Technology Maturity for Massively Parallel Adaptive Computing

A Workshop to Discuss Readiness of Technology Ingredients to Realize MPAC Platforms

Fayé A Briggs
Intel Fellow, Director Scalable Server Architecture
Digital Enterprise Group
Intel Corporation

March 2-3, 2009
Portland, Oregon

Acknowledgement: Shekhar Borkar, Jim Held, and Shu-Ling Garver

Legal Disclaimer

LEGAL INFORMATION:

Agenda

Welcome

Goal of the workshop

Process technology roadmap

Spectrum of silicon solutions

Challenges in silicon architectures

Emerging Applications

Challenging the Workshop

Silicon Technology Follows Moore's Law

90nm
2003

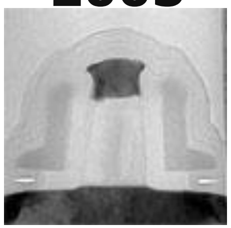

65nm
2005

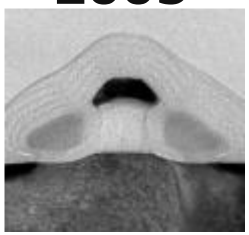

45nm
2007

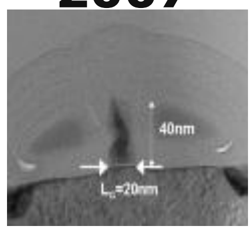

32nm
2009

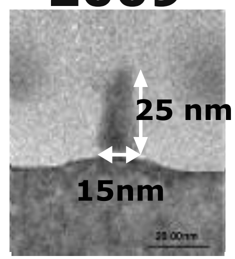

22nm
2011

2005 – 2012

16nm

11nm

8nm

2013 –201

New Intel technology generation every 2 years
Intel R&D technologies drive this pace well into the next decade

Roadmap

Research

All products, computer systems, dates, and figures specified are preliminary based on current expectations, and are subject to change without notice.

intel

Recent 15 Years of Processor Core Evolution

Cores

Pentium

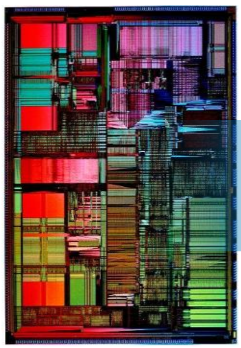

Pentium
In-Order
66 MHz
3.1 M transistors
800 nm

Pentium II & III

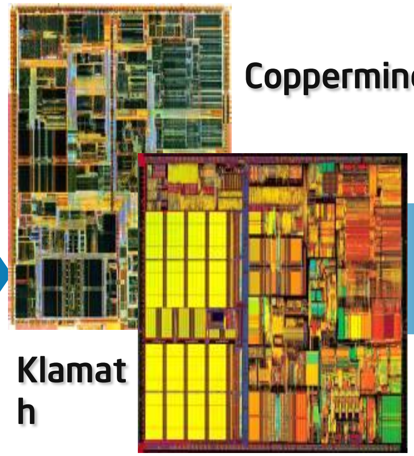

Pentium II & Pentium III
In-Order
300-450MHz, 1GHz
7.5, 28 M transistors
250, 180 nm

Netburst Architecture

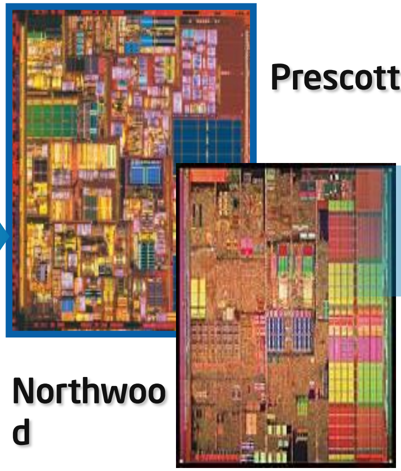

Intel Pentium 4 Processor
Out-Of-Order
2-3+ GHz frequency
55, 120 M transistors
130nm, 90 nm

Core µArchitecture

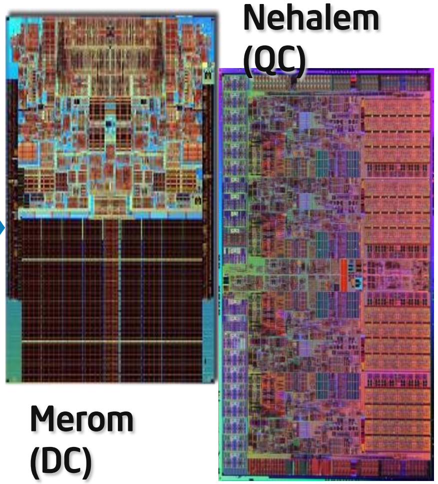

Intel Core 2 Processor
Out-Of-Order
3+ GHz frequency
291, 731 M transistors
65, 45 nm

Note: Just a representative list, not exhaustive

Accelerating Multi- and Many-core

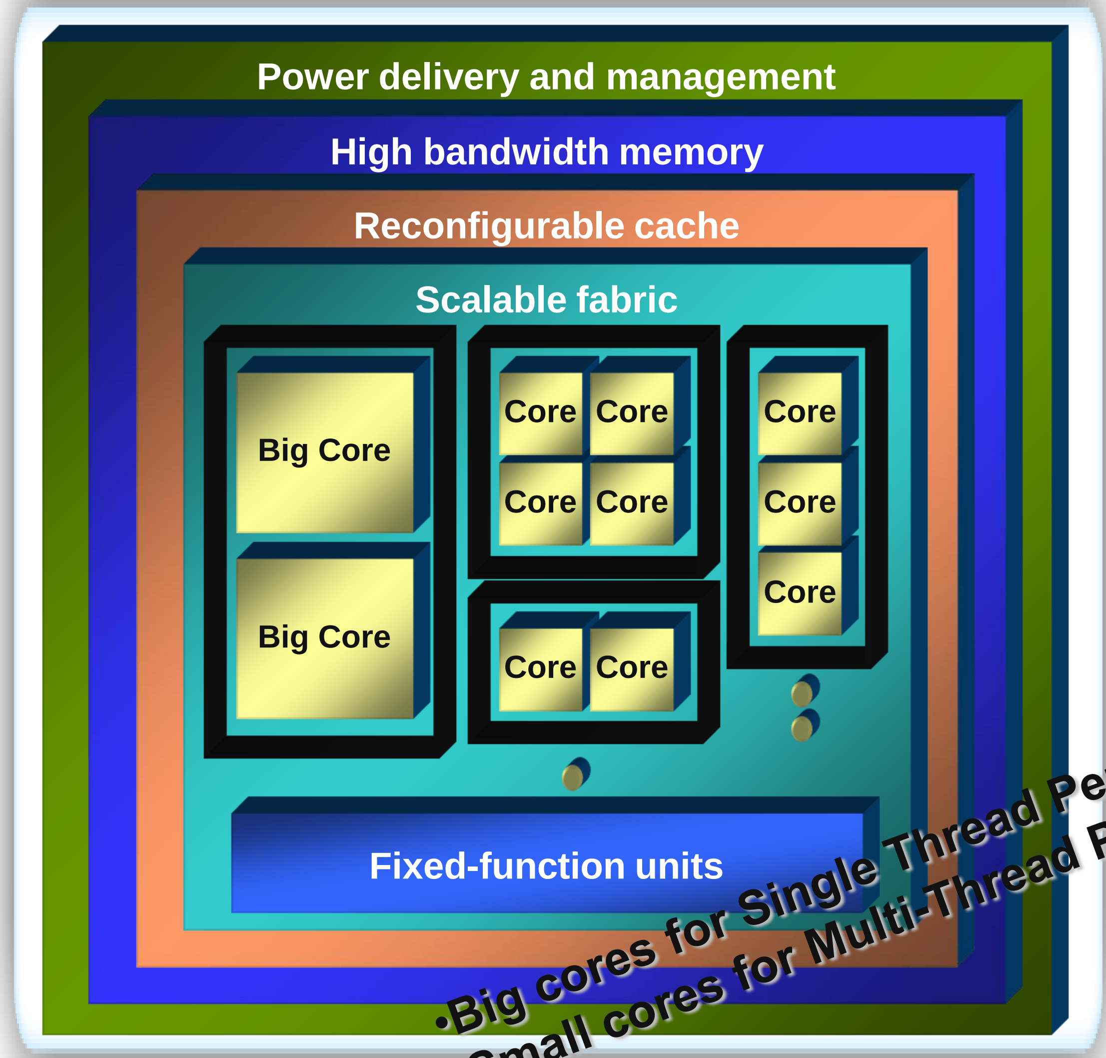

- Big cores for Single Thread Performance
- Small cores for Multi-Thread Performance

Performance Through Parallelism

Intel Developer FORUM

Exploit Potential Extreme Variations

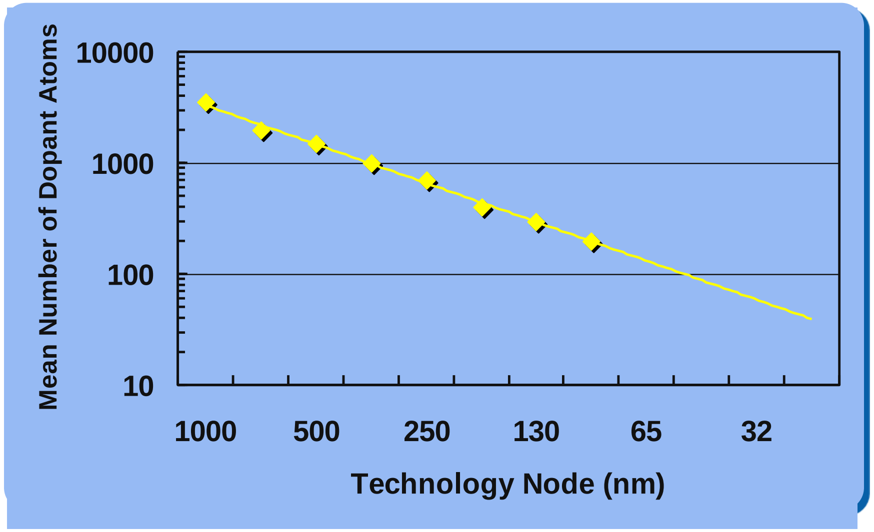

Random Dopant Fluctuations

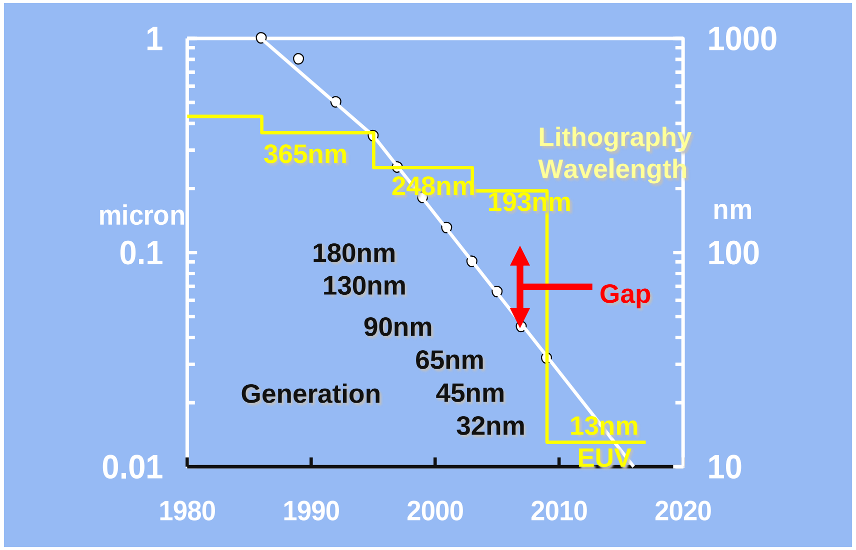

Source: Mark Bohr, Intel

Sub-wavelength Lithography

Exploit Potential Reliability Issues

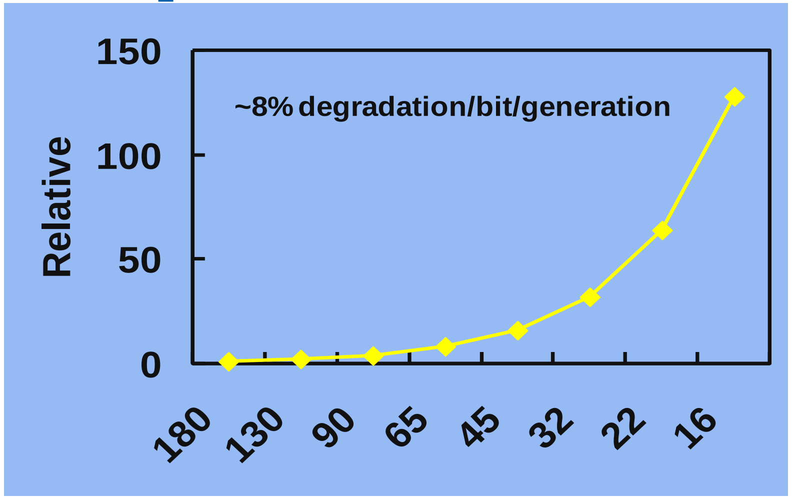

Soft Error FIT/Chip (Logic & Mem)

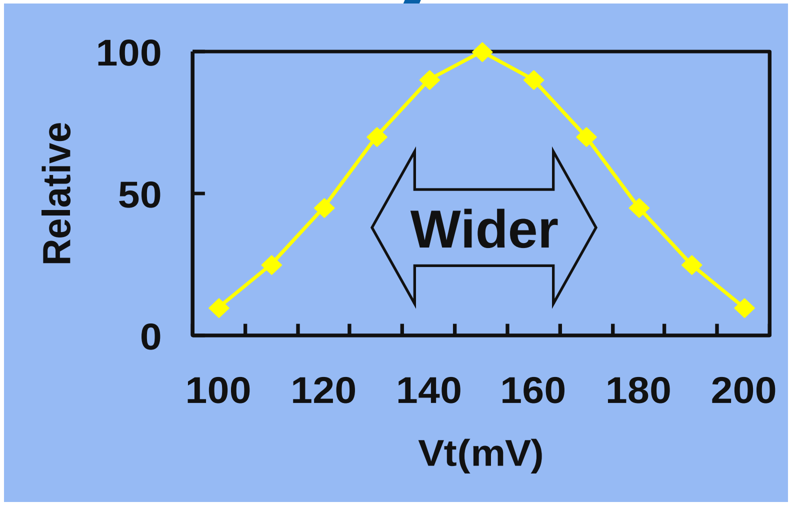

Extreme device variations

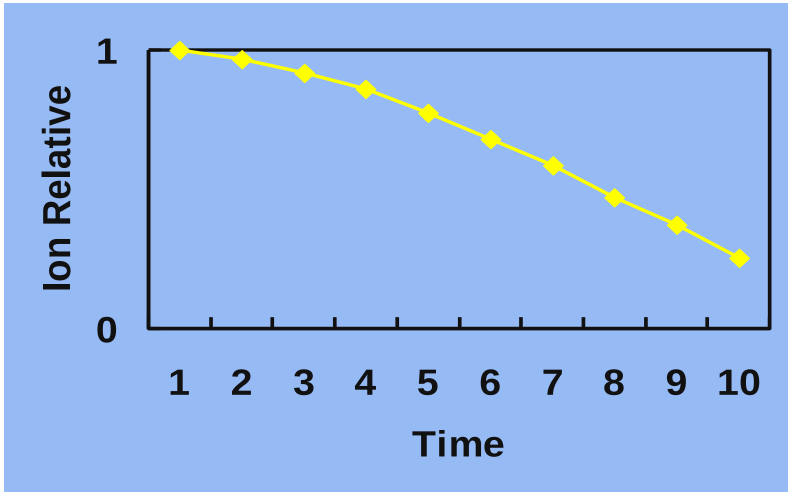

Time dependent device degradation

Properties of Architecture

- Massively Parallel through many-core and multithreading
- Extremely tolerant of noisy environments
- Built from relatively slow, low-power, transistors that can tolerate variability
- Manufacturing defect tolerant
- Fault-tolerant (core, network elements)
- Distributed memory with “clustered” coherence
- Partitionable into virtual machines
- Be programmable with tools for extreme parallelism
- Connected set of networks
- How much precision is desired in computation?
- Capable of solving real world problems that go well beyond today’s platforms

Many-Cores with Resiliency to Adapt to Variability

intel

Intel Tera-scale Research

100+ Research Projects Worldwide

Microprocessor

Examples:
Scalable memory
Multi-core architectures
Specialized cores
Scalable fabrics
Energy efficient circuits

Platform

Examples:
3D Stacked Memory
Cache Hierarchy
Virtualization/Partitioning
Scaleable OS's
I/O & Networking

Programming

Examples:
Speculative Multithreading
Transactional memory
Workload analysis
Compilers & Libraries
Tools

ACCELERATE TRANSITION TO PARALLEL PROGRAMMING

Software Products

www.intel.com/software/products

University Outreach
Intel ® Press
Intel® Software College

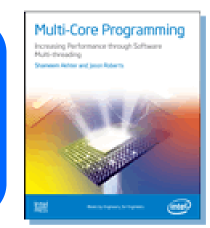

Teraflops and Beyond on a Chip

Architecture
- Compute element
- Data Communication (type of N-o-C)
- Memory hierarchy & topology

Design
- High performance FPMAC
- High Bandwidth & low latency 2D Mesh
- Tiled methodology
- Power management
- Resiliency of CM and N-o-C

Future could deliver PFLOP performance under 100W

Many-Core Cluster

On-Die Fabric Supporting Resiliency & Scalability

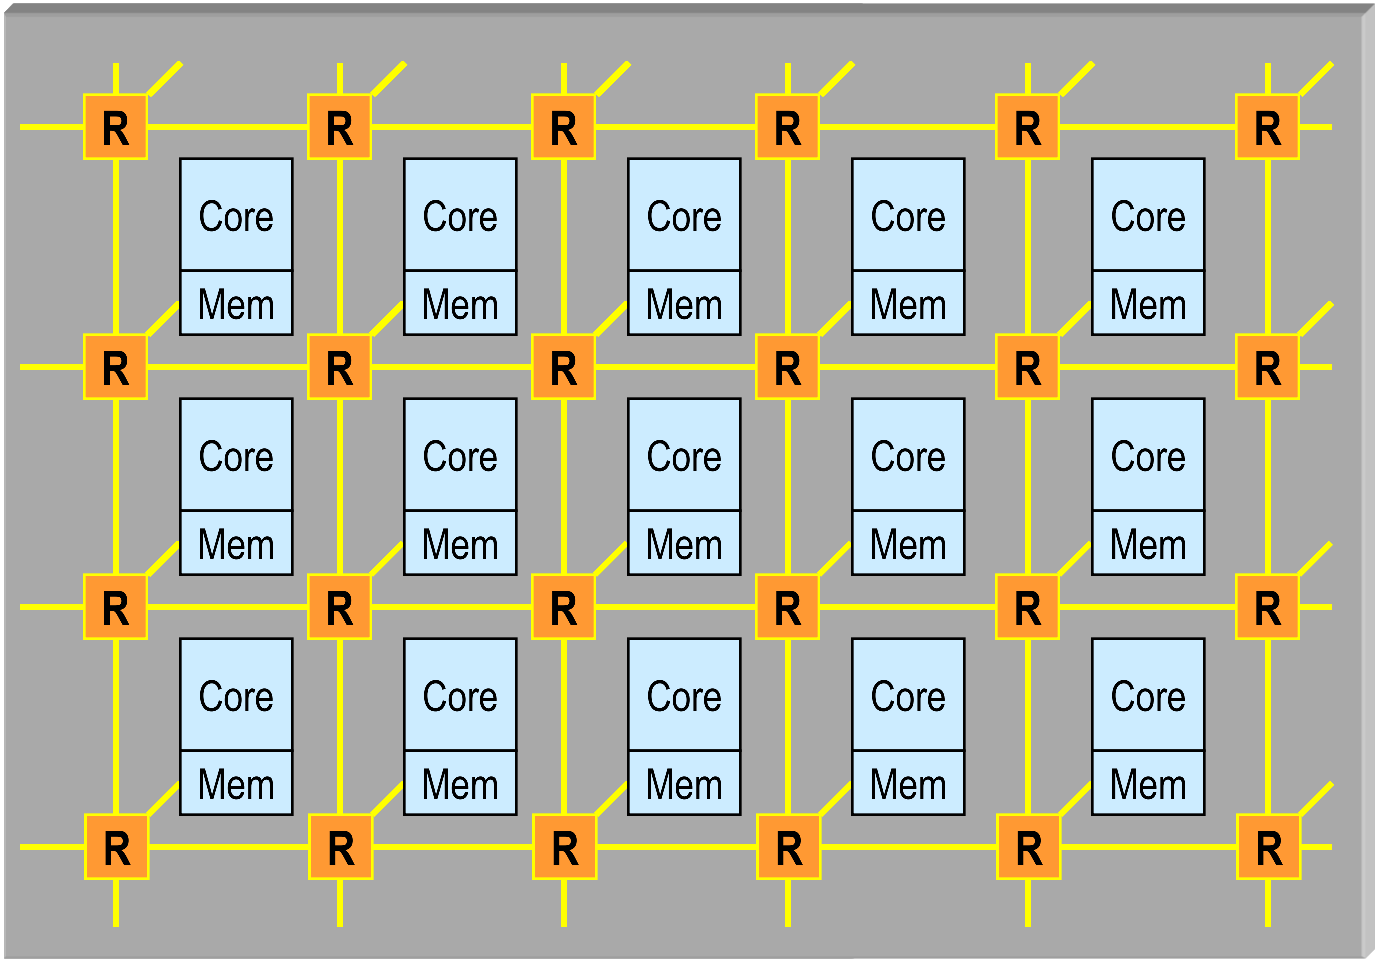

- ✓ Multi-generation scalability
- ✓ Modular design and validation
- ✓ Resiliency
- ✓ Do we need coherency?

Many-Core Platform Vision

Adapting the Hardware

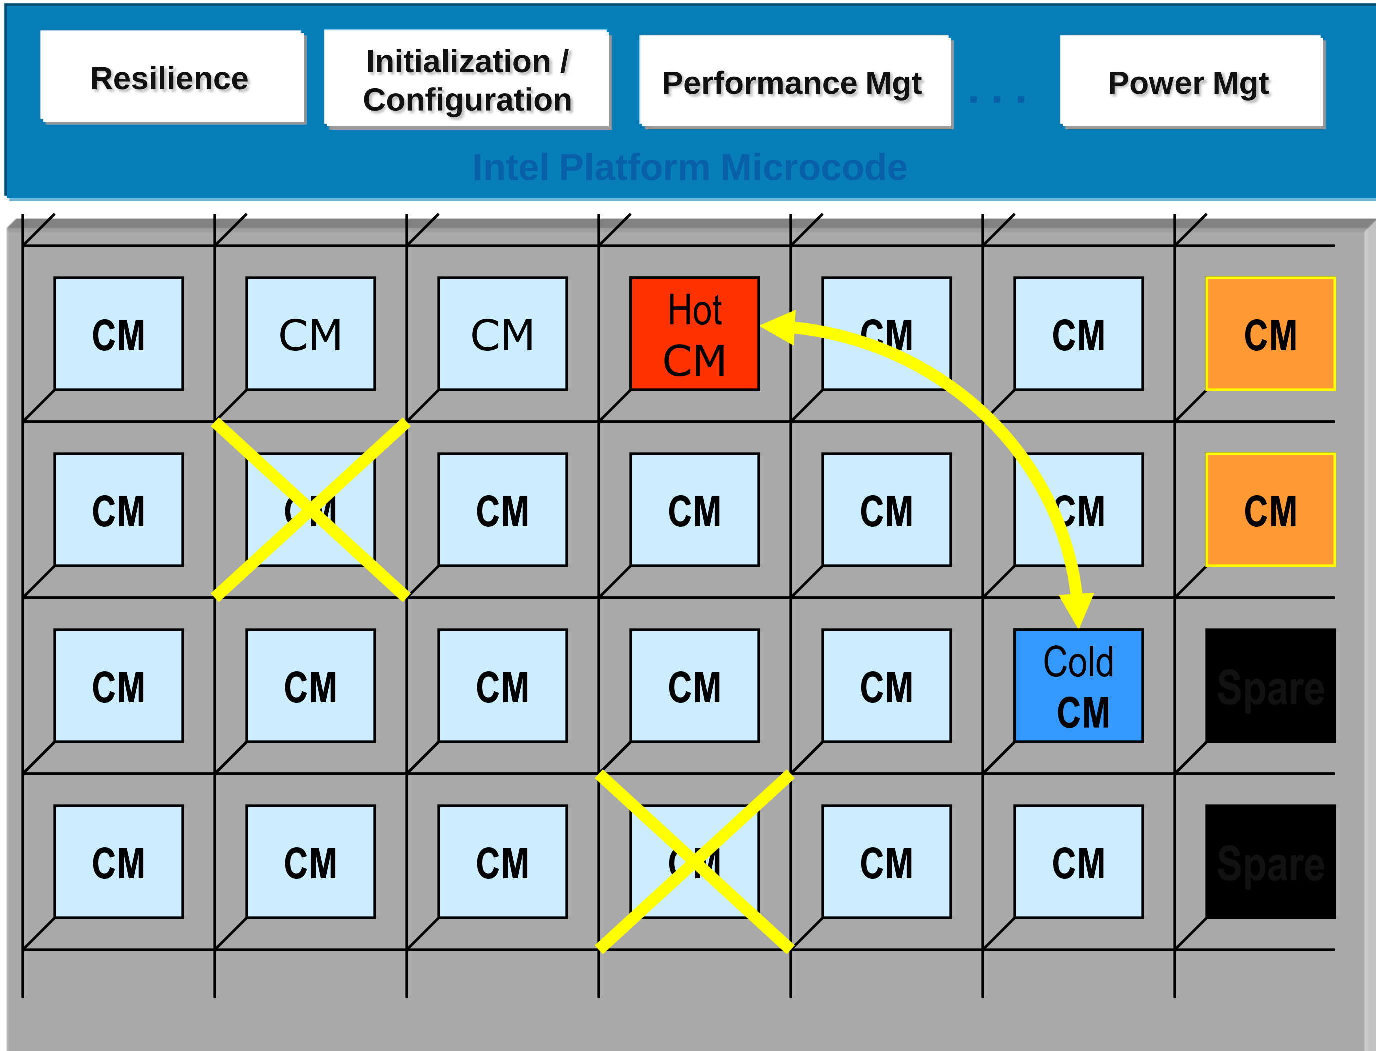

- ✓ Addressing variability and reliability issues
- ✓ Performance Mgt – “binning”
- ✓ Thermal & Power Mgt

Key ingredients for Teraflops on a chip

- Core
- Communication
- Technology
- Clocking
- Power saving techniques

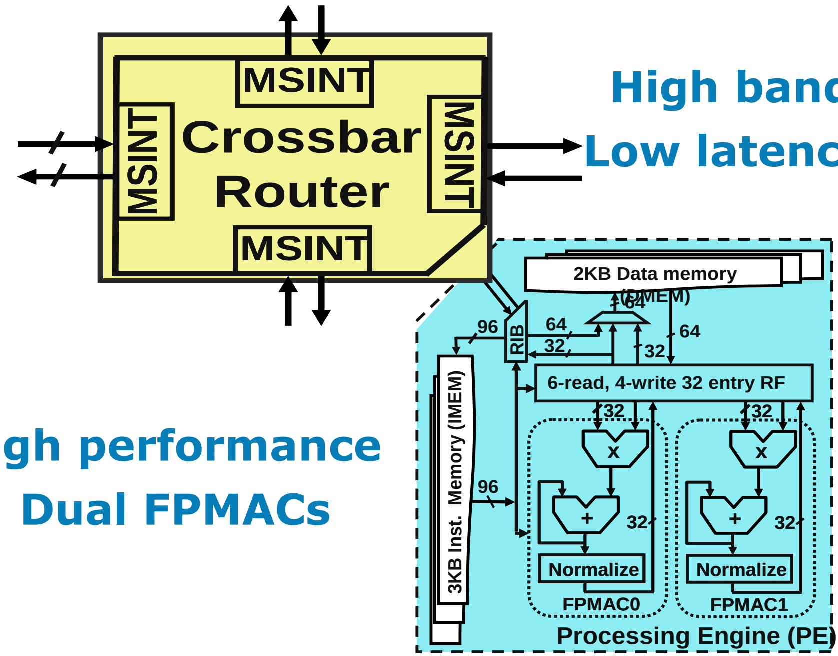

High bandwidth
Low latency XBAR

High performance
Dual FPMACs

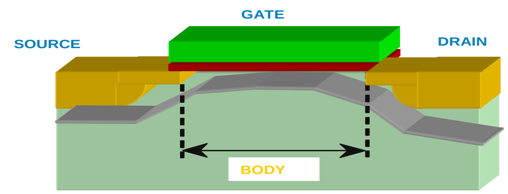

65nm
eight metal
CMOS

Goal: Demonstrate TFLOP performance under 100W

First monolithic programmable Teraflops processor

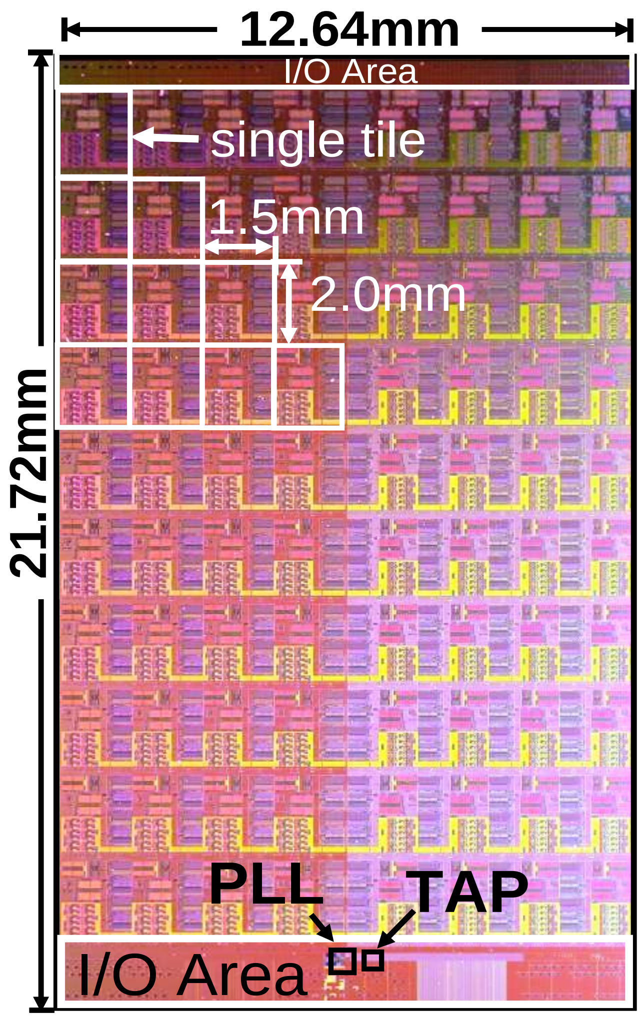

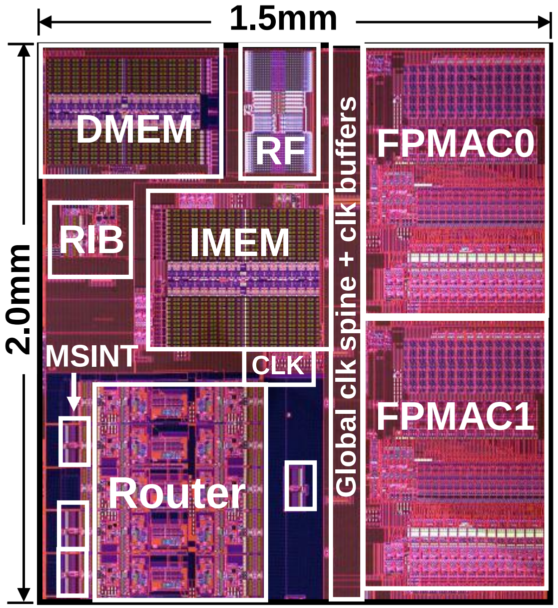

Technology	65nm, 1 poly, 8 metal (Cu)
Transistors	100 Million (full-chip) 1.2 Million (tile)
Die Area	$275mm^2$ (full-chip) $3mm^2$ (tile)
C4 bumps #	8390

Final result

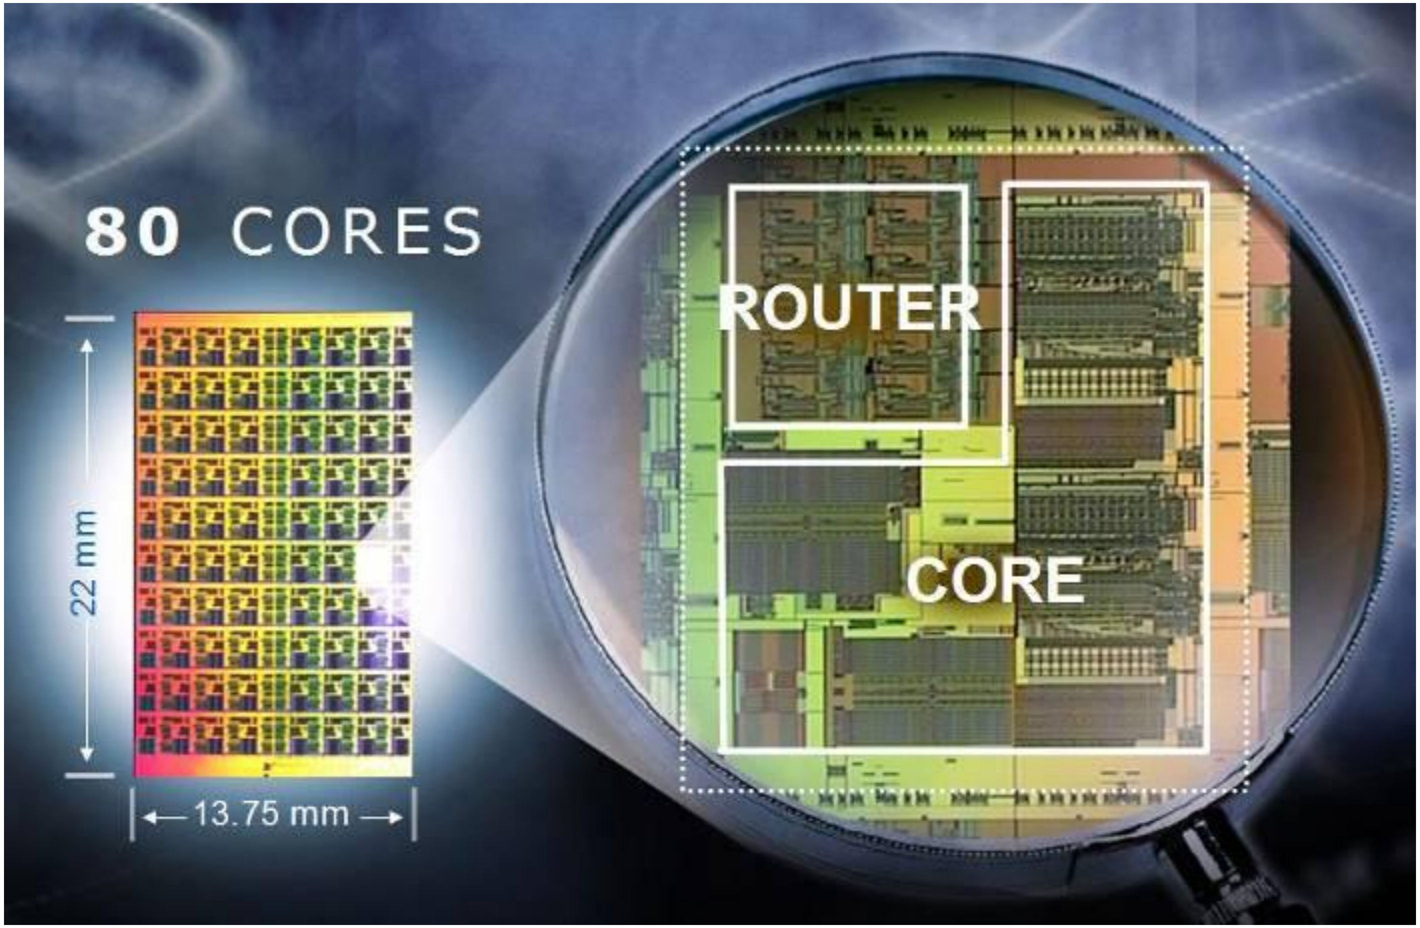

Tera-scale Memory Bandwidth

Tera Bytes: 3D Stacked Memory

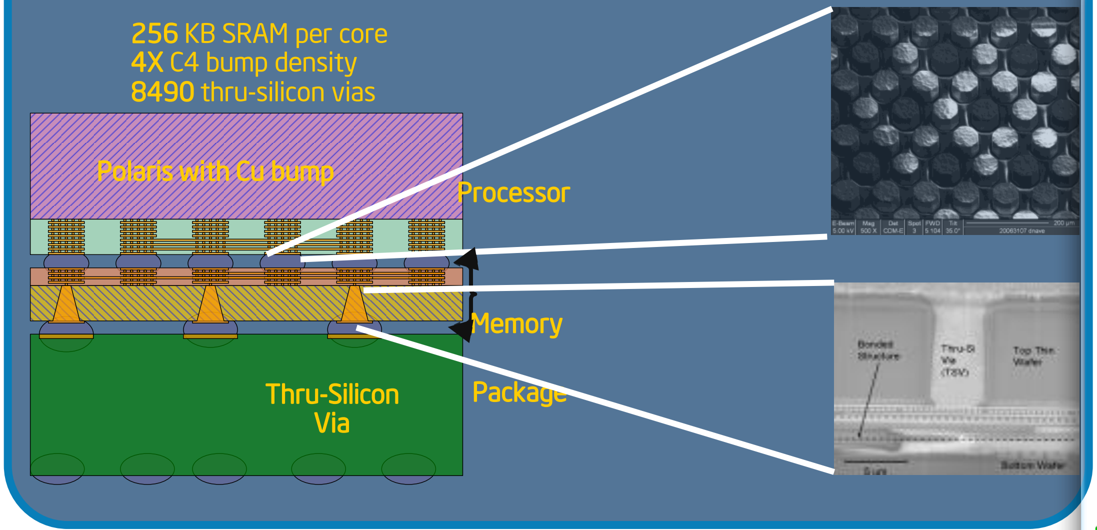

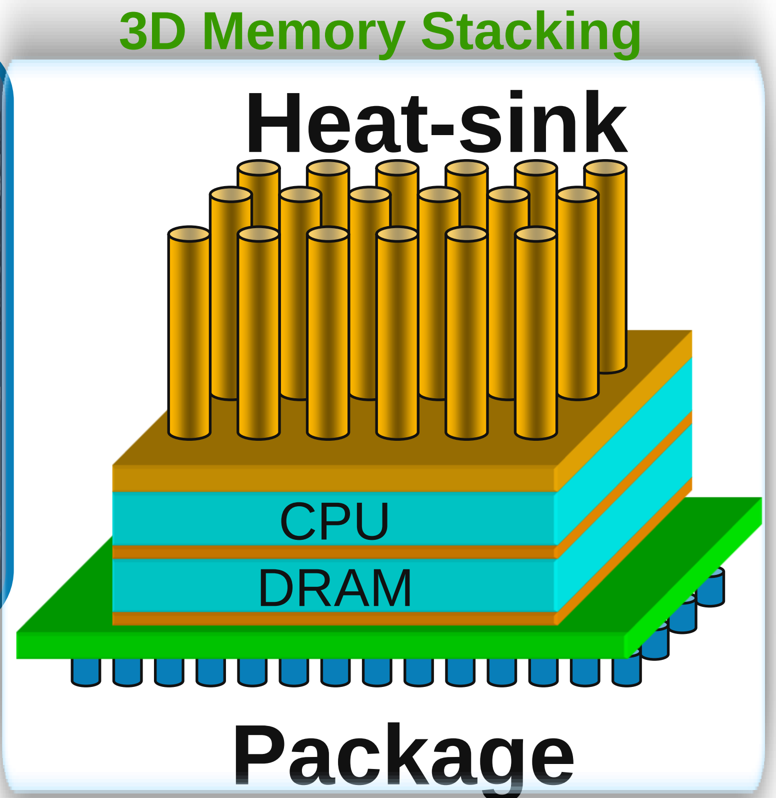

The 3D stacking technology is proven in small volumes but there is more to be done to support HVM capability

Software Tools for Many-Cores

Intel® Software Development Products	● = Currently Available	Operating Systems		Operating Systems		
		Windows*	Linux*	Windows	Linux	Mac OS*
		Development Environments		*Development Environments*		
		Visual Studio*	GCC*	Visual Studio	GCC	Xcode*
Compilers	C++	●	●	●	●	●
	Fortran	●	●	●	●	●
Performance Analyzers	VTune® Performance Analyzer	●	●	●	●	
Performance Libraries	Integrated Performance Primitives	●	●	●	●	●
	Math Kernel Library	●	●	●	●	●
	Mobile Platform SDK			●		
Threading Analysis Tools	Thread Checker			●	●	
	Thread Profiler			●		
Cluster Tools	MPI Library	●	●	●	●	
	Trace Analyzer and Collector	●	●	●	●	
	Math Kernel Library Cluster Edition	●	●	●	●	
	Cluster Toolkit	●	●	●	●	
XML Tools**	XML Software Suite 1.0		●	●	●	

Comprehensive Set of Intel® Software Development Products Enable Multithreaded Application Development Across Intel® Platforms

** Additional XML tools information can be found at www.intel.com/software/xml

Emerging Real-World Apps

Image/Picture recognition, classification, & ***search***

Static & dynamic object recognition

Speech recognition, translation

Security

Autonomous navigation

Many More…..

Short Response Time Expected

Goal of MPAC Workshop

Review key ongoing research activities in architectures, algorithms & applications that exploit massive parallel adaptive computing

- Understand state of the art & approaches

Identify technology gaps in realizing platforms capable of cost-effectively & energy-efficiently processing bio-inspired algorithms & apps

- What's holding us back?

Breakout sessions to identify problem statements, known technological shortfalls, and where applicable, possible solutions

Thanks

Backup

BEYOND TERA-SCALE COMPUTING

Microprocessor

Essential
Scalable memory
Many-core architectures
Specialized cores?
Scalable fabrics
Energy efficient circuits

Complementary
Si Process Technology
Resiliency

Platform

Essential
3D Stacked Memory
Cache Hierarchy
Distributed Memory
Virtualization/Partitioning
Scalable Distributed OS's
I/O & Networking

Complementary

Programming

Essential
Speculative Multithreading
Data parallel operations
Compilers & Libraries
Tools

Complementary
Domain specific languages
Usage Models

Take away

Object Recognition Progress

LARRABEE – TERA-SCALE SOLUTION

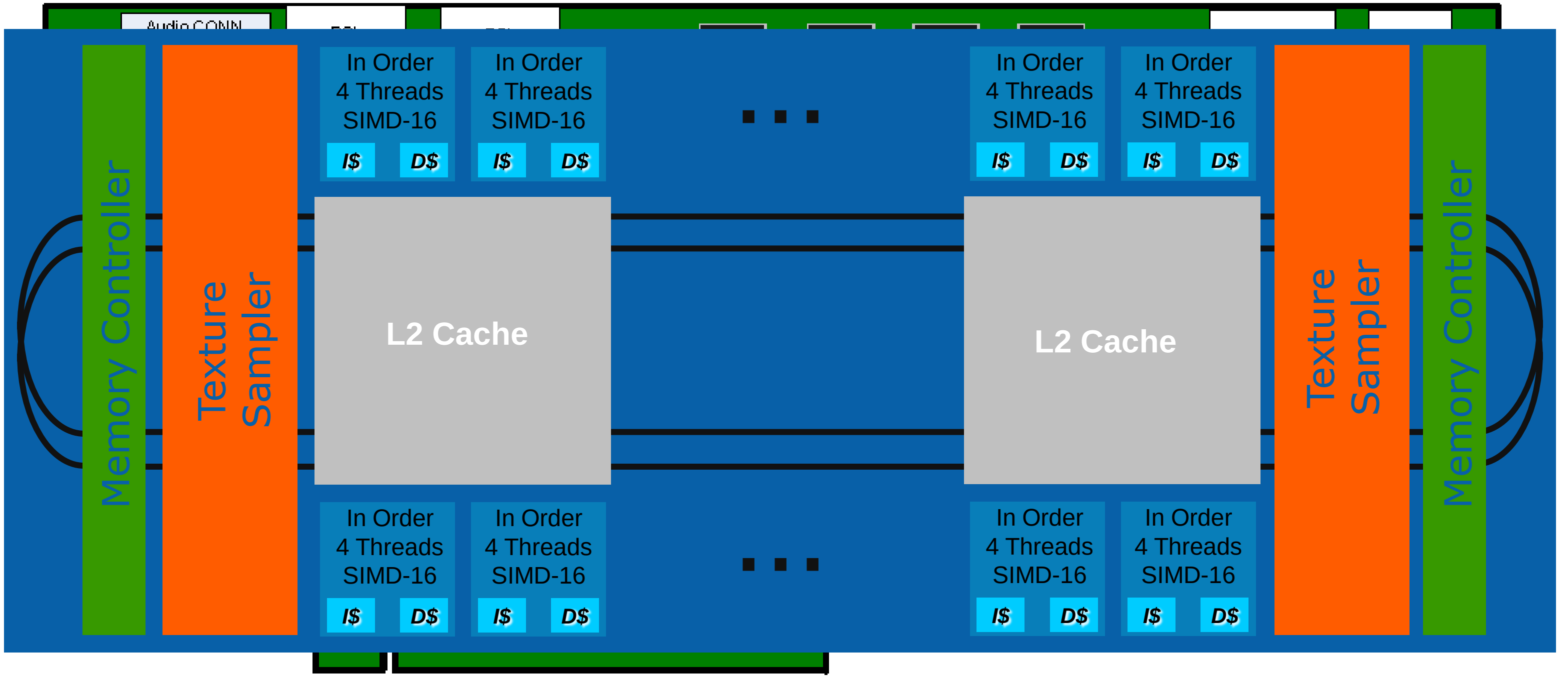

- Discrete high end GPU on general purpose platform
- TeraFlops of fully programmable performance
- **GPU** - >16 cores @ ~2.0GHz, >150W
- JPEG textures, physics acceleration, anti-aliasing, enhanced AI, Ray Tracing etc.