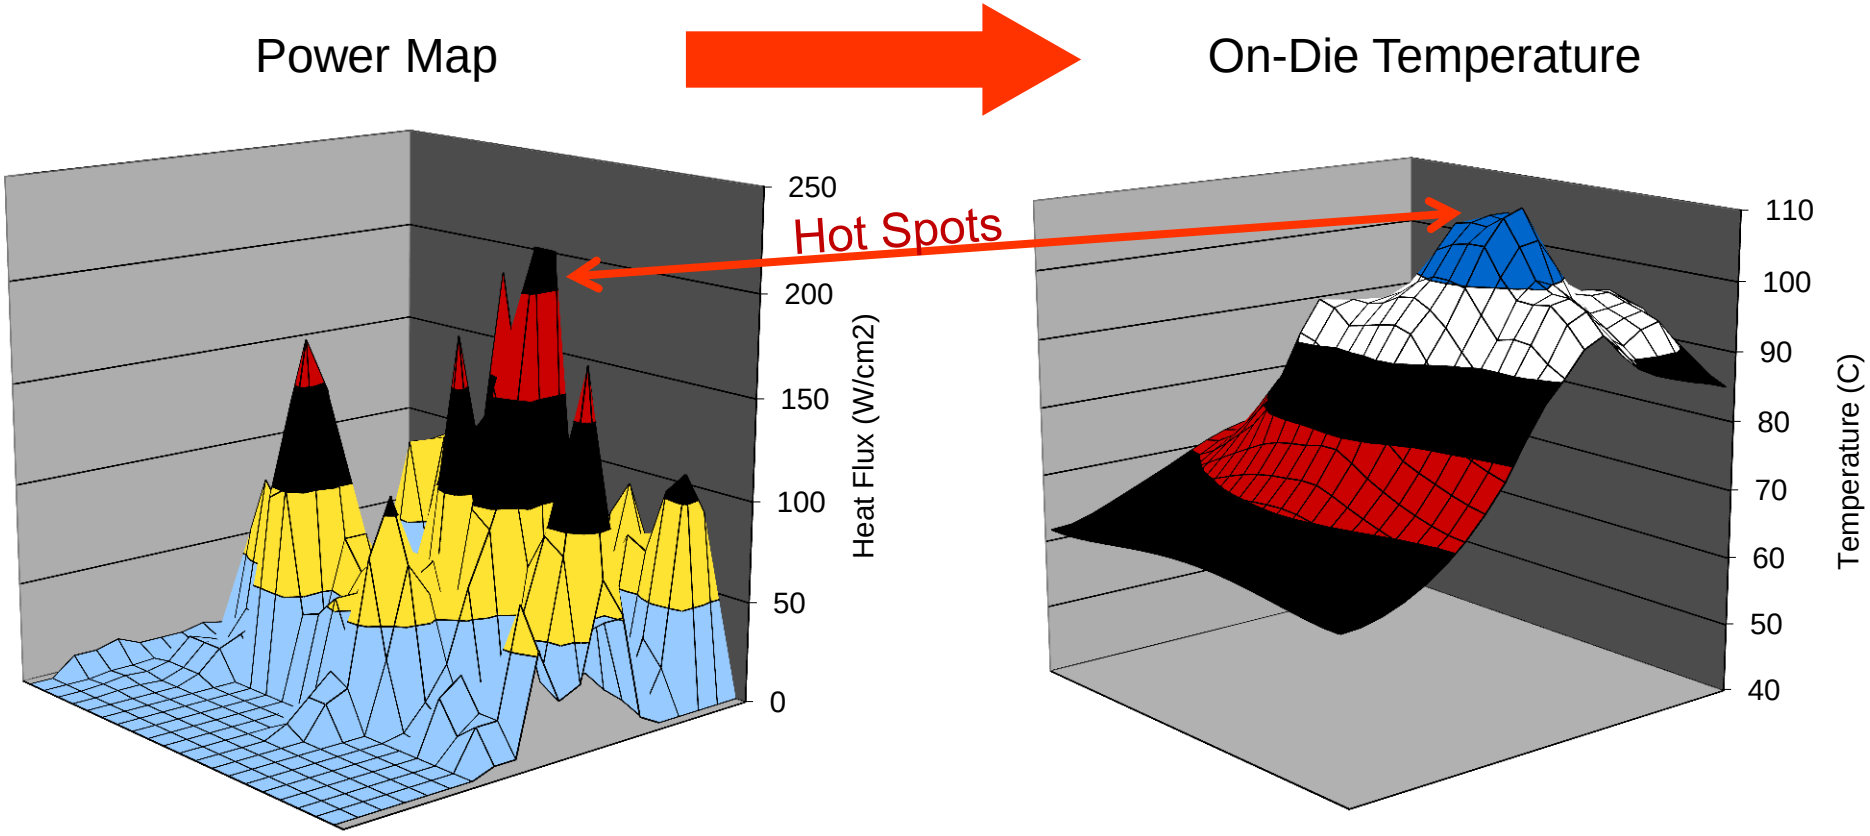

With increasing power density and large on-die caches, detailed, non-uniform power models are required

Thermal Modeling

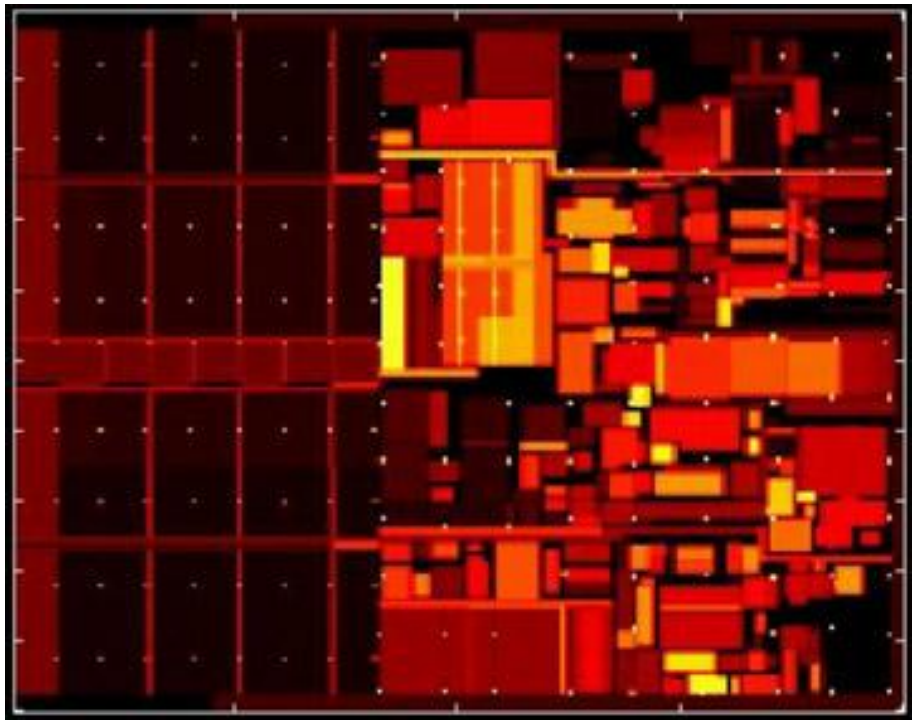

Simulated power density

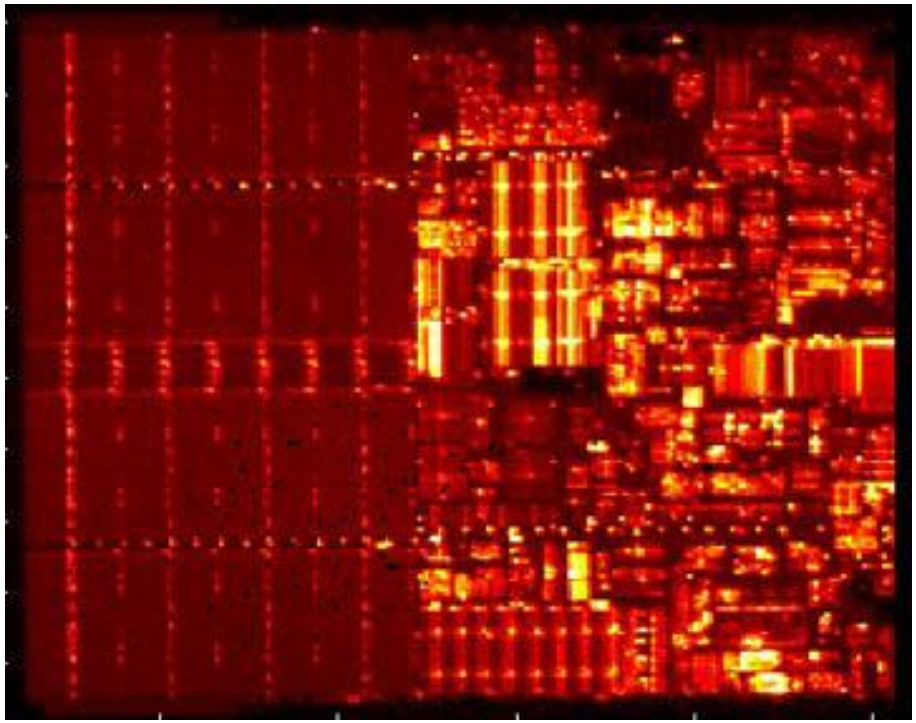

Infrared emission microscope measurement

D. Genossar and N. Shamir "Intel® Pentium® M Processor Power Estimation, Budgeting, Optimization and Validation", Intel Technology Journal 5/2003

Thermal Sensors

- Multiple temperature sensors
  - One in each core hot spot
  - One in the die center
- Temperature information is available through PECI bus for system fan management

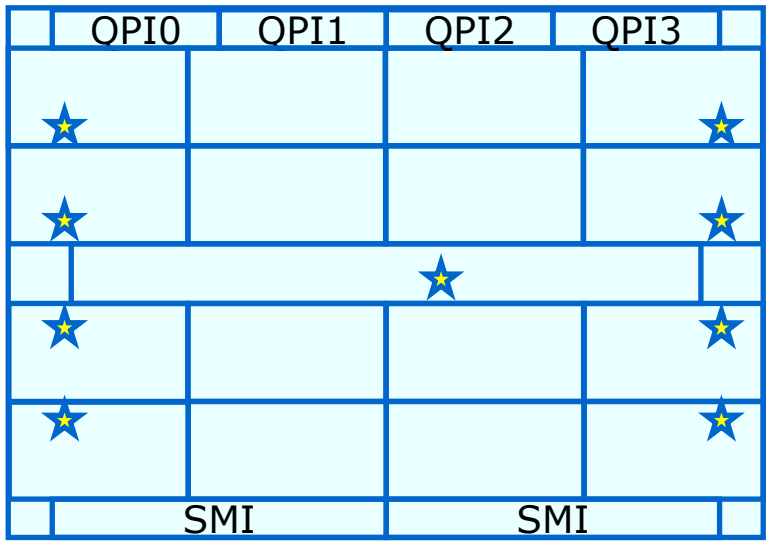

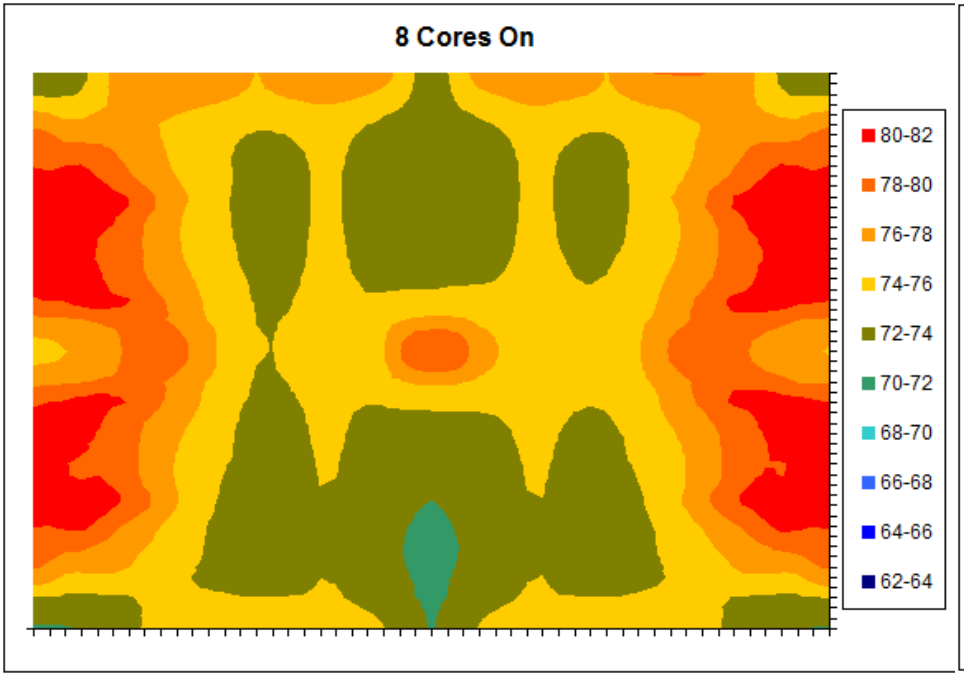

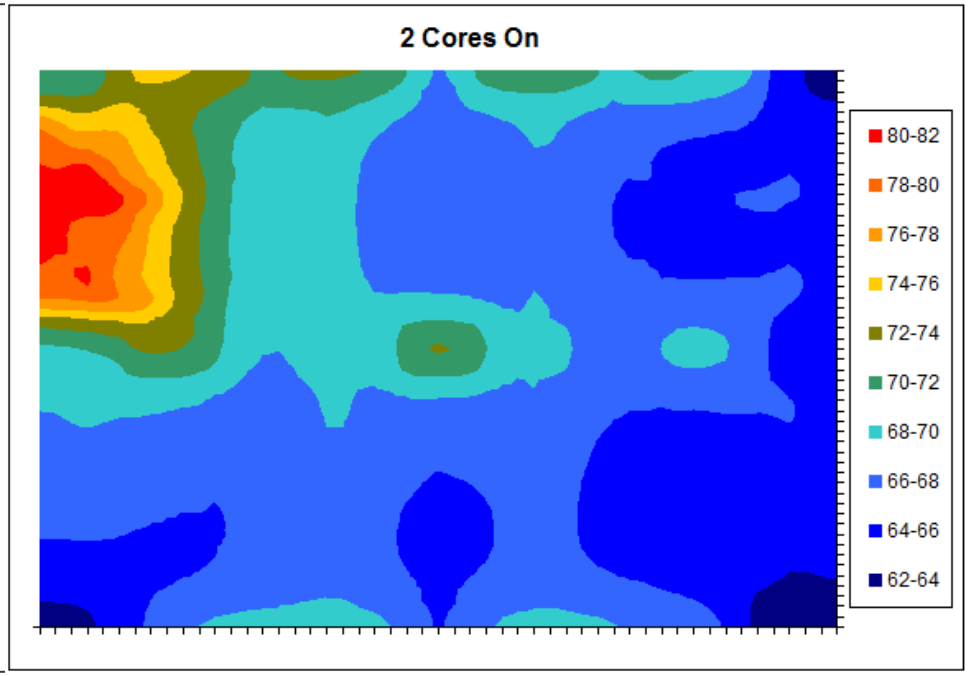

Power Management Unit

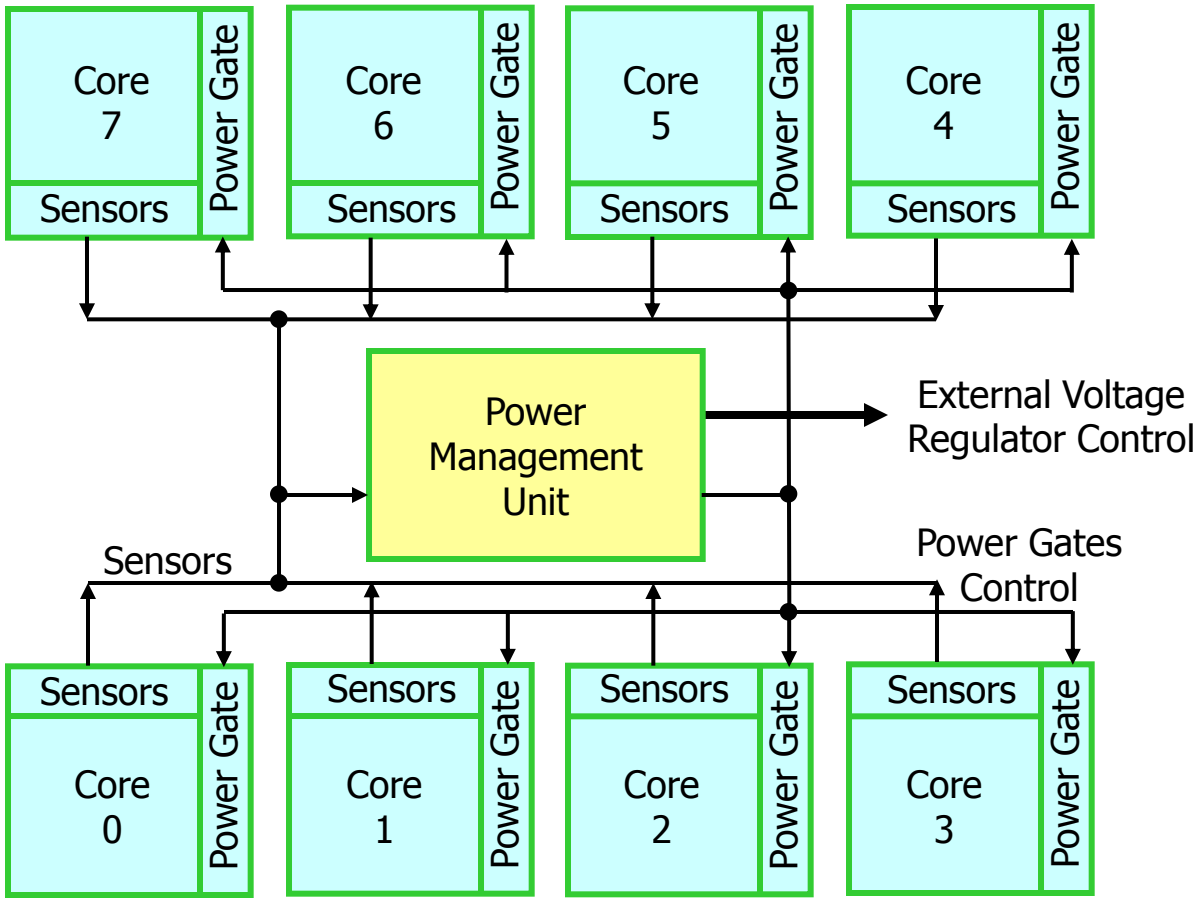

PMU controls processor voltage and frequency based on compute loading and thermal data

Agenda

- Microprocessor Design Trends
- Process Technology Directions
- Active Power Management
- Leakage Reduction Techniques
- Packaging and Thermal Modeling
- Future Directions and Summary

Future Directions

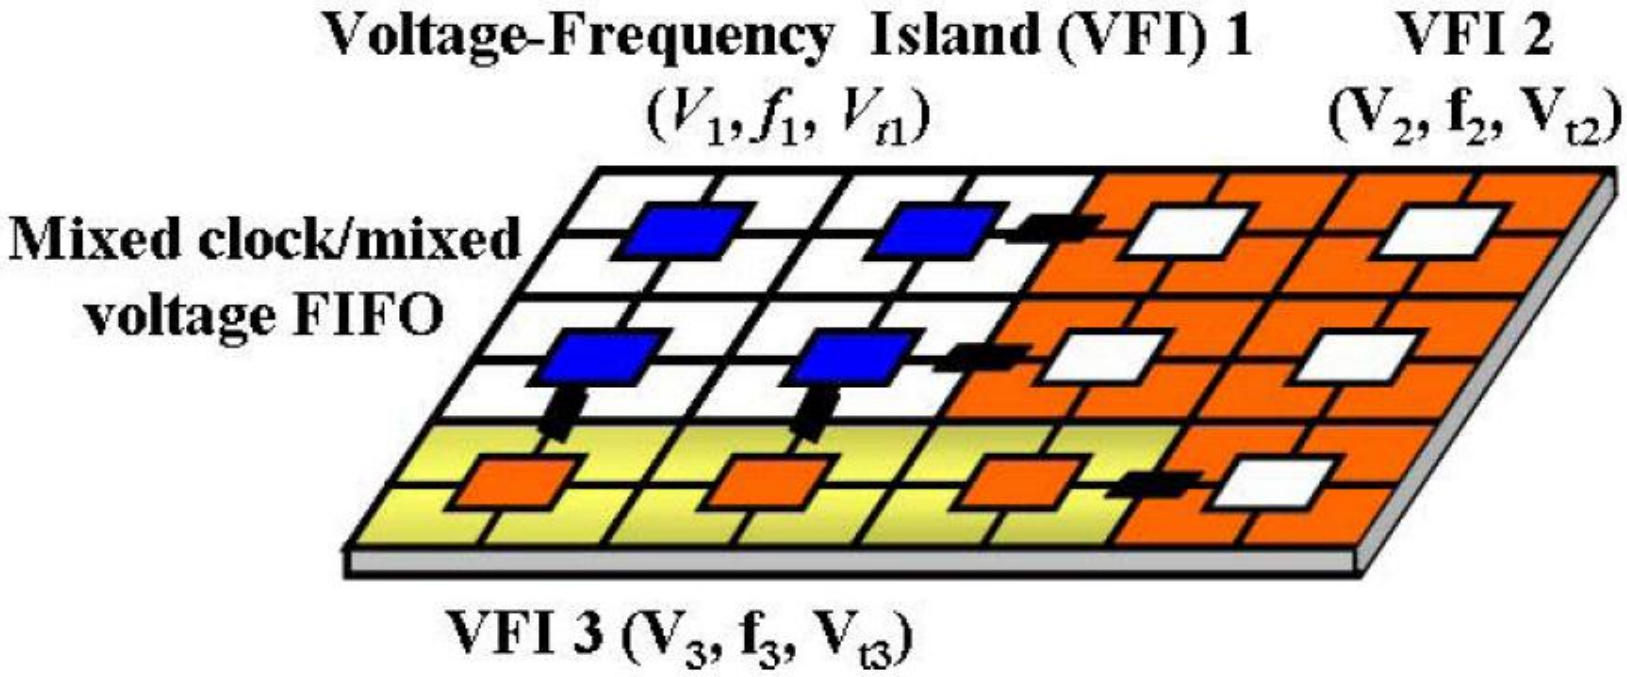

- 2D mesh network with multiple Voltage / Frequency islands
- Communication across islands achieved through FIFOs

Ogras (CMU), DAC 2007

Fine Grain Power Management

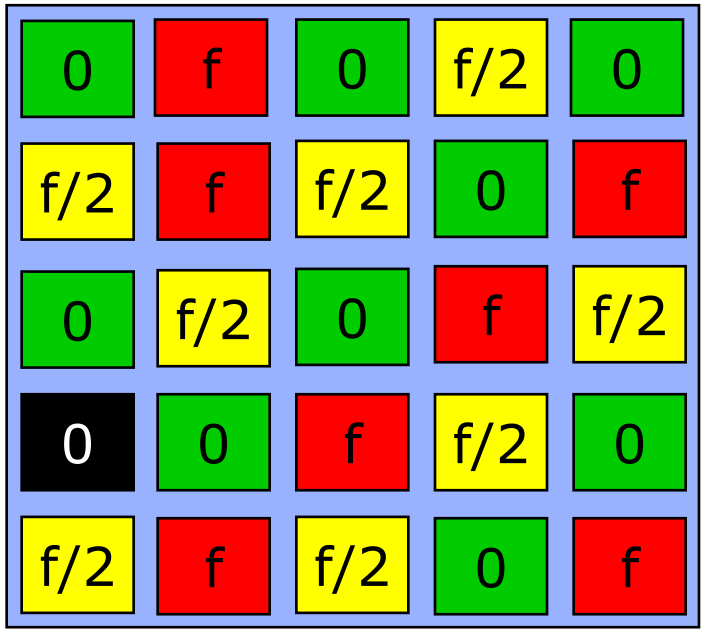

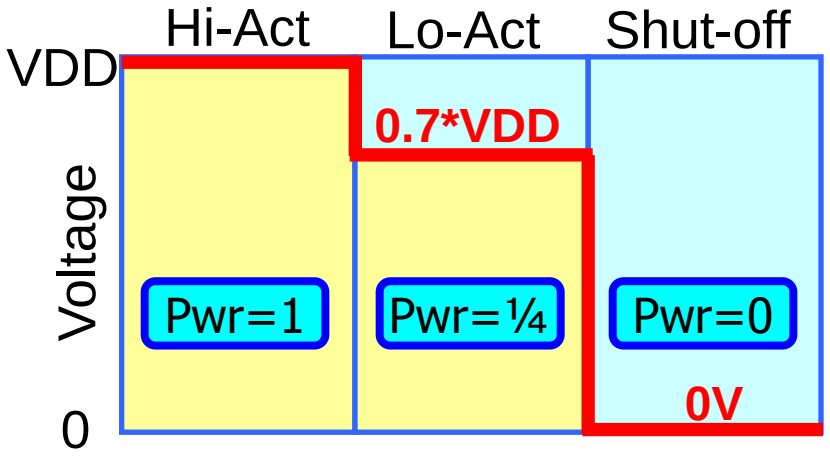

25-core processor example:

f — Cores with critical tasks
Freq = f, at Vdd
TPT = 1, Power = 1

f/2 — Non-critical cores
Freq = f/2, at 0.7*Vdd
TPT = 0.5, Power = 0.25

0 — Temporarily shut down
TPT = 0, Power = 0

0 — Permanently disabled
TPT = 0, Power = 0

Summary

- Moore's Law has fueled the worldwide technology revolution for over 40 years and will continue for at least another decade
  - 0.7x transistor dimension scaling every two years
  - Tri-gate devices provide significant benefits
- Continued microprocessor performance improvement depends on our ability to manage active power and leakage
  - Clock and power gate un-used or disabled blocks
  - Multiple voltage and clock domains
  - Dynamic voltage and frequency adjustment
- Core and cache recovery enables multiple product options
  - Disabled cores and cache slices are clock and power gated