Making the Right Hand Turn to Power Efficient Computing

Justin Rattner

Intel Fellow and Director
Microprocessor Research Labs
Intel Corporation

Outline

- Technology scaling
- Types of efficiency
- Making the right hand turn

Historic Perspective

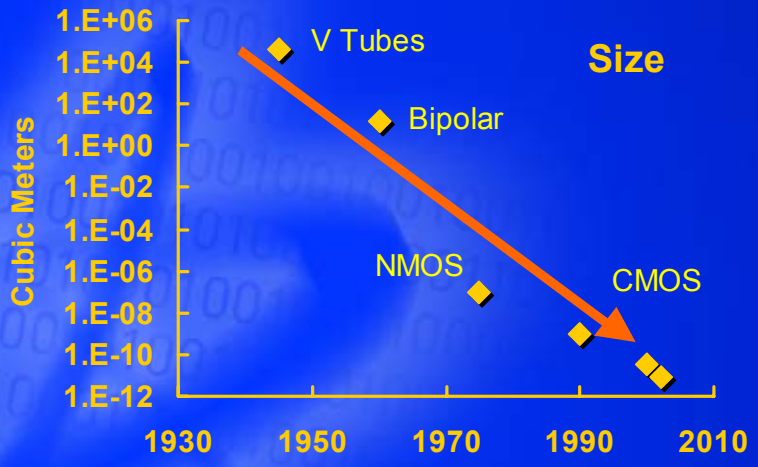

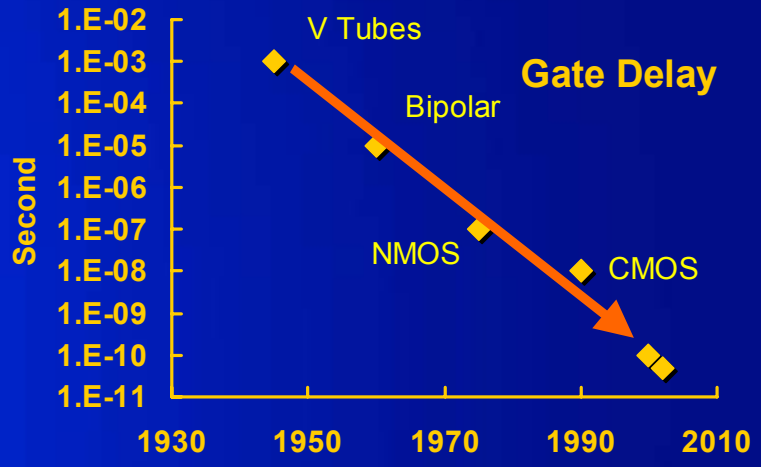

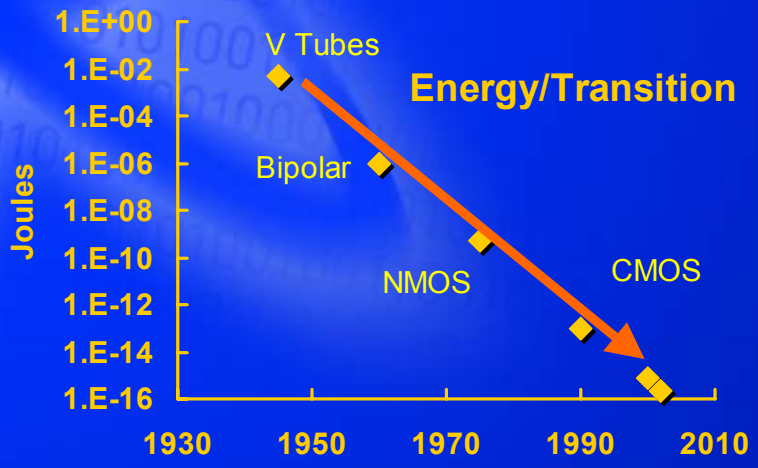

V Tubes ⇨ Bipolar
Bipolar ⇨ NMOS
NMOS ⇨ CMOS
CMOS ⇨ ?

Scaling will continue

Technology Scaling

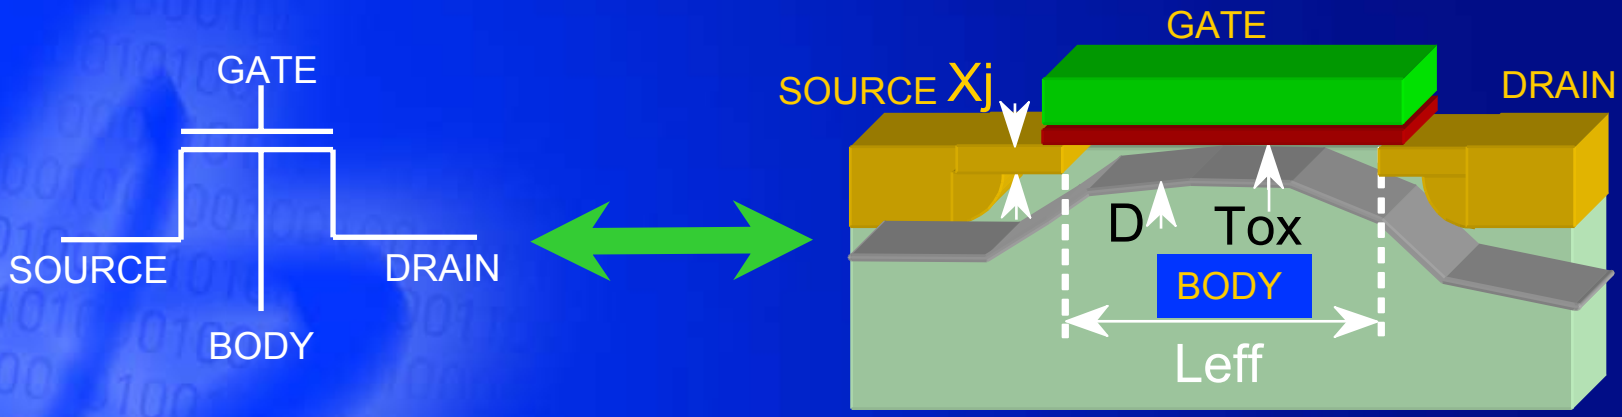

X & Y Dimensions scale down by 30%	Doubles transistor density
Z-Oxide thickness scales down	Faster transistor, higher performance
Vcc & Vt scaling	Lower active power

Technology scaling is a great thing

Supply Voltage Scaling

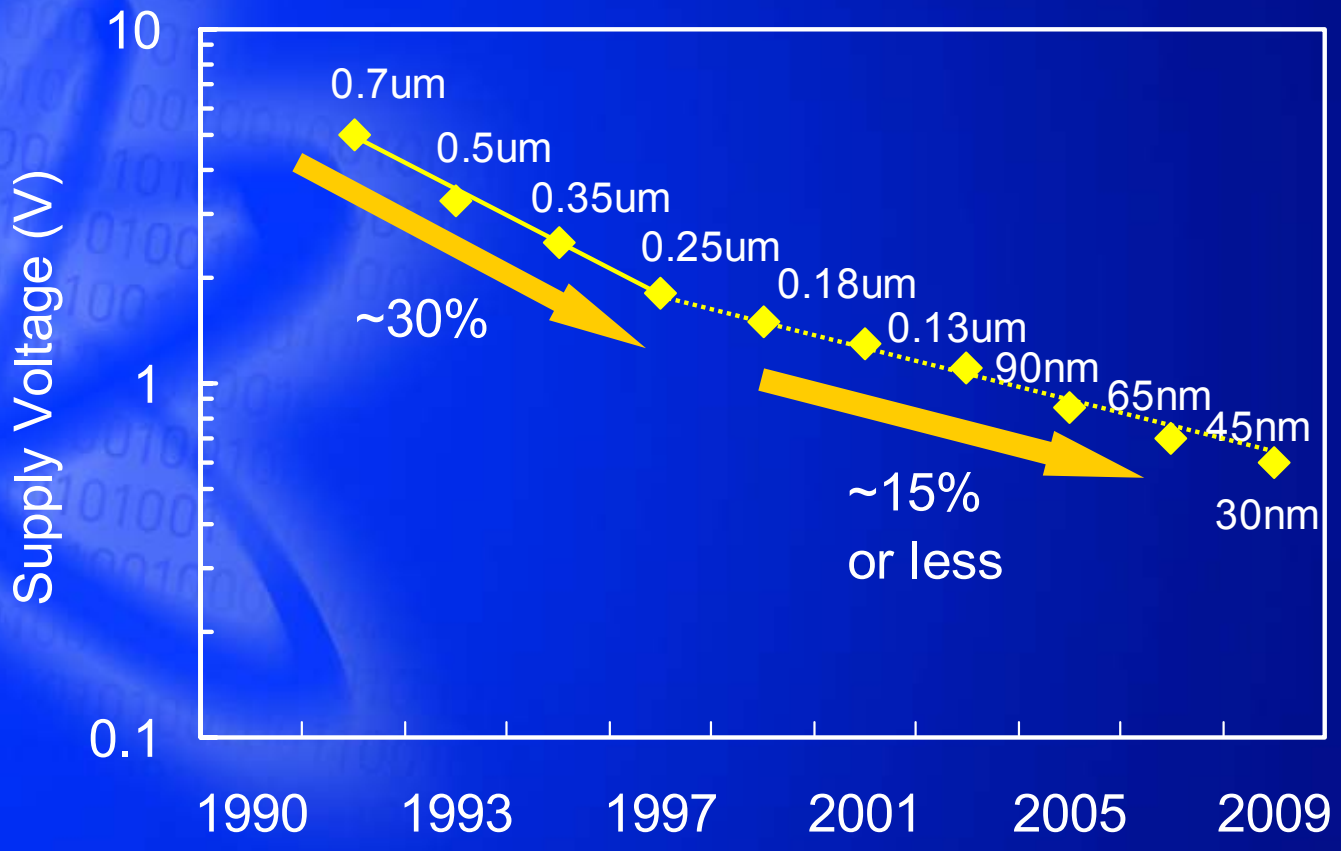

Supply voltage scaling has slowed

Active Power Projection

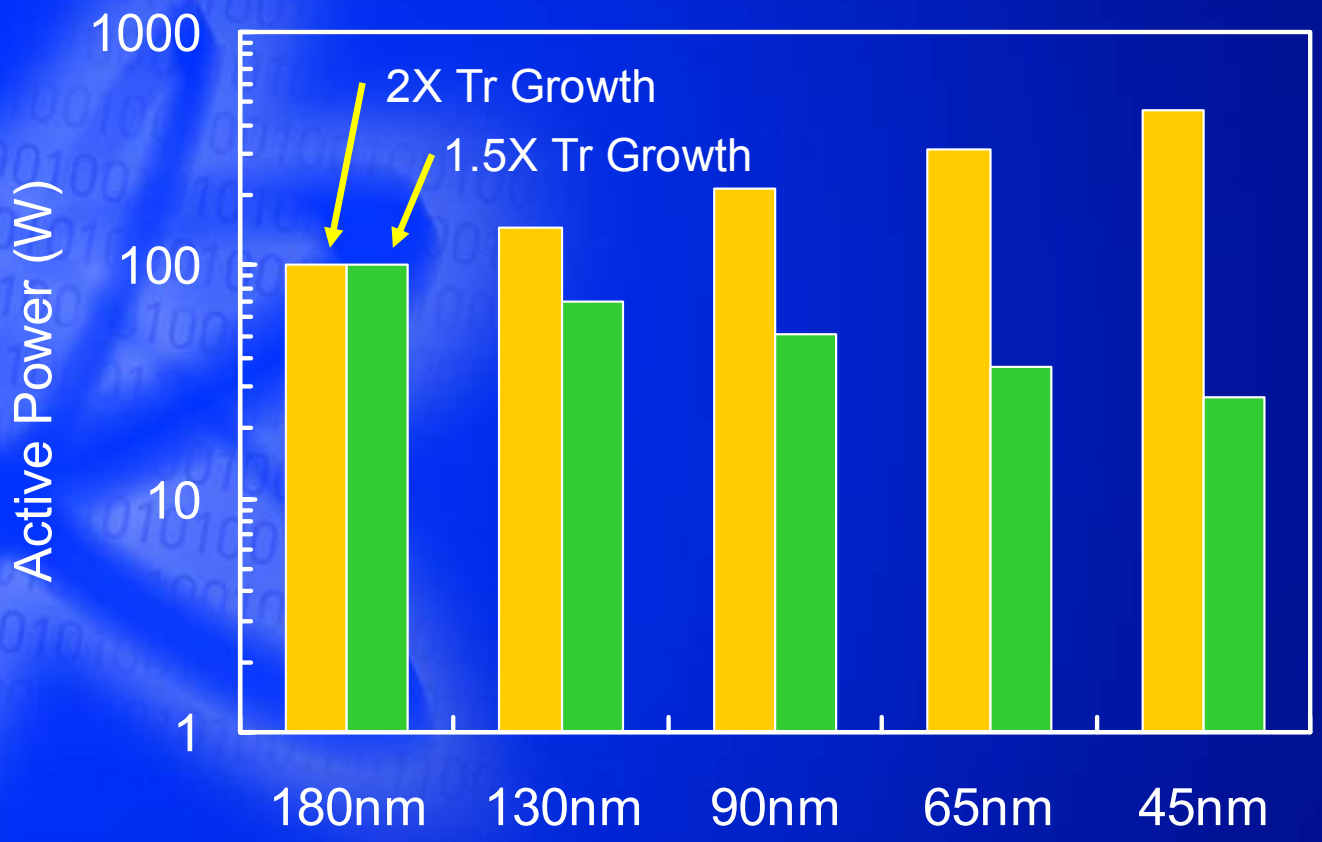

Assumptions:
- 15% Vdd scaling
- 50% Freq scaling
  (Per generation)

Power will limit transistor integration

Sub-threshold Leakage

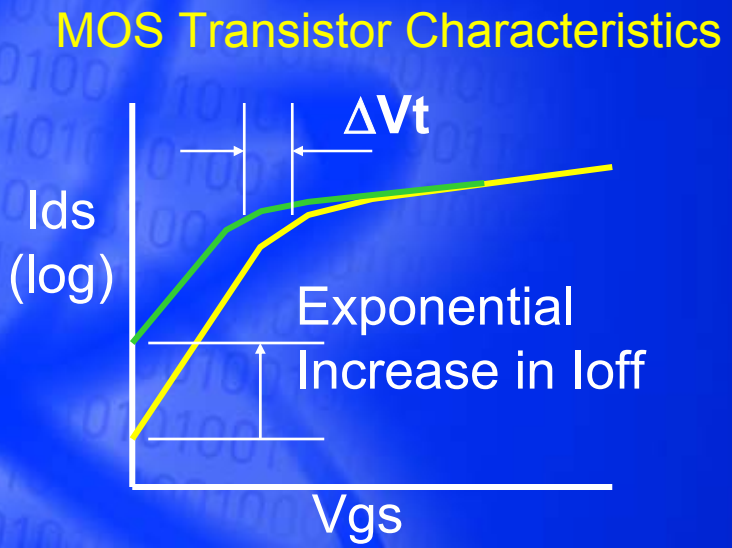

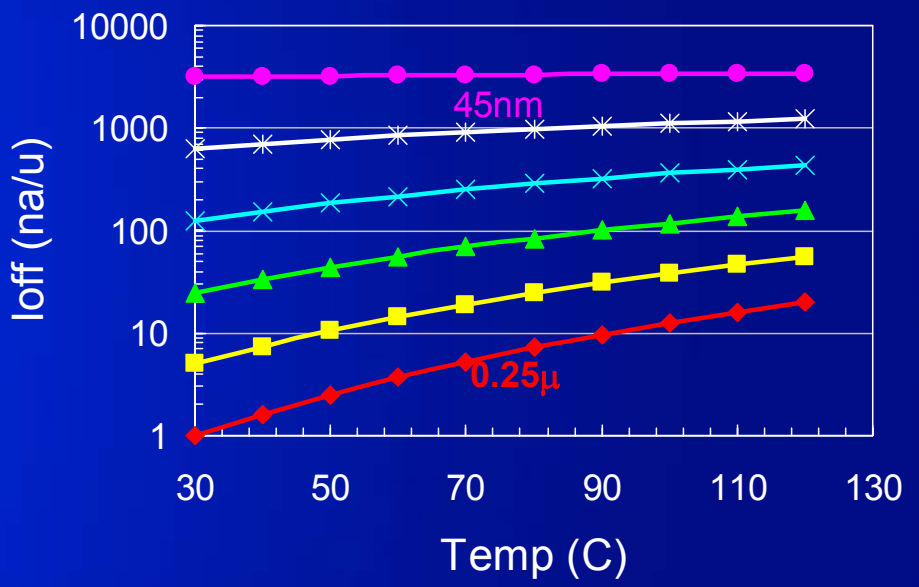

Transistors will be *dimmers*, not *switches*

SD Leakage Power

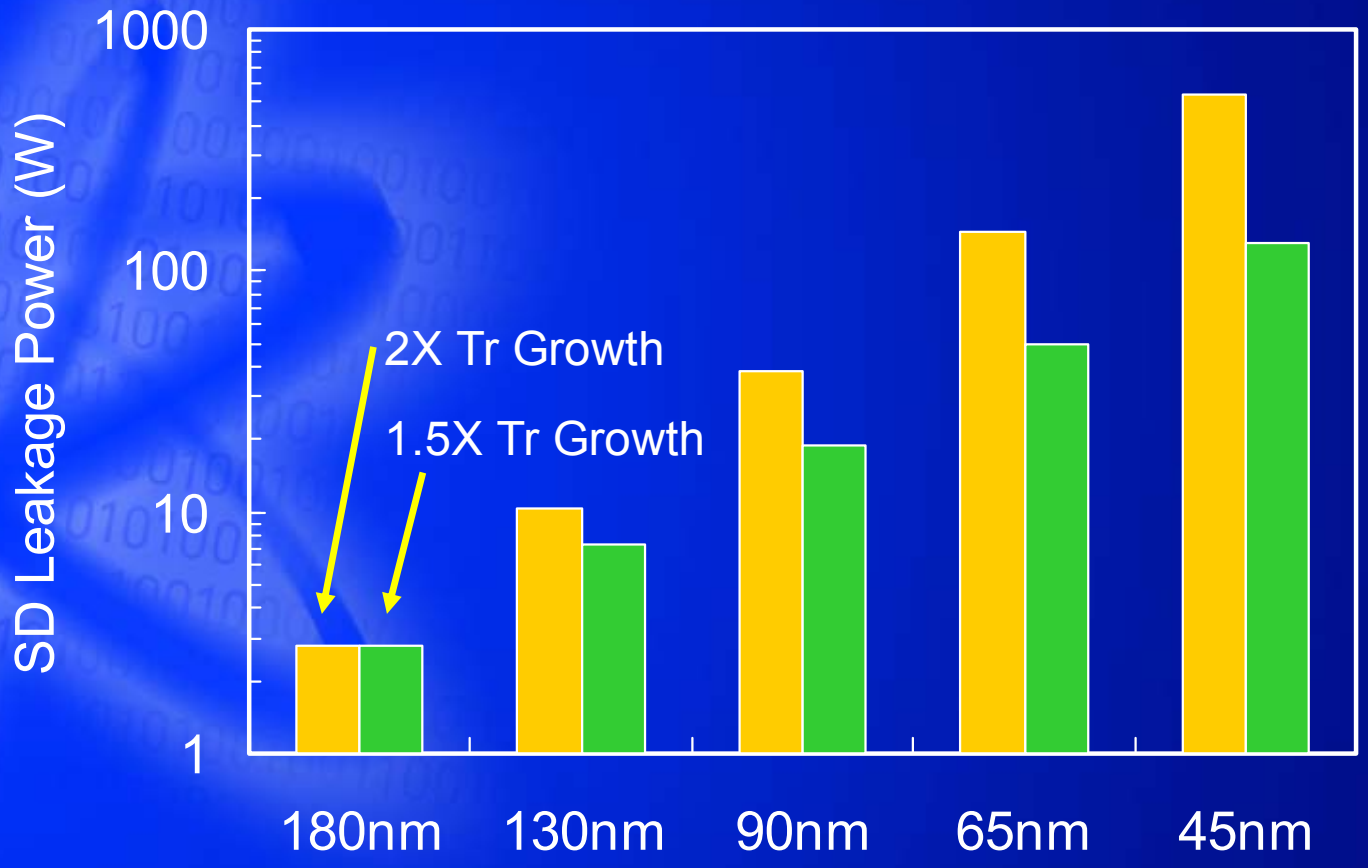

Assumptions:
- 15% Vdd scaling
- 5X Ioff scaling
  (Per generation)

Excessive sub-threshold leakage power

Variations in P, V, and T

Process

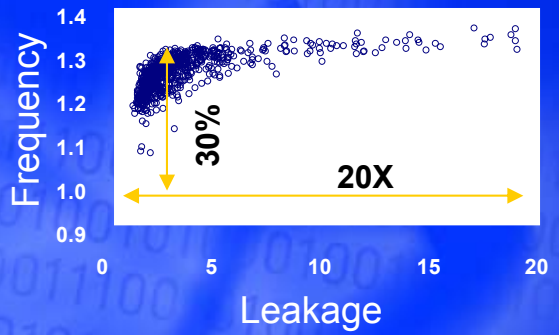

- Die-to-die variation
- Within-die variation
- Static for each die

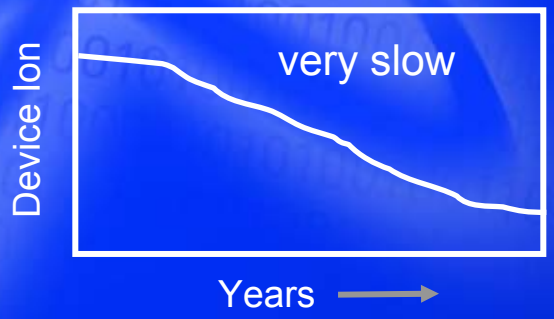

Time dependent degradation

Voltage

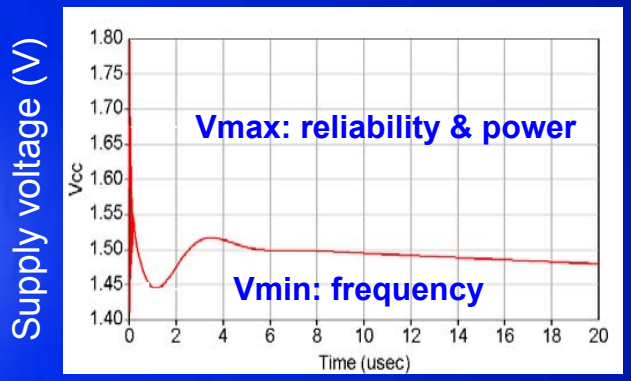

Time (usec)

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

Temperature

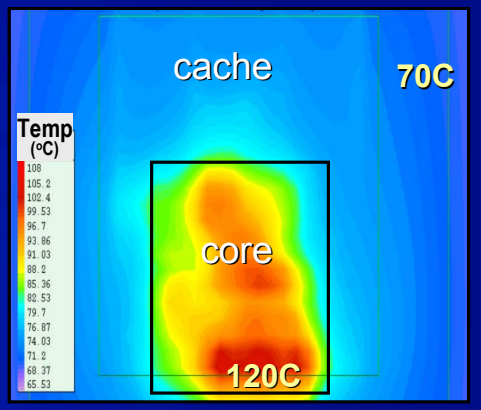

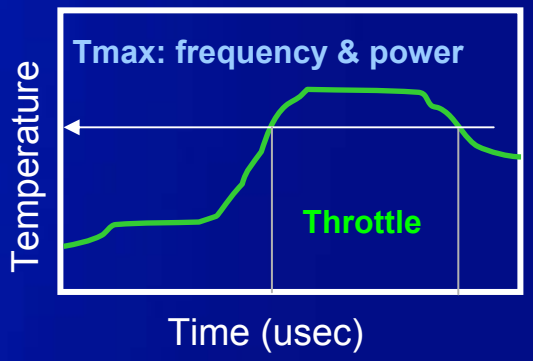

- Activity & ambient change
- Dynamic: 100-1000us
- Within-die variation

Impact of Critical Paths

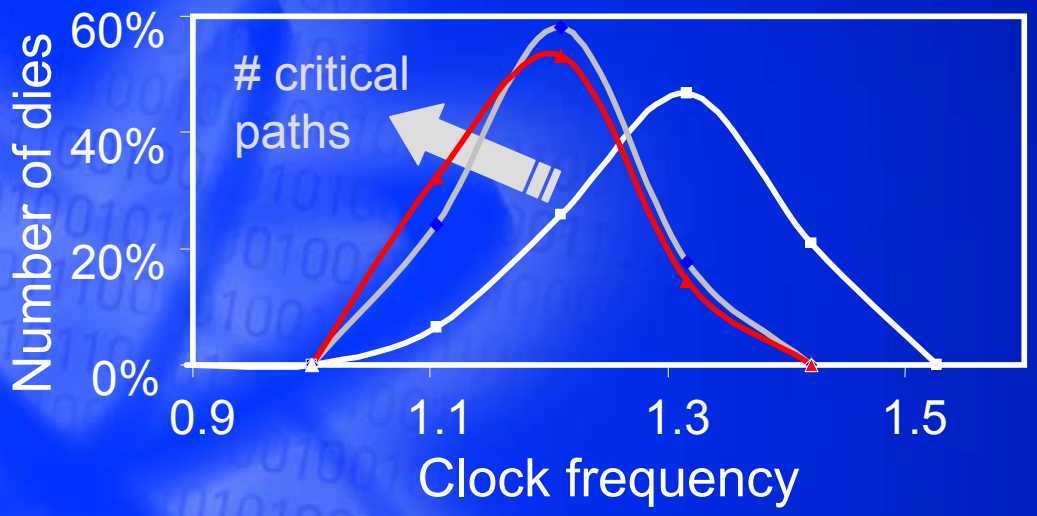

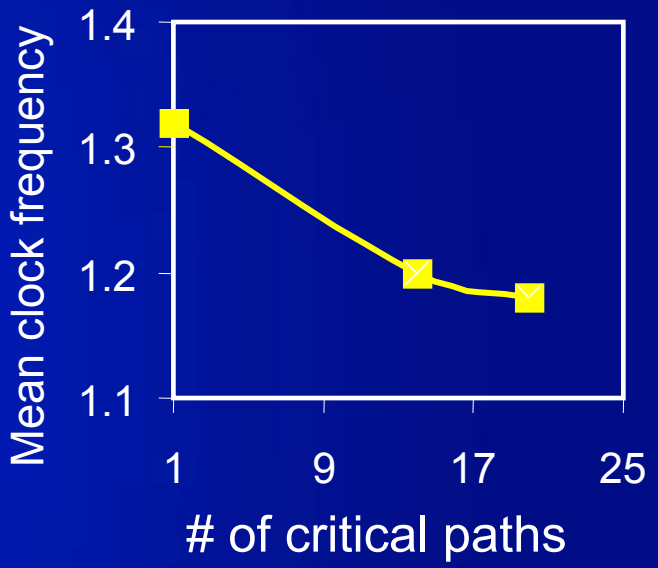

Impact of transistor parameter variations:

- Wide distribution of circuit frequency
- Lower mean freq with # of critical paths
- Encourages more localized, clustered designs

Tough Platform Demands

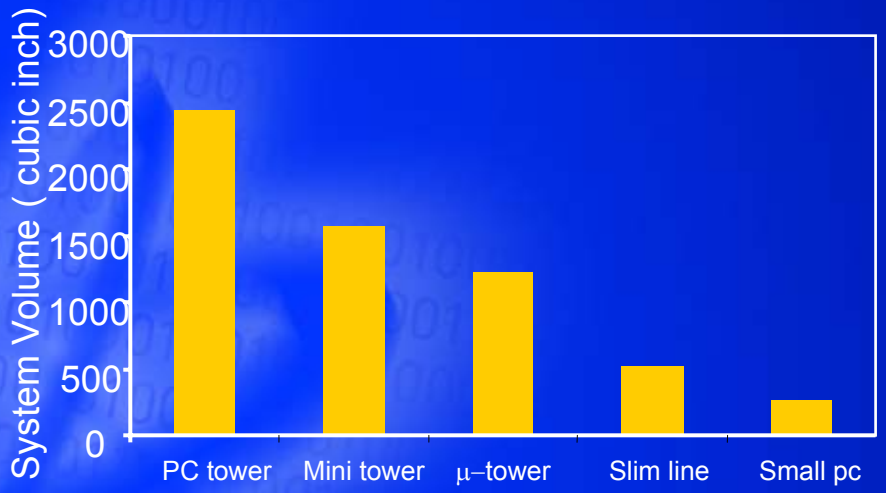

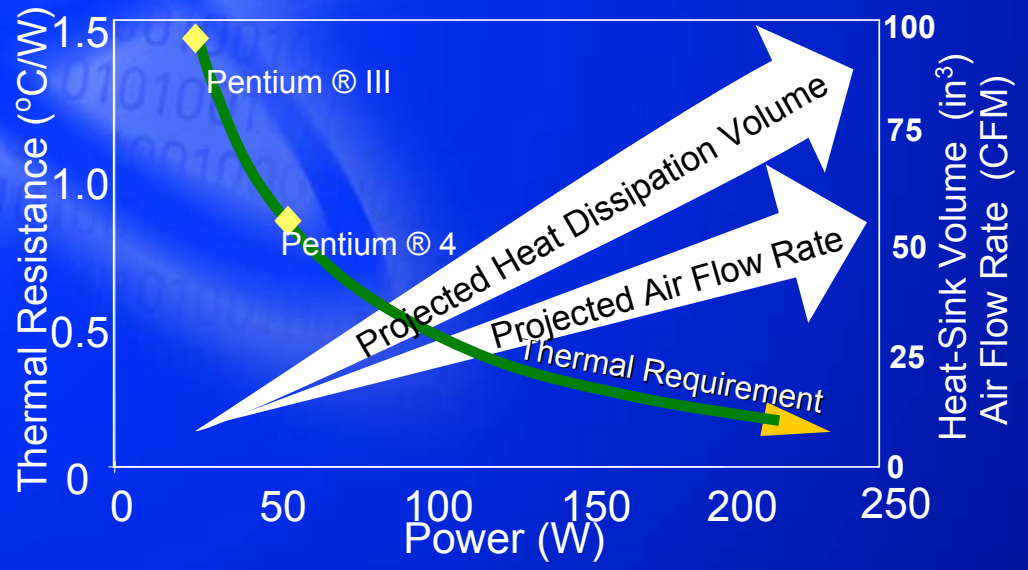

Shrinking volume

Reduced Noise

Yet, Higher Performance

At Odds

Reduced thermal budget

Higher heat sink volume

Higher air flow rate

BOM Cost Squeeze

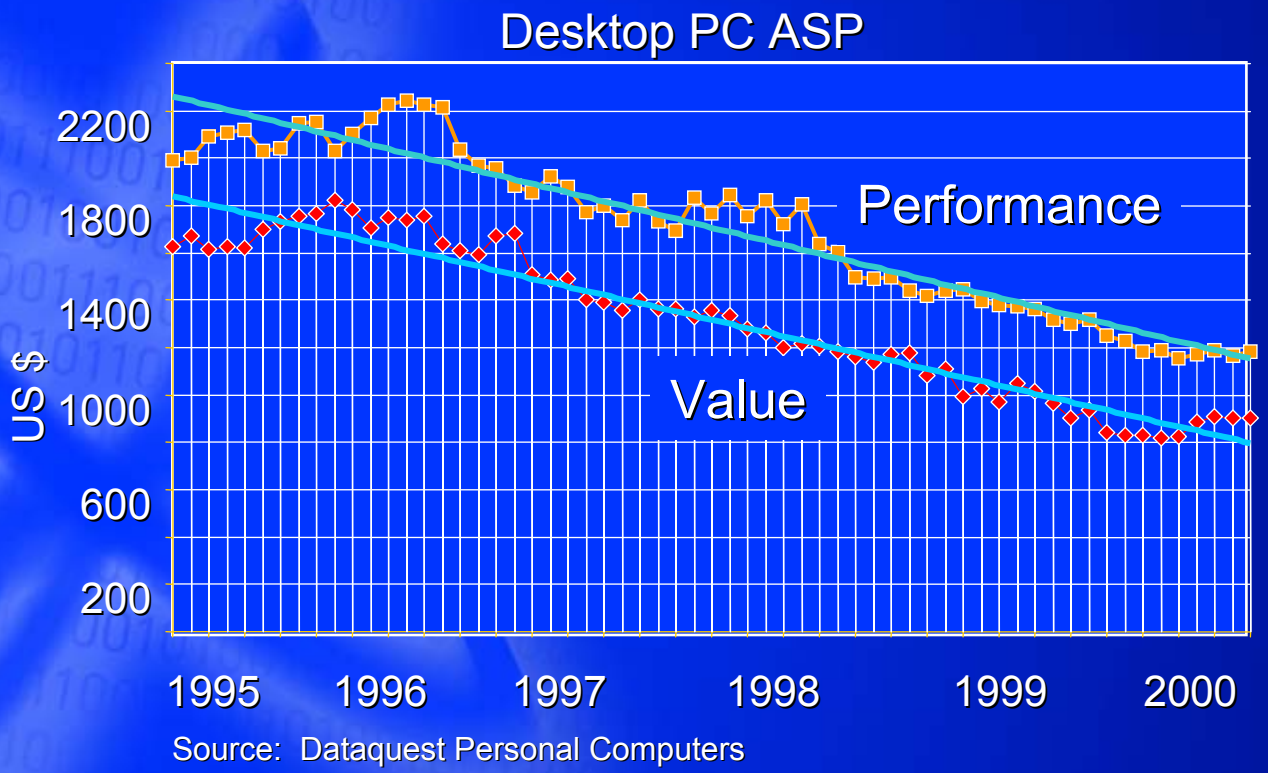

$2000 PC cost ('97)

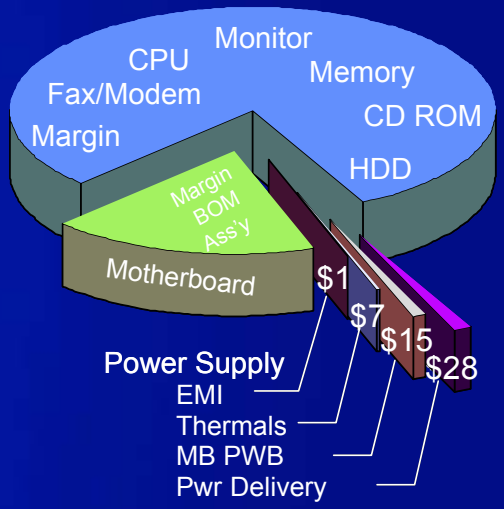

Budget for power and cooling is shrinking

Data Centers: Rack Mount Limits

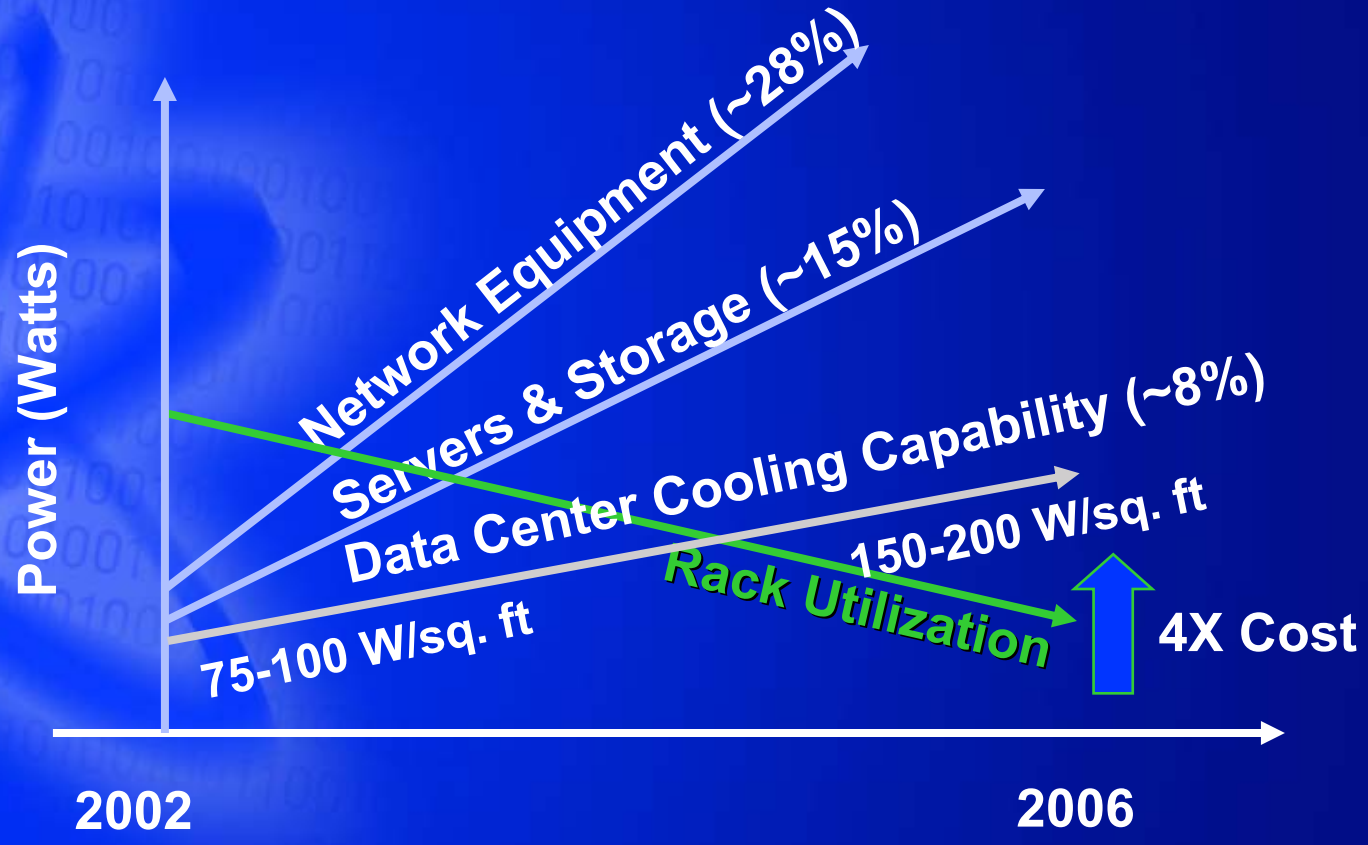

Wasted space and higher cost in future

Outline

- Technology scaling
- Looking at efficiency
- Making a right hand turn

Power Efficiency

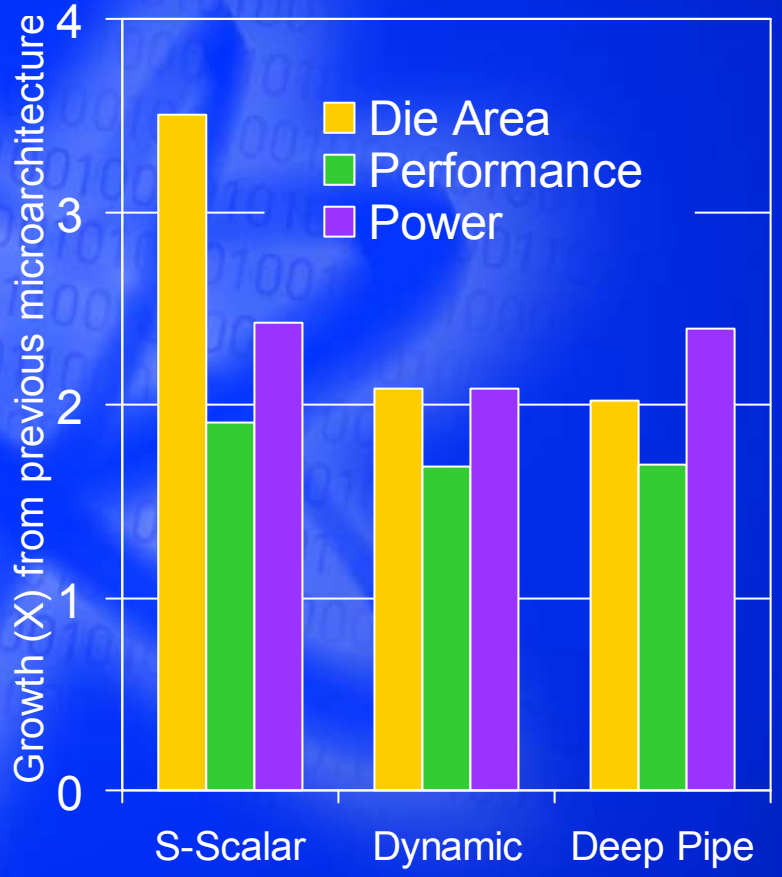

In the same process technology, compare:

Scalar ⇨ Super-scalar

⇨ Dynamic

⇨ Deep pipe

2-3X Growth in area

~1.4X Growth in Integer Performance

~1.7X Growth in Total Performance

2-2.5X Growth in Power

2-2.5x growth in power / generation

Energy Efficiency

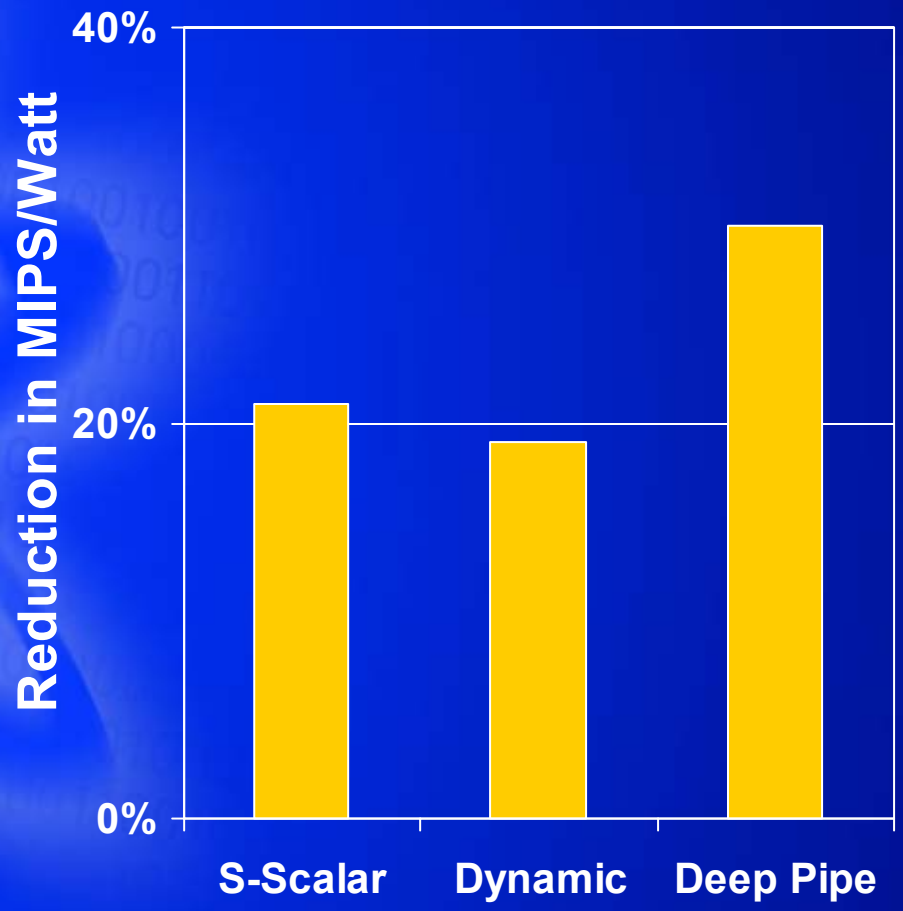

20-30% drop in energy efficiency / generation

Circuit Efficiency

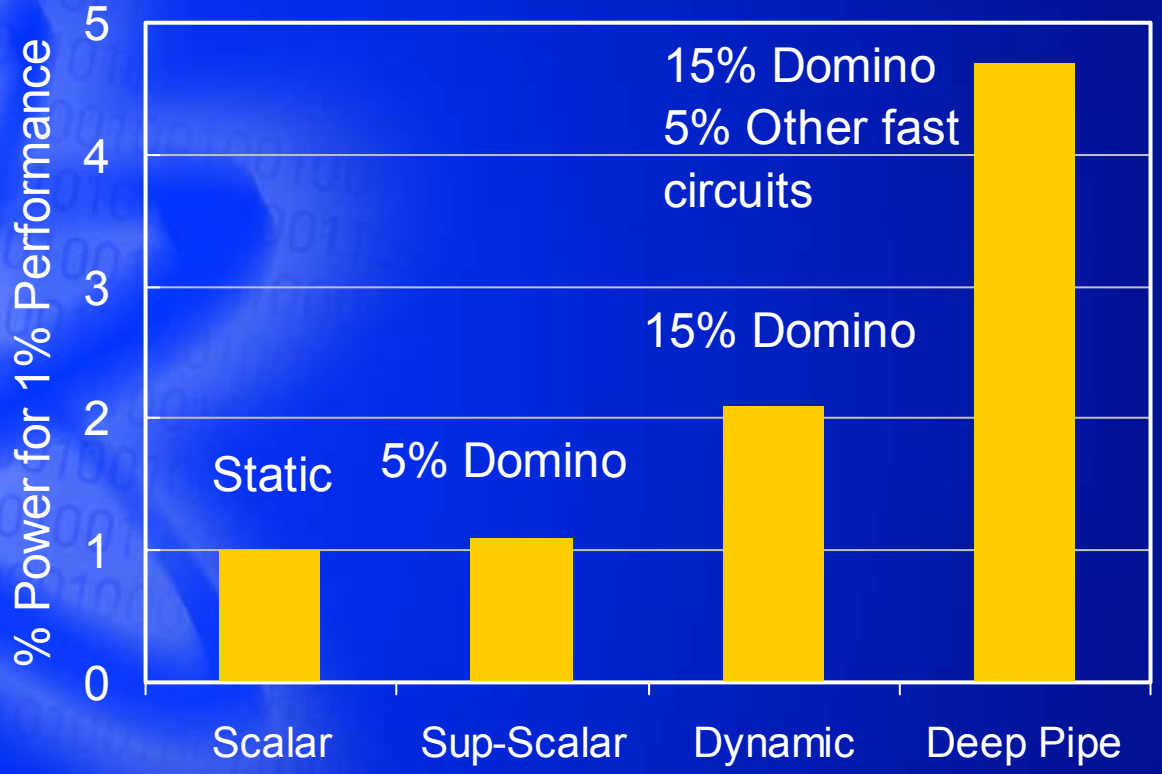

Assumptions:
Activity: Static = 0.2, Domino = 0.5
Clock consumes 40% of full chip power

Faster circuits contribute to power inefficiency

Outline

- Technology scaling
- Types of efficiency
- Making a right hand turn

Power Comes First

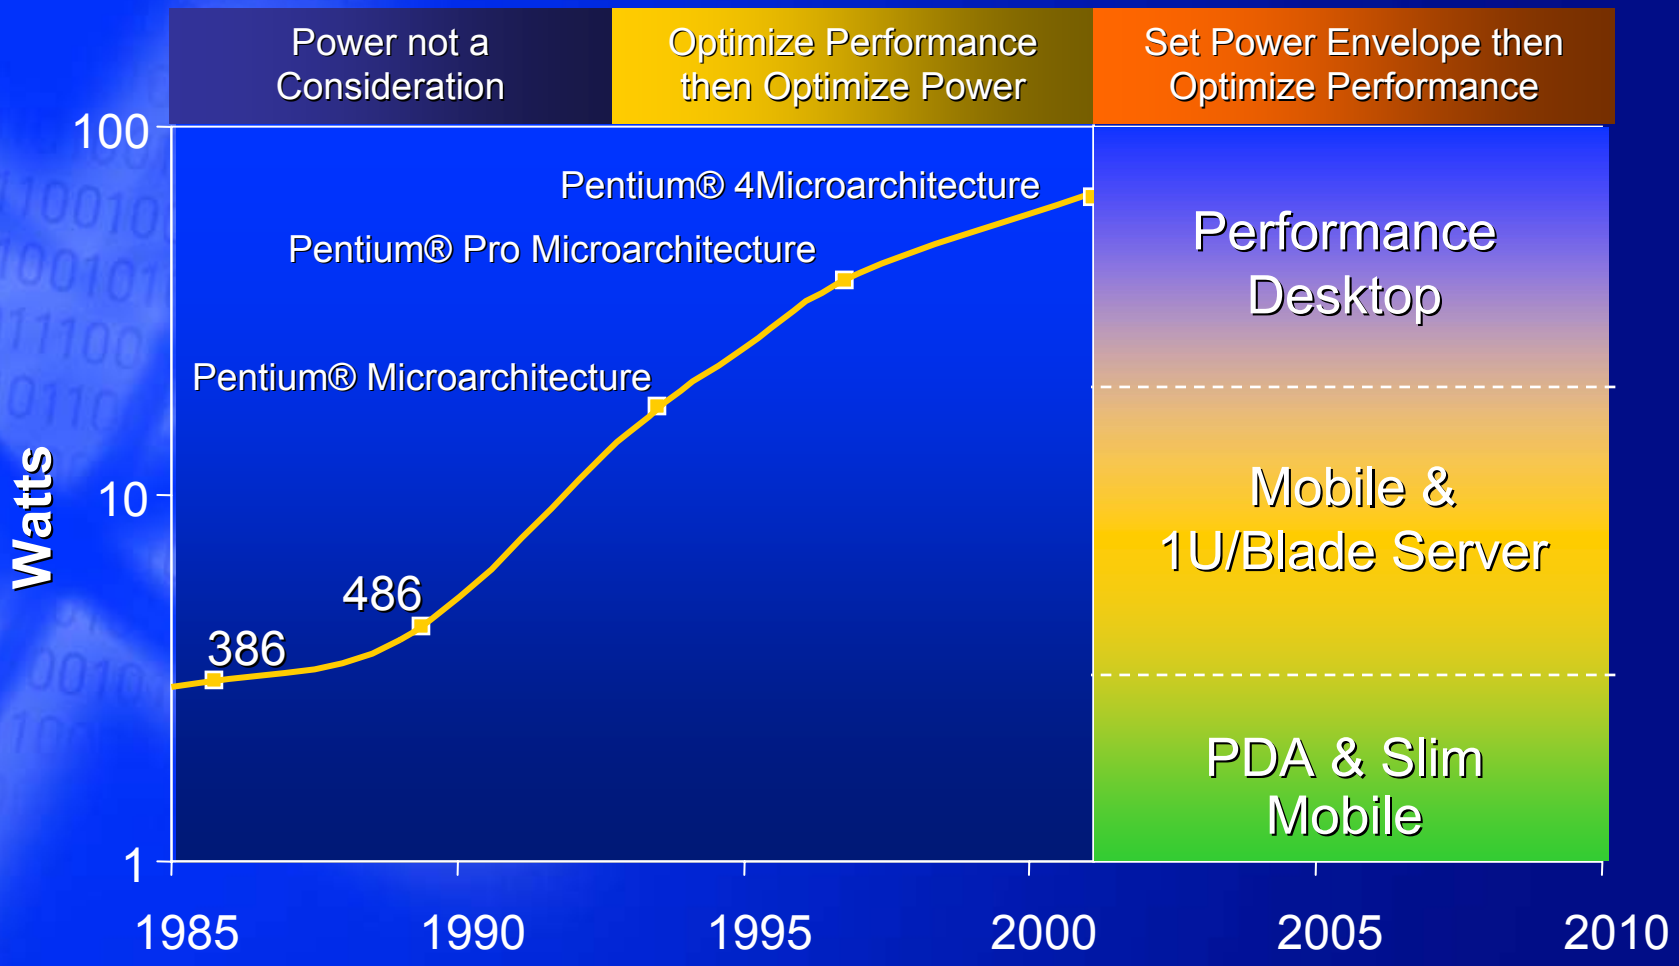

Business as usual is not an option

Low Power and High Performance

- ***Maximize battery life (fixed energy)***

$$Energy = Texec * Power \approx (1/Perf) * Power$$

Increasing the Performance by 10% and the Power by 10% will end up with same battery life

- ***Maximize performance within a given power envelope (Thermal constrains)***

$$f \approx K*V$$

$$Power = \alpha*C*V^2*f \approx \alpha*C*f^3$$

$$\Delta power/Power = ((f+\Delta f)^3 - f^3)/f^3 \approx 3\Delta f/f$$

$$Perf = IPC*f$$

➔ ***The right trade off between Performance and Power***

ΔIPC < 3ΔPower is the metric

"Less is More"

- Strive to accomplish the same task in less energy and less time
  - Higher performance at lower energy can always be traded with same performance at lower power
- Methodology works at all levels
  - Aggressive clock gating
  - Caching - dumb and smart
  - Better branch predictors
  - Smart work reduction
  - Prioritize useful over speculated work
  - Fixed functions

"Less is More" in Banias

- Improved branch prediction
  - Over 20% fewer branch mispredictions
- Dedicated stack manager
  - Over 5% uop reduction
- Uop fusion
  - Over 10% uop reduction
- Big L2 cache

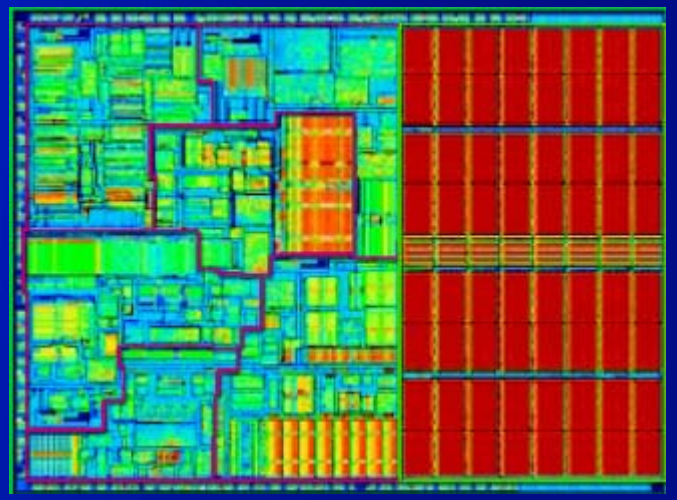

Achieving Higher Performance at Lower Power

Reducing Active Power

Multiple Supply Voltages

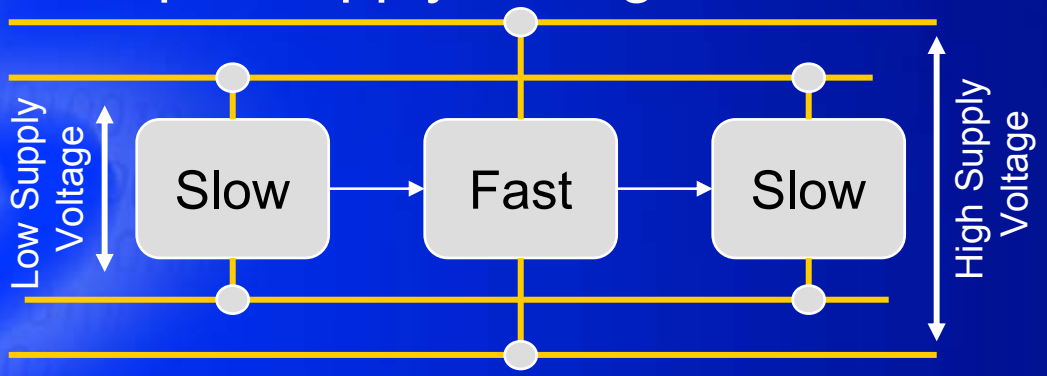

Throughput Oriented Design

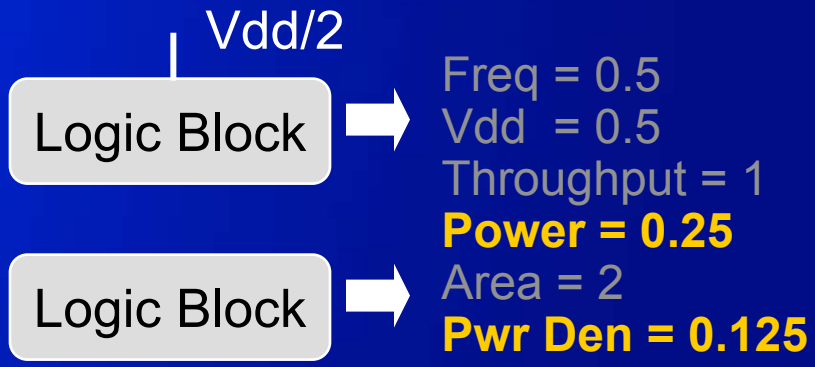

Critical Scheduling

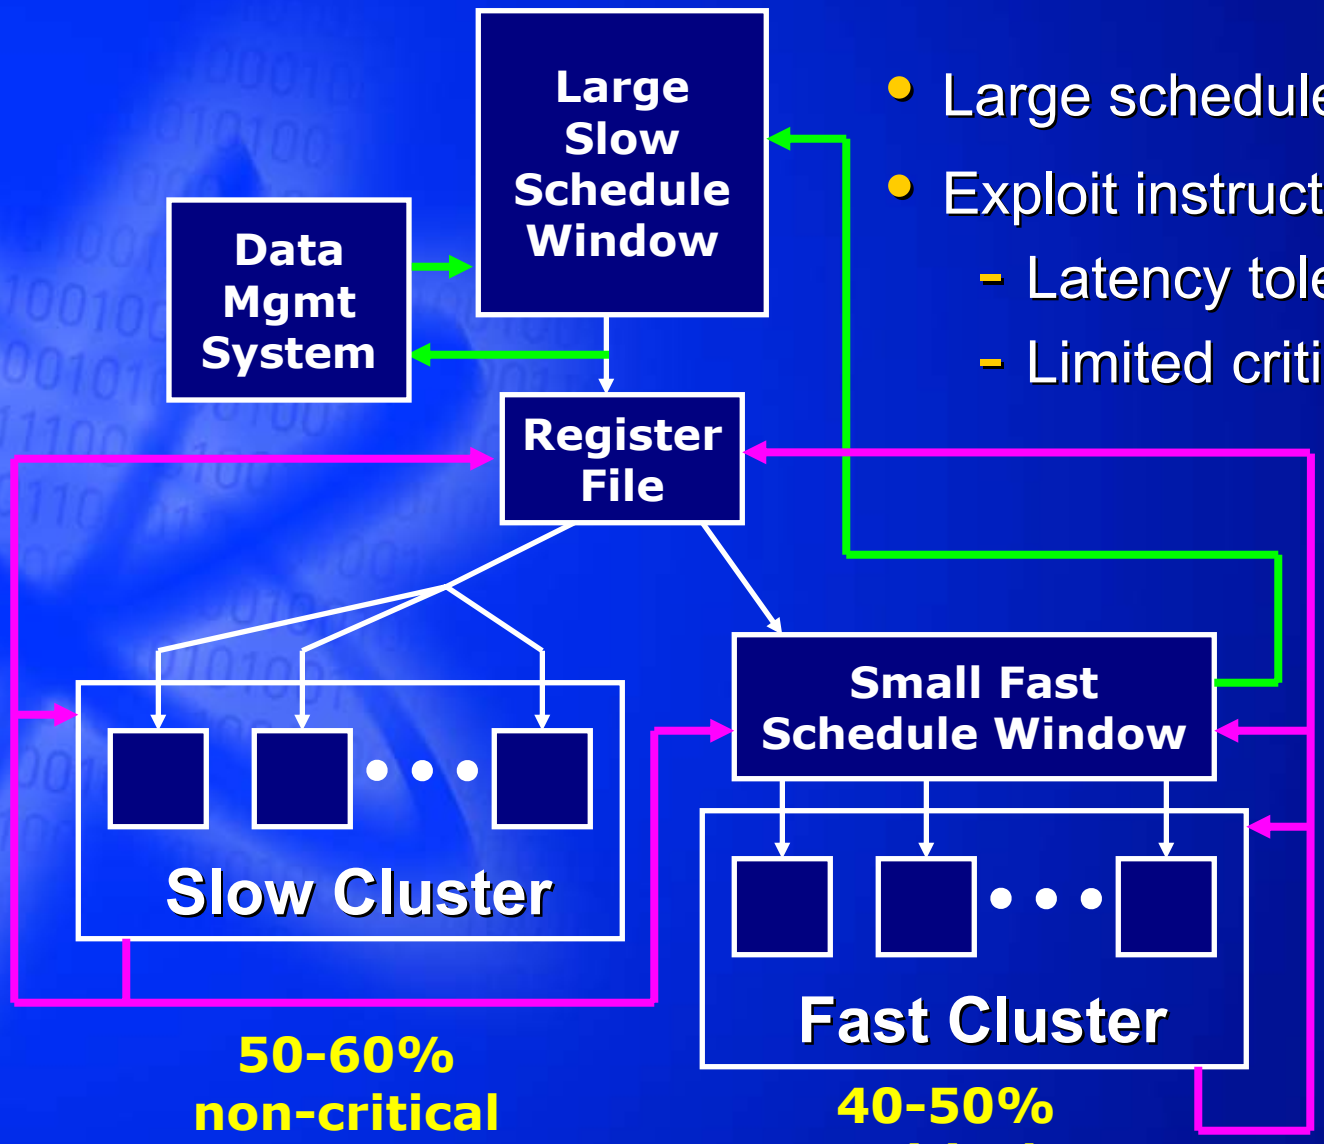

- Large schedule window (ILP)
- Exploit instruction criticality
  - Latency tolerant bypass
  - Limited critical resources

Recycling Waste

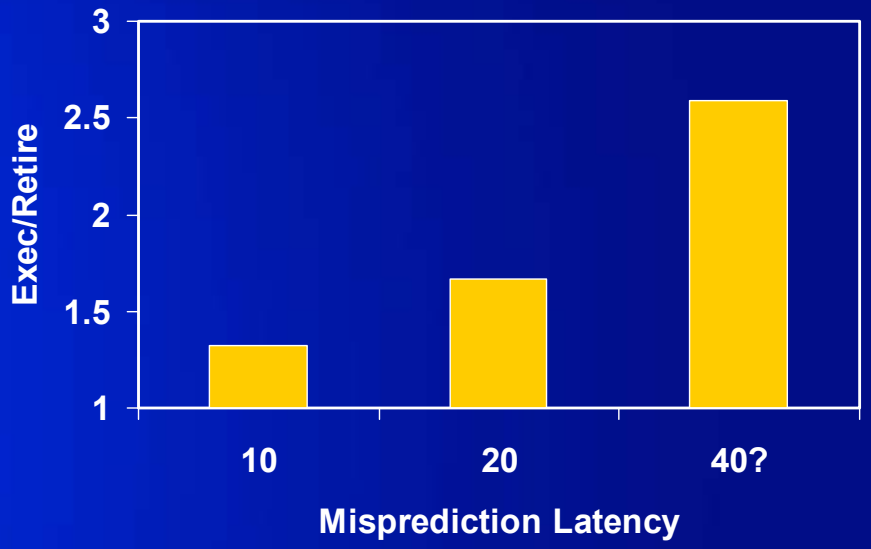

- Wasted execution
  - Spec Exec vs Retired
  - ~30% in 1st gen OOO
  - ~60% in 2nd gen OOO
  - ~160% in future
- Leverage info from wasted execution
- Improve branch Prediction
  - 30% reduction in mispredication rate
  - 18% to 48% less wasted execution
- Can we reuse the some of the execution result too?

Efficiency Through xMT

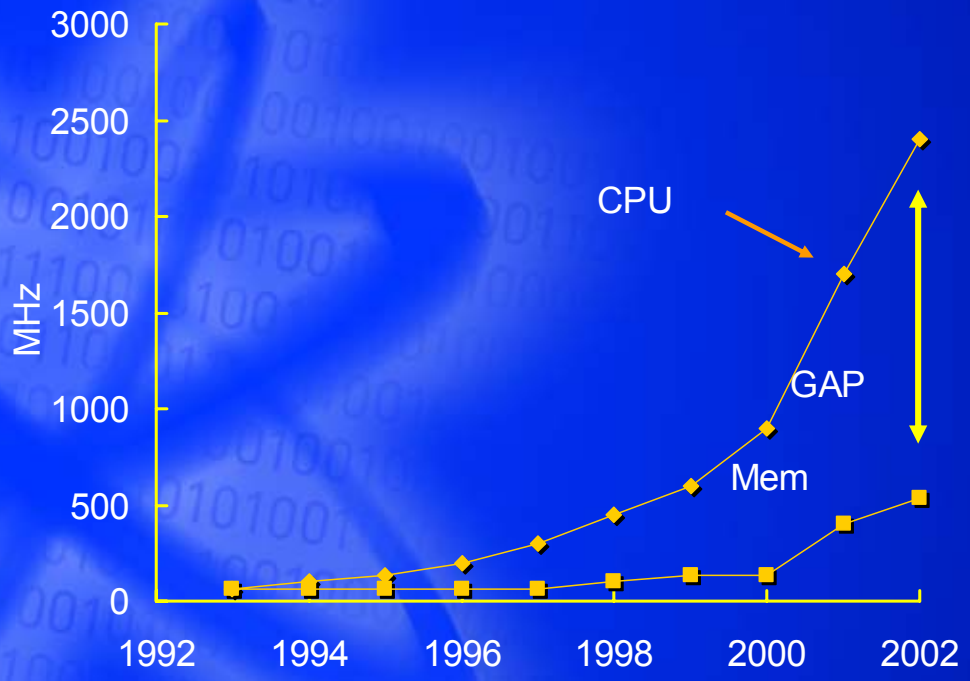

Thermals & power delivery designed for full HW utilization

Single Thread

Multi-thread

Multithreading improves performance without impacting thermals & power delivery

Power Efficient Asymmetric Threads

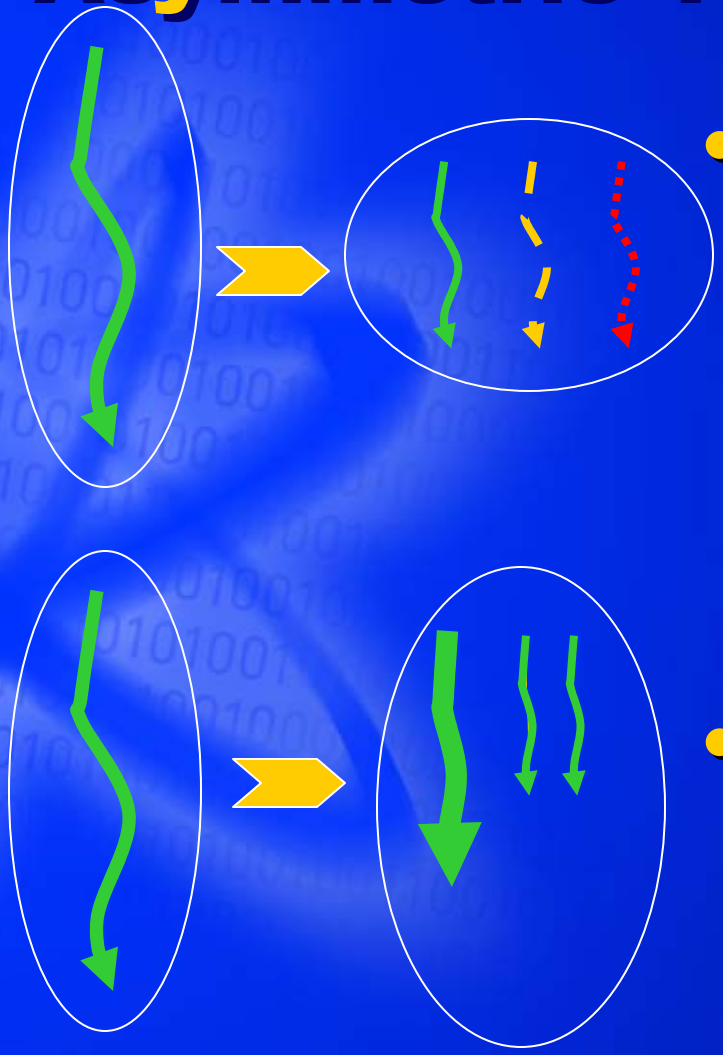

- Function AsymmetricThreading (w/co-processor ISA)
  - Partition single thread into a main thread with special function threads
  - Special function unit is more area and power efficient
- Performance Asymmetric Threading (ISA compatible)
  - Serial code on heavy core
  - Parallel code on smaller and power efficient core

Chip Multi-Processing

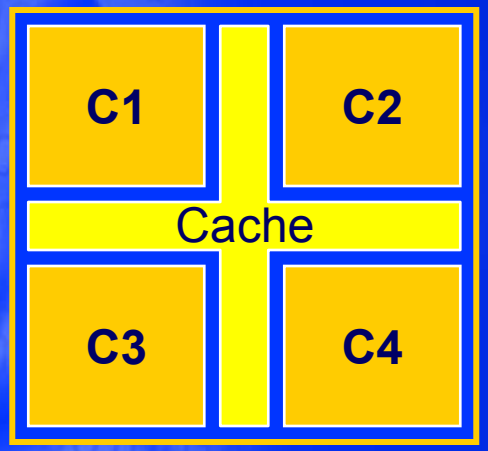

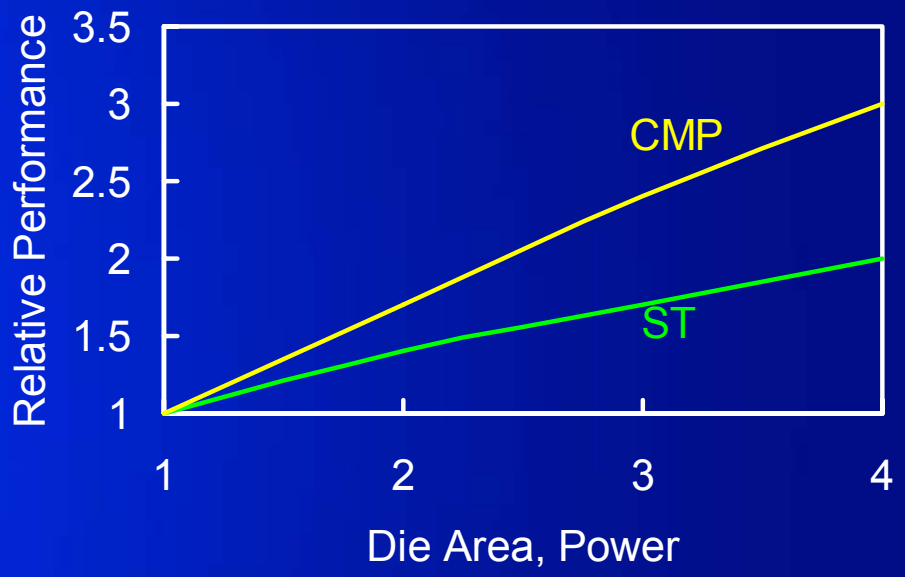

- Multi-core, each core MT
- Shared cache and front side bus
- Each core has different Vdd & Freq
- Core re-cycling to spread hot spots
- Lower junction temperature

Summary

- Technology scaling can and will continue
- Challenges to power and energy efficiency are real but surmountable..
- ..through evolutionary approaches to circuits and microarchitecture