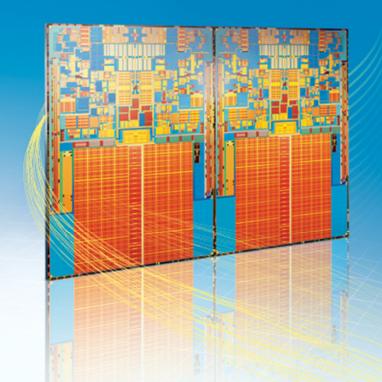


CAD challenges in a mobility-driven era





Acknowledgments

- Chirayu Amin
- Shekhar Borkar
- Steve Burns
- Mike Kishinevsky
- Umit Ogras
- Nikolai Ryzhenko

- Emily Shriver
- Krishnamurthy Soumyanath
- Sriram Vangal
- Jianping Xu
- Raj Yavatkar



It is not the strongest of the species that survives, nor the most intelligent that survives. It is the one that is the most adaptable to change.

- Charles Darwin

October 30, 2009

Caveat

- This is a workshop
- This is a personal perspective
- I have the problems not the solutions!

BACKGROUND & CONTEXT

October 30, 2009

5

Innovation Enabled Technology Pipeline Our Visibility Continues to Go Out ~10 Years

32nm 2009 22nm 2011

14nm 2013

<u>10nm</u>

7nm 2015+ 5nm

Manufacturing

Development

Research





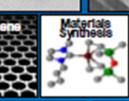




















Future options subject to change

INVESTOR MEETING 20



The End of Scaling is Near?

(Decades of Predictions)

"Optical lithography can't do sub-micron"

'Optical lithography will reach its limits in the range of 0.75-0.50 microns'

"Optical lithography should reach its limits in the 1990-1994 period" "X-ray lithography will be needed below 1 micron"

"Minimum geometries will saturate in the range of 0.3 to 0.5 microns"

"Channel lengths can be reduced to approximately 0.2 microns" "Minimum gate oxide thickness is limited to ~2 nm" "Oxide reliability may limit oxide scaling to 2.2 nm"

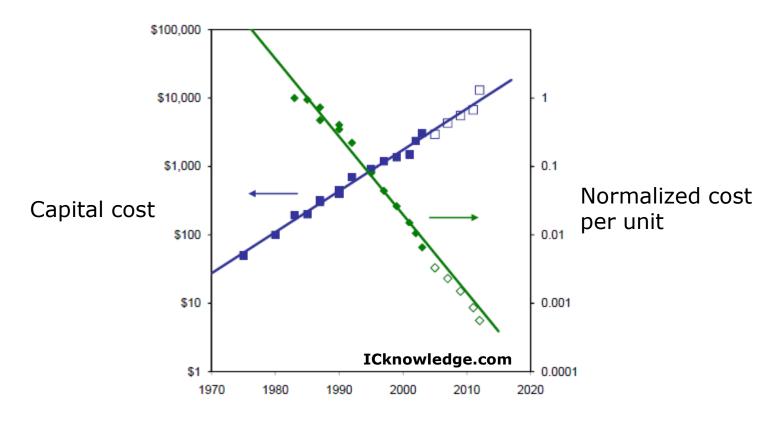
"Copper interconnects will never work" "Plasma etched aluminum will not happen in our lifetime"

"Scaling will end in ~10 years"

1/

INVESTOR MEETING 2012 (intel

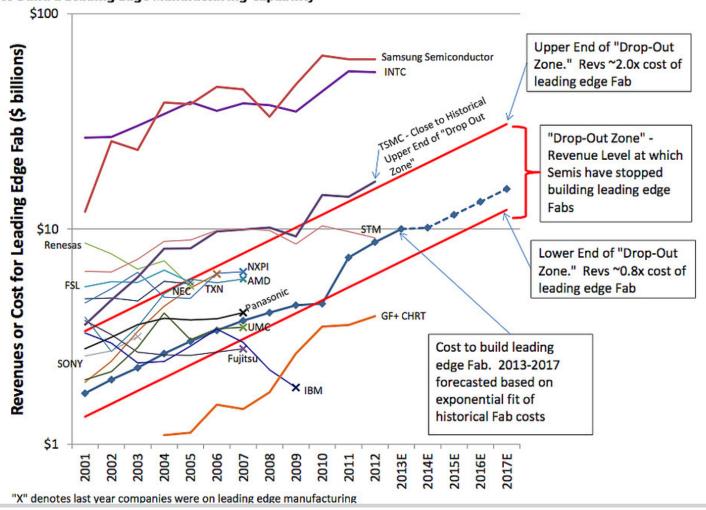
Rock's Law



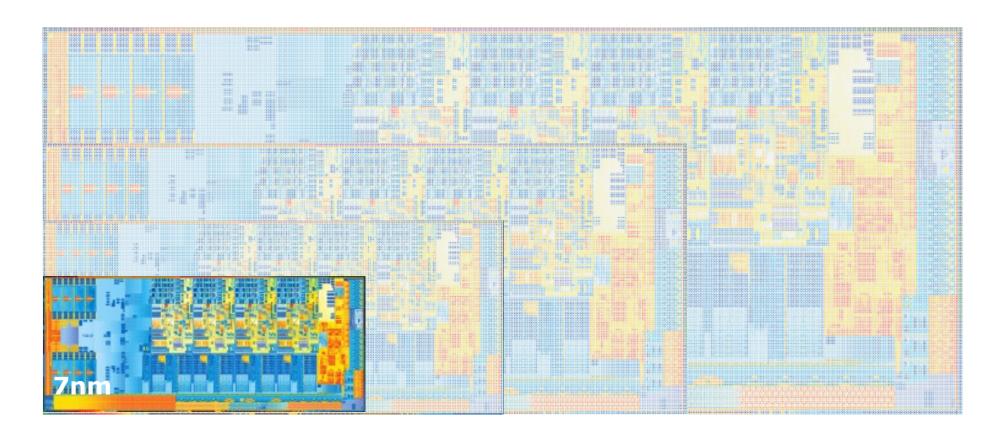
Fab costs double every four years

Revenue for leading edge fab

Chart 8: Semiconductor Manufacturer Revenues vs. Leading Edge "Drop-Out Zone" Represented by 0.8x-to-2.0x the Cost to Build a Leading Edge Manufacturing Capability

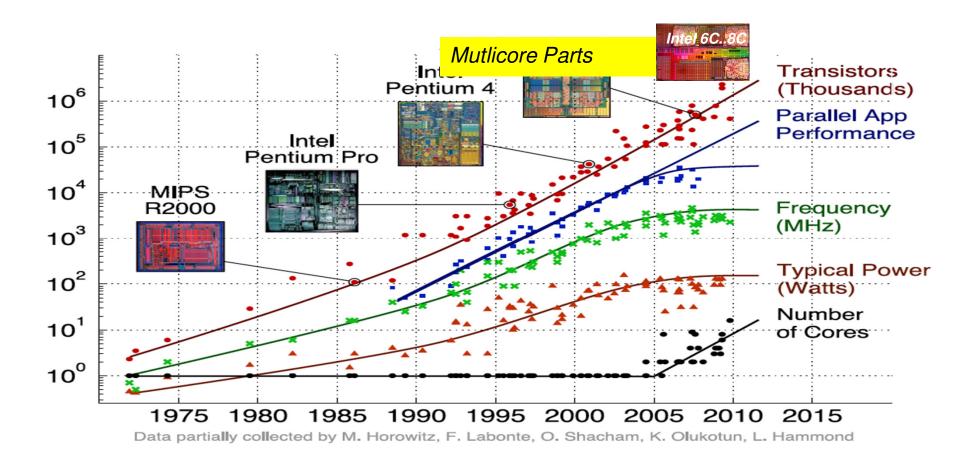


Scaling

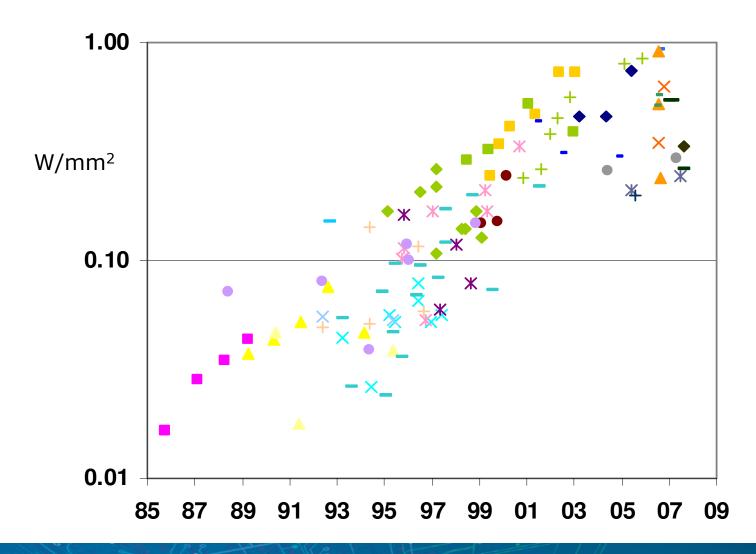


IVB, 160mm², 1.4B transistors

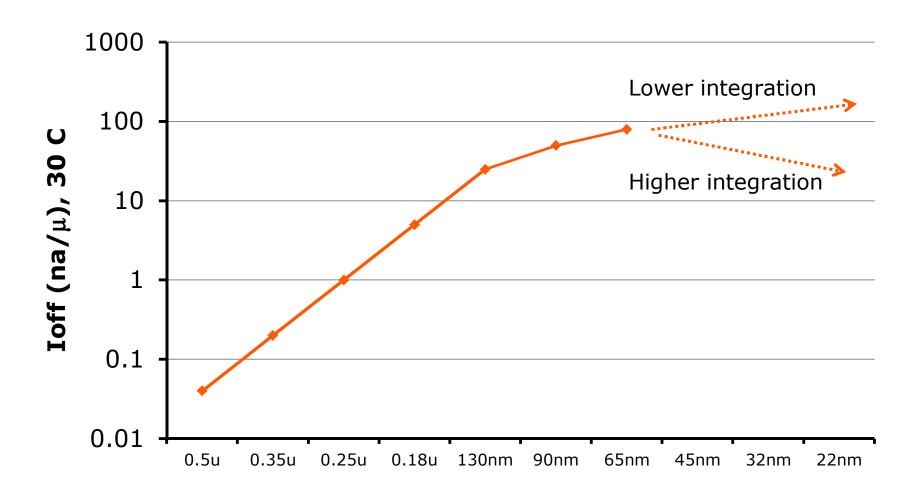
Transistor count



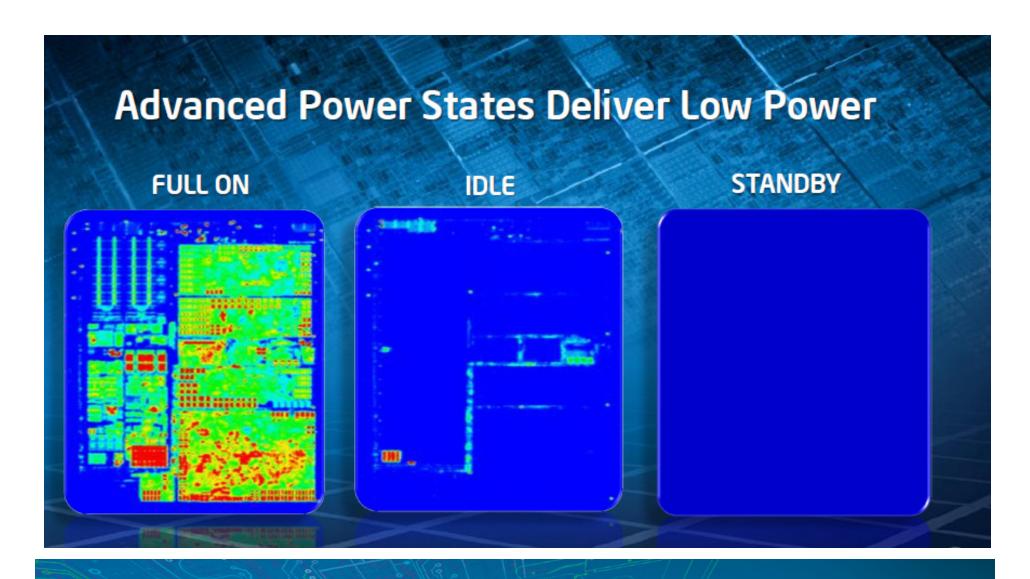
Power density



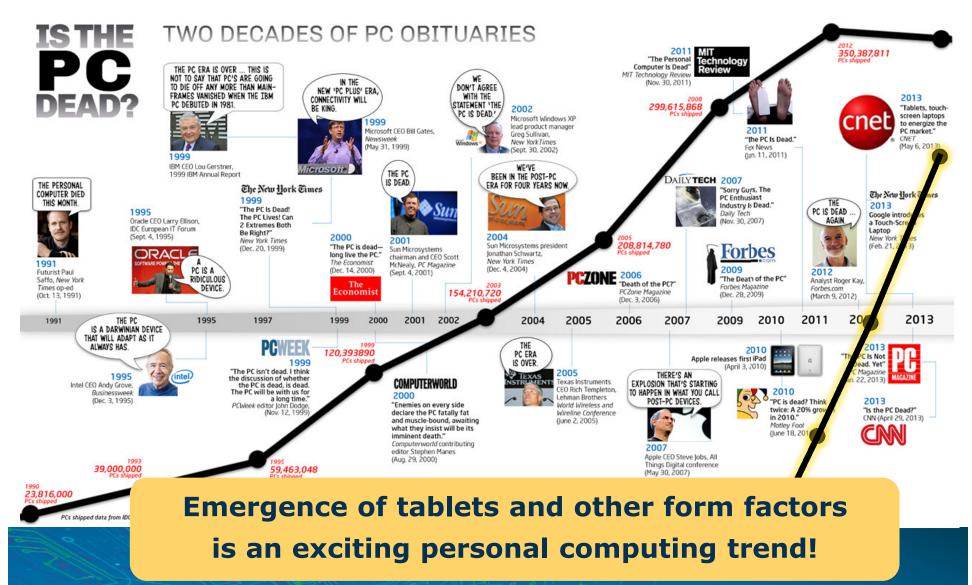
Leakage



Dark silicon

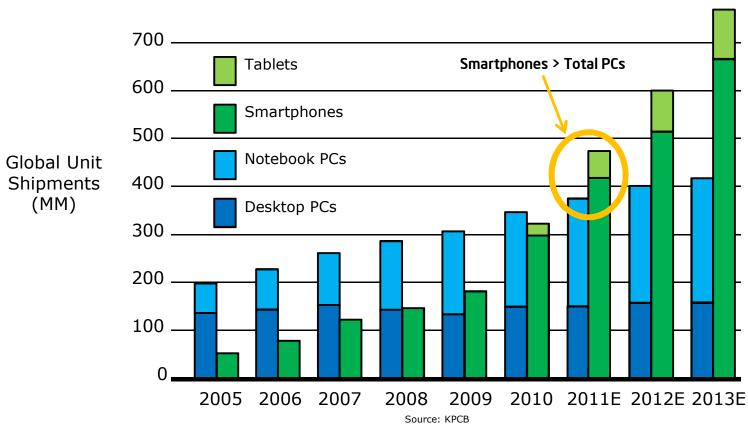


Personal computing: PCs...tablets... phones...convertibles...phablets...



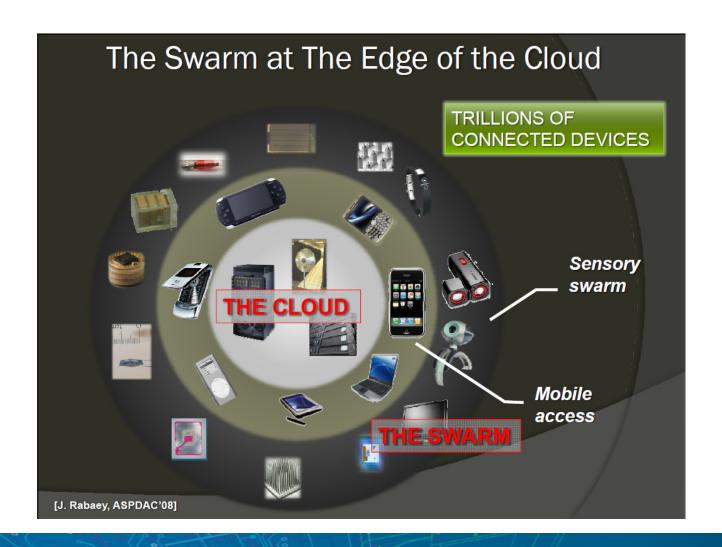
Mobility trends

Global Unit Shipments of Desktop PCs + Notebook PCs vs Smartphones + Tablets

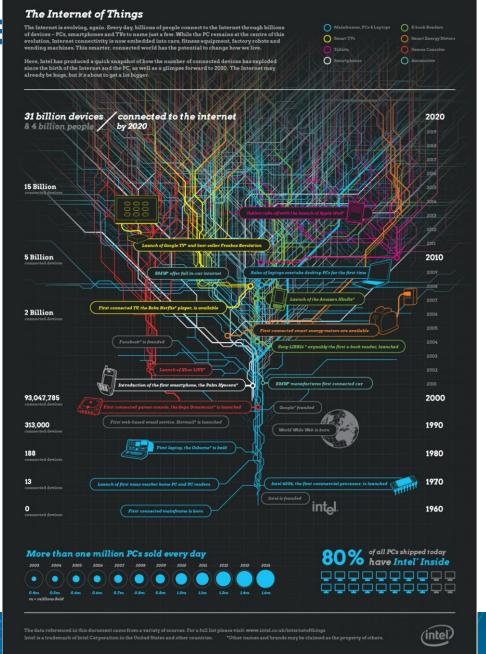


Source: Katy Huberty, Ehud Gelblum, Norgan Stanley Research Data and Estimates as of 2/11

Internet Of Things

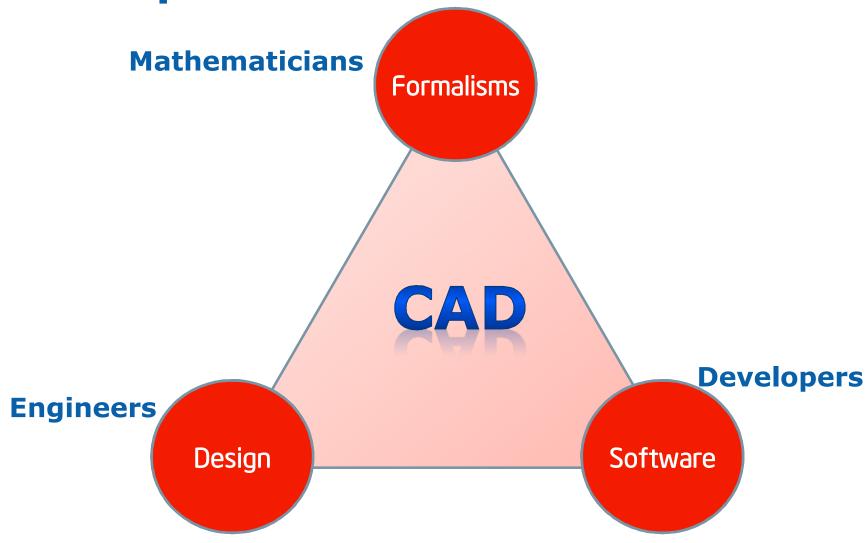


Interne

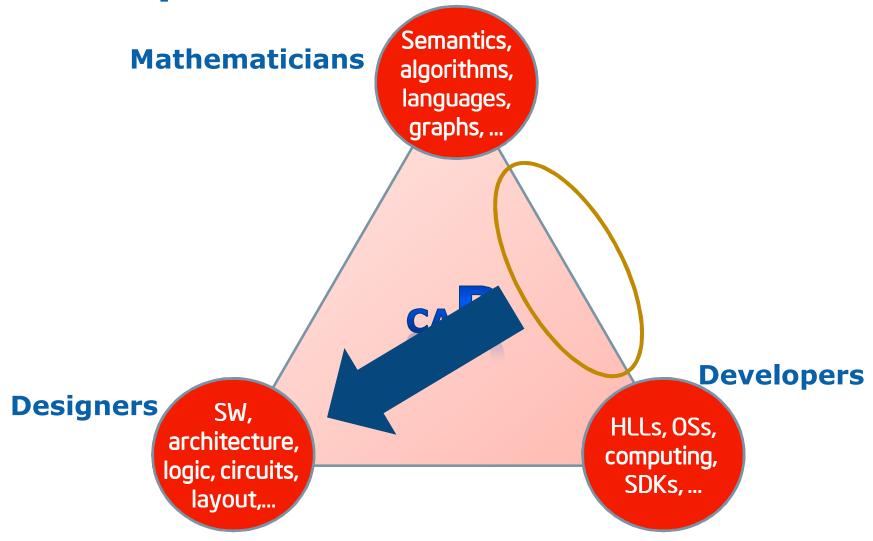


What does this have to do with CAD research?

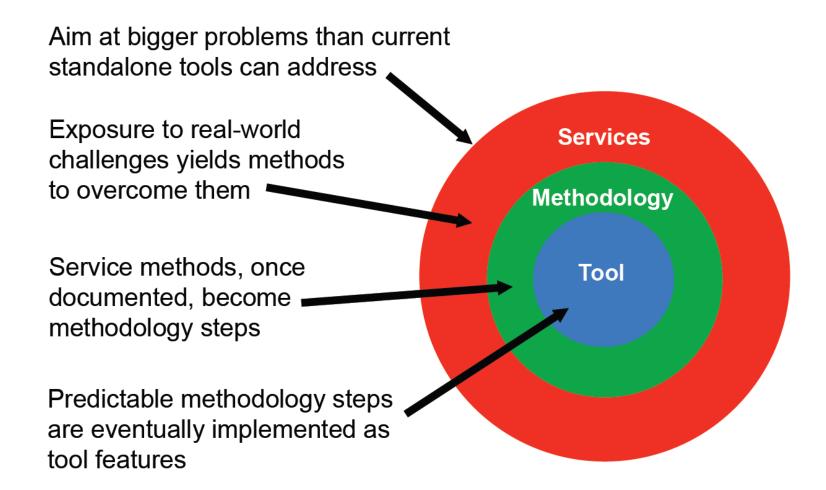
Three pillars of CAD



Three pillars of CAD

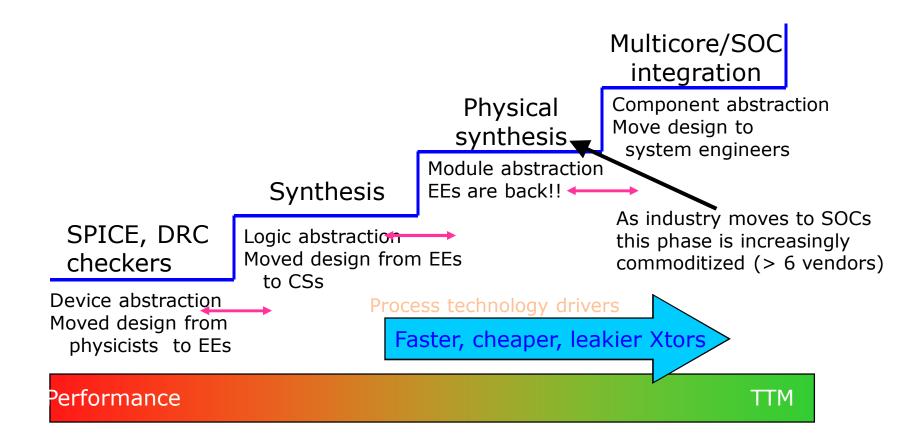


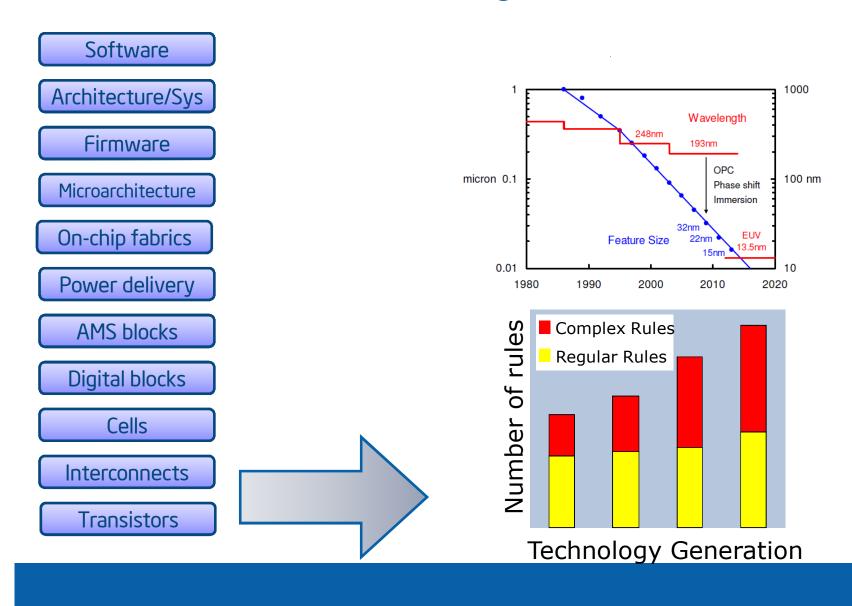
Bull's eye strategy



"Pioneering the Future of Verification A Spiral of Technological and Business Innovation" Kathryn Kranen, CEO, Jasper Design Automation, Haifa Verification Conference, Dec. 2011

A perspective on (implementation) CAD





Software

Architecture/Sys

Firmware

Microarchitecture

On-chip fabrics

Power delivery

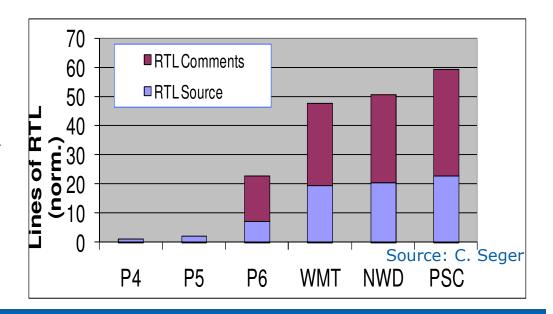
AMS blocks

Digital blocks

Cells

Interconnects

Transistors



Software

Architecture/Sys

Firmware

Microarchitecture

On-chip fabrics

Power delivery

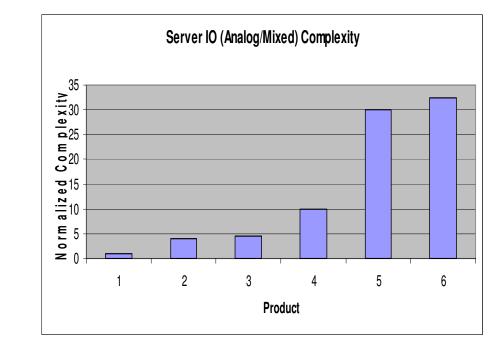
AMS blocks

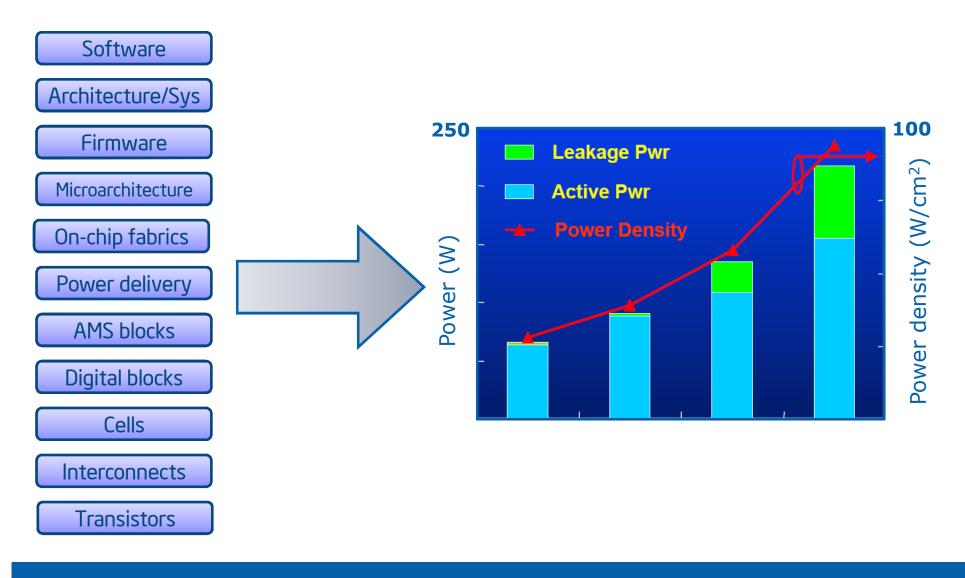
Digital blocks

Cells

Interconnects

Transistors





Software

Architecture/Sys

Firmware

Microarchitecture

On-chip fabrics

Power delivery

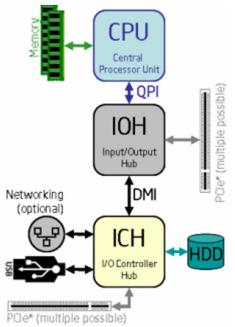
AMS blocks

Digital blocks

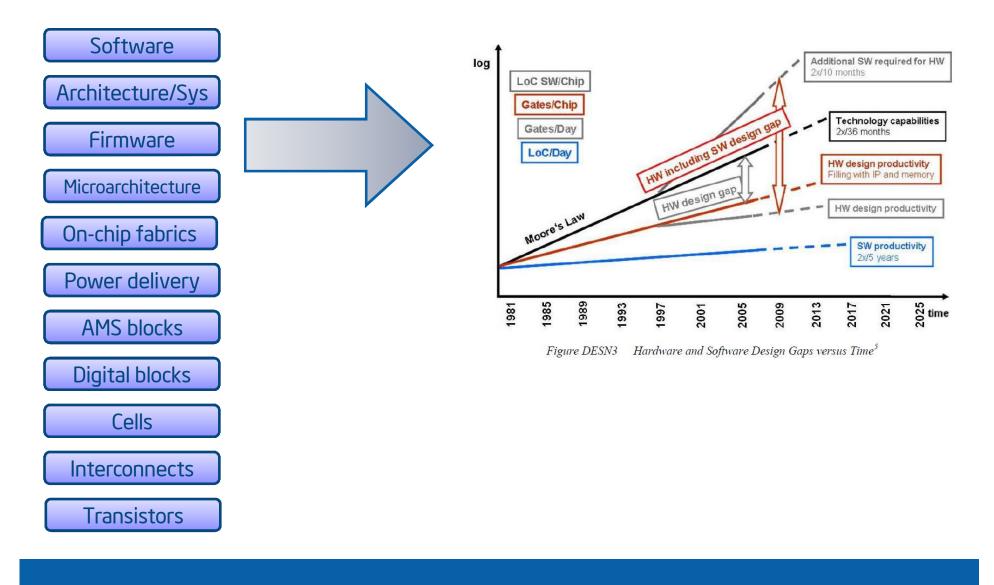
Cells

Interconnects

Transistors









Architecture/Sys

Firmware

Microarchitecture

On-chip fabrics

Power delivery

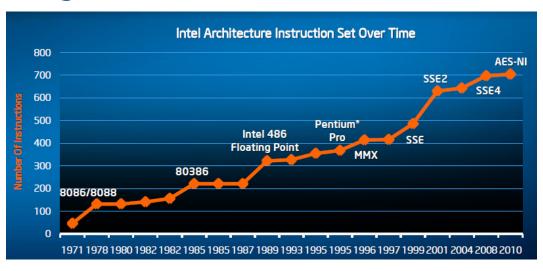
AMS blocks

Digital blocks

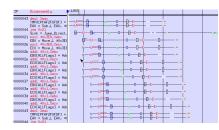
Cells

Interconnects

Transistors

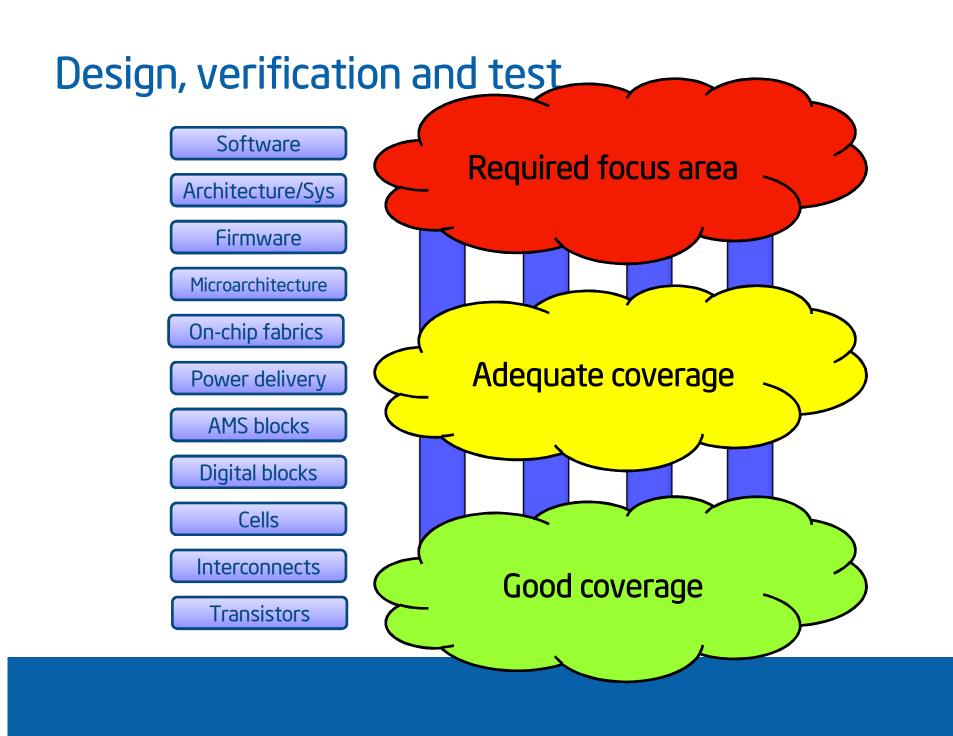


Pentium® Processor: In-order

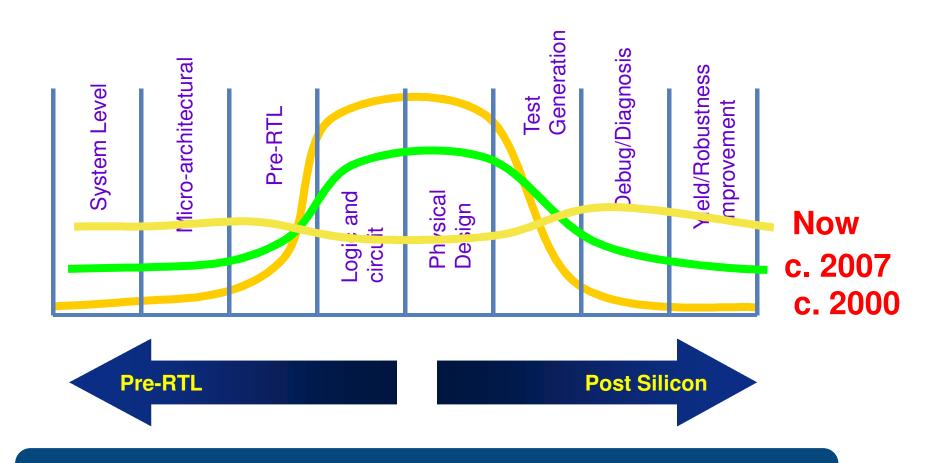


Pentium® 4 Processor: OO, TraceCache, ...



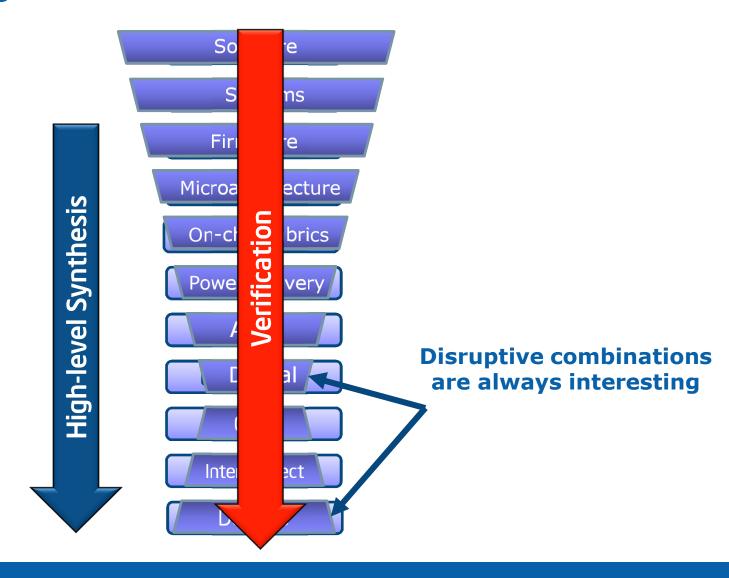


Shifting design R&D focus: Transition from the middle



Challenges shifting from middle to front- and back-end

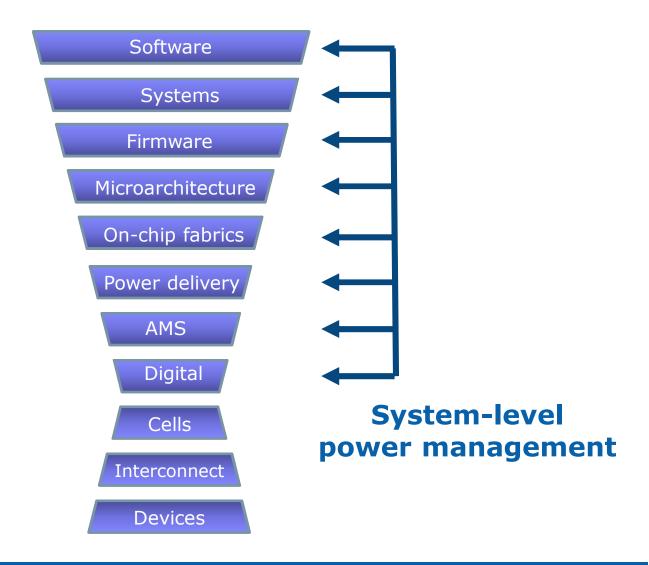
The wineglass model for CAD research



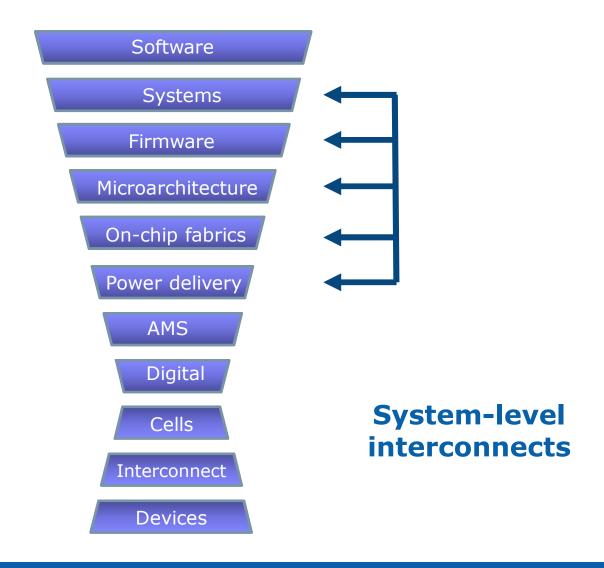
Example CAD Top 10 for 2020

- Mixed-signal verification
- System-level modeling and verification (esp. new devices)
- Power management
- System-level interconnects
- Embedded software verification (esp. security)
- Reconfigurable design
- Reliability impact on performance
- Design rule complexity
- Integration of novel devices
- Physical synthesis for large blocks (100M+ cells)

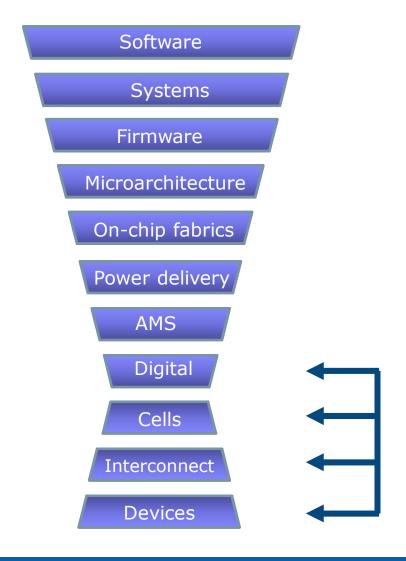
Example 1



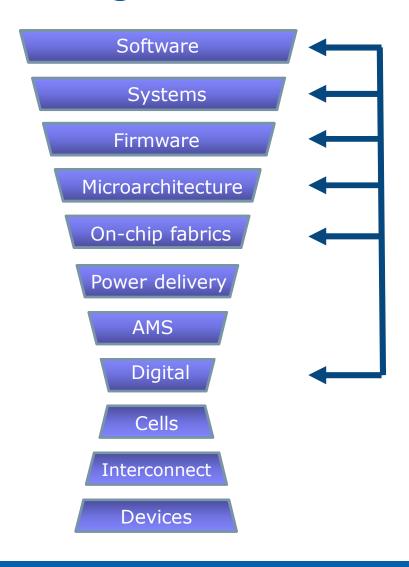
Example 2



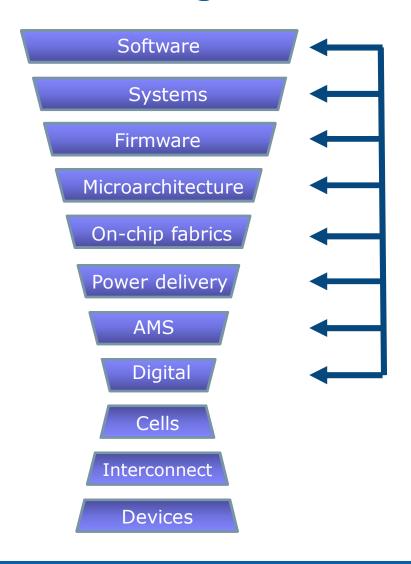
Integration of novel devices



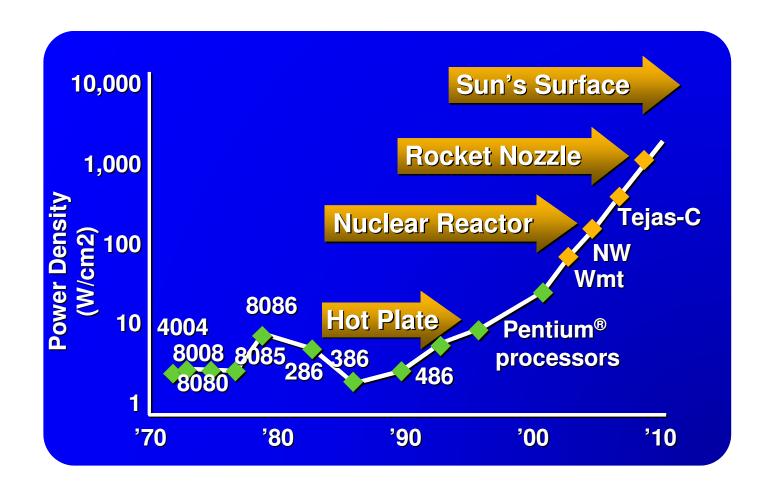
System-level modeling and verification



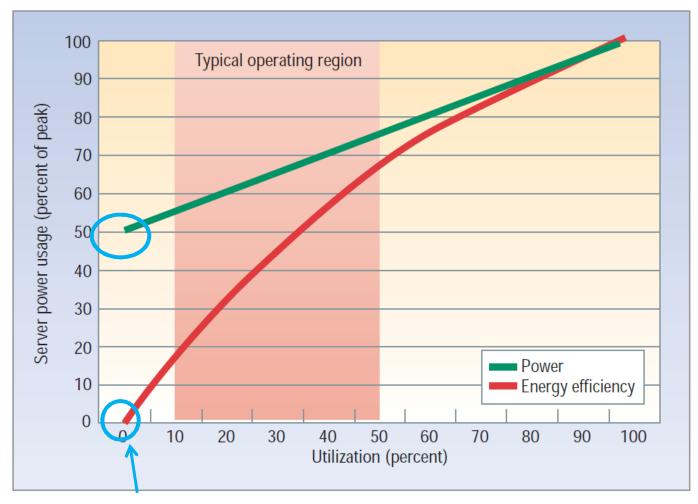
System-level power management



Power density trends

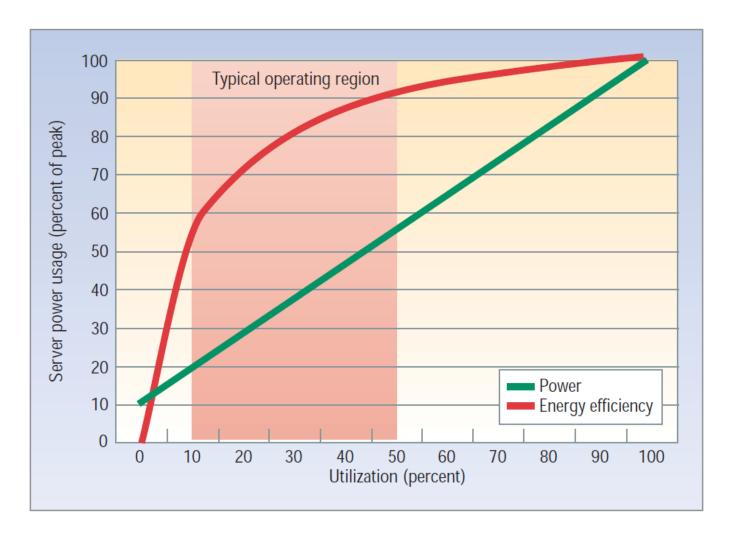


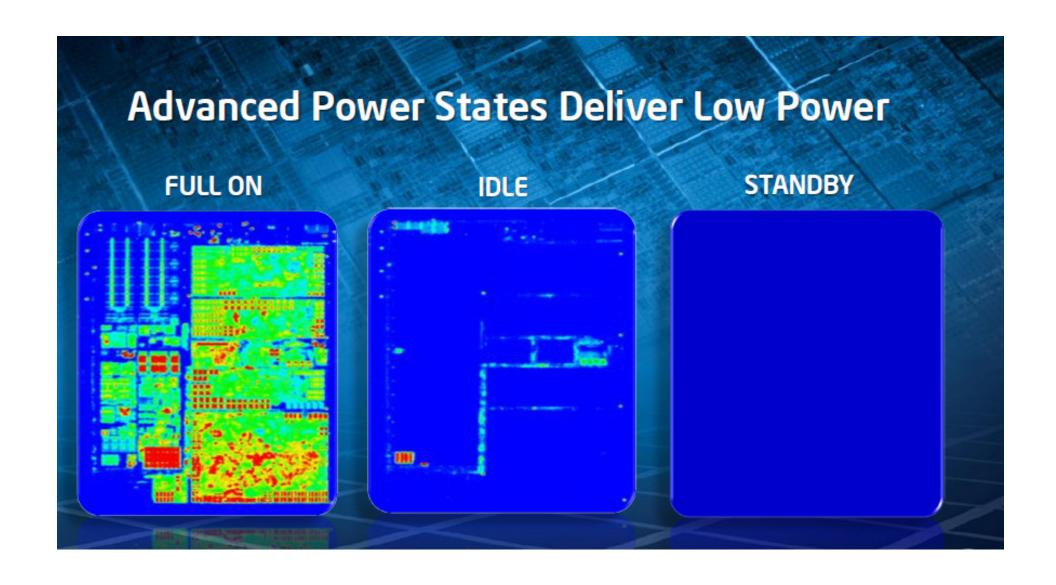
Energy proportional computing



Server is doing no work, but consumes half its peak power!

Do nothing well!





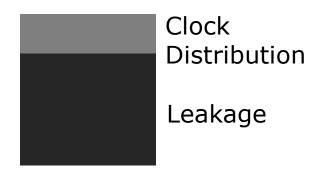
CPU Core Power Consumption

- High frequency processes are leaky
 - Reduced via high-K metal gate process, design technologies, manufacturing optimizations



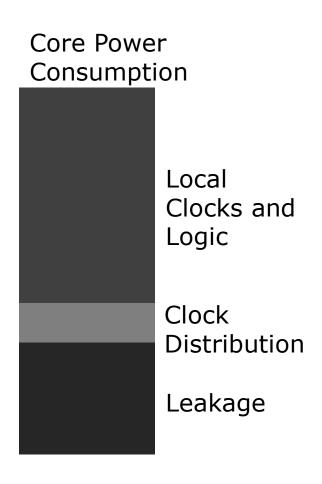
CPU Core Power Consumption

- High frequency designs require high performance global clock distribution
- High frequency processes are leaky
 - Reduced via high-K metal gate process, design technologies, manufacturing optimizations

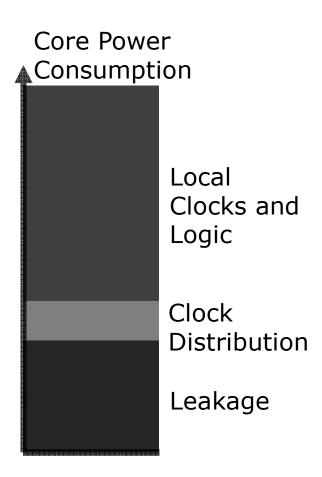


CPU Core Power Consumption

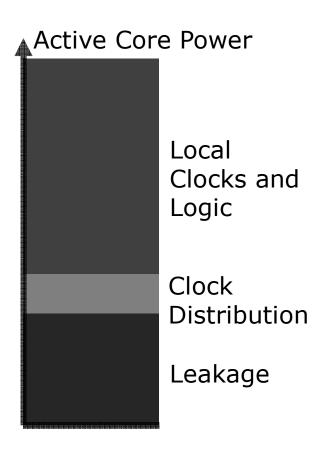
- Remaining power in logic, local clocks
 - Power efficient microarchitecture, good clock gating minimize waste
- High frequency designs require high performance global clock distribution
- High frequency processes are leaky
 - Reduced via high-K metal gate process, design technologies, manufacturing optimizations



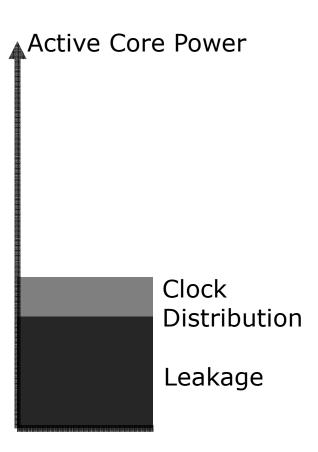
• C0: CPU active state



- C0: CPU active state
- C1, C2 states (early 1990s):
 - Stop core pipeline
 - Stop most core clocks

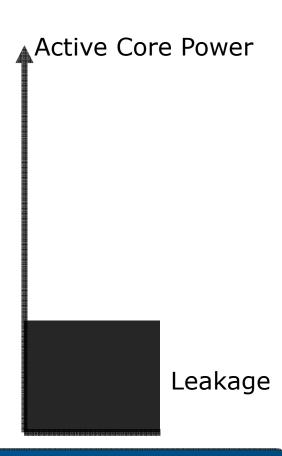


- C0: CPU active state
- C1, C2 states (early 1990s):
 - Stop core pipeline
 - Stop most core clocks
- C3 state (mid 1990s+):
 - Stop remaining core clocks
 - Voltage reduced to retention levels



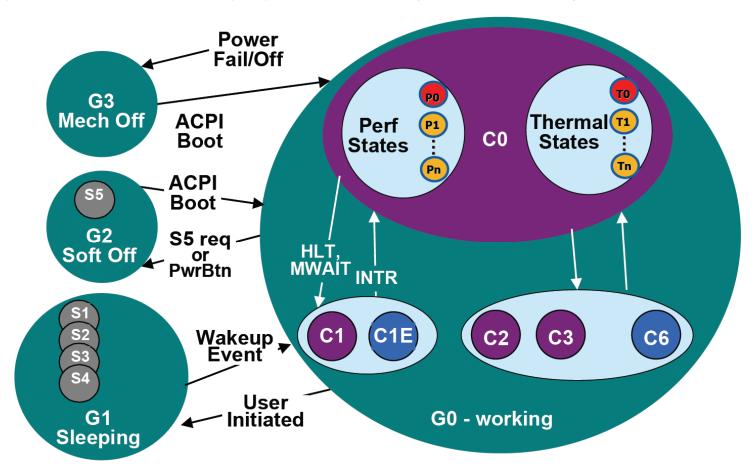
- C0: CPU active state
- C1, C2 states (early 1990s):
 - Stop core pipeline
 - Stop most core clocks
- C3 state (mid 1990s+):
 - Stop remaining core clocks
 - Voltage reduced to retention levels
- C6 state (2008):
 - Processor saves architectural state
 - Turn off power, eliminating leakage

Independent clock and voltage domains allow very good core power management today



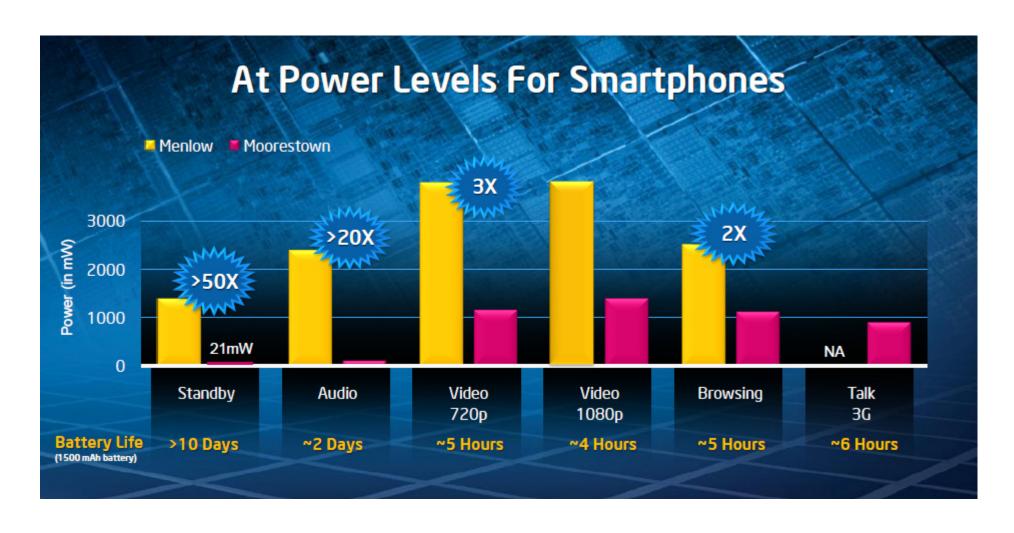
Hierarchical Power State Model

ACPI = Adv. Configuration & Power Interface Spec Gx = global state, Sx = sleeping st, Cx = CPU power st, Px = perf st, Tx = thermal st



Source: "Platform Power Management Opportunities for Virtualization", IDF Fall '08

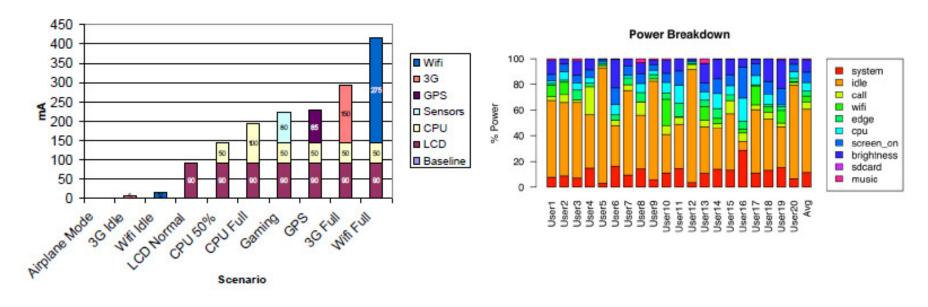
Power management & modeling



Application power scenarios

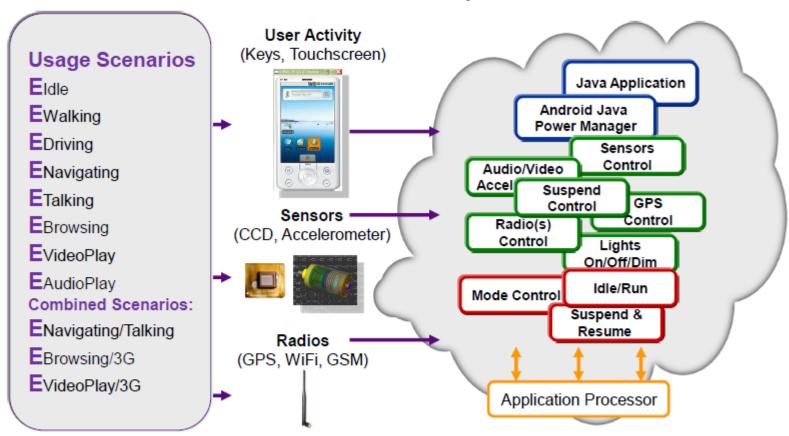
Scenario Profiles

User Profiles



Usage Scenarios

Highly dynamic operation of multiple interacting hardware and software components



(C) Synopsys, ARM Gibbons, Nohl, Flynn, *System-Level Design and Software Development for Energy Efficient Platforms*, DAC 11 tutorial

Power management stack

Application

Operating system

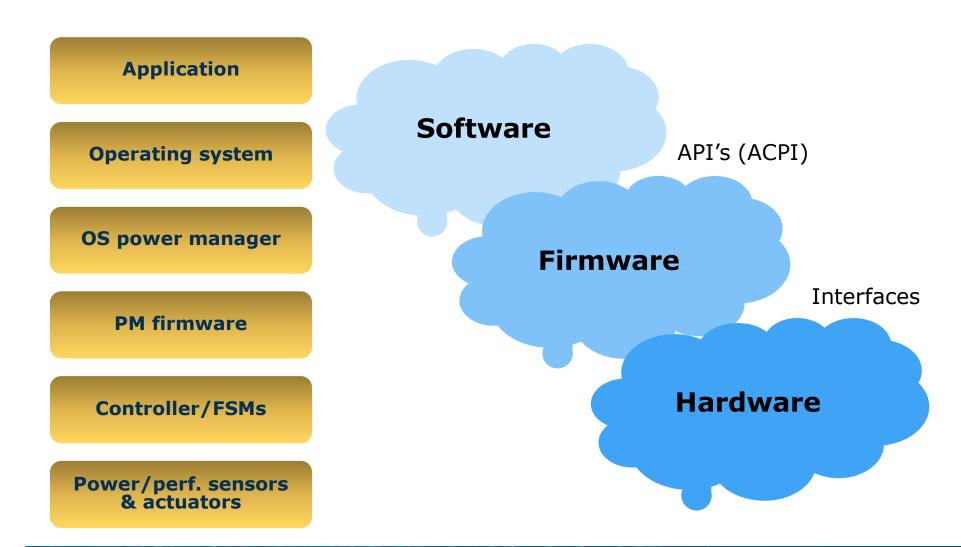
OS power manager

PM firmware

Controller/FSMs

Power/perf. sensors & actuators

PM software-hardware co-design



Power management: Ideal state

Application

Operating system

OS power manager

PM firmware

Controller/FSMs

Power/perf. sensors & actuators

Software

Well-characterized applications
Well-defined interfaces
Power management states
Si behavior accurately modeled
Entry/exit latencies well
calibrated

Hardware

Power management: Reality

Application

Applications' power/perf behavior is hardware specific

Operating system

Optimization complexity

/are/hardware generations

ried hardware models

OS power mar

PM firmware

Controller/FSMs

Power/perf. sensors & actuators

Fir vara

Validation

complexity

Power

....

Controller / Jugg

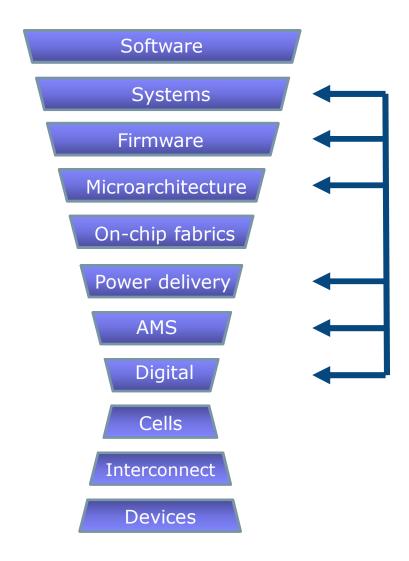
The right set of Si looks are not incorporated

The Si behavior of existing hooks has excursions

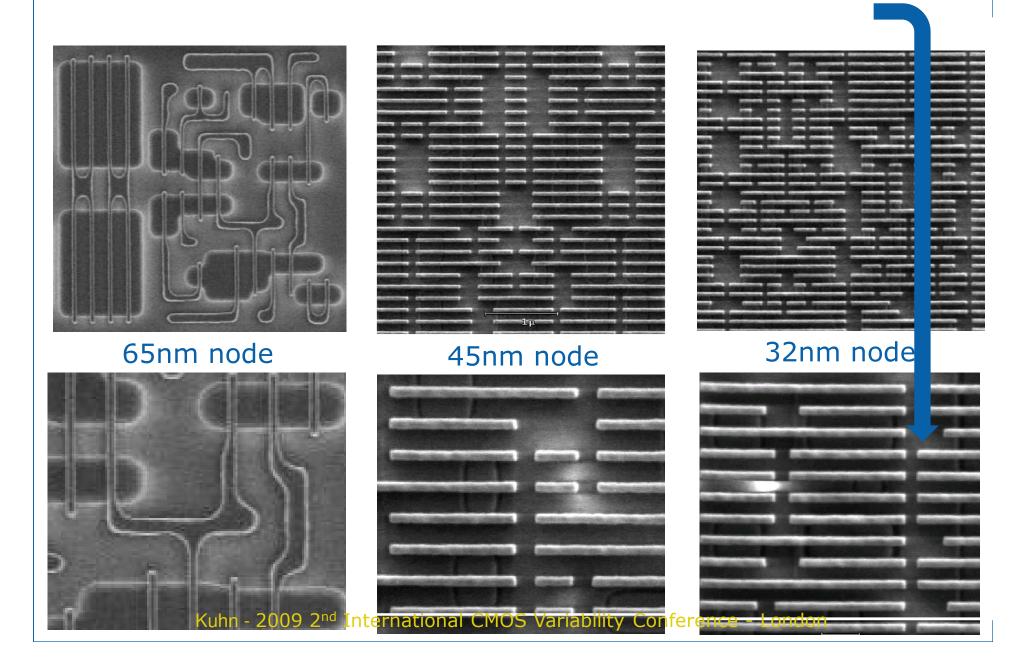
Power Management: Research Vectors

- Application-level power behavioral modeling of software, platform, and components
- Calibrated system-level power models
- Power management algorithms and implementations
- Design discovery of power state options
- Circuit techniques for power management
- Validation ingenuity

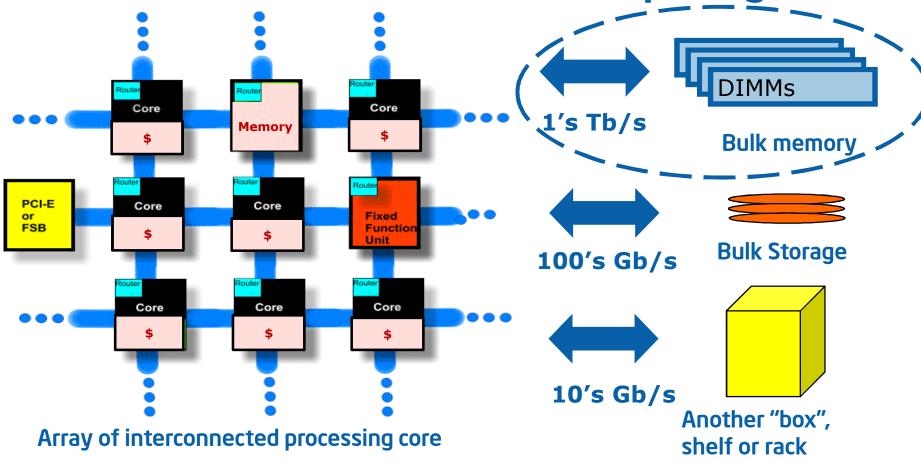
AMS validation



Soumya's challenge: Build a radio that looks like this



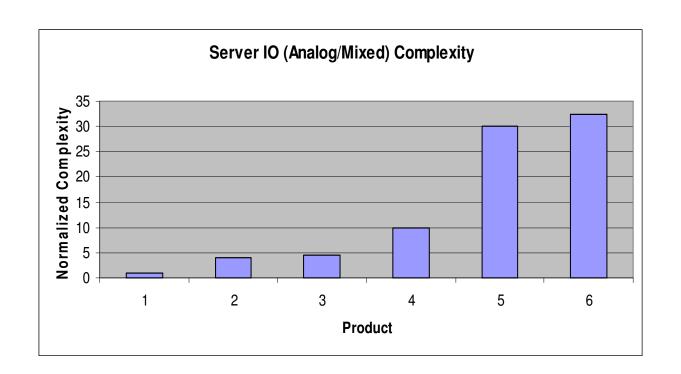
10 Bandwidth for Tera-scale Computing



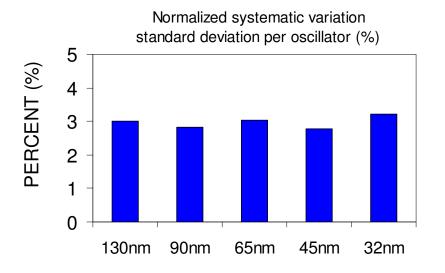
- Tera-Flops implies tera-bytes/sec of bandwidth
- Limited off-chip intercon. density scaling implies high pin speed

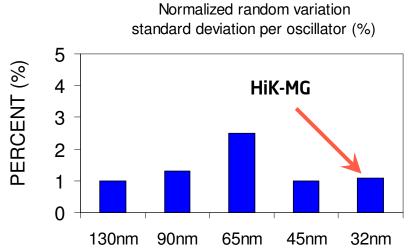
Source: Randy Mooney

Server IO (analog/mixed signal) complexity

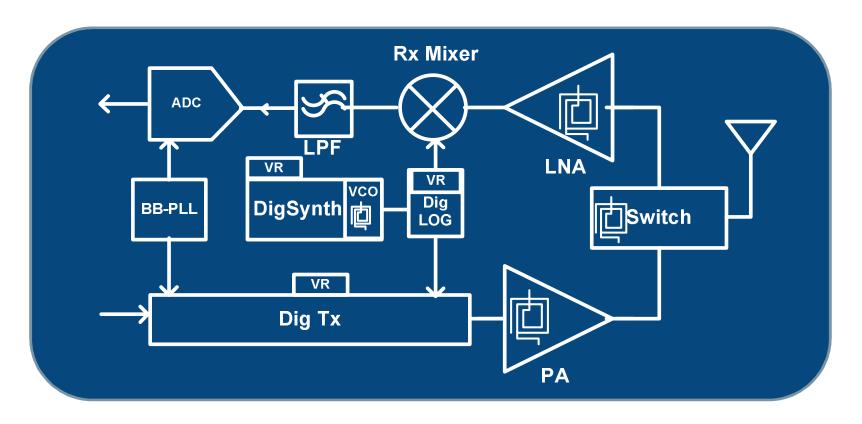


Variation trends





Wrapping it up

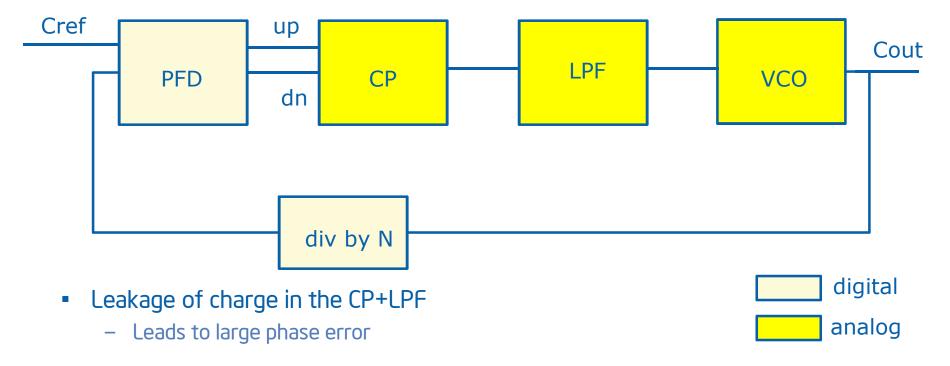


Mix of charge, voltage and time domain coded analog blocks with VR's

"Big" (D)digital around "small" (a) analog (some of which is "digital")

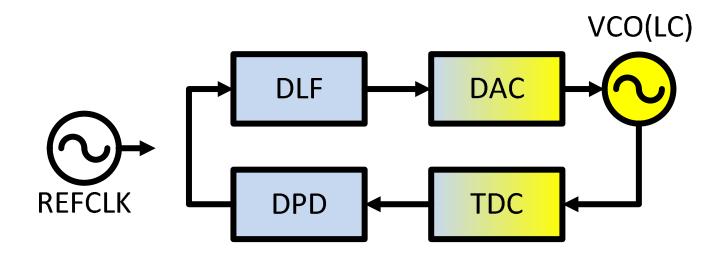
Krishnamurthy Soumyanath, 2010 Intel AMS Verification Workshop

The analog PLL



- Poor area scaling of analog content
- Reducing voltage headroom/generation
- Sensitivity to process variations

Property verification of a digital PLL



- Will the PLL lock for arbitrary initial conditions and across different PVT?
- Will it always lock within x us?
- Will it stay locked once locked?
- Jitter tolerance
 - How much jitter on input clock can it tolerate
 - What is the rms and P2P jitter on output

A. Ravi et al, "A 9.2–12GHz, 90nm digital fractional-N synthesizer with stochastic TDC calibration and -35/-41dBc integrated phase noise in the 5/2.5GHz bands," *IEEE VLSI Symposium Digest*, pp.143 - 144, 2010

Analog Property Verification

Delay Lock Loop (DLL):

- Will it lock properly? (i.e. phase delays between consecutive outputs are good)
- Will it enter harmonic lock? Under what startup conditions?

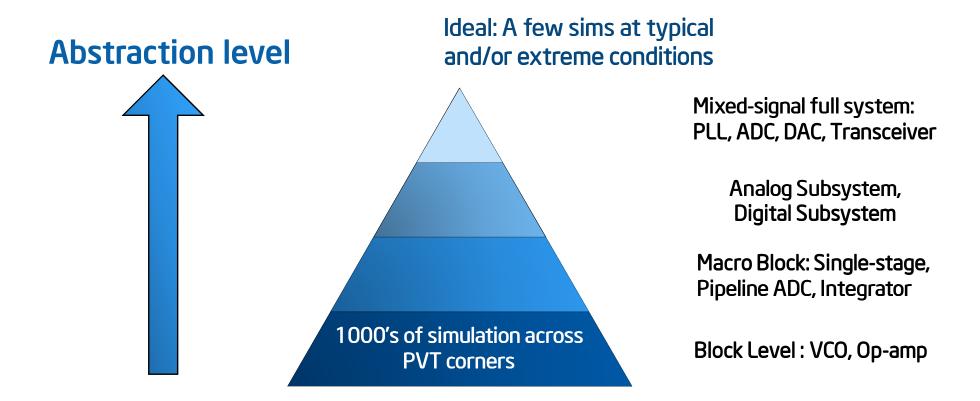
Phase Lock Loop (PLL):

- Will it lock properly? (i.e. generated output clock is good)
- Under what control input, startup, and PVT conditions will it fail to lock?

I/O Systems:

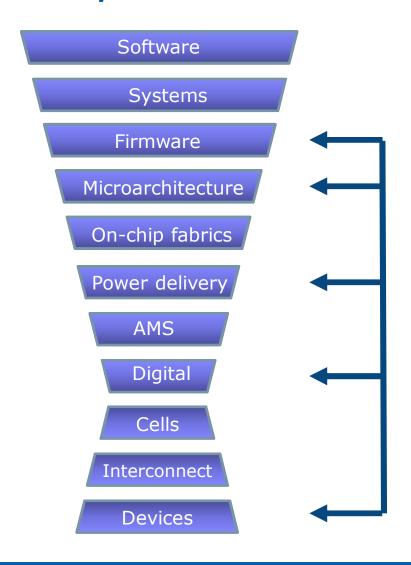
- Will I/O adaptation algorithms converge to good values for preemphasis, equalization, and termination settings? Under what conditions will we reach saturation for control settings?
- Will we get desired timing and voltage margins?

Mixed-Signal Validation Pyramid

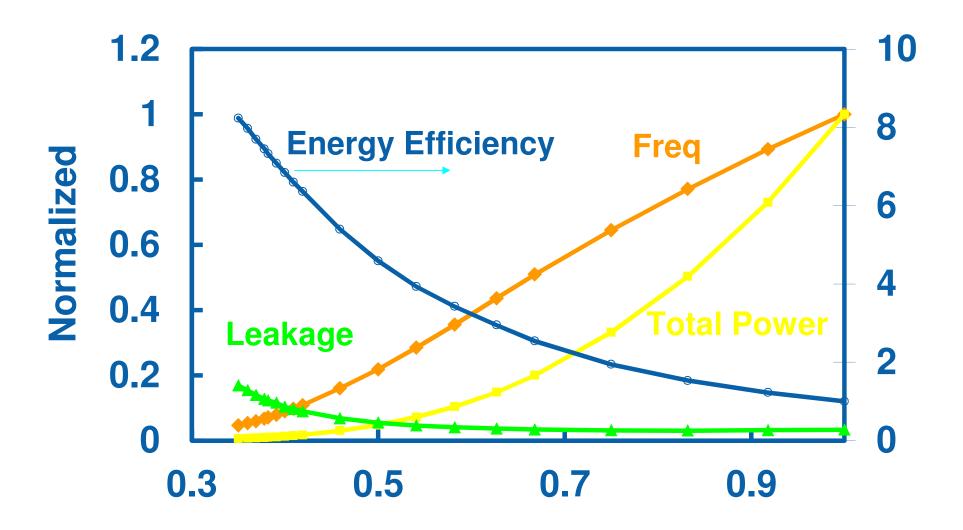


Raise the abstraction level for AMS verification

Reliability impact on performance

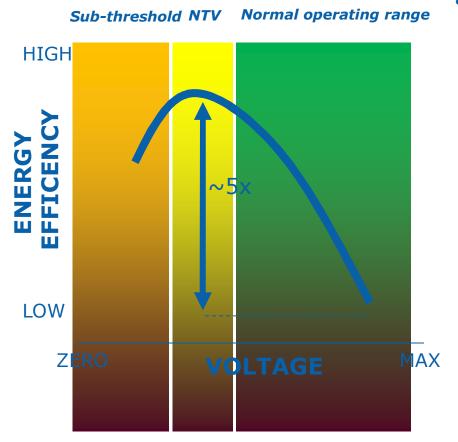


Voltage scaling for energy efficiency



Source: Shekhar Borkar

Benefits of Near Threshold Voltage



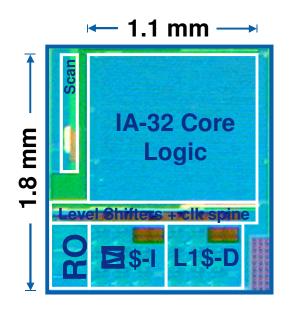
- Peak energy efficiencies at NTV
- Fine-grain power management

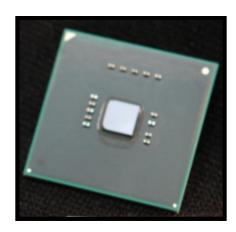
POTENTIAL FOR...

- More always-on / instant wake devices
- Intelligent everyday devices with battery/solar powered CPUs
- Longer battery lives for mobile computing
- Scalable many-core chips for the datacenter
- Meeting Extreme-scale Compute Challenges

Source: Sriram Vangal

Experimental NTV processor





951 Pin FCBGA Package



Custom Interposer

Technology	32nm High-K Metal Gate
Interconnect	1 Poly, 9 Metal (Cu)
Transistors	6 Million (Core)
Core Area	2mm ²



Legacy Socket-7 Motherboard

A Complete Vmin Solution

- NTV circuits proven and technically feasible
- Building product requires addressing challenges and additional technologies

Process variations

Power Delivery

Dynamic variations

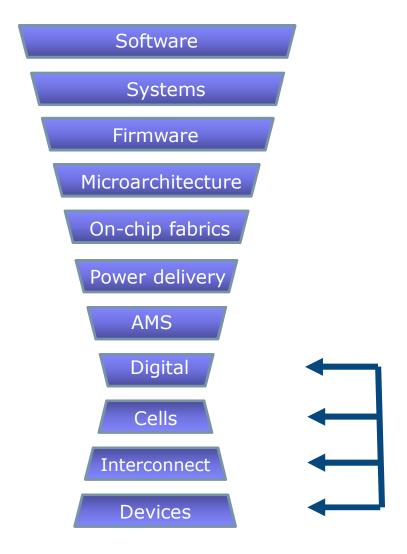
CAD tools for timing closure

Libraries

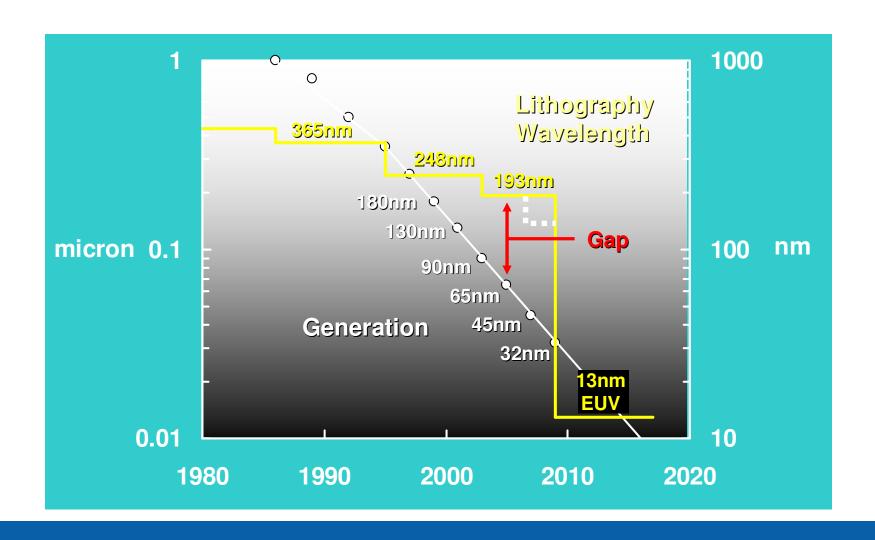
Area and power overhead

Source: Sriram Vangal

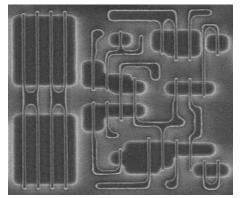
Design rule complexity



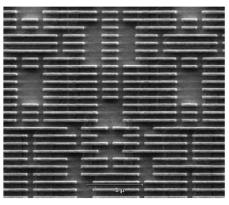
Gridded design



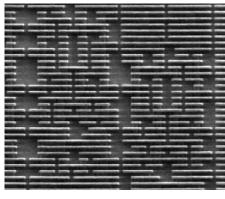
Gridded design



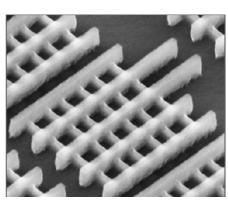




45nm node

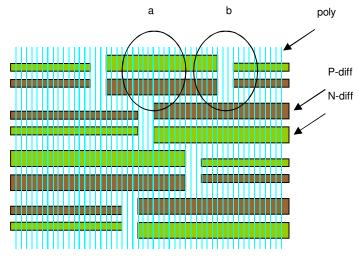


32nm node

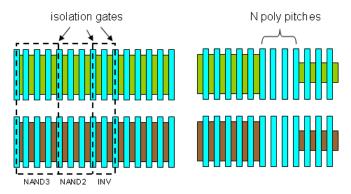


22nm node

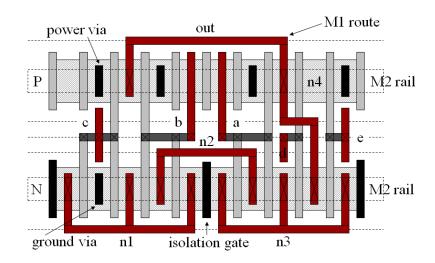
Our approach to regular layout fabrics

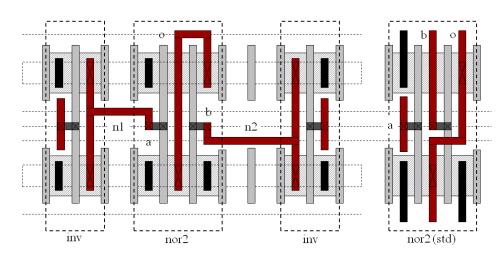


Islands of cells with the same diffusion width

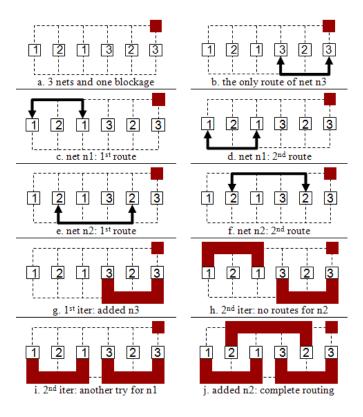


Isolation gates separate cells within a diffusion block. Diffusion blocks are separated by empty space.





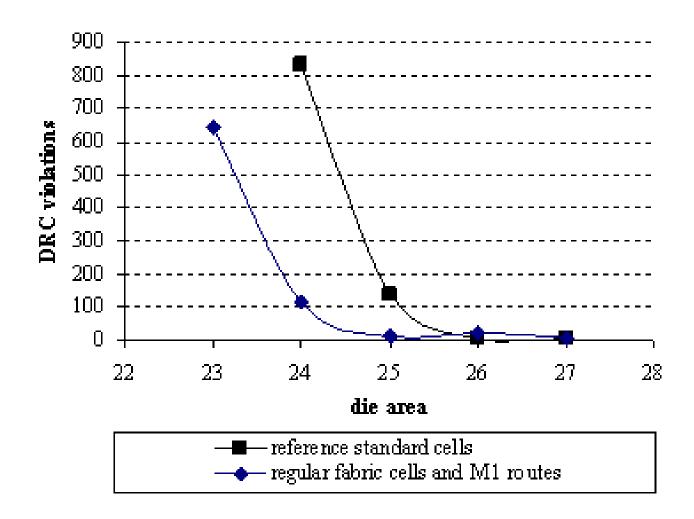
Auto route and cell synthesis



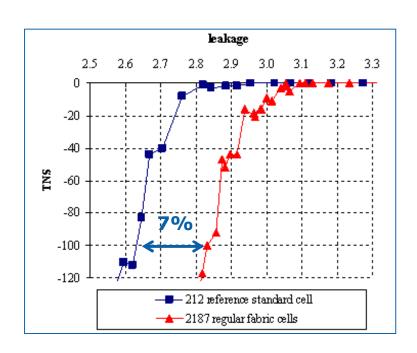
- Input is a sized transistor netlist.
- Placement of transistors is done.
- Nets are split onto 2-end connections.
- Possible layout patterns are created on a grid.
- Unfeasible layout patters are pruned.
- Patterns are enumerated to get a complete routing.

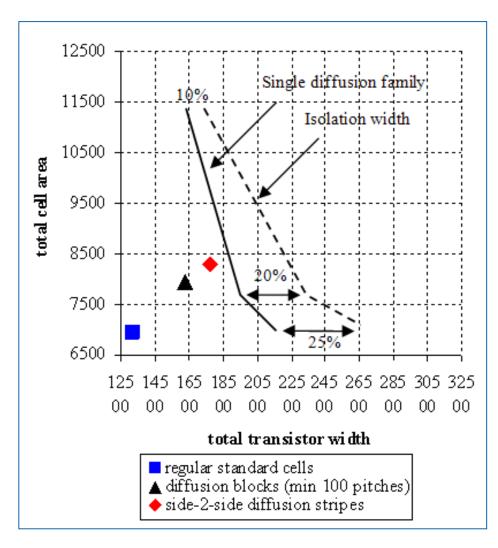
- Regular layout fabric allows automatic cell synthesis for individual cells....
 - ... which is combined with intra-cell routing for better pin placement

Experimental results



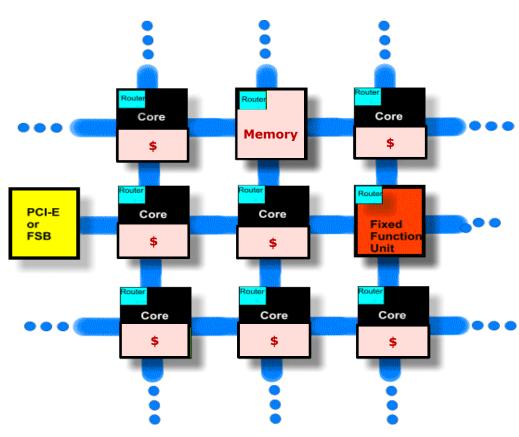
Flows with Long Diffusion Stripes





On-chip Interconnects

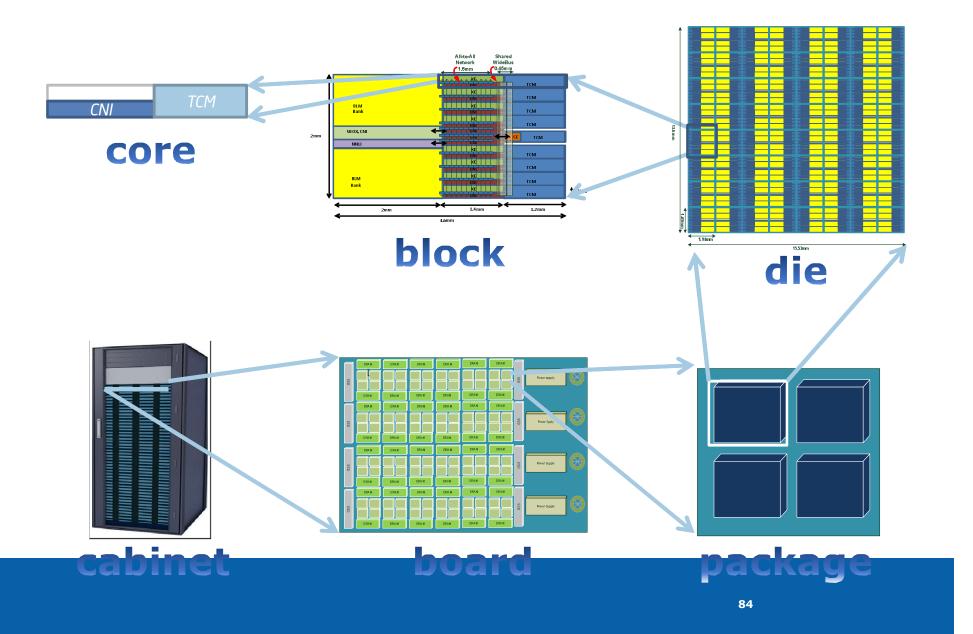




SOCs

Many Core Architectures

System Interconnects



Modular SOC Fabrics

Fabric Convergence

- Single interface?
- Packetized approach?
- Hierarchical NoCs
- Reusable components

Fabric Intelligence/Services

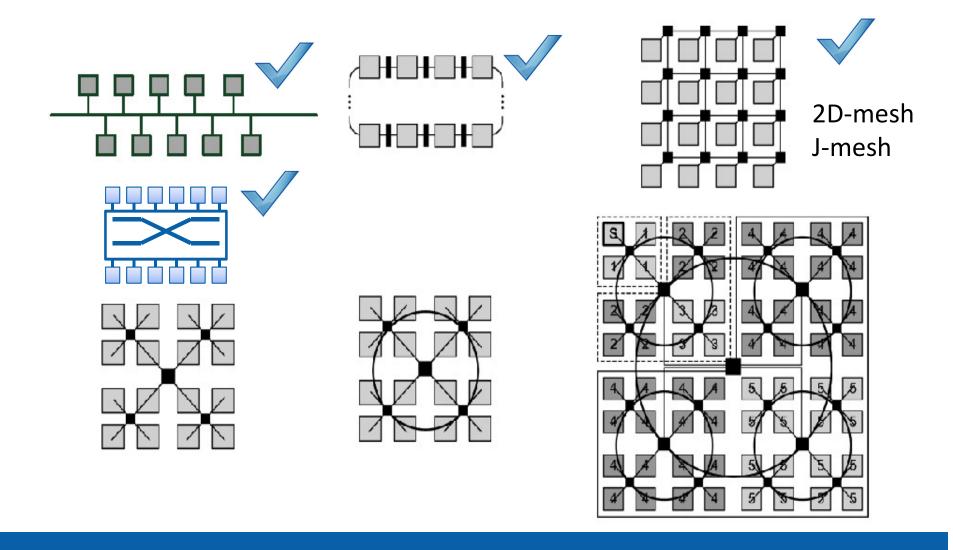


Fabric Automation

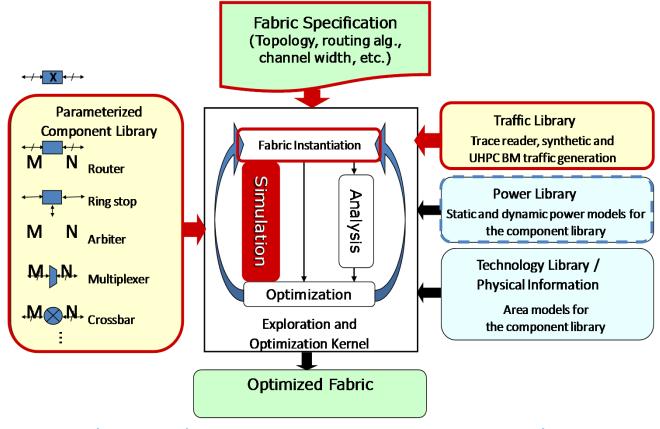
- Automation Tool Chain
 - -- Coh & Non-Coh
 - -- QoS support
 - -- Hierarchies, etc

Source: Raj Yavatkar

Topology options

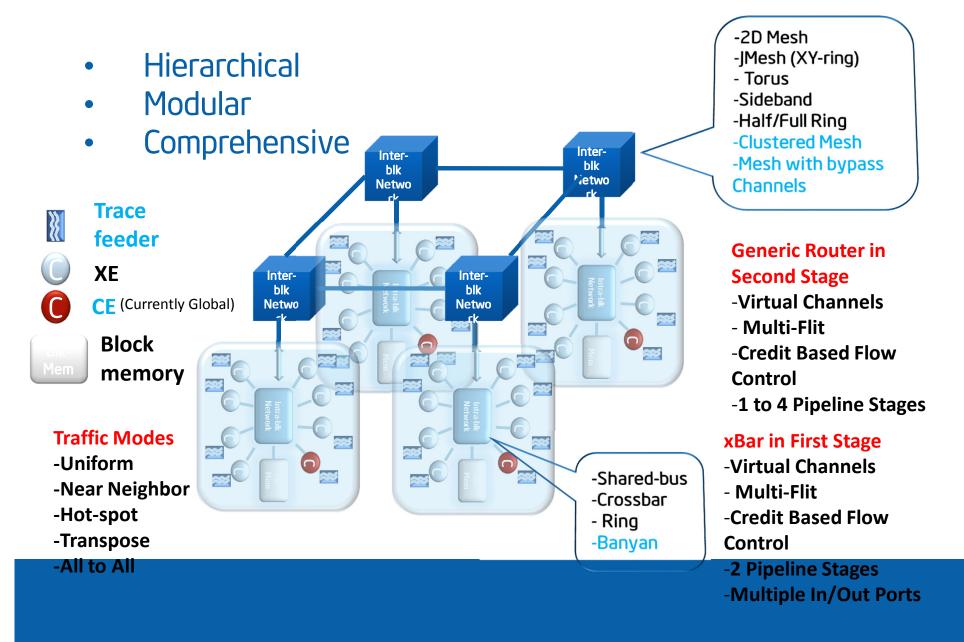


xPLORE UHPC simulator

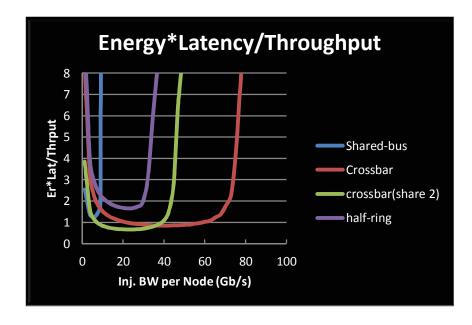


- Cycle accurate simulator in System C Currently, syr
- Cover on-chip fabric (up to 1 k cores) •
- Supports hierarchical interconnects
- Currently, synthetic and trace-driven
- 720 Agents (6x10x12) 10^6 cycles ~ 45 min.
- In future, integrate with core and memory simulation tools

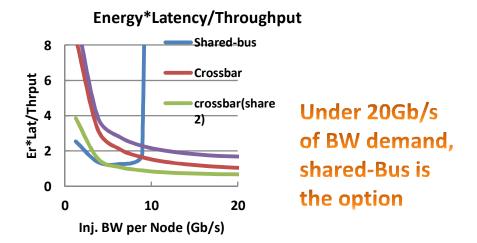
Simulator status



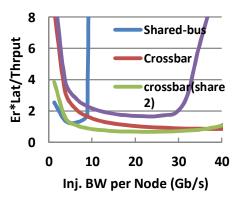
Putting it all together



For high BW demand, crossbar is the option.

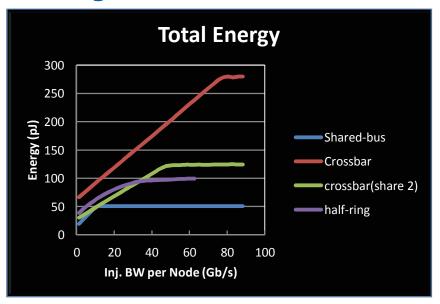


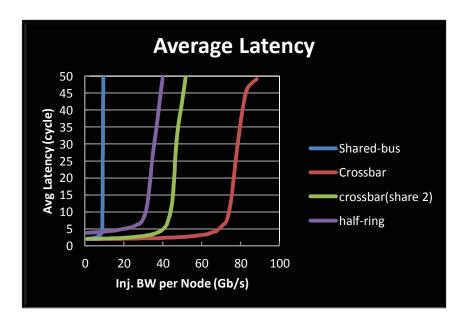


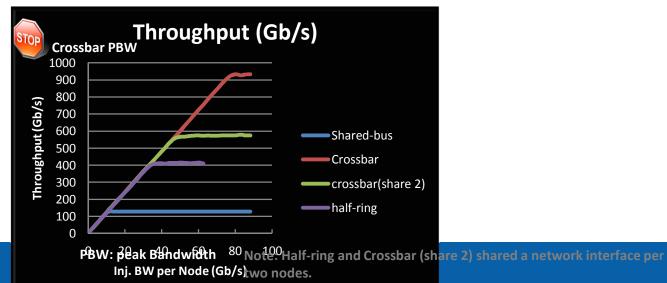


Under 40Gb/sof BW demand,crossbar (share2) is the option

Early results







Specific research problems

- Physically-aware performance & area optimization for comm. fabrics
 - Performance and floorplan characterization for early design decisions
- Quality of service analysis and optimization
 - Fast analytical estimation of quality of service with 10-20% accuracy
 - Analytical optimization methods (buffer sizing, arbitration function selection)
- Traffic models
 - Algorithms for compact workload characterization including bursty traffic
- Functional correctness
 - Deadlock-freedom, memory consistency
 - Scalable formal verification of critical functional properties
- Uncore, interconnect, and system power management
 - DVFS- and AVS-aware exploration techniques