VLSI Design Challenges for Gigascale Integration

Shekhar Borkar

Intel Corp.

October 25, 2005

Outline

- Technology scaling challenges
- Circuit and design solutions
- Microarchitecture advances
- Multi-everywhere
- Summary

Goal: 10 TIPS by 2015

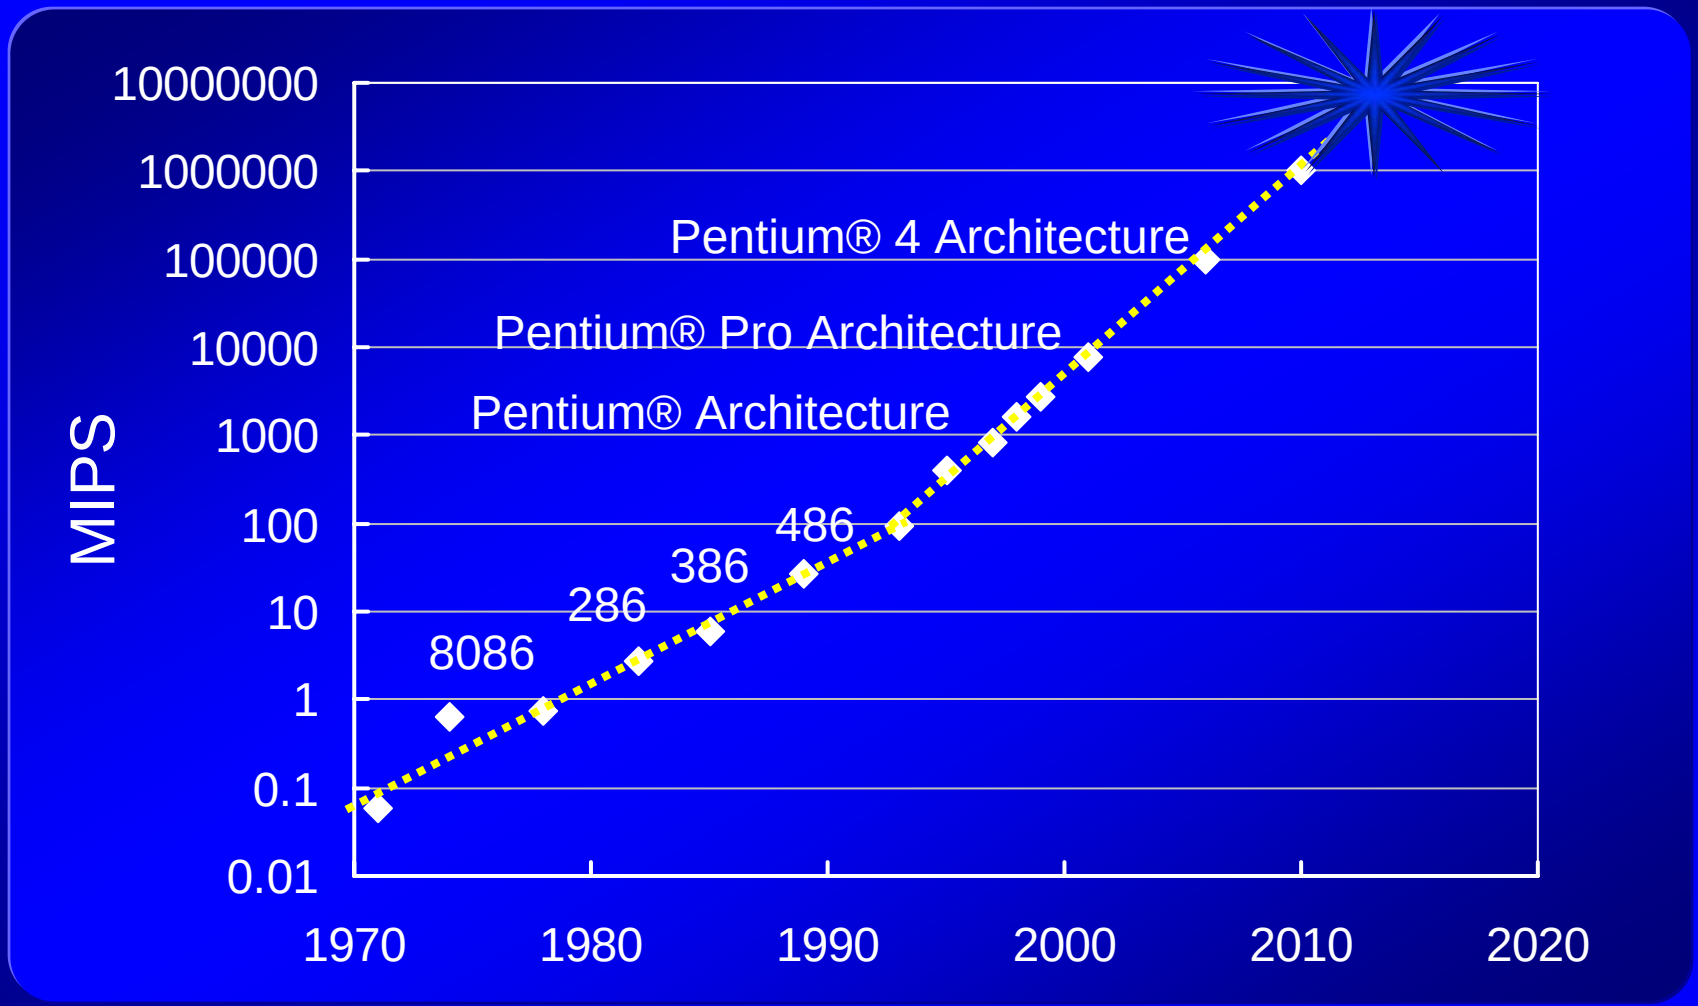

How do you get there?

Technology Scaling

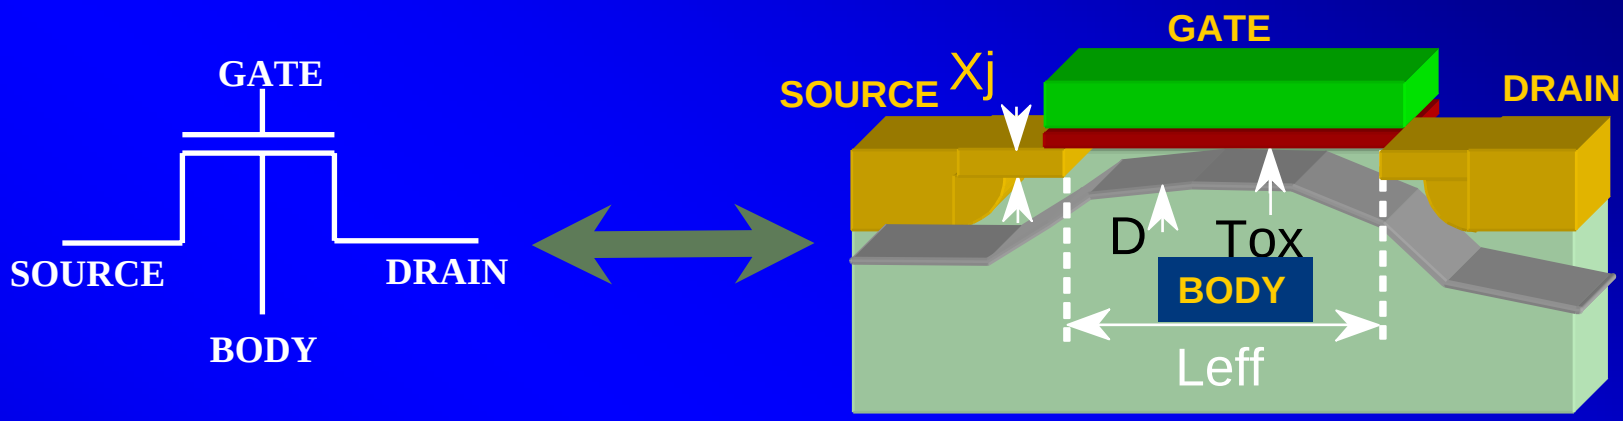

Dimensions scale down by 30%	Doubles transistor density
Oxide thickness scales down	Faster transistor, higher performance
Vdd & Vt scaling	Lower active power

Scaling will continue, but with challenges!

Technology Outlook

High Volume Manufacturing	2004	2006	2008	2010	2012	2014	2016	2018
Technology Node (nm)	90	65	45	32	22	16	11	8
Integration Capacity (BT)	2	4	8	16	32	64	128	256
Delay = CV/I scaling	0.7	~0.7	>0.7	Delay scaling will slow down				
Energy/Logic Op scaling	>0.35	>0.5	>0.5	Energy scaling will slow down				
Bulk Planar CMOS	High Probability						Low Probability	
Alternate, 3G etc	Low Probability						High Probability	
Variability	Medium			High			Very High	
ILD (K)	~3	<3	Reduce slowly towards 2-2.5					
RC Delay	1	1	1	1	1	1	1	1
Metal Layers	6-7	7-8	8-9	0.5 to 1 layer per generation				

The Leakage(s)...

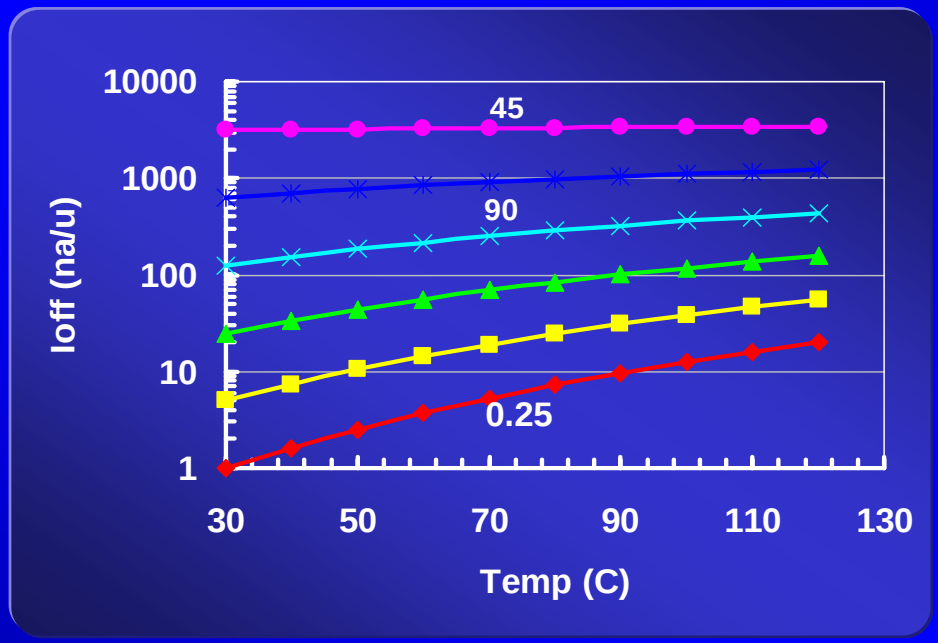

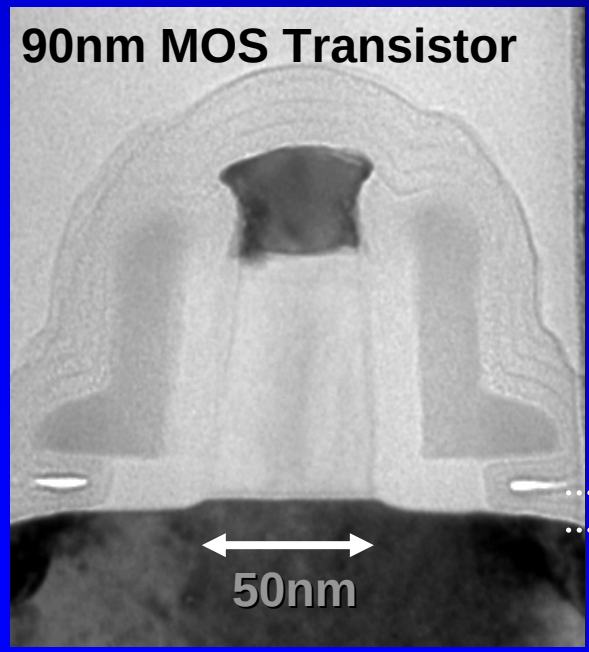

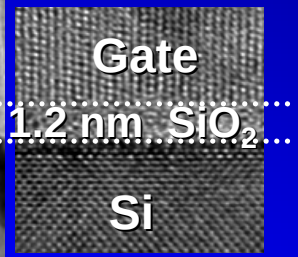

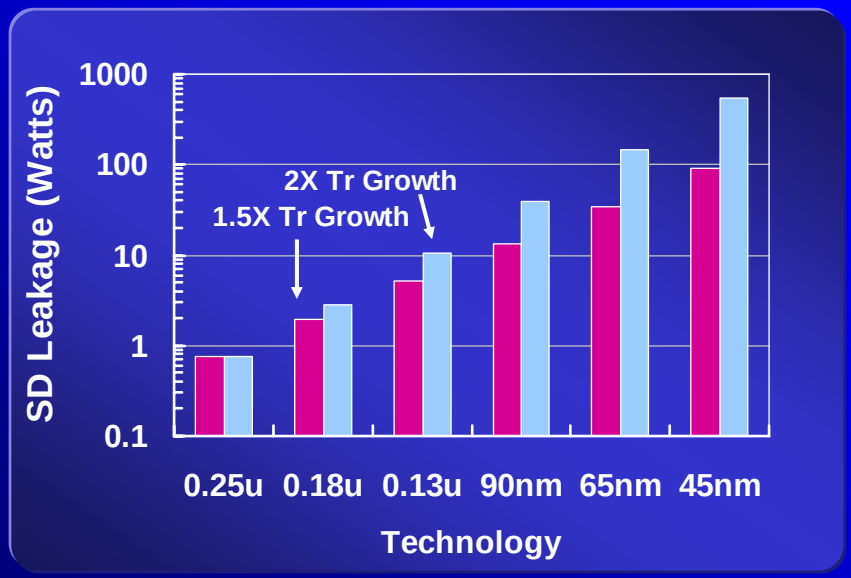

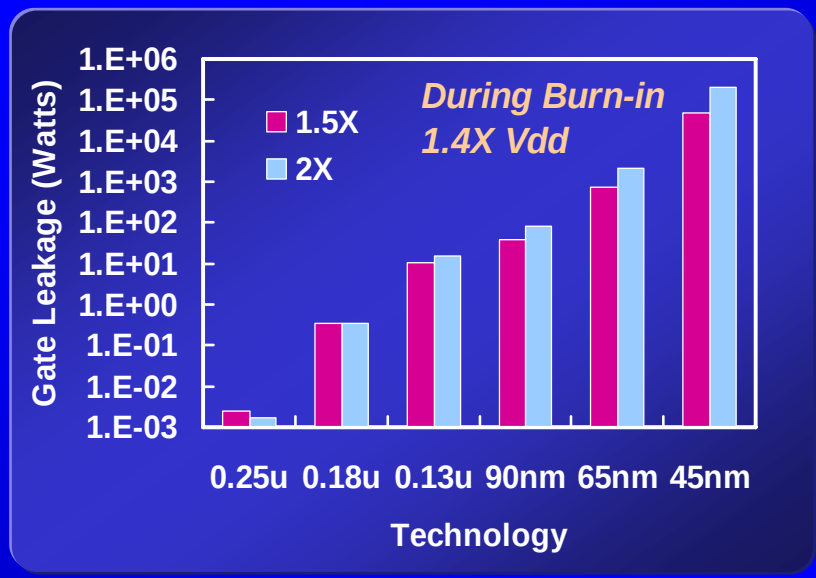

Must Fit in Power Envelope

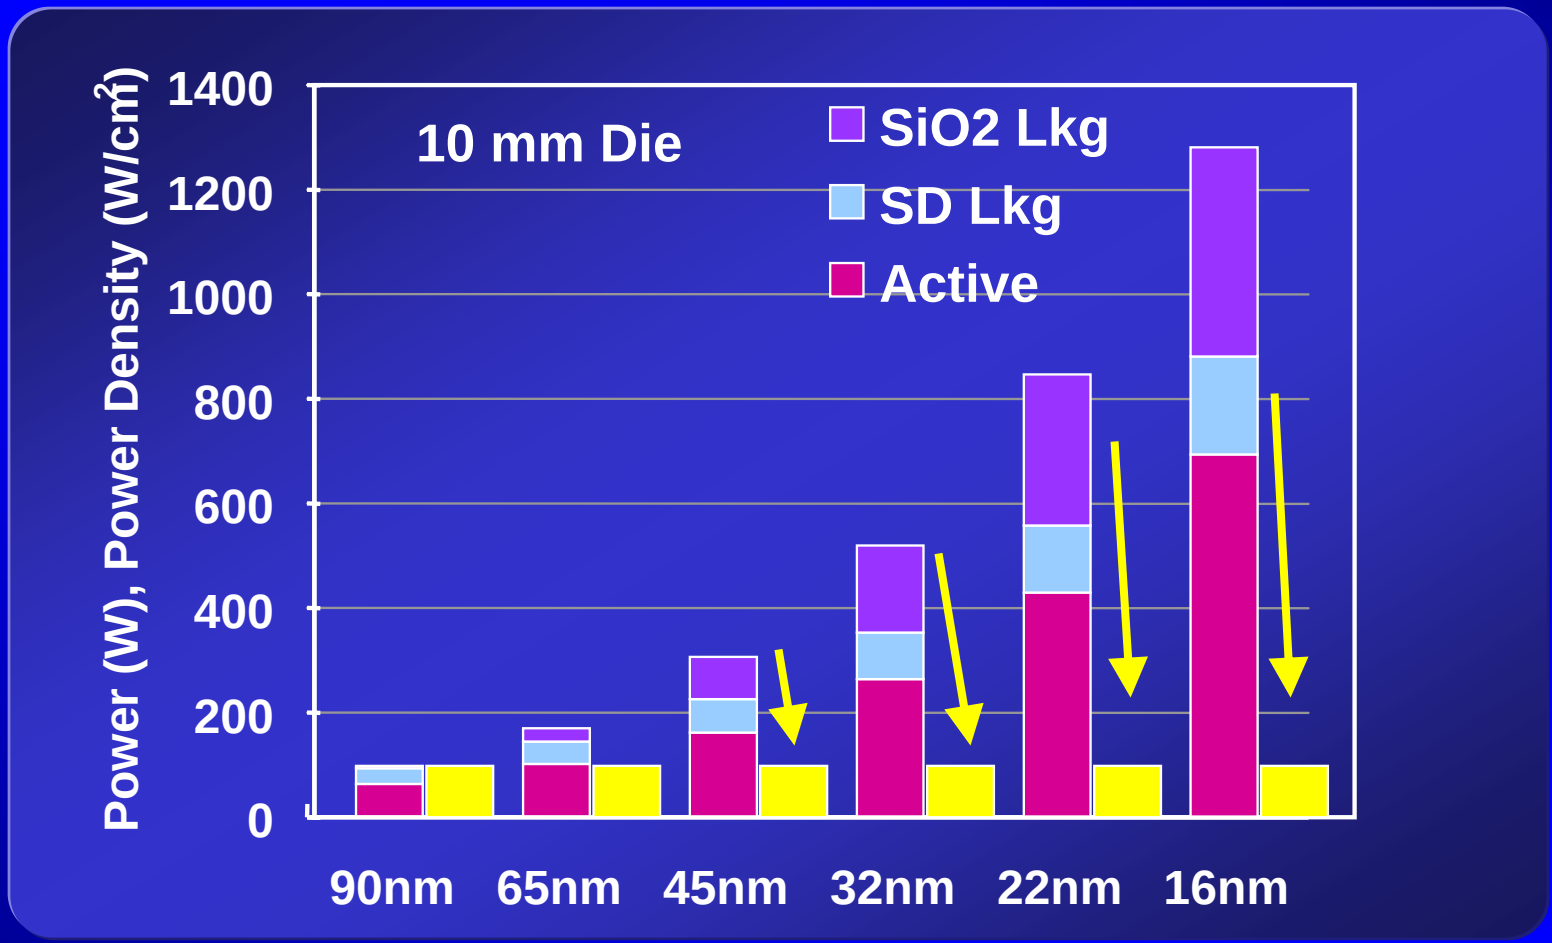

Technology, Circuits, and Architecture to constrain the power

Solutions

- Move away from Frequency alone to deliver performance
- More on-die memory
- Multi-everywhere
  - Multi-threading
  - Chip level multi-processing
- Throughput oriented designs
- Valued performance by higher level of integration
  - Monolithic & Polylithic

Leakage Solutions

For a few generations, then what?

Planar Transistor

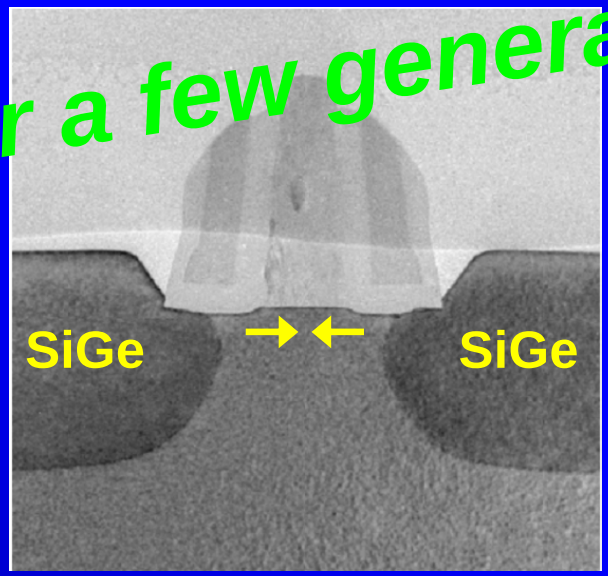

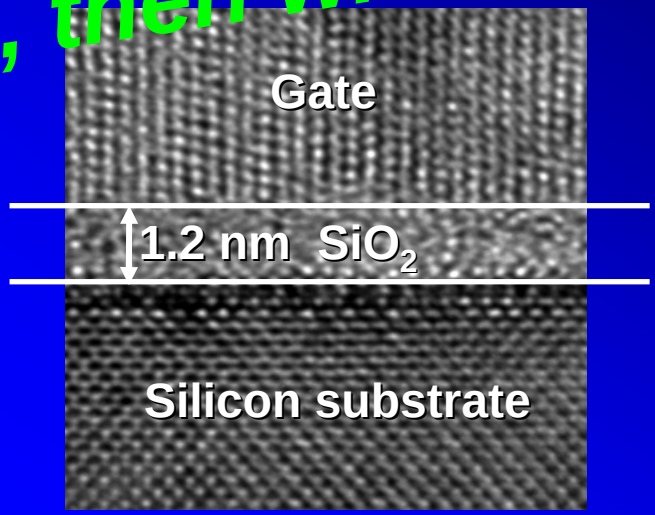

Tri-gate Transistor

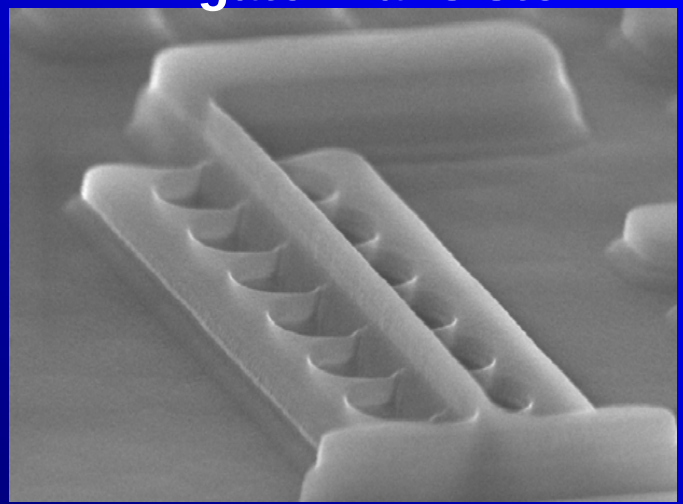

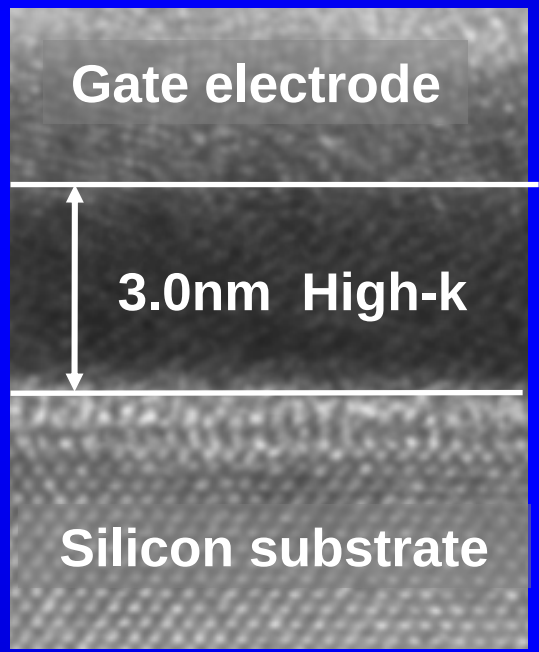

Active Power Reduction

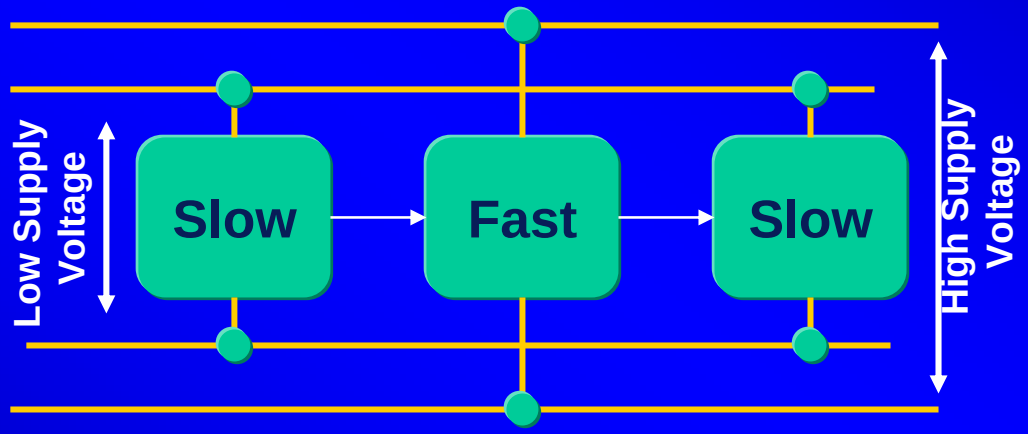

Multiple Supply Voltages

Throughput Oriented Designs

Vdd

Logic Block

Freq = 1
Vdd = 1
Throughput = 1
Power = 1
Area = 1
Pwr Den = 1

Vdd/2

Logic Block

Logic Block

Freq = 0.5
Vdd = 0.5
Throughput = 1
Power = 0.25
Area = 2
Pwr Den = 0.125

Leakage Control

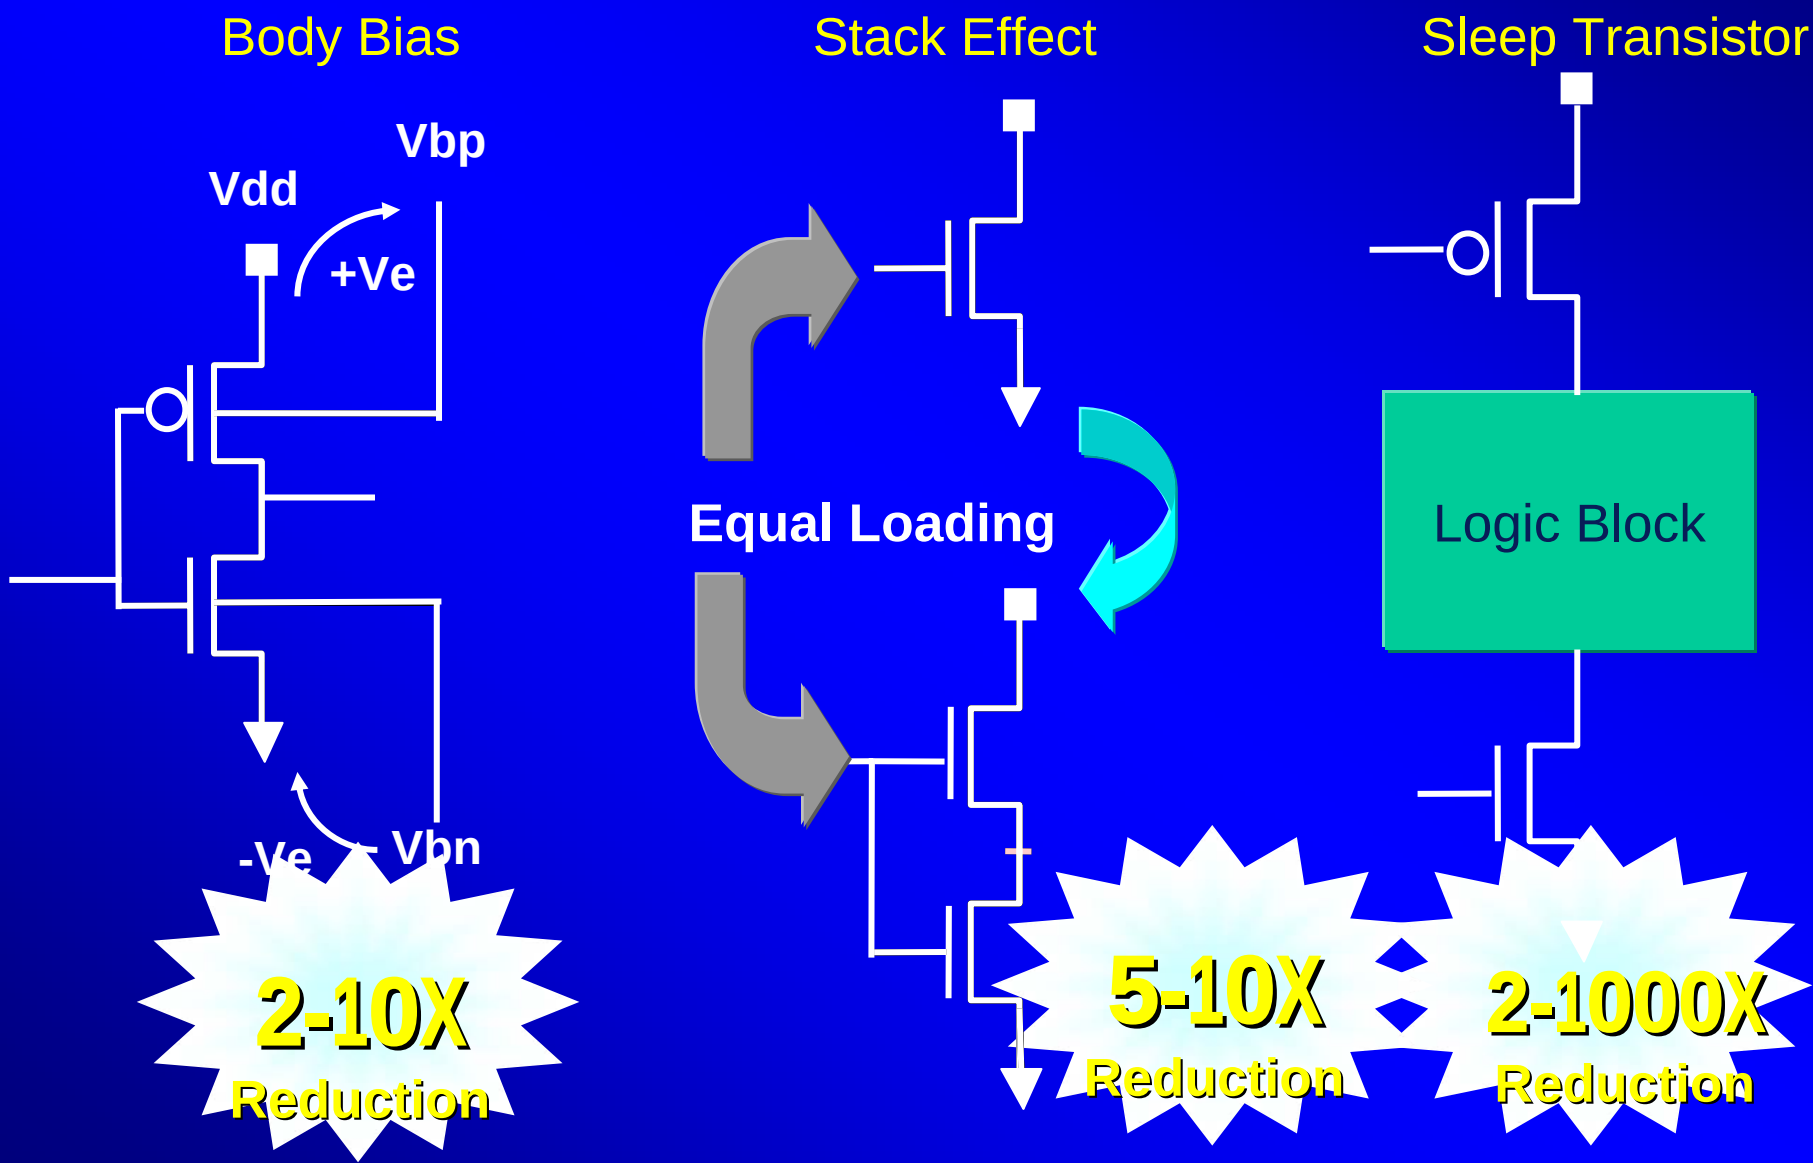

Optimum Frequency

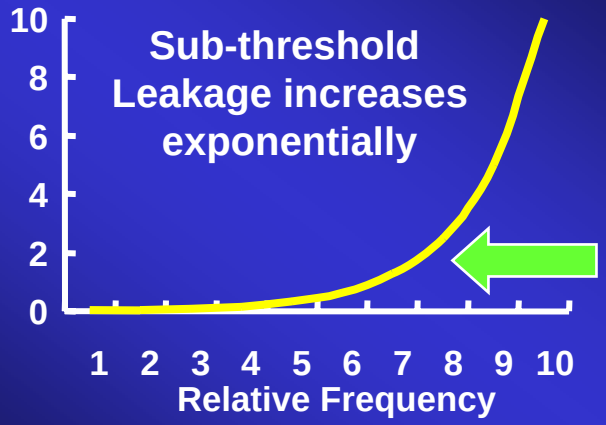

Process Technology

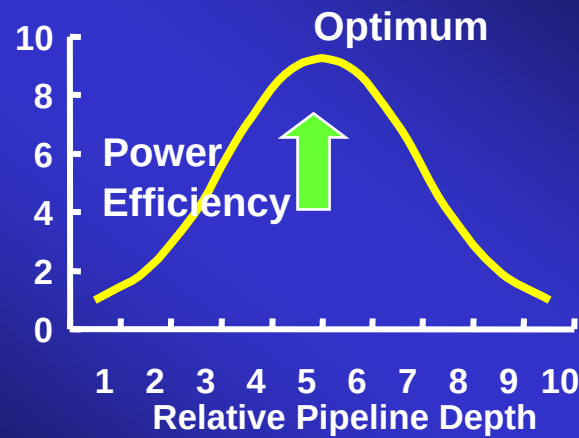

Pipeline Depth

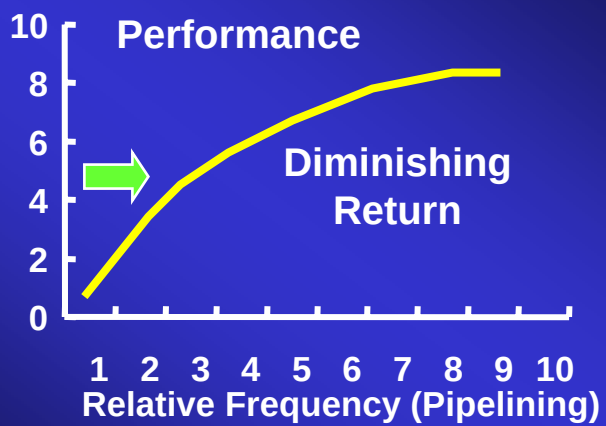

Pipeline & Performance

Maximum performance with
- **Optimum pipeline depth**
- **Optimum frequency**

Memory Latency

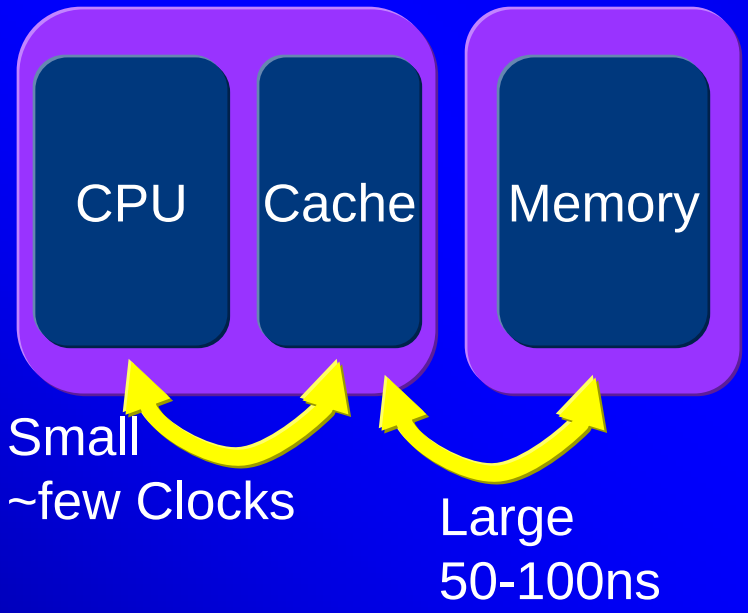

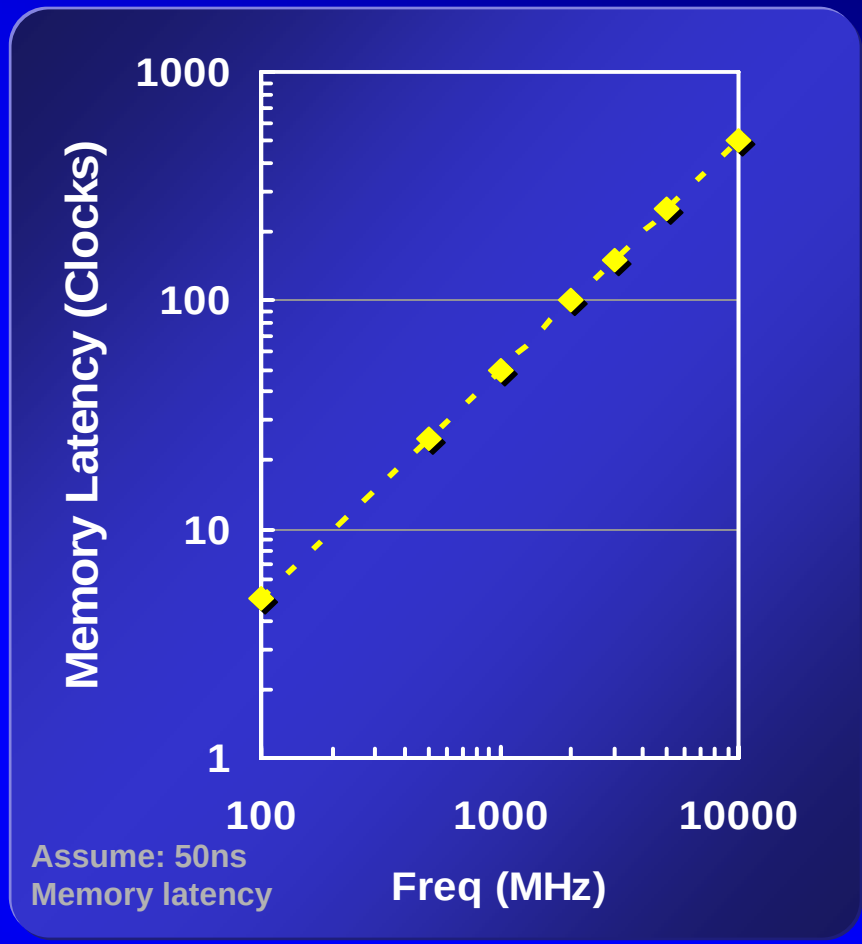

Cache miss hurts performance
Worse at higher frequency

Increase on-die Memory

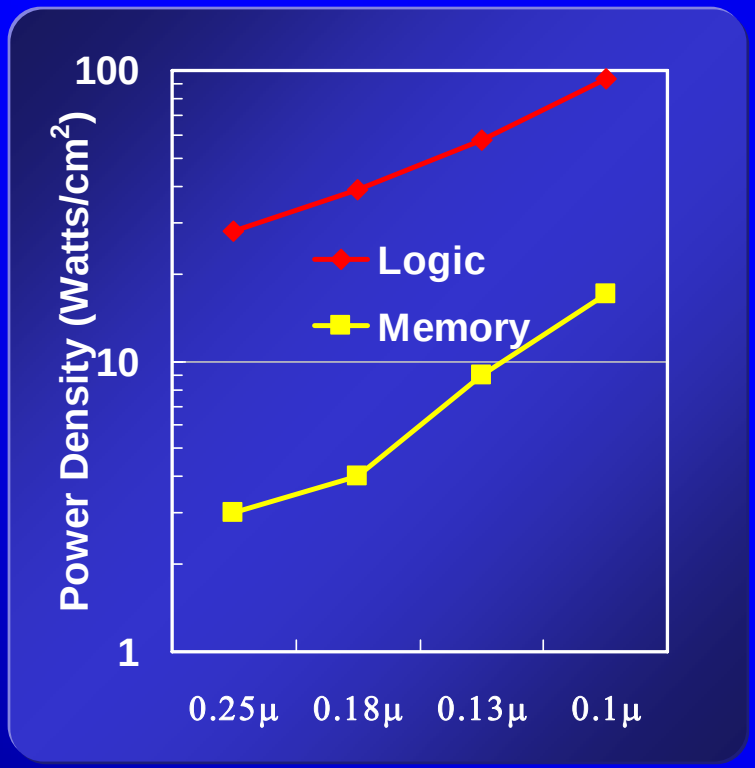

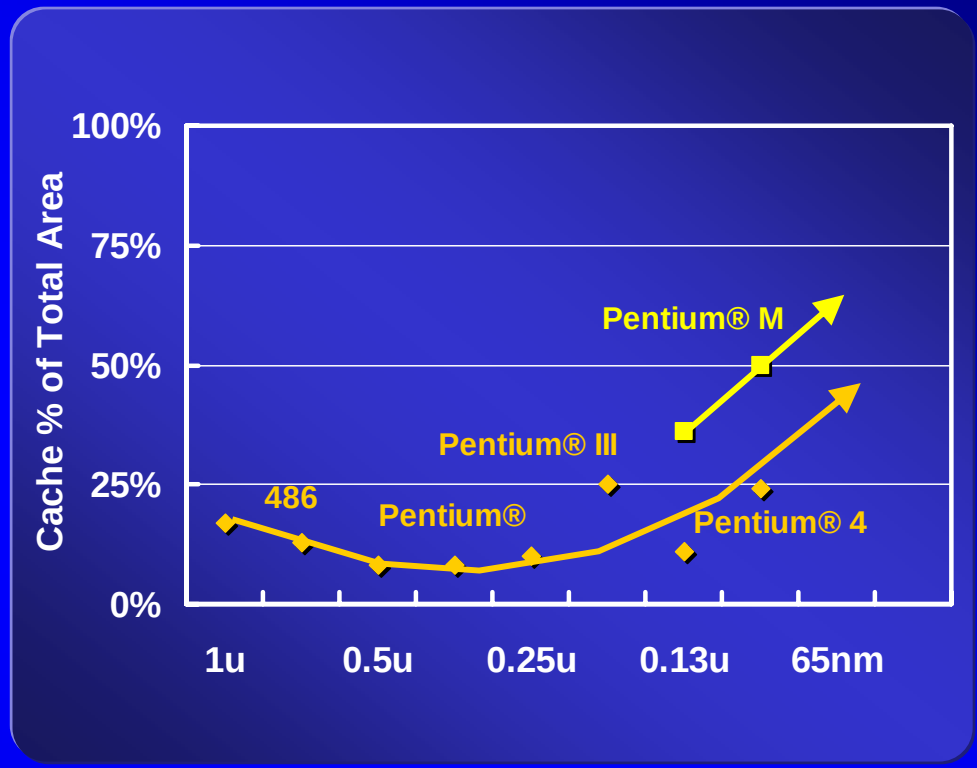

Large on die memory provides:

1. Increased Data Bandwidth & Reduced Latency
2. Hence, higher performance for much lower power

Multi-threading

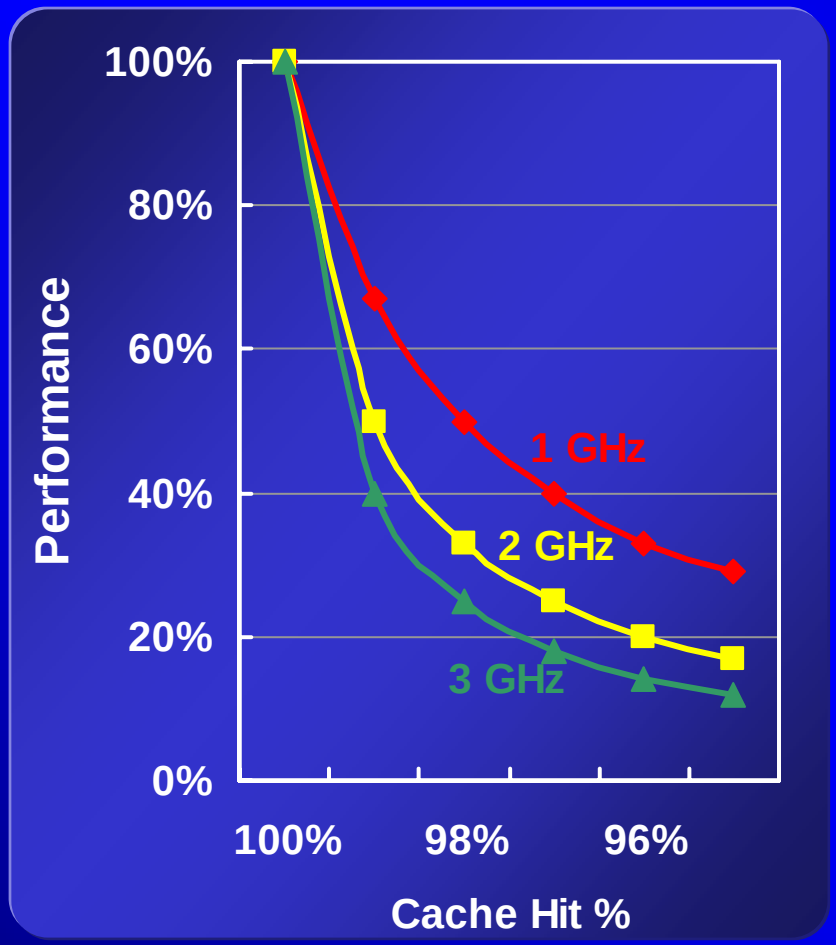

Thermals & Power Delivery designed for full HW utilization

Single Thread

Full HW Utilization

ST | **Wait for Mem**

Multi-Threading

MT1 | **Wait for Mem**

MT2 | **Wait**

MT3

Multi-threading improves performance without impacting thermals & power delivery

Single Core Power/Performance

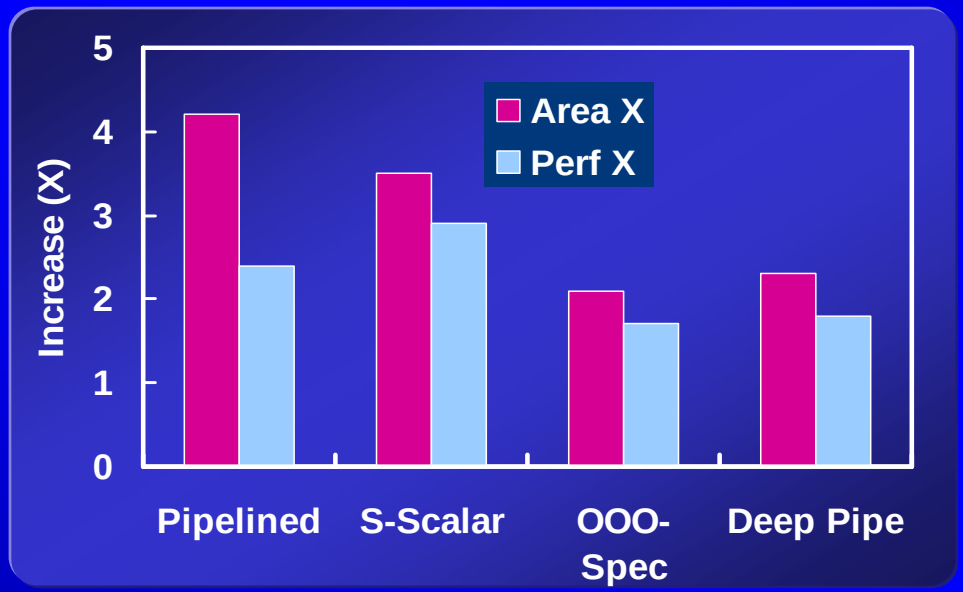

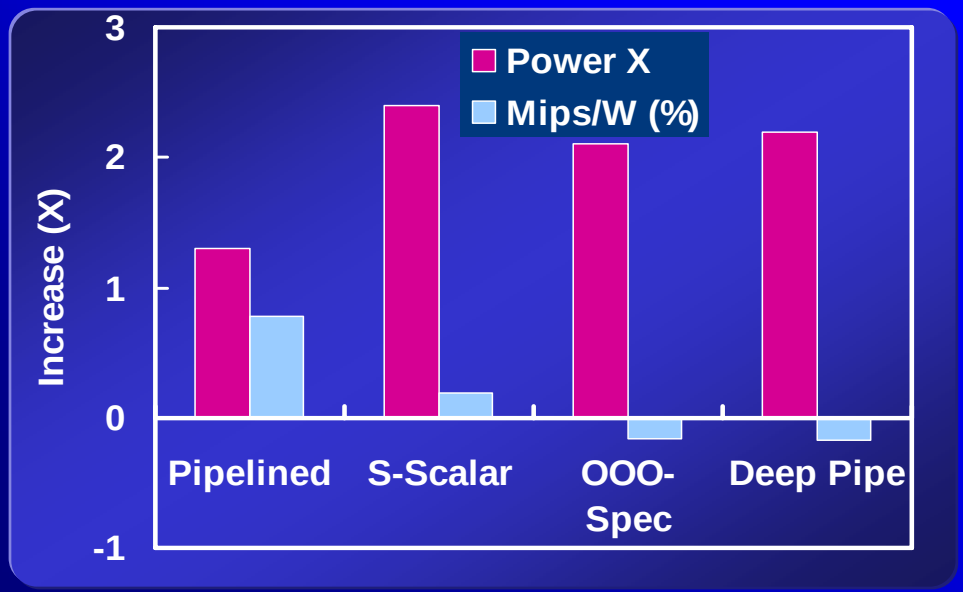

Moore's Law ⇨ more transistors for advanced architectures

Delivers higher peak performance

But…

Lower power efficiency

Chip Multi-Processing

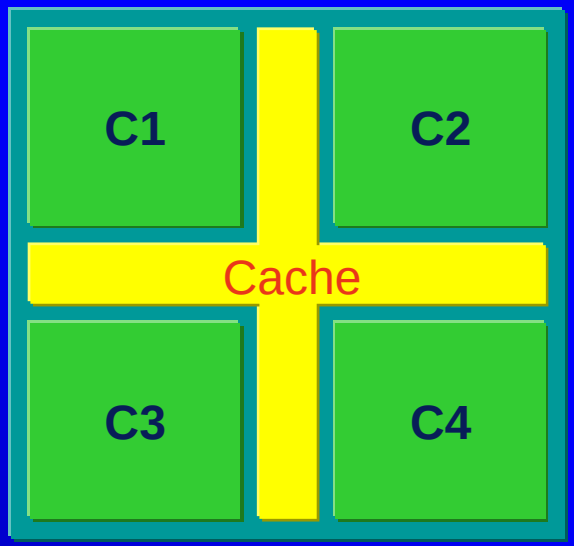

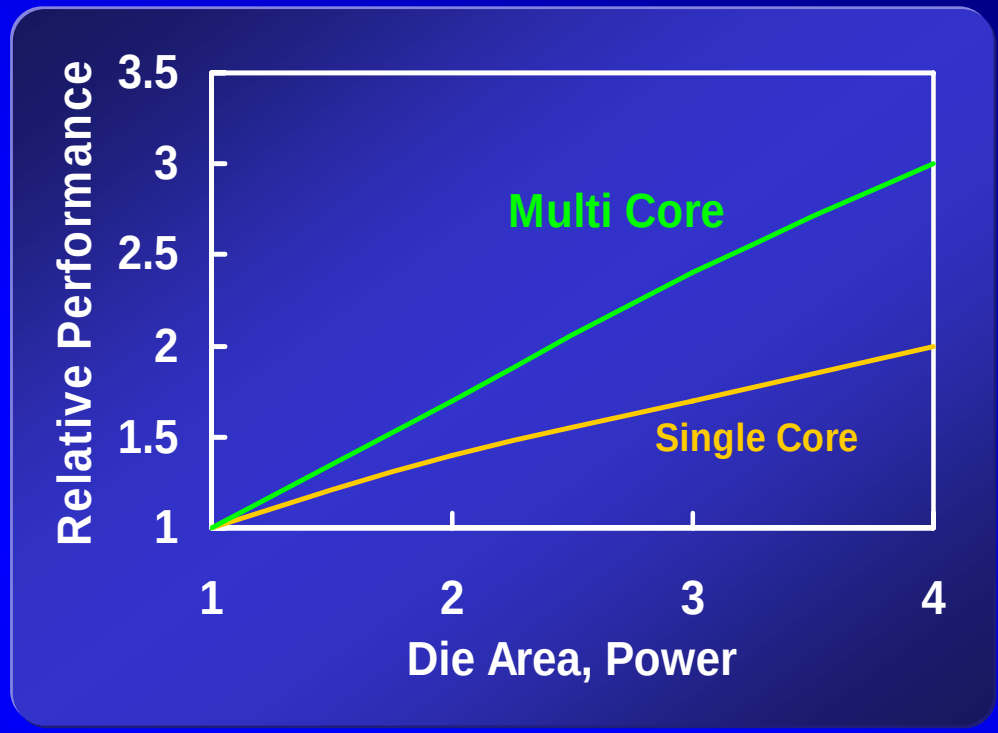

- Multi-core, each core Multi-threaded
- Shared cache and front side bus
- Each core has different Vdd & Freq
- Core hopping to spread hot spots
- Lower junction temperature

Dual Core

Rule of thumb

Voltage	Frequency	Power	Performance
1%	1%	3%	0.66%

In the same process technology…

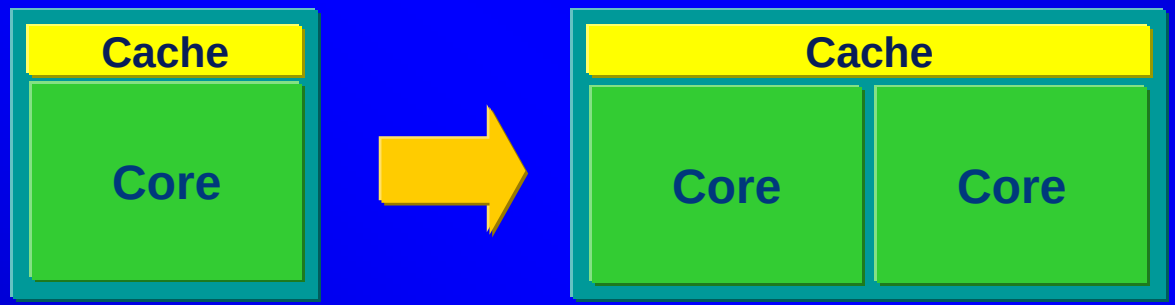

Voltage = 1
Freq = 1
Area = 1
Power = 1
Perf = 1

Voltage = -15%
Freq = -15%
Area = 2
Power = 1
Perf = ~1.8

Multi-Core

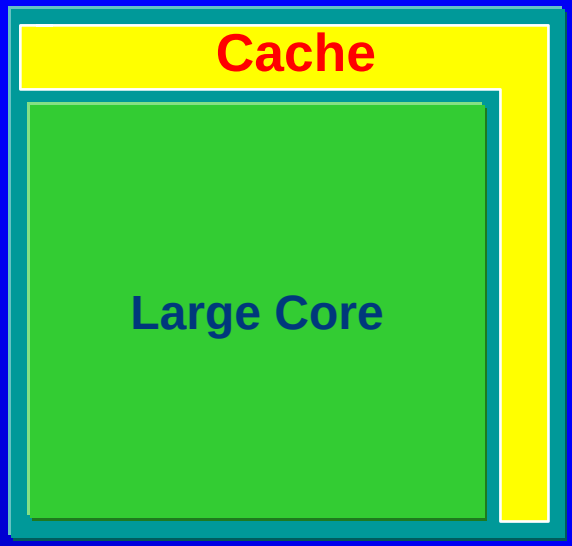

Power

Performance

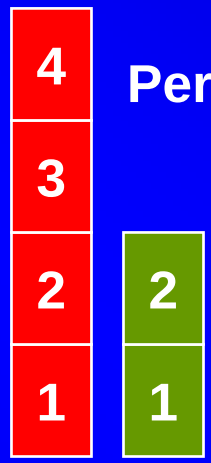

Power = 1/4

Performance = 1/2

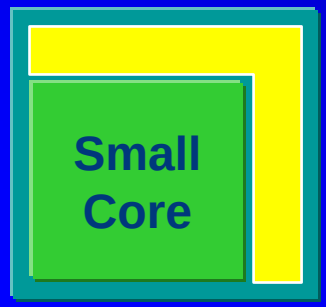

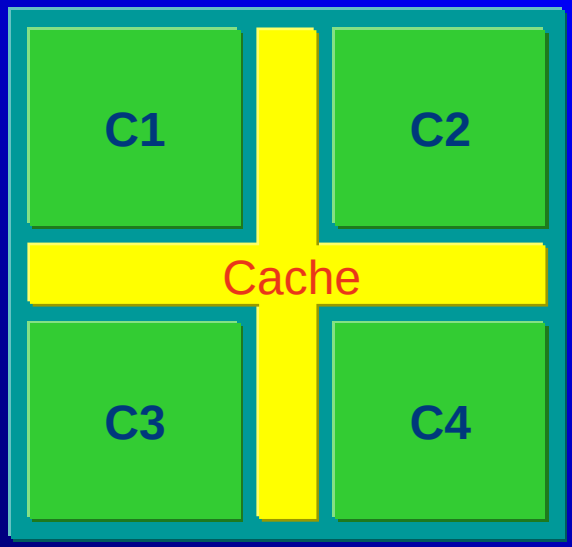

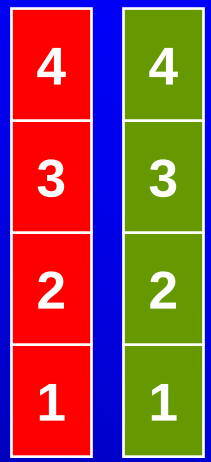

Multi-Core:
Power efficient
Better power and thermal management

Special Purpose Hardware

TCP/IP Offload Engine

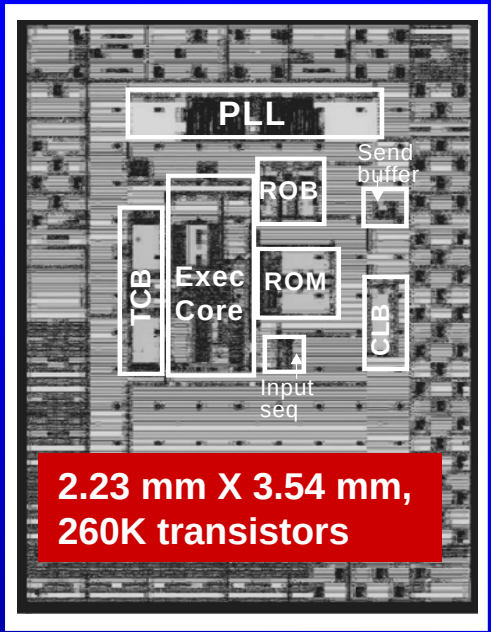

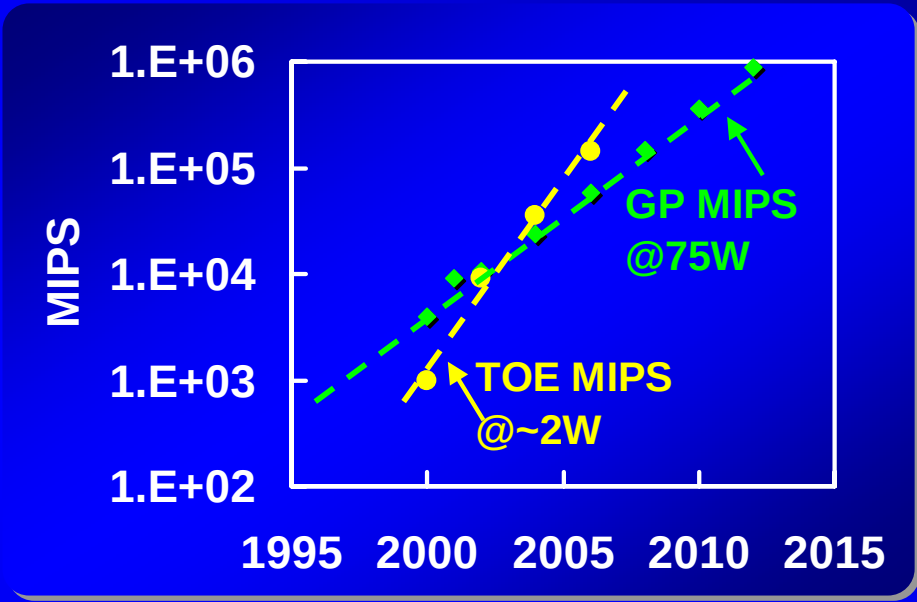

Opportunities: Network processing engines
MPEG Encode/Decode engines, Speech engines

Special purpose HW provides best Mips/Watt

Performance Scaling

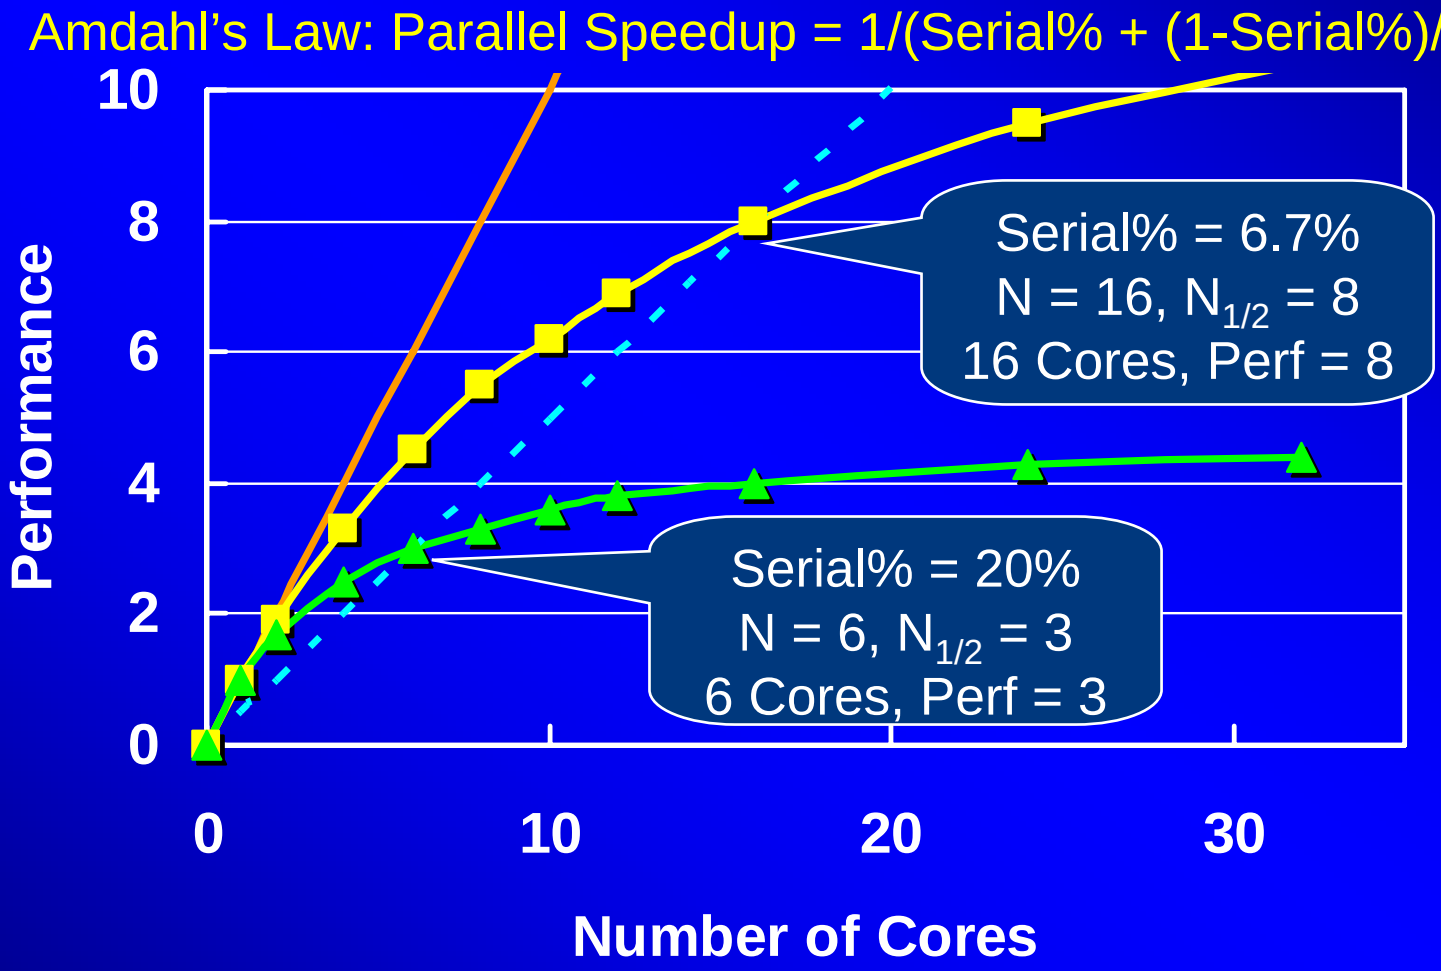

Parallel software key to Multi-core success

From Multi to Many…

13mm, 100W, 48MB Cache, 4B Transistors, in 22nm

12 Cores

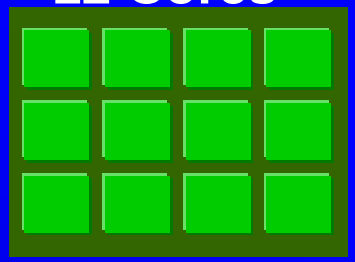

24 Cores

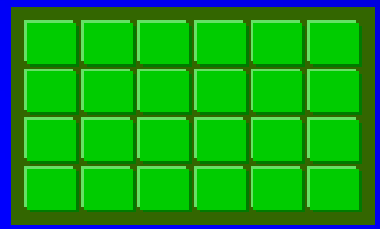

144 Cores

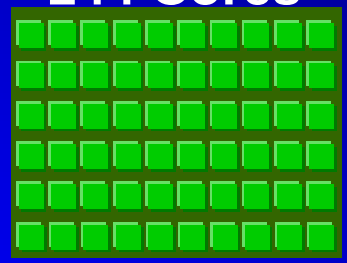

Single Core Performance

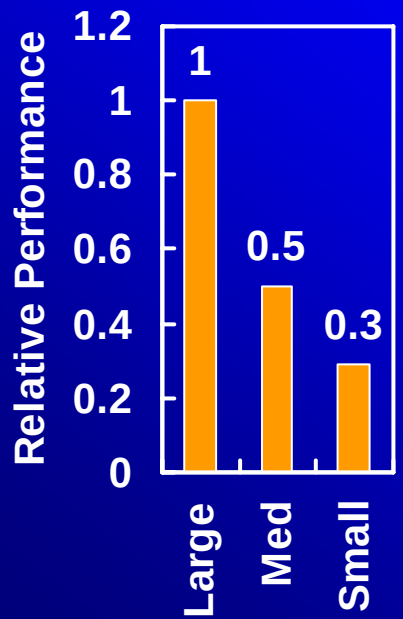

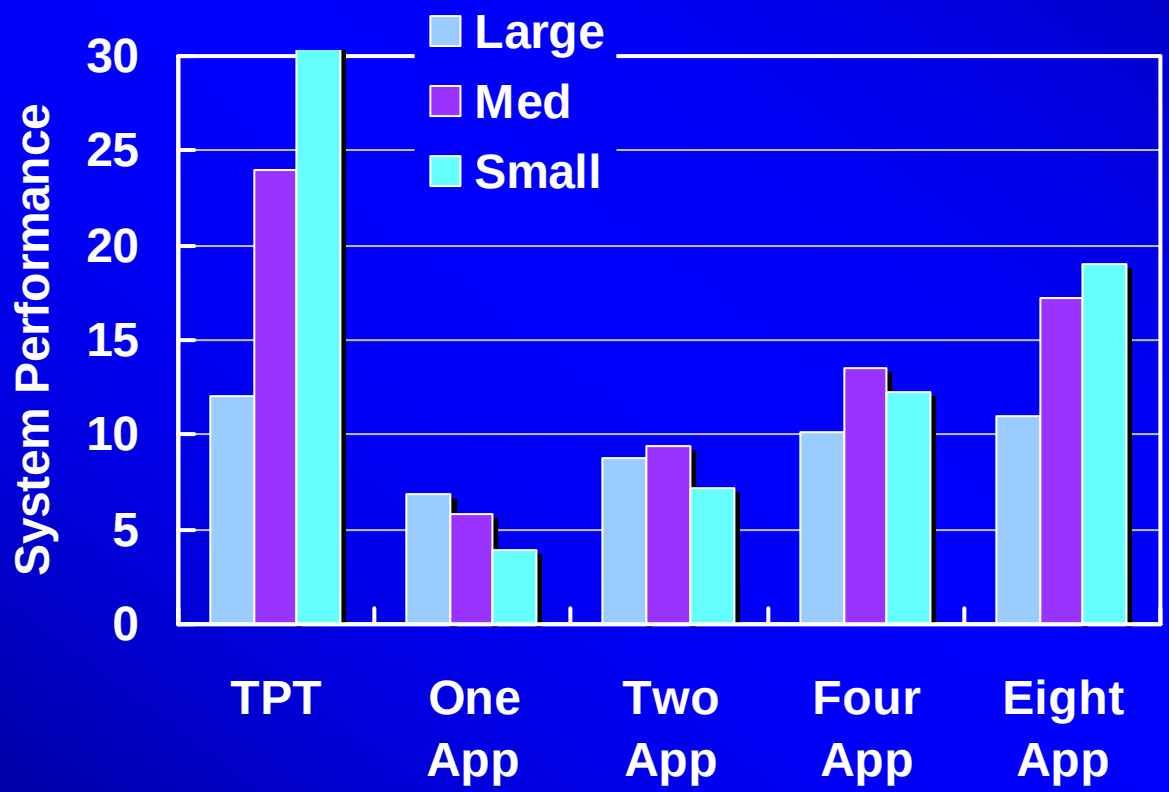

Future Multi-core Platform

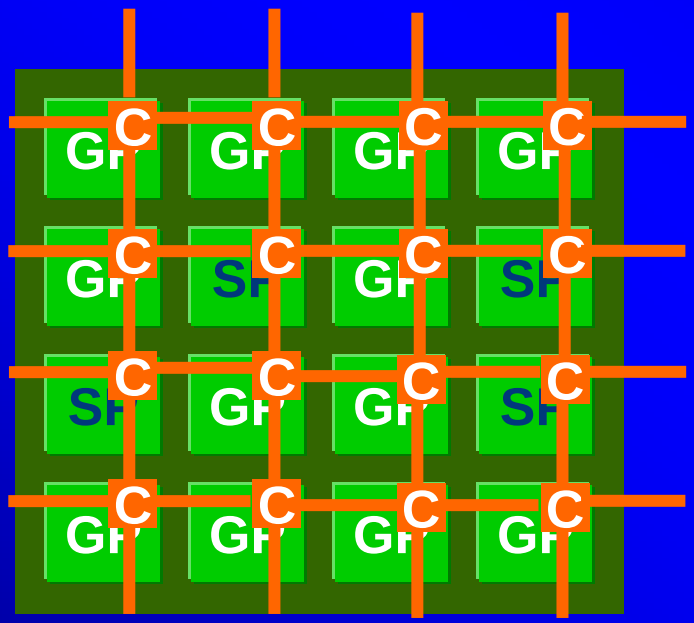

General Purpose Cores

Special Purpose HW

Interconnect fabric

Heterogeneous Multi-Core Platform

The New Era of Computing

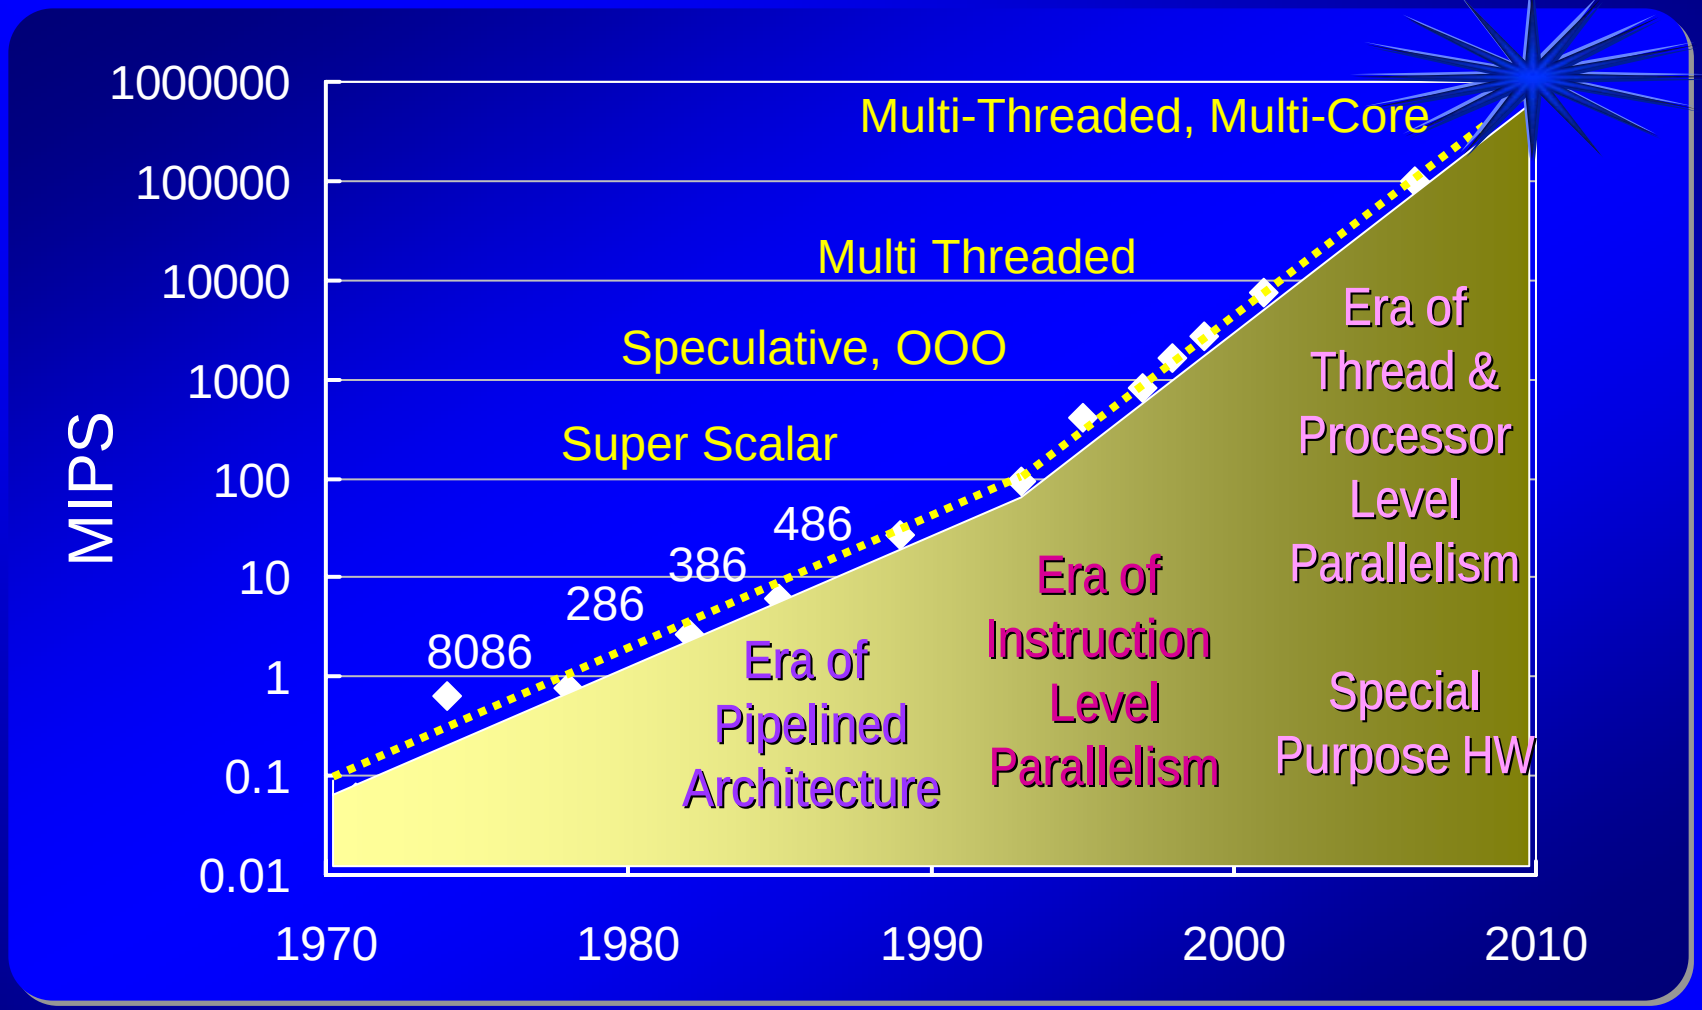

Multi-everywhere: MT, CMP

Summary

- ***Business as usual*** **is not an option**
  - **Performance at any cost is history**
- **Must make a *Right Hand Turn (RHT)***
  - **Move away from frequency alone**
- **Future μArchitectures and designs**
  - **More memory (larger caches)**
  - **Multi-threading**
  - **Multi-processing**
  - **Special purpose hardware**
  - **Valued performance with higher integration**