



The Future of Multi-core: Intel's Tera-scale Computing Research

Jim Held
Intel Fellow & Director,
Tera-scale Computing Research
Intel Corporation




Agenda

- Tera-scale Computing
 - Motivation
 - Platform Vision
- Tera-scale Computing Research Agenda
- Teraflops Research Processor
 - Key Ingredients
 - Power Management
 - Performance
 - Programming
 - Key Learnings
- Work in progress
- Summary



Multi-core for Energy-Efficient Performance



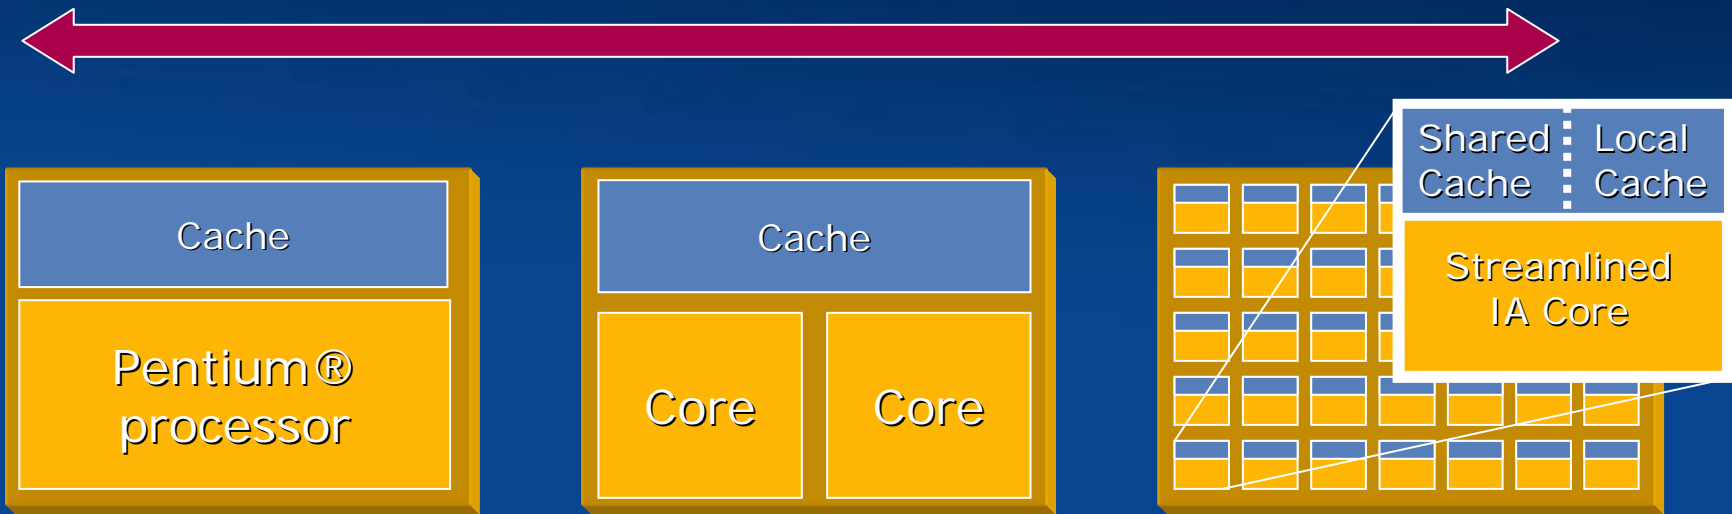
Relative single-core frequency and Vcc



Many Cores for Tera-scale Performance

Optimized for speed

Optimized for performance/watt



Pentium® processor era chips optimized for raw speed on single threads

Today's chips use cores which balance single threaded and multi-threaded performance

Future: 10s-100s of energy efficient, IA cores optimized for multithreading

Tera-scale Application Areas



Personal Media Creation and Management

- Search for and edit photos and videos based on image; no tagging
- Easily create videos with animation



Entertainment

- Watch yourself star in a movie or game
- Hold and interact with objects in the virtual world
- Control with speech and gesture
- Immersive: 3D, interactive

Health

- Virtual health worker monitors and assists elders/patients living alone
- Real-time realistic 3D visualization of body systems
- Effects of changes in diet, exercise and disease on body



Source:
<http://vhp.med.umich.edu/Surgical-Simulation.jpg>

Learning and Travel

- Surround yourself with sights and sounds of far-away places
- Practice new languages and customs



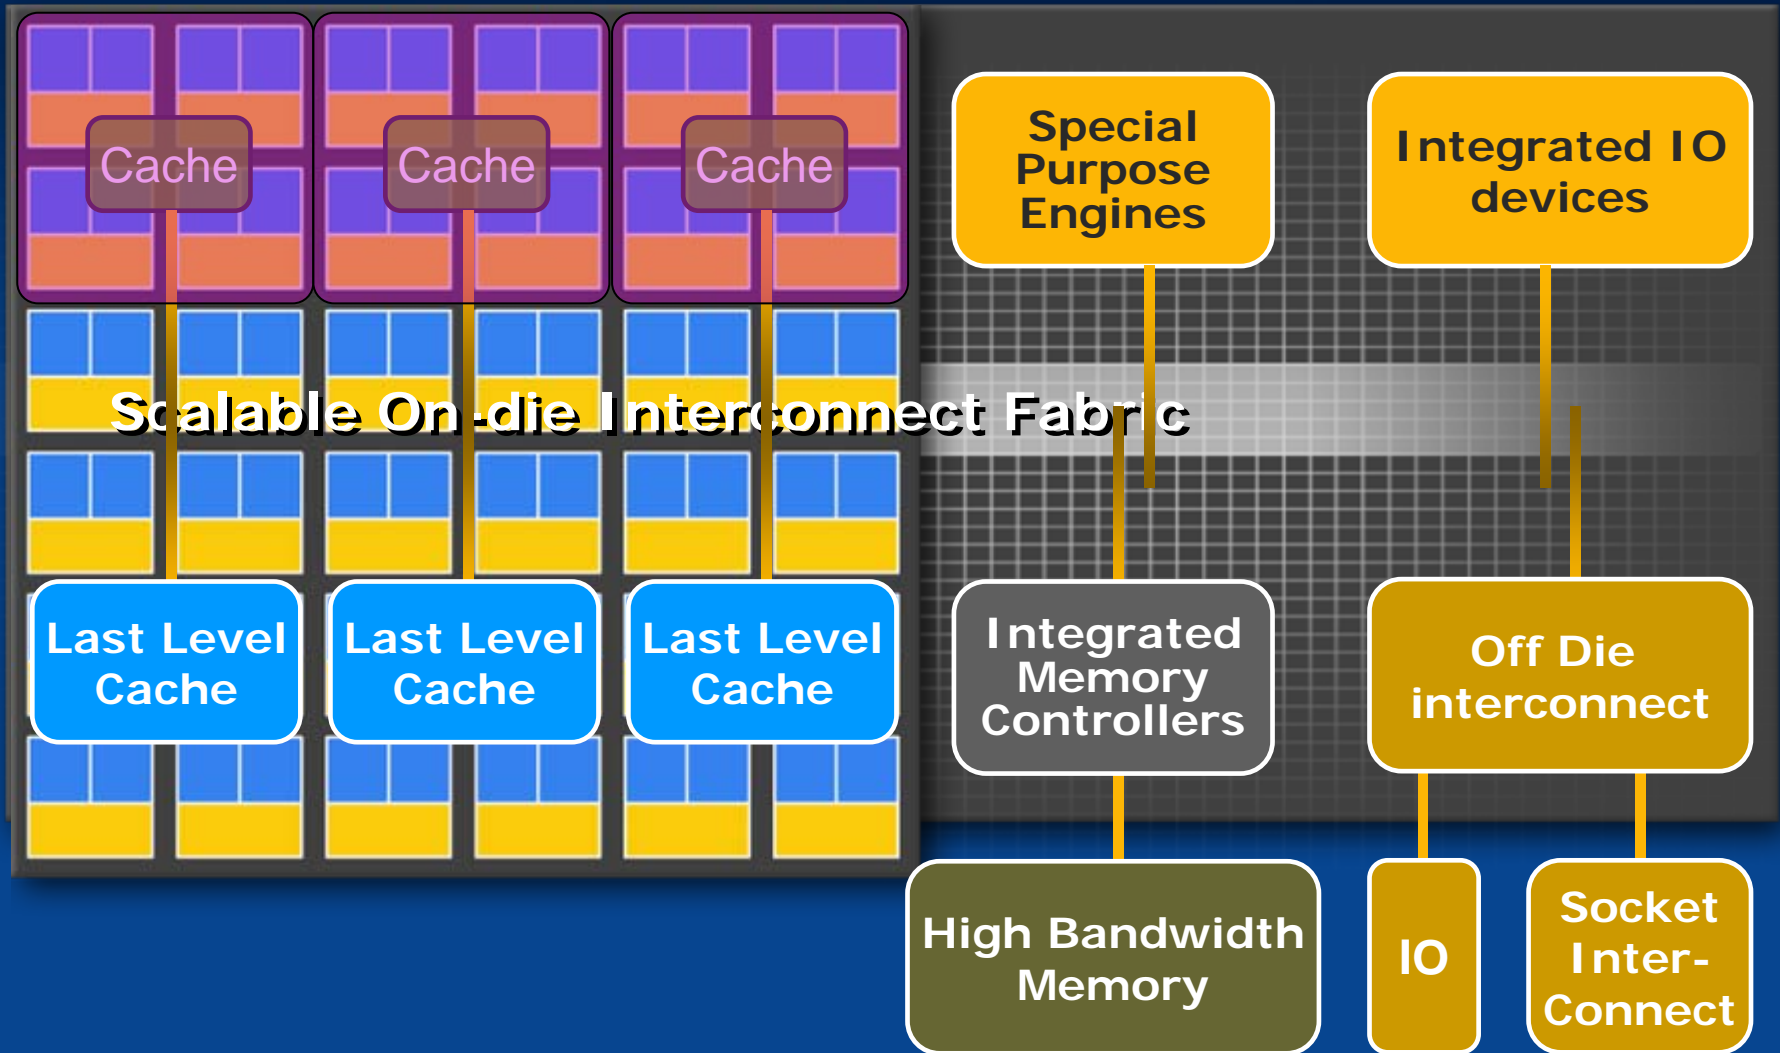
Source: Steven K. Feiner, Columbia University

Telepresence & Collaboration

- As if you are in the same place with family and friends—without the travel
- Appointments with doctors, teachers, leaders
- Develop and perform art with those far away



A Tera-scale Platform Vision



Tera-scale Computing Research

Applications – Identify, characterize & optimize

Programming – Empower the mainstream

System Software – Scalable services

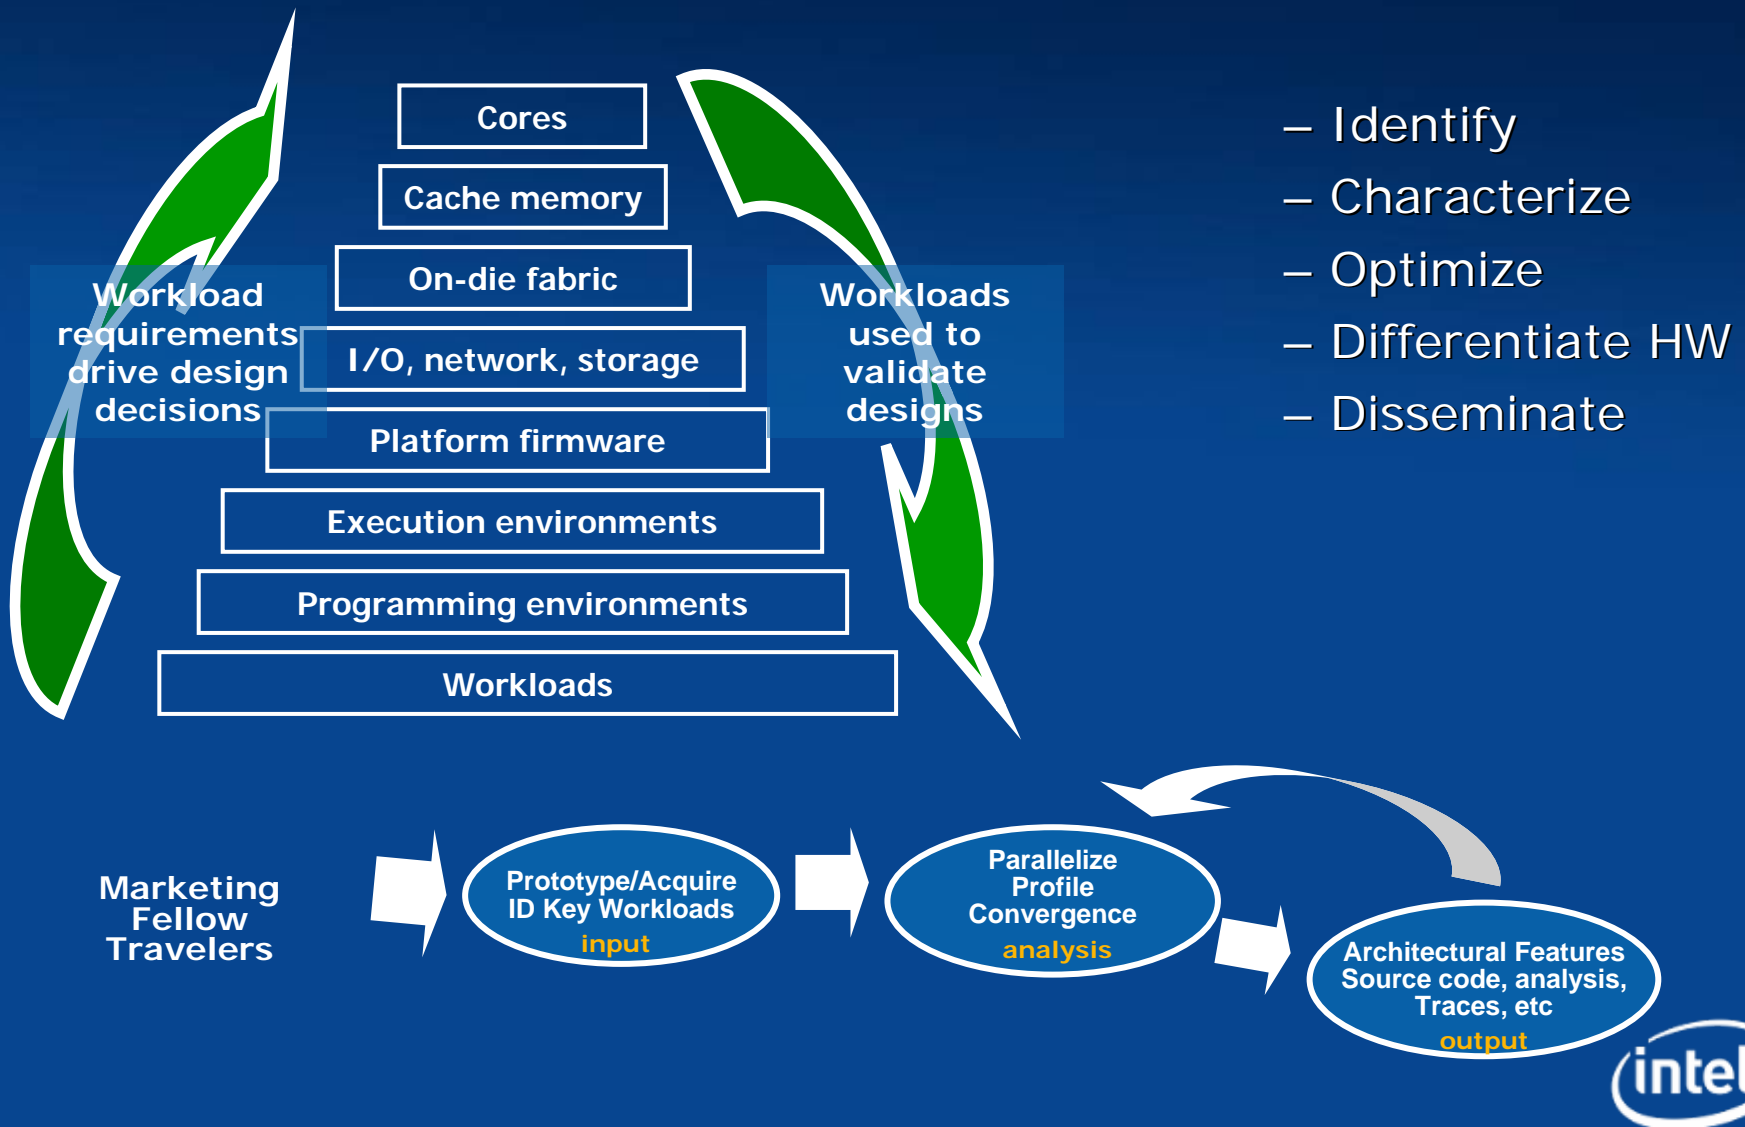
Memory Hierarchy – Feed the compute engine

On-Die Interconnect – High bandwidth, low latency

Cores – power efficient general & special function




Emerging Application Research




Emerging Applications will demand Tera-scale performance


Computer Vision



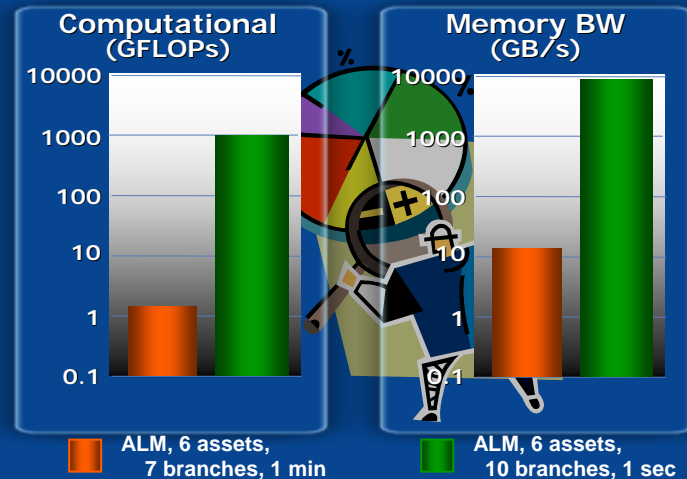
Ray-Tracing

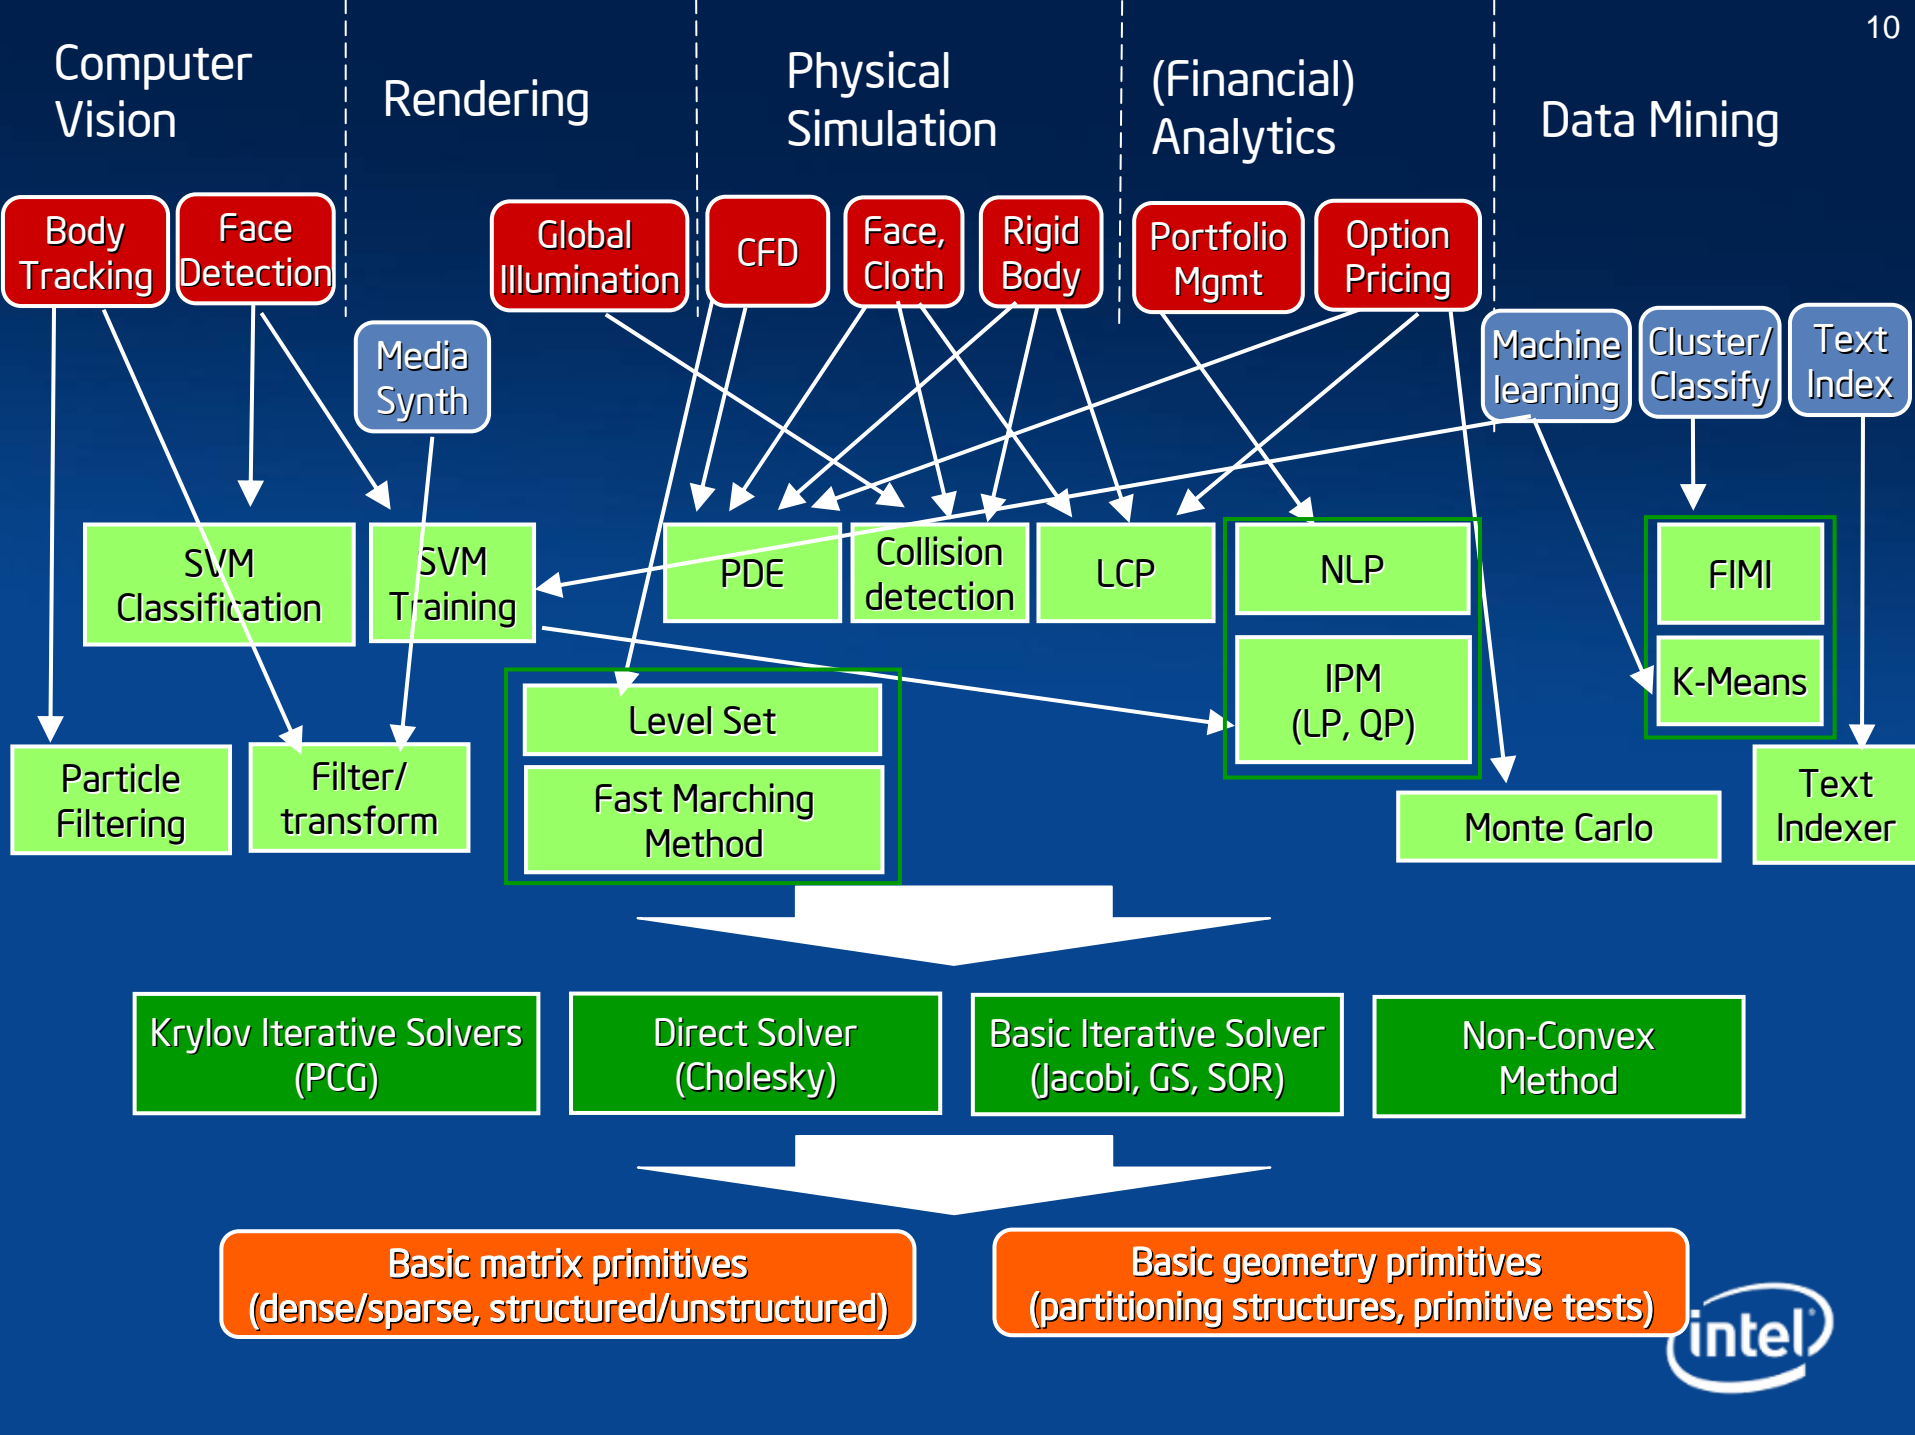


Physical Simulation

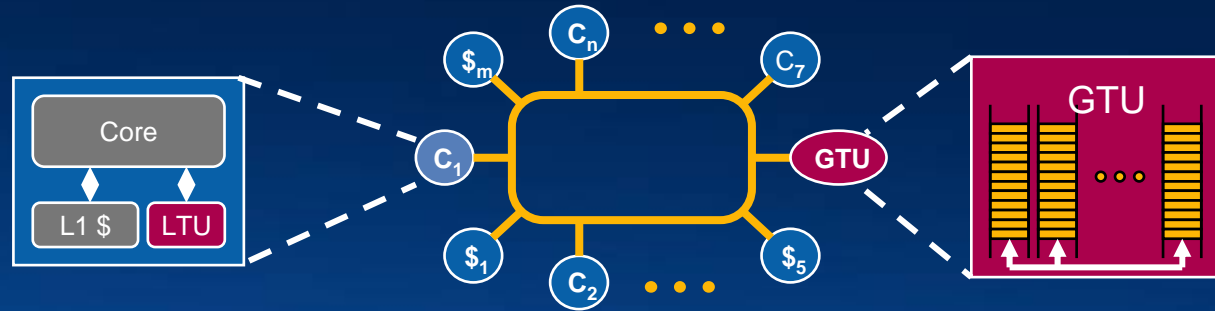



Financial Analytics





Application Acceleration: HW Task Queues



Task Queues

- scale effectively to many cores
- deal with asymmetry
- supports task & loop parallelism

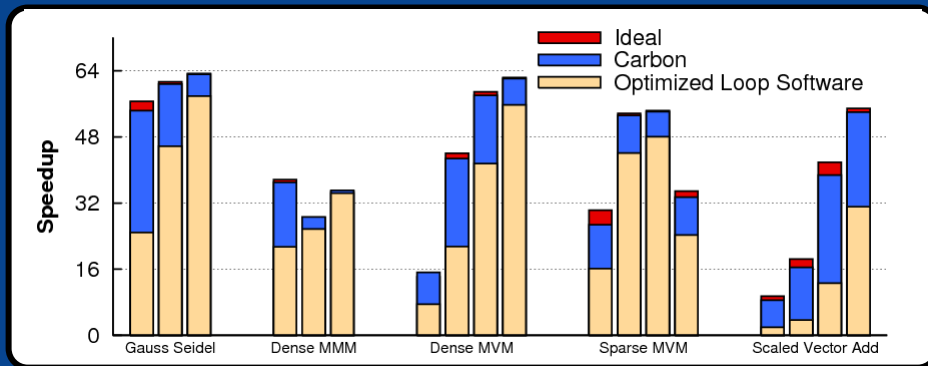
Local Task Unit (LTU)

Prefetches and buffers tasks

Global Task Unit (GTU)

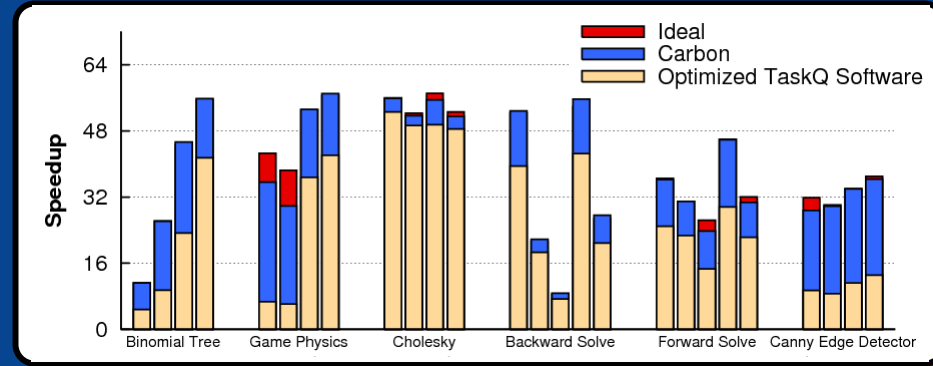
Caches the task pool
Uses distributed task stealing

Loop Level Parallelism



88% benefit optimized S/W

Task Level Parallelism



98% benefit over optimized S/W



Programming Environment Research

- **Languages & programming abstractions**

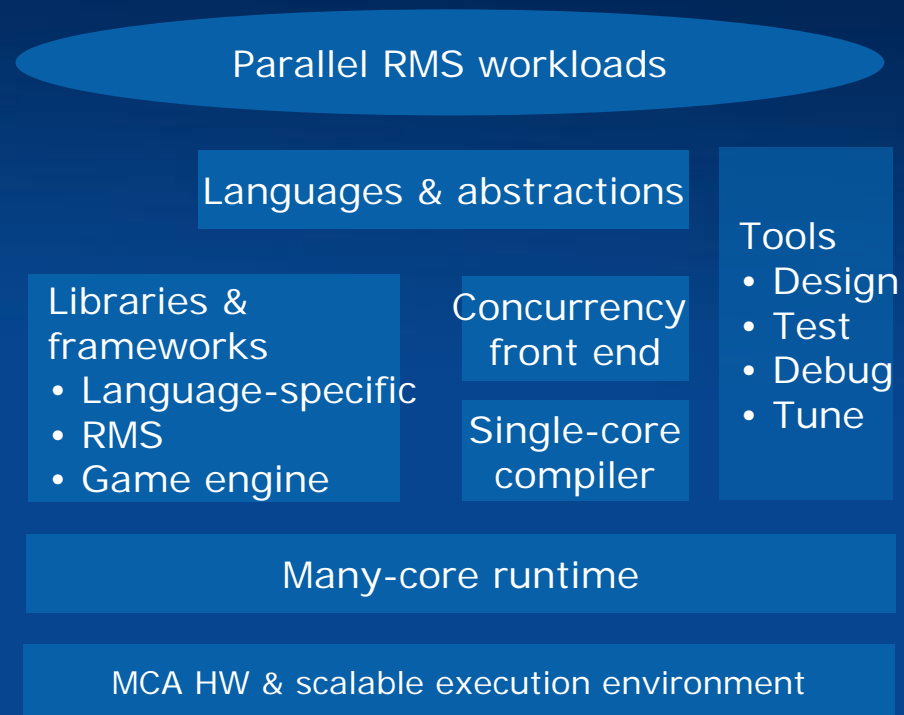
- Transactional Memory
- Data parallel operations
- Lightweight tasks
- Fine-grain synchronization
- Message passing

- **Compilers**

- Multi-language support
- Optimizations for Prog. Abst.
- Dynamic & static compilation
- Speculative multithreading

- **Many-core runtime scalability**

- Efficient scalable support for Programming Abstractions



Data parallel vector sum

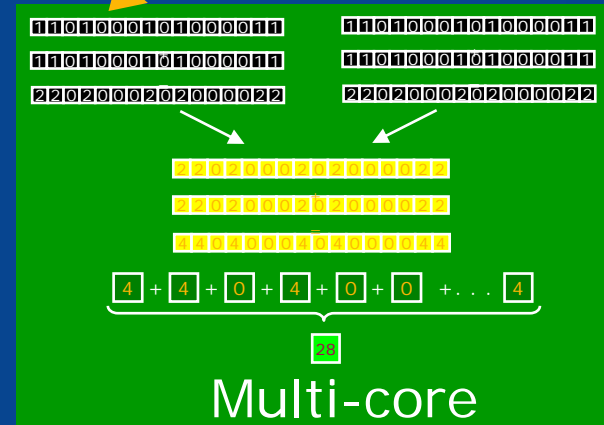
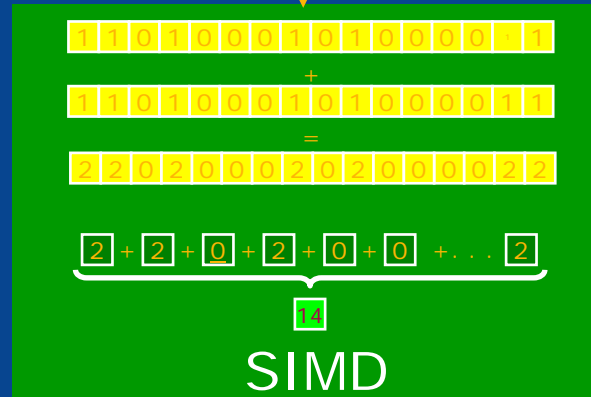

Data parallel
program

```
VEC B;
. . .
x = Sum_Of_Elements(B);
```

Data parallel compiler
& runtime

Compiler understands
and optimizes data
parallel operations

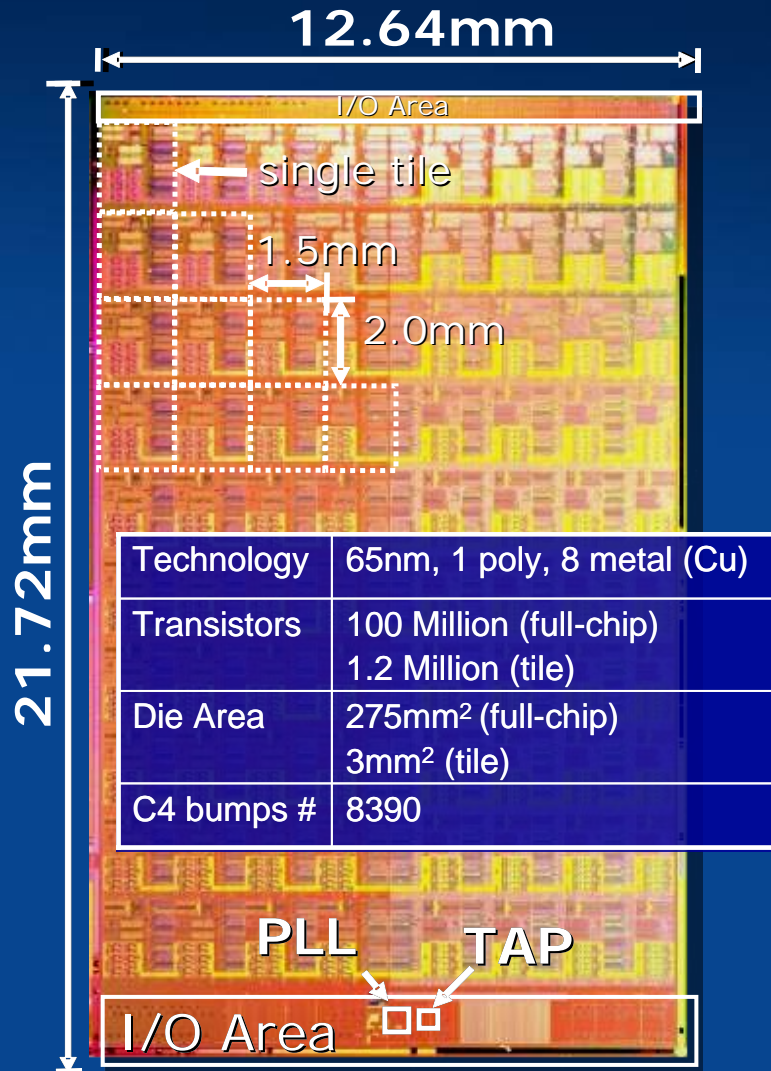
Automatic mapping
based on vector size
and architecture



Ct Manages Parallelism Without Exposing Threads



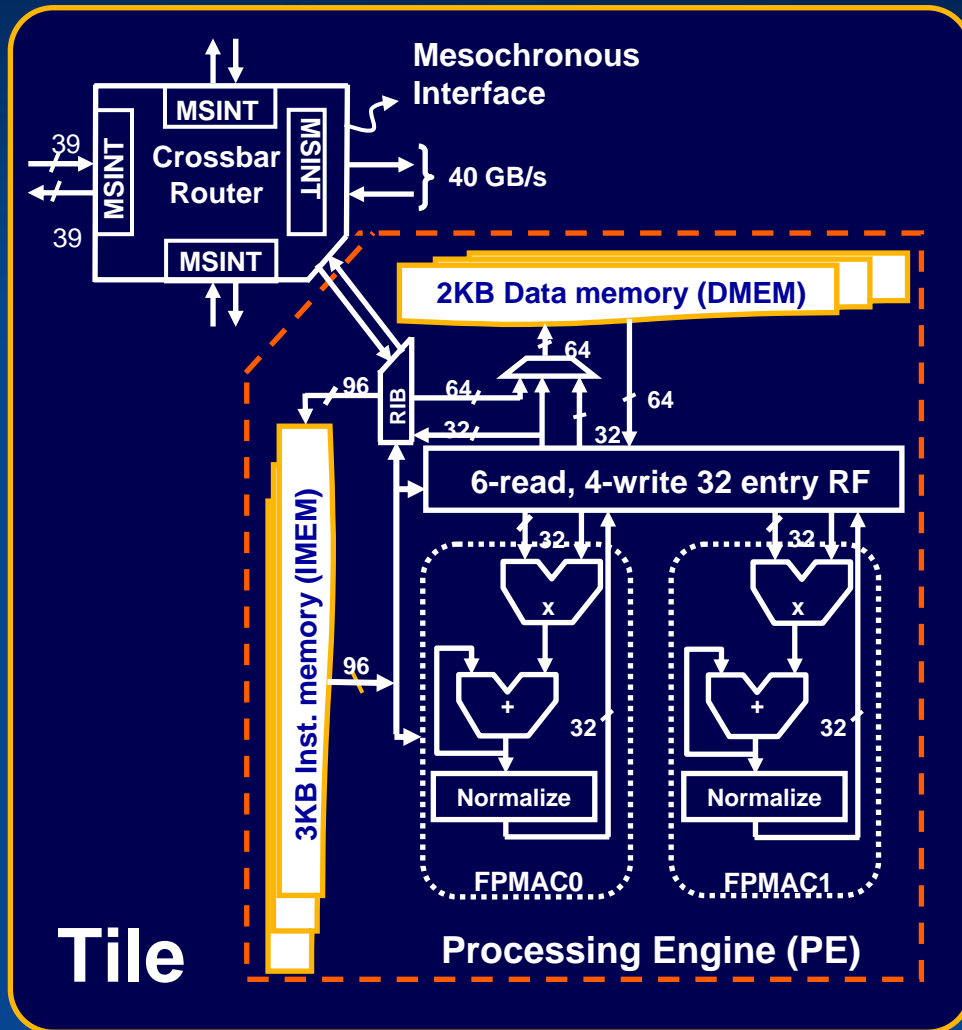
Teraflops Research Processor



Goals:

- Deliver Tera-scale performance
 - Single precision TFLOP at desktop power
 - Frequency target 5GHz
 - Bi-section B/W order of Terabits/s
 - Link bandwidth in hundreds of GB/s
- Prototype two key technologies
 - On-die interconnect fabric
 - 3D stacked memory
- Develop a scalable design methodology
 - Tiled design approach
 - Mesochronous clocking
 - Power-aware capability

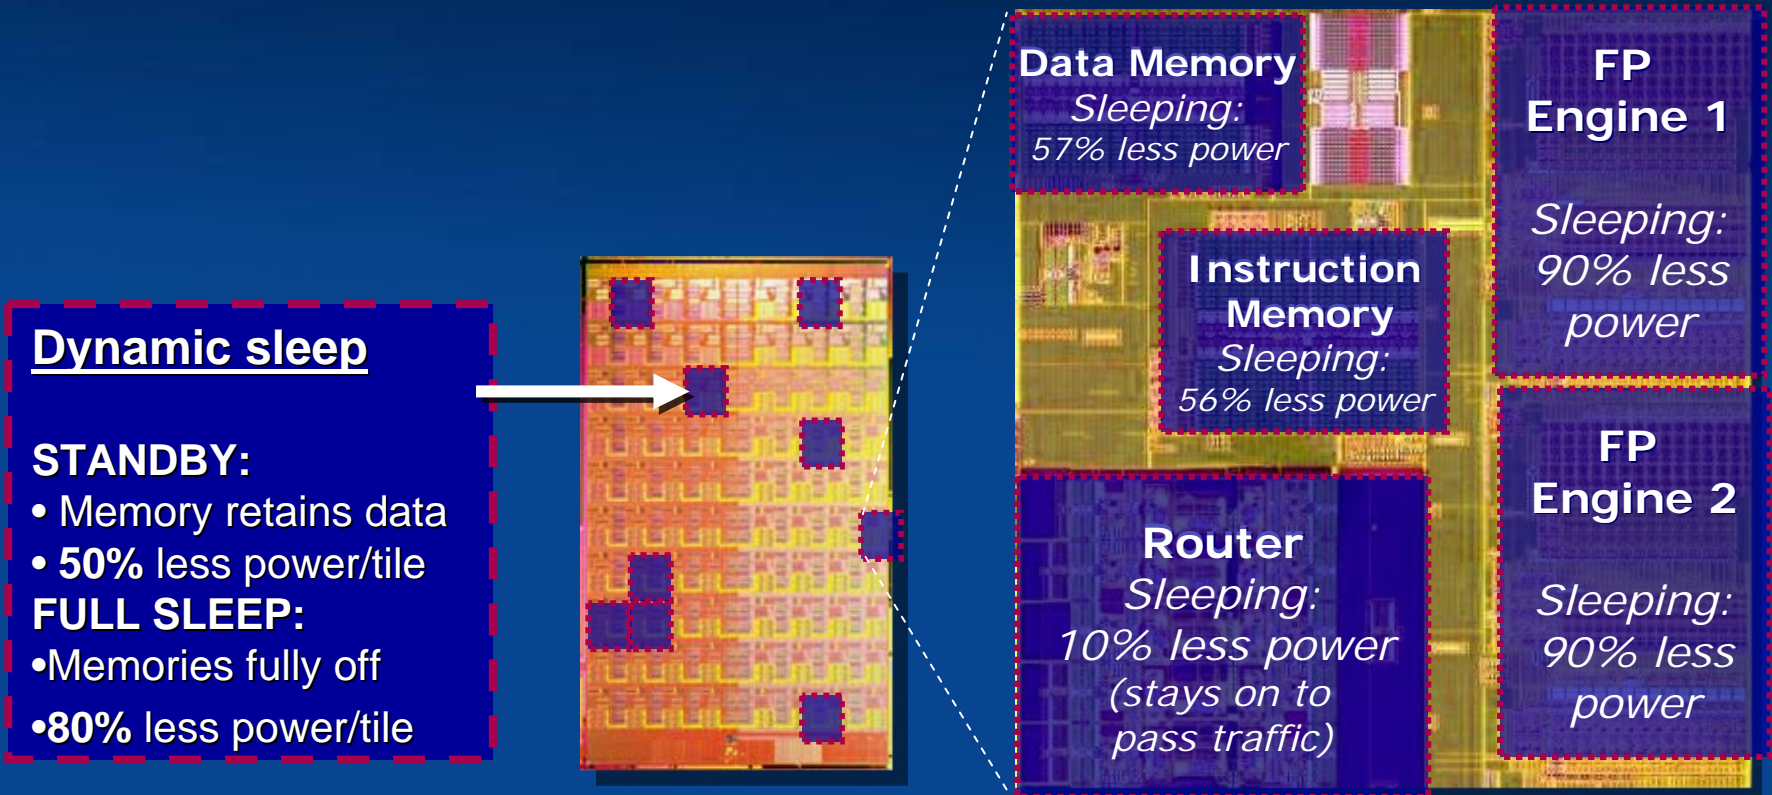
Key Ingredients



- Special Purpose Cores
 - High performance Dual FPMACs
- 2D Mesh Interconnect
 - High bandwidth low latency router
 - Phase-tolerant tile to tile communication
- Mesochronous Clocking
 - Modular & scalable
 - Lower power
- Workload-aware Power Management
 - Sleep instructions and Packets
 - Chip voltage & freq. control

Fine Grain Power Management

21 sleep regions per tile (not all shown)




Scalable power to match workload demands

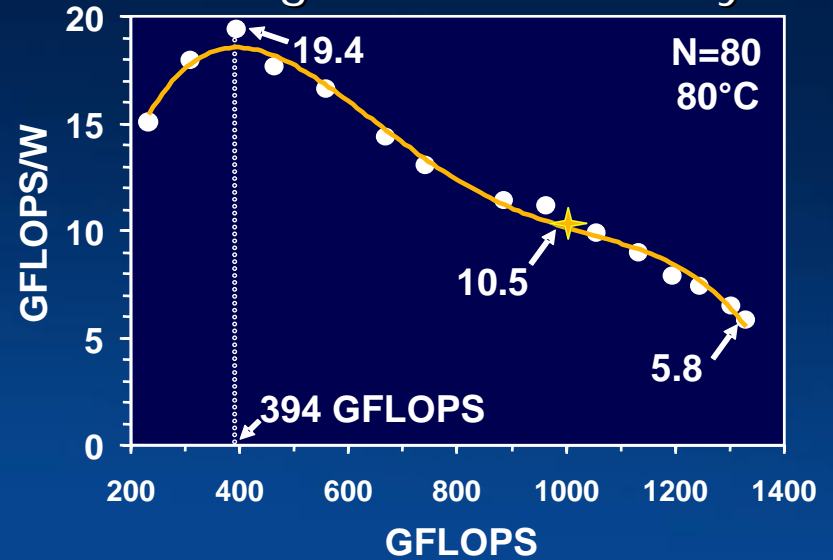


Power Performance Results

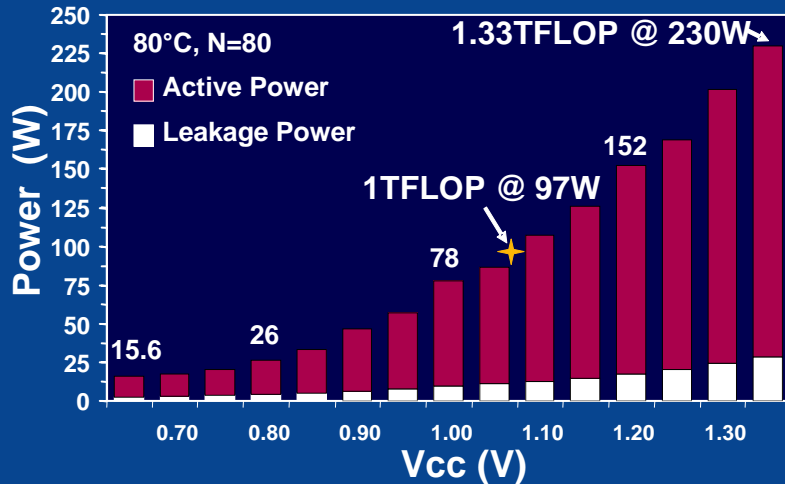
Peak Performance



Average Power Efficiency

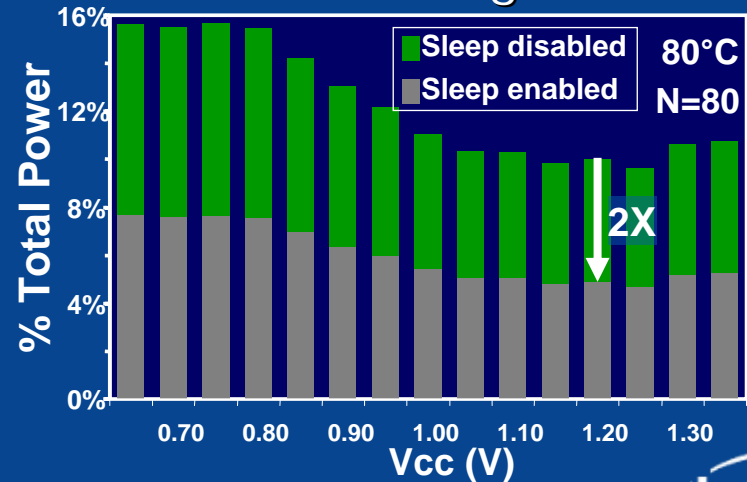


Measured Power



Stencil: 1TFLOP @ 97W, 1.07V;

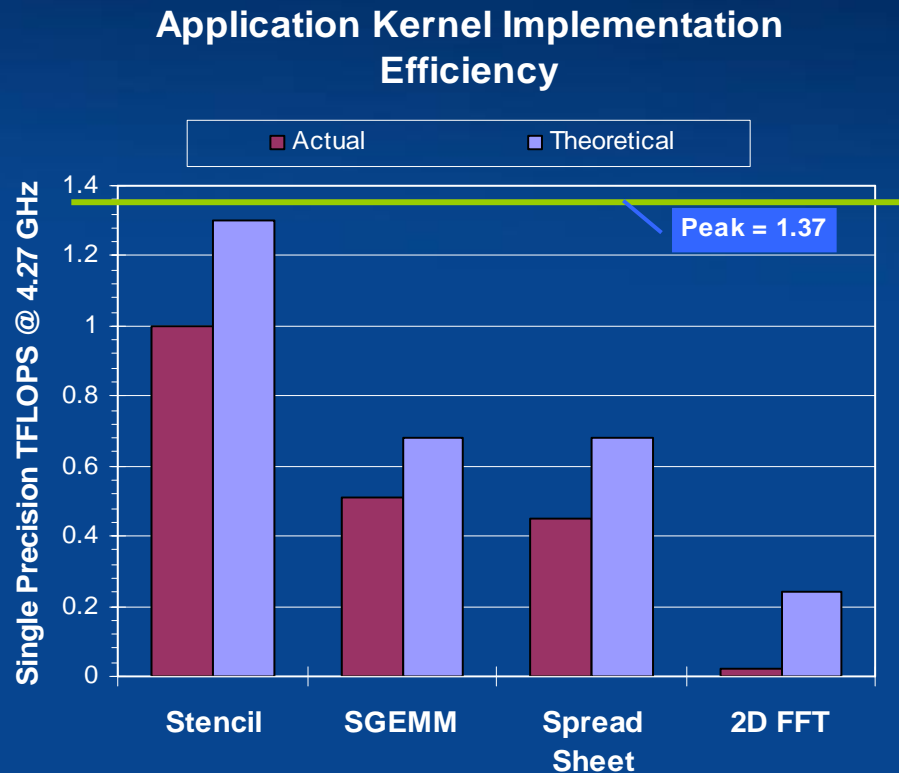
Leakage



All tiles awake/asleep



Programming Results



- Not designed as a general Software Development Vehicle
 - Small memory
 - ISA limitations
 - Limited data ports
- Four kernels hand-coded to explore delivered performance:
 - Stencil 2D heat diffusion equation
 - SGEMM for 100x100 matrices
 - Spreadsheet doing weighted sums
 - 64 point 2D FFT (w 64 tiles)
- Demonstrated utility and high scalability of message passing programming models on many core






Key Learnings

- Teraflop performance is possible within a mainstream power envelope
 - Peak of 1.01 Teraflops at 62 watts
 - Measured peak power efficiency of 19.4 GFLOPS/Watt
- Tile-based methodology fulfilled its promise
 - Design possible with ½ the team in ½ the time
 - Pre & Post-Si debug reduced – fully functional on A0
- Fine-grained power management pays off
 - Hierarchical clock gating and sleep transistor techniques
 - Up to 3X measured reduction in standby leakage power
 - Scalable low-power mesochronous clocking
- Excellent SW performance possible in this message-based architecture
 - Further improvements possible with additional instructions, larger memory, wider data ports



Work in Progress: Stacked Memory Prototype

256 KB SRAM per core
4X C4 bump density
3200 thru-silicon vias



Memory access to match the compute power



Summary

- Emerging applications will demand teraflop performance
- Teraflop performance is possible within a mainstream power envelope
- Intel is developing technologies to enable Tera-scale computing



Questions?

