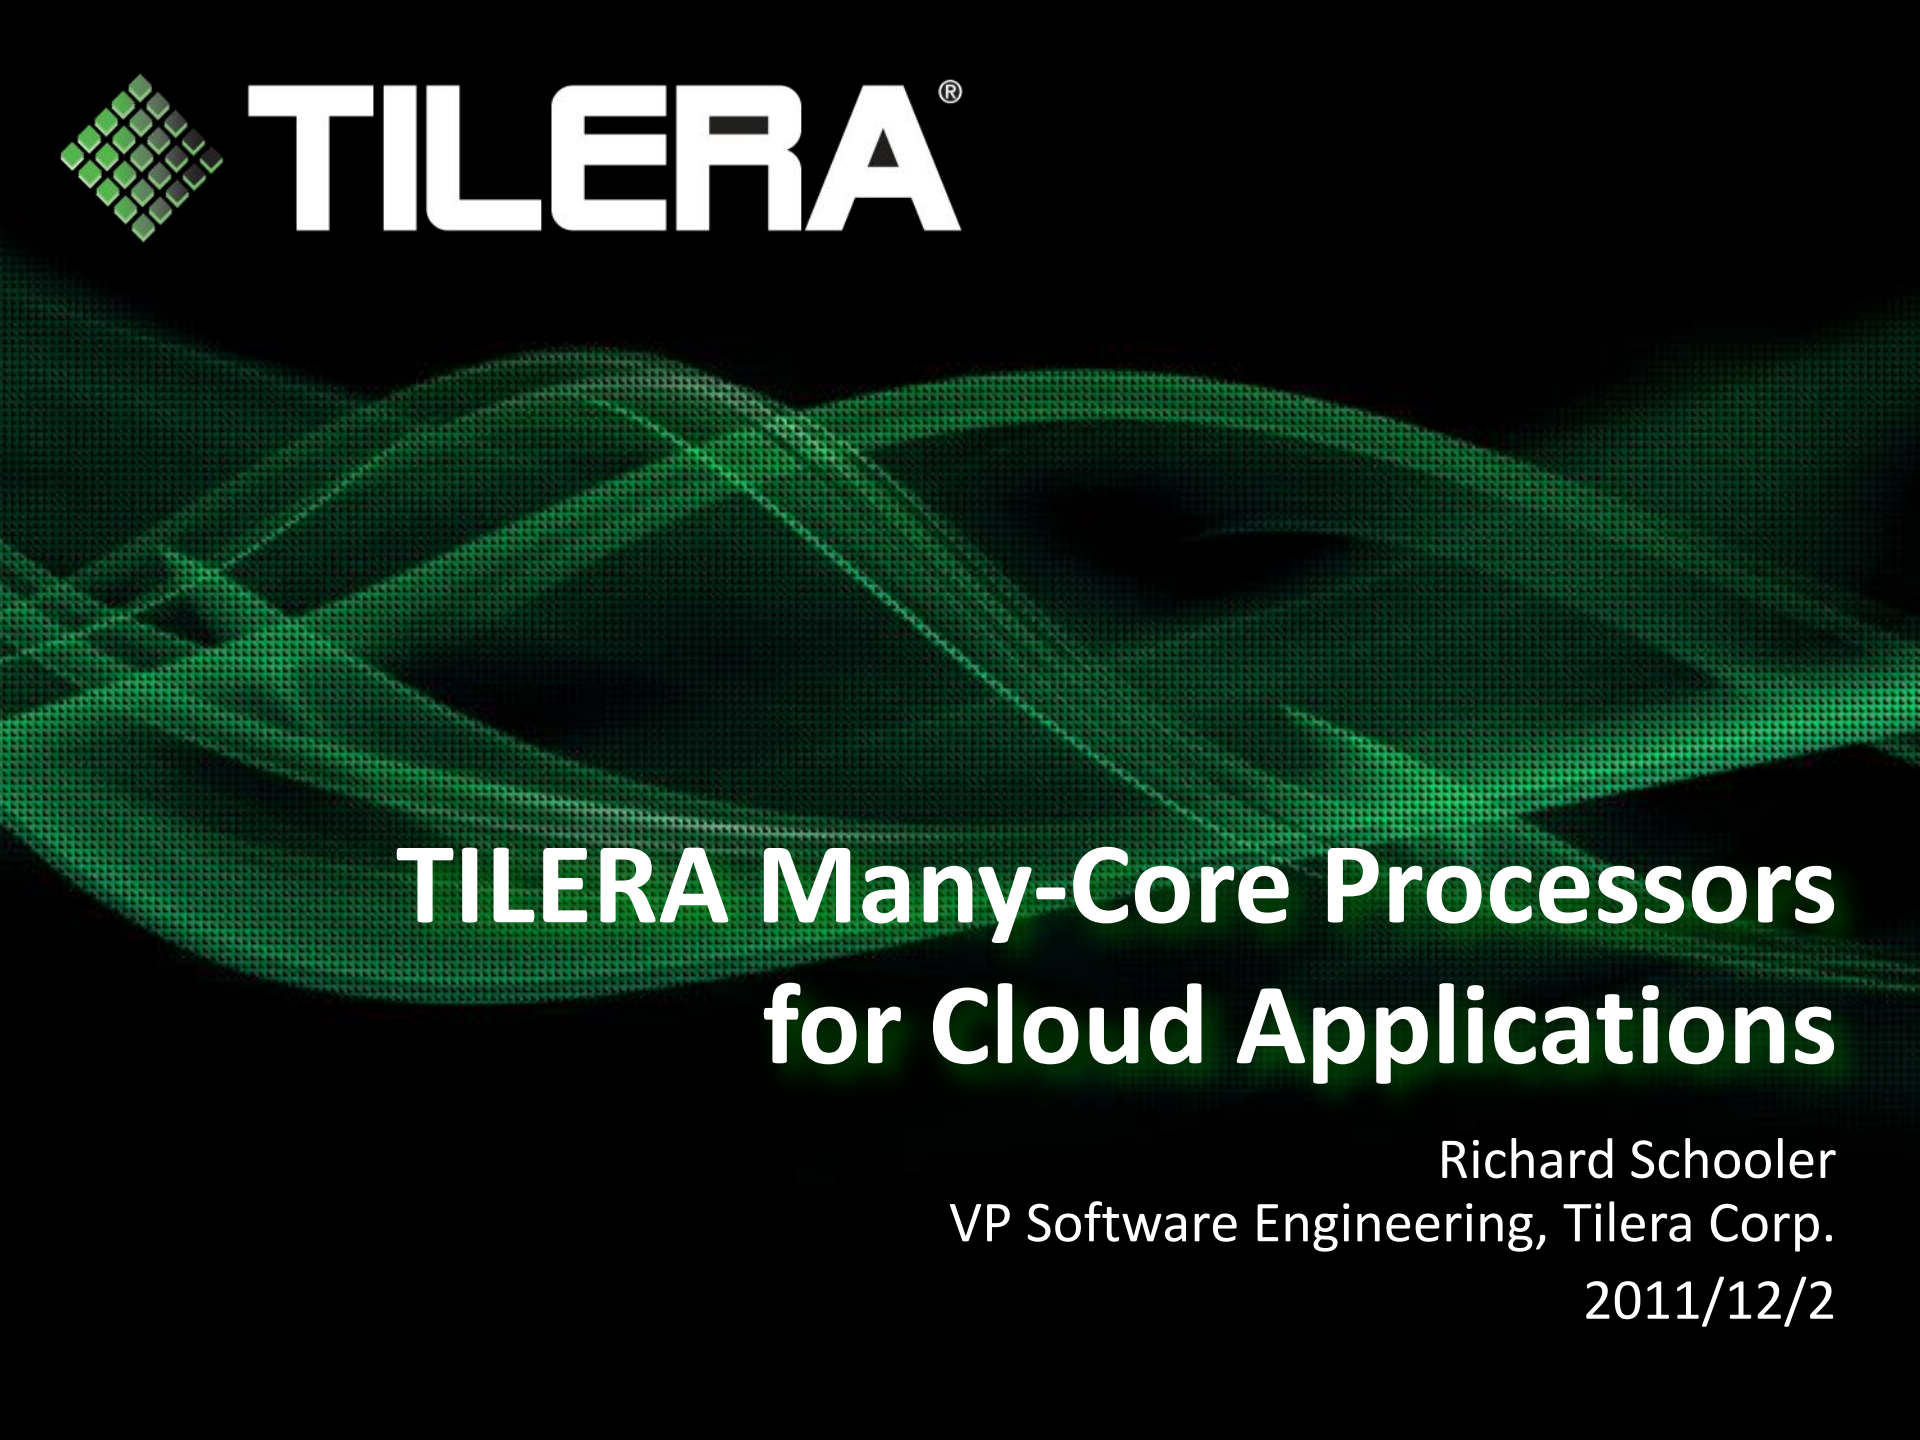

TILERA Many-Core Processors for Cloud Applications

Richard Schooler
VP Software Engineering, Tilera Corp.

2011/12/2

Agenda

- Company Overview
- Tilera Many-Core Technology
- Cloud Applications
  - Memcached, Media, Hadoop, etc.

Our Focus

- Compute performance
- Power efficiency
- Ease of programming

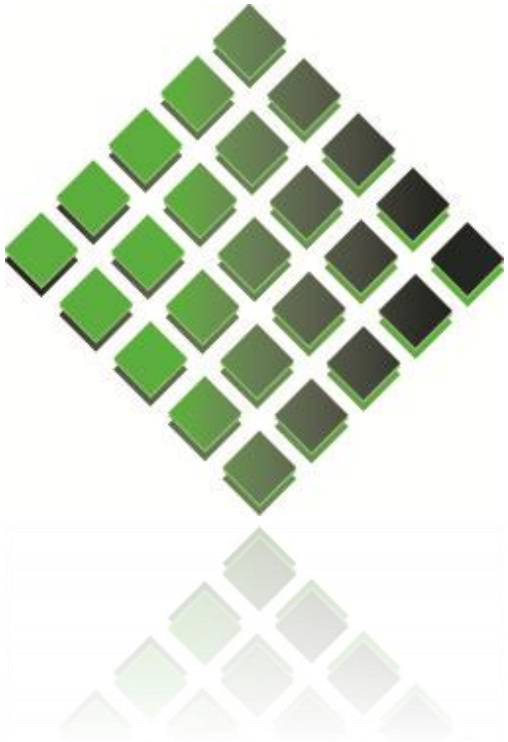

Company Timeline

A rich heritage in developing world's leading multicore

1994

32-node cache coherent multiprocessor based on a mesh

2002

First 16-core mesh-based multi-core processor (MIT RAW processor)

2004

Tilera founded and funded by Bessemer, Walden

2007

First commercial 64-core processor shipping

2009

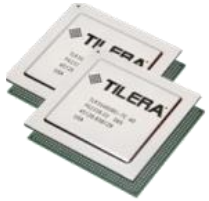

Second generation commercial 36/64-core processors shipping

2011

Third generation processors with 16, 36, 64 and 100 cores.

16 years of world-class research and development

Why Do We Need Many-Core?

Applications require ever-increasing performance

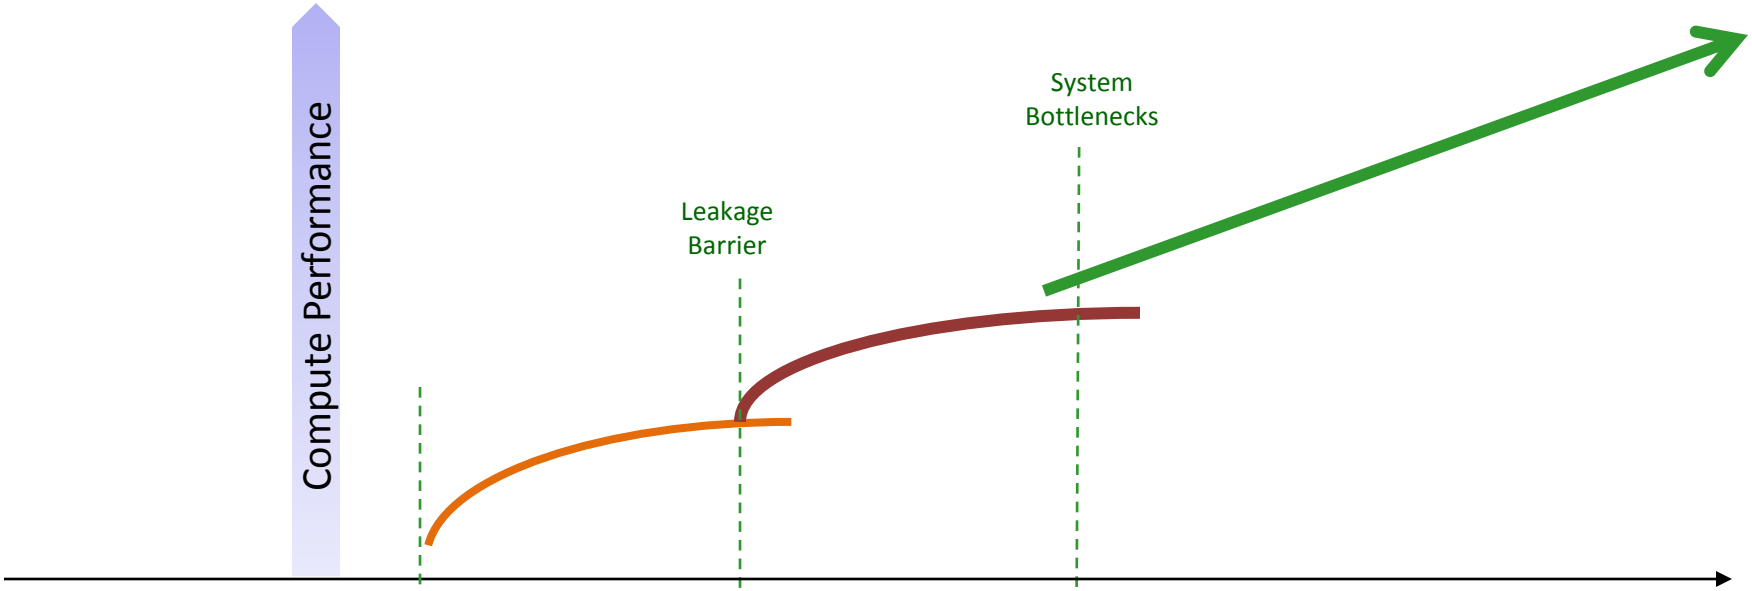

Applications	1990	2000	2010	2015
Networking	100Base-TX	1 GbE	10 GbE	100 GbE
Multimedia	SVGA	SD	HD	H.265 Super HD
Internet	World Wide Web	Web 2.0	Cloud	Media Cloud RAM Cloud Security Cloud ?

Current and future performance requires scalable many-core

Scalability, the Promise of Many-Core

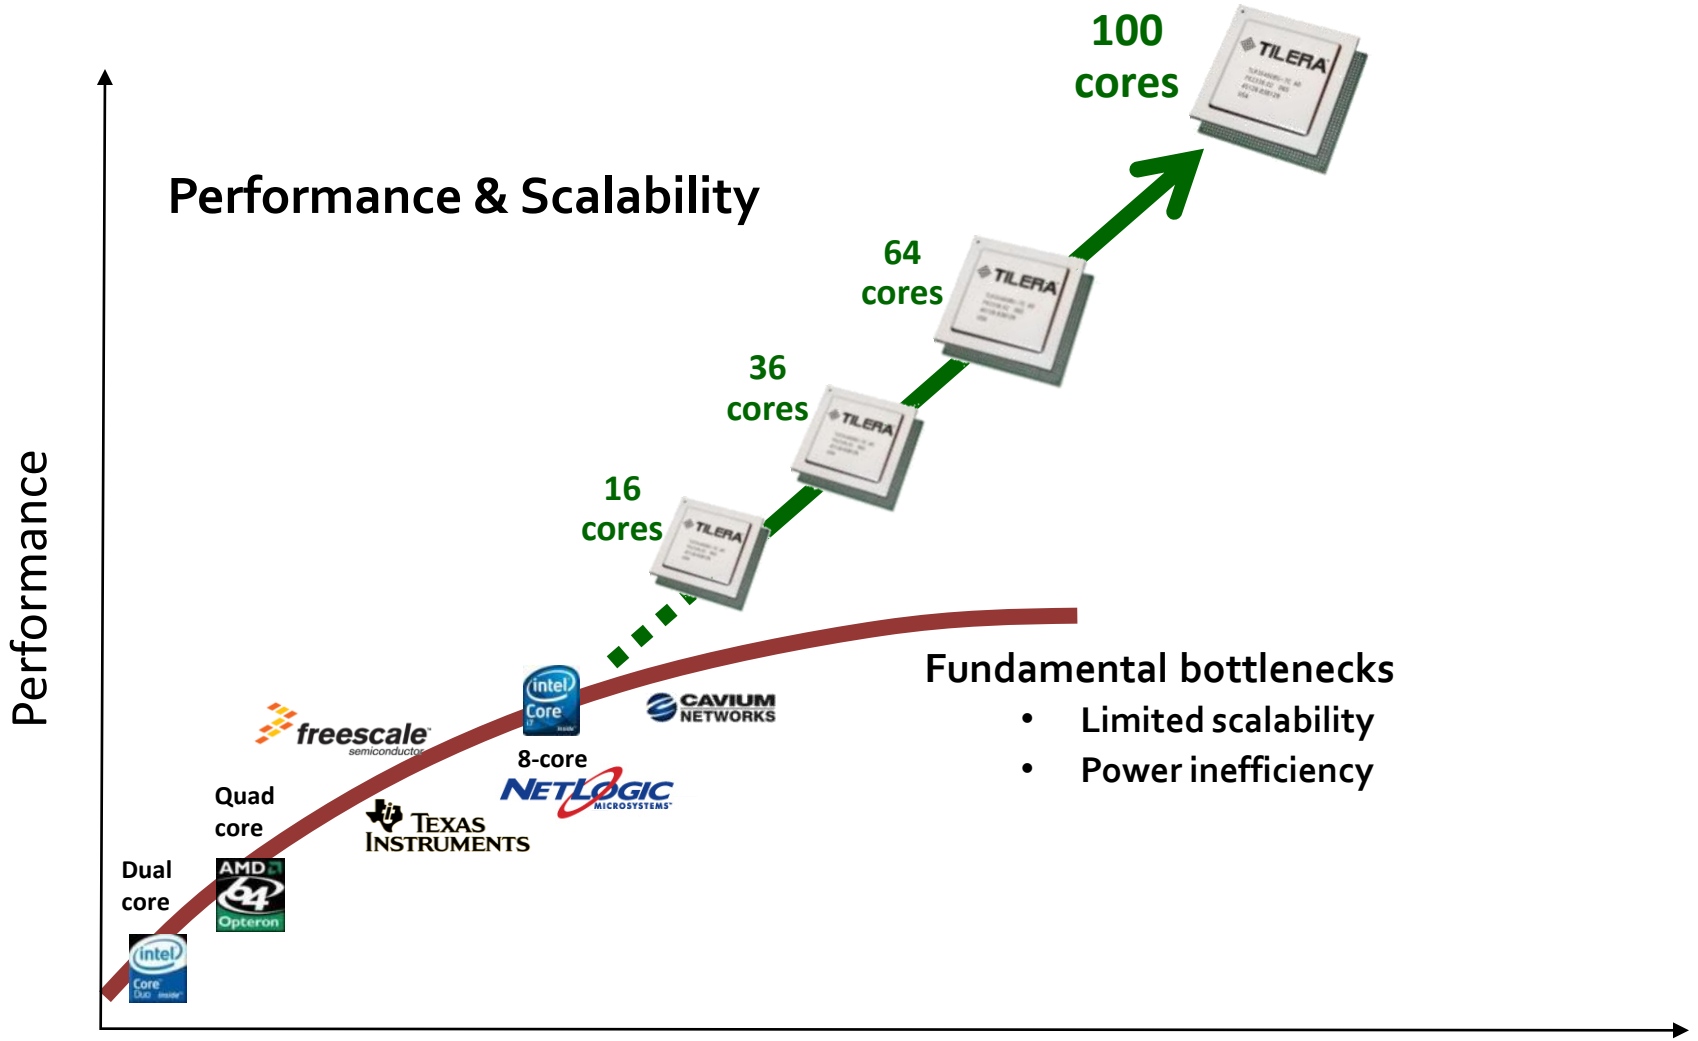

Only Tilera's architecture scales to hundreds of cores

Many-Core Architecture Efficiency

Better Performance/Watt/$

x86 Technology

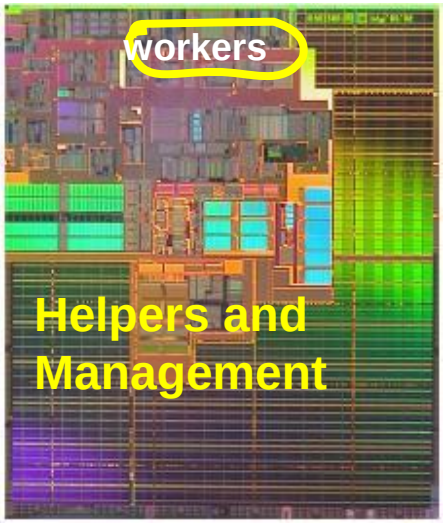

Madison Itanium2

- Most of x86 die will not be used for web 2.0 tasks
- Wasted cost and power

Tilera processors

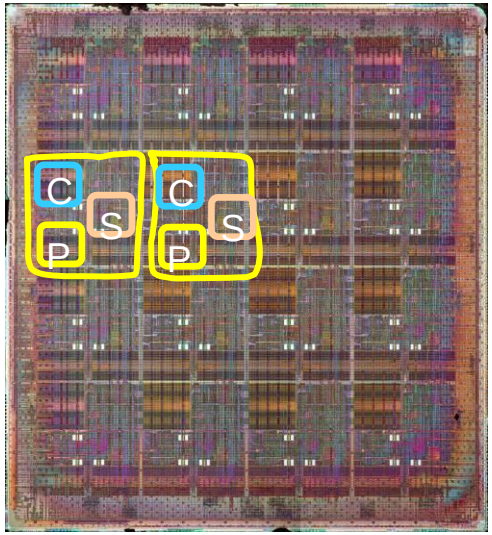

- Web 2.0 tasks are small and many
- Many efficient cores fit the requirements

SPECint_rate2006: 96.4 @ 25W = 3.9/W, 4x Sandy Bridge (GCC)

Tilera Many-Core Architecture

TILE Architecture

Complete 64-bit cores with integrated cache

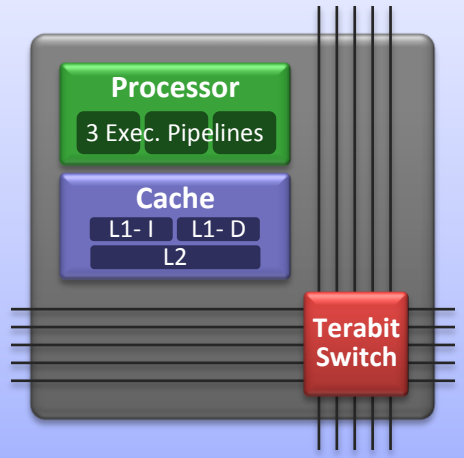

< 400 mW Power

2 Dimensional on-chip mesh network

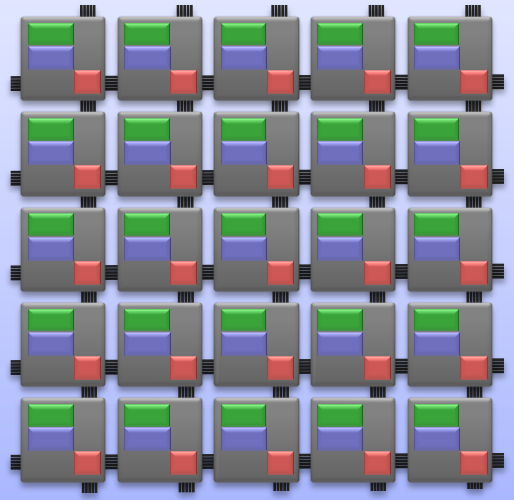

iMesh: Terabits of on-chip BW

Powerful SoC features

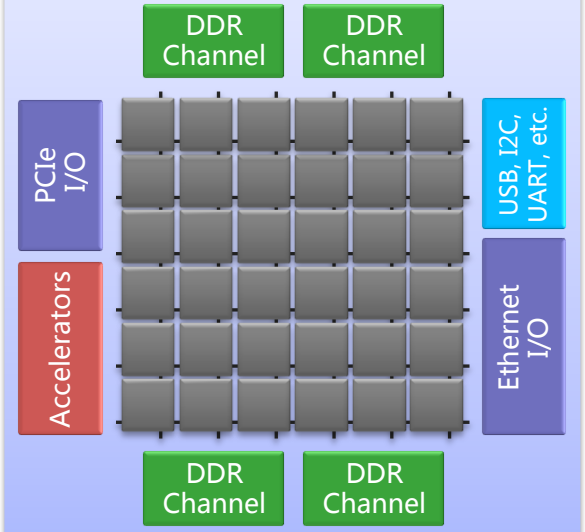

Low-latency mesh interconnect to RAM & I/O

Server on Chip Architecture

Highly Integrated & Low Power Many-Core

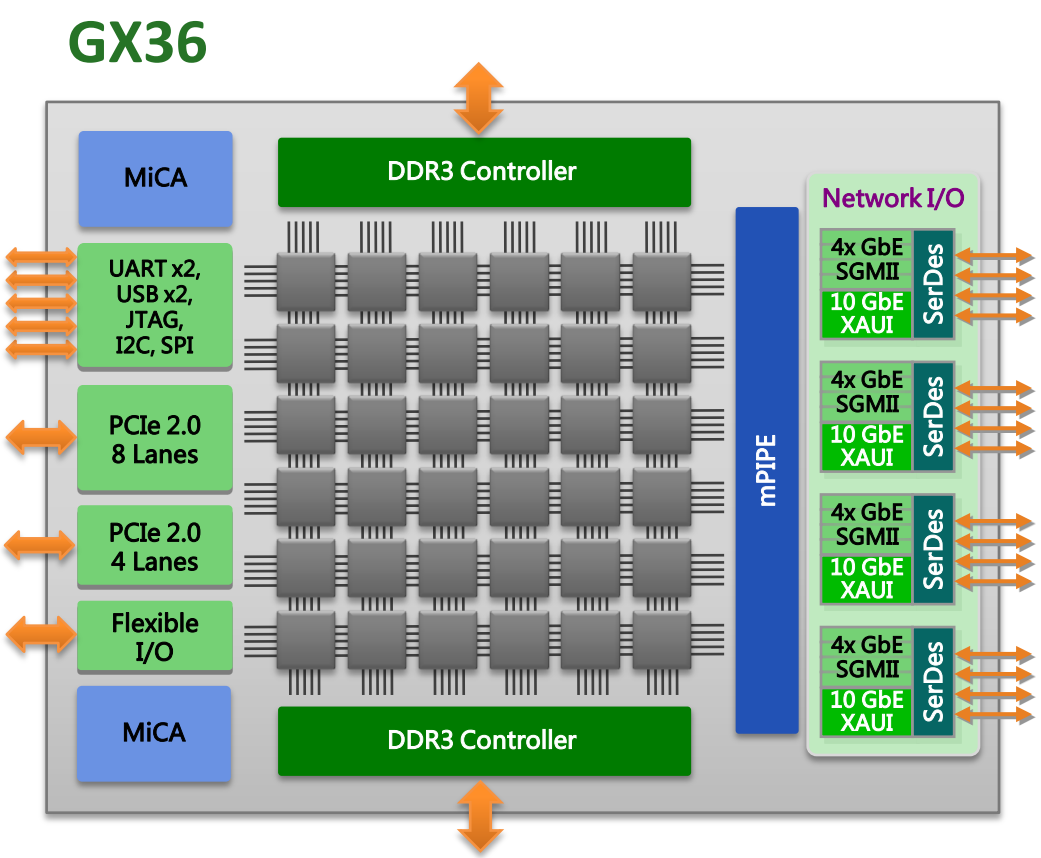

Tiles

- 36x 64-bit Processor Cores
- 1.2GHz
- 12 MBytes total cache
- 66 Tbps iMesh BW

DDR3 RAM

- 200 Gbps memory BW
- Address up to 512GB DRAM

I/O

- 40 Gbps packet I/O
  - 4 ports XAUI / 16 GbE (SGMII)
- Programmable packet engine
  - 60 Mpps throughput
- 48 Gbps PCIe I/O

Acceleration

- MiCA™ Engines:
  - 20-40 Gbps crypto (IPsec, SSL…)
  - 10 Gbps compress & decompress

TILE-Gx Covers Server Performance Spectrum

First Generation
Tile Server Processor
TILEPro, 90nm

Second Generation
Tile Server Processor
TILE-Gx, 40nm

28 nm

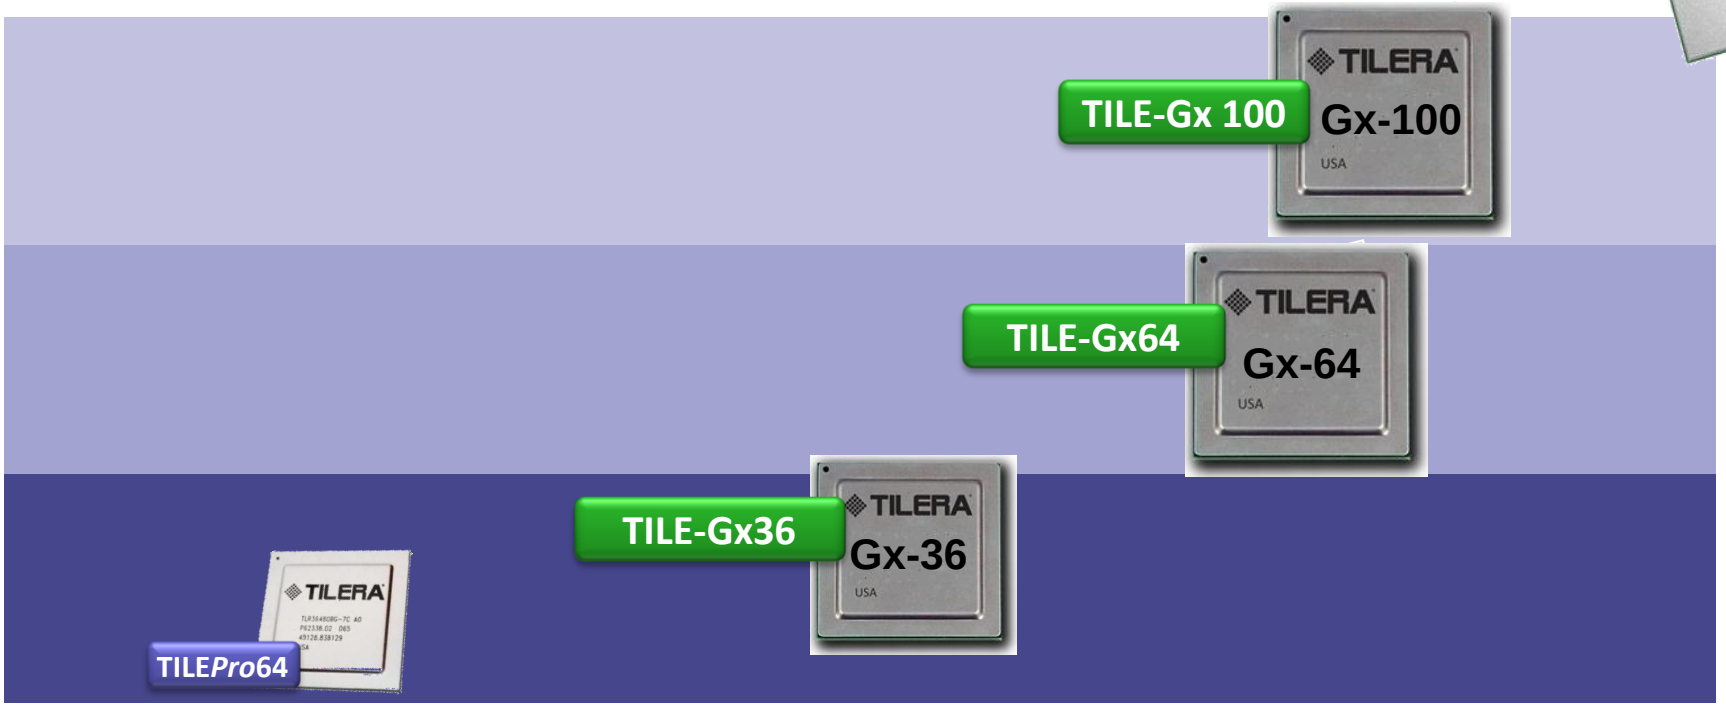

2009 **2011** **2012** **2013**

Tilera System Software

Familiar, open source, software stack

Multicore Development Environment (Runtime Software)

Standard application stack

Application Layer
- Standard C/C++ libs
- Open JDK

Operating System
- SMP Linux
- Bare metal

Hypervisor Layer
- H/W abstraction (KVM)
- I/O device drivers

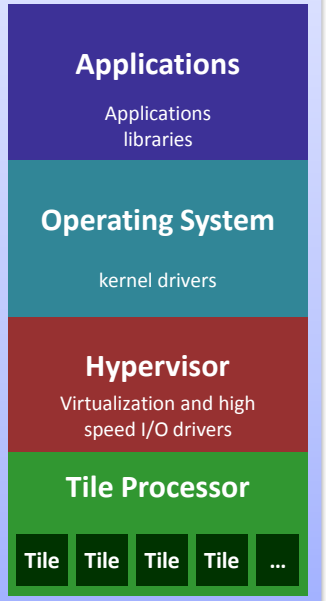

Innovative Stack Components

Zero Overhead Linux
- Real-time performance
- Mutes all Linux interrupts

Tile Multicore Components
- Parallel programming
- Performance tuning

Network interface drivers
- Many core load balancing
- Network stack bypass

Standard stack to run existing applications

Tilera Software Tools

Best-in-class application development tools

Multicore Development Environment (Tools)

Standards-based tools

Standard Programming

- GNU ANSI C/C++
- Java, Erlang, TBB
- PHP, Perl, Python

Integrated Tools

- GCC V4.4 compiler
- Gdb, mudflap
- gprof, oprofile
- perf_events
- Eclipse IDE

Innovative Multicore Tools

Multicore Debugger

- Aggregate control
- State display

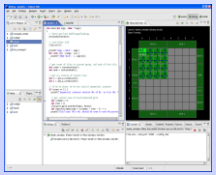

Multicore Profiler

- Architectural state
- Application states
- Source-level tracking

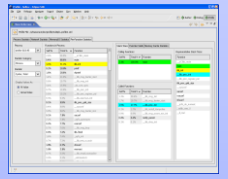

Full Chip Simulator

- Timing & functional accuracy

Standard Programming for Porting Simplicity

Tilera: The Solution for Cloud Requirements

- First generation servers in production
  - 1U 64 Core Server: TILEmpower
  - 2U 512 core server: Quanta S2Q
  - Security Gateway , Media Processing, Memcached

- Second generation server processor
  - TILE-Gx 64-bit many-core processor
  - Integrated functions: I/O, crypto, etc.
  - Sampling now

- Proven value proposition
  - Higher performance, lower power
  - Facebook published 4x better perf/watt results

facebook

Memcached Study

- Facebook paper: International Green Computing Conference, 2011

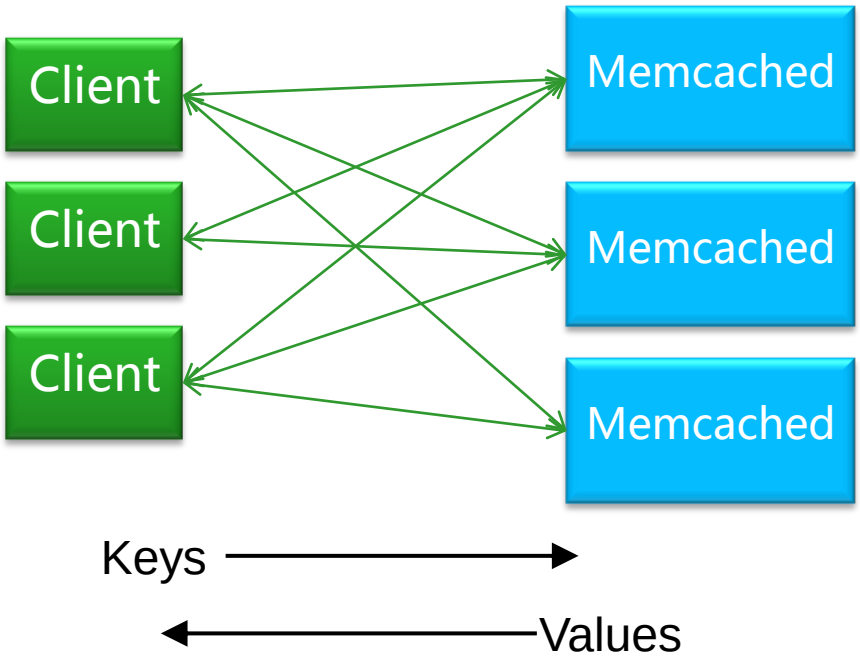

- Key-value store
- In-memory database
- Distributed hash-table
- Static key hashing
- Extremely stable: thousands of clients, memcached servers

- Compare Xeon, Opteron, TILEPro64

Memcached Study - Results

- PRO64 Processor 4X performance per watt over x86
- “Low-power many-core processors are well suited to Kvstore workloads”
- “TILEPro64 can yield at least 67% higher throughput than low-power x86 servers at comparable latency.”
- “S2Q server with 8 processors handles at least three times as many transactions per second per Watt as the x86-based servers”
- “…even within a single [x86] socket with 4 cores, performance scales poorly. … TILEPro64 implementation can easily take advantage of more cores for higher performance.”
- “…This trade-off (using Wimpy nodes like Atom) can mean higher costs for hardware, system administration, and fault management of very large clusters.”

Tilera: Best Performance/Watt

Best Performance for Media Cloud

- Tilera TILEPro provides 4x performance per Watt advantage over Westmere and Sandybridge
- For content providers real-time video delivery
- 10,000 streams of SD transcoding case study
  - Intel solution: 500 servers, 125 KW
  - Tilera solution: 80 servers, 35 KW!!

Video	Intel Westmere 6 cores 2.8 GHz	Tilera TILEPro64 866 MHz	Performance Advantage	Power Advantage
1080P Frames	166	270	**1.62 Advantage**	**4x Advantage**

Video	Intel Sandybridge 4 cores 3.4 GHz	Tilera TILEPro64 866 MHz	Performance Advantage	Power Advantage
1080P Frames	168	270	**1.61 Advantage**	**4x Advantage**

Cloud Server Processing Offload

Make your x86 work more efficient

Offload Application:

Network Packet Fast Path forwarding , DPI, TCP/IP

Crypto Protocol offload:

TCP/IP, IPSec, OpenSSL

AES/DES/3DES/ Kasumi/Snow3G

RSA/SHA-1

GZIP offload

Computing Offload

Snort

Video, Picture processing

GX8036 PCIe Offload Card

- **CPU GX8036**
  - 36x 64bit 1Ghz Core offer **1080 BOPS**
  - 12MB Cache
- **Memory**
  - Dual Channel SODIMM of 64bit 1333 DDR3 with ECC
  - Up to 21.3GB/s Memory bandwidth
  - Up to 8GB ECC DDR3 Size
- **Hardware Accelerator**
  - Crypto: 20Gbps Symmetric Crypto
    - RSA 1024/1024 CRT 40K/s
    - True Random Number Generator
  - GZIP: 10Gbps full duplex
- **Networking Interface**
  - Four line rate 1/10G Ethernet ports
  - Powerful Hardware packet processing engine
  - Hardware Time stamp/ IEEE1588v2
- **PCIe**
  - 8 lane PCIe 2.0 40Gbps bandwidth
  - Strong DMA engine

Hadoop on Many-Core

- Java: porting HotSpot C2 (server) JIT compiler
- Hadoop tuning:
  - Exploit many-core parallelism
  - JVM tuning, etc.
- Are you interested in working with us?

High Performance Offload Card for Hadoop Cluster

- Offload GZIP Processing
- Network Access and Protocol Stack
- Crypto
- Compute Processing

Make Your Hadoop Cluster more efficient

Conclusion

- The conference is over; time to look to the future
- The many-core era is coming
- Let’s work on green Cloud infrastructure together!

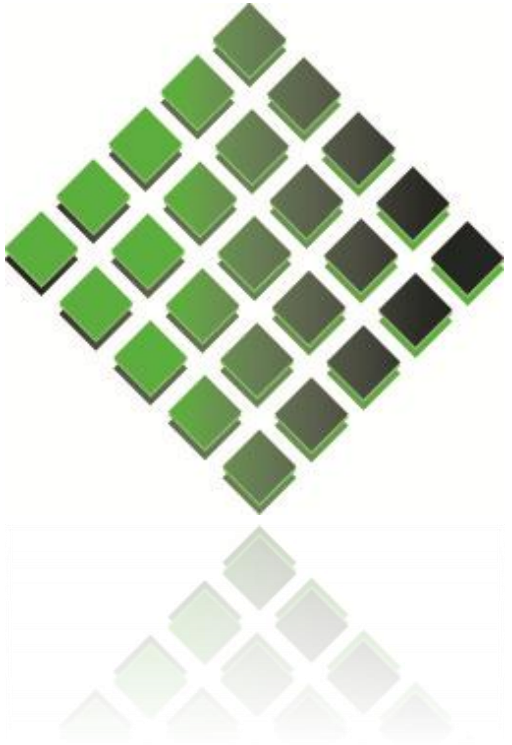

SPECint_rate2006 Performance/W

Benchmark	Gx 36 36 cores 1.2GHz w/GCC 4.4 (25 W)	Sandy Bridge 4 cores/8T 3.4GHz w/GCC 4.4 (95W)
perlbench	108.63	138.00
bzip2	66.53	90.70
gcc	74.14	59.60
mcf	60.91	61.40
gobmk	146.25	124.00
hmmer	95.40	85.50
sjeng	147.05	123.00
libquantum	68.56	73.50
h264ref	246.25	194.00
omnetpp	76.38	61.10
astar	67.13	67.80
xalancbmk	111.38	94.10
SPECint Rate	**96.4**	**91.2**
SPECint Rate/W	**3.9**	**0.96**

Higher performance than high-end Xeon

4x performance per Watt