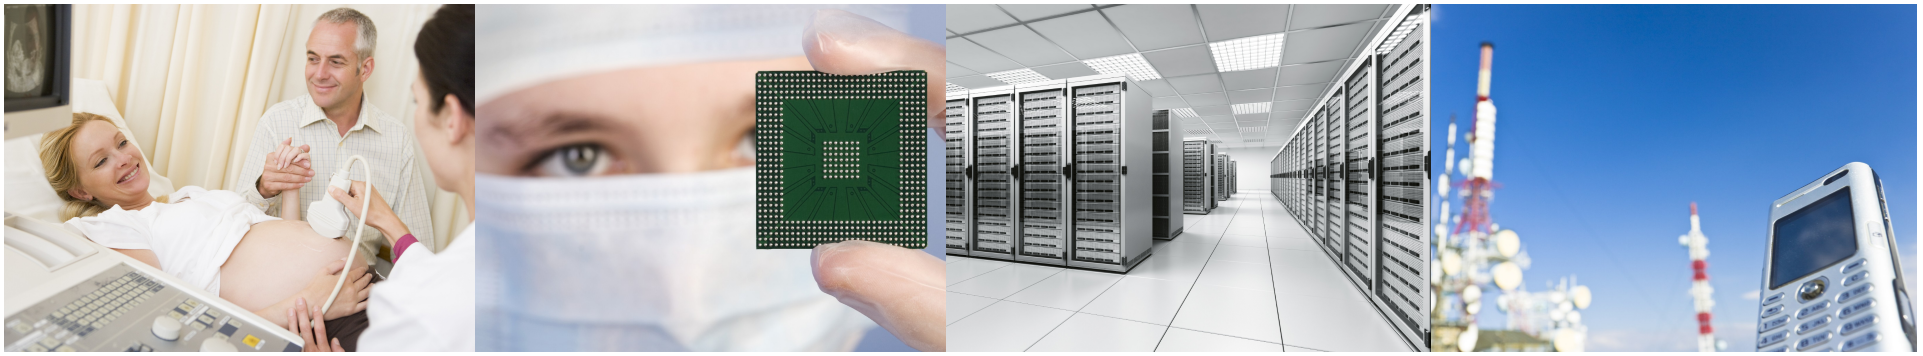

Hybrid System Design: The Only Practical Way

September 2010

adapteva

Company Introduction

- Adapteva has developed a ground breaking signal processing technology with energy efficiency of 50 GFLOPS/Watt
- Veteran SOC design team with 15 successful low power product tape-outs and over $100M in product related revenue
- Founded 2008. $1.5M raised in fall 2009 through strategic Series A strategic investments
- Now sampling first product to lead customers and strategic partners

Hard Times for Semi. Companies

- $10M development cost per chip
- ARM RISC Core up to 120MHz
- USB 2.0 , Ethernet, CAN, GPIO, etc
- 512KB flash, 65KB SRAM
- 12 bit A/D, 10 bit D/A

$2.36 in 10ku

vs.

- Coffee beans
- Water..

$4

adapteva

More Bad News for Semi. Companies

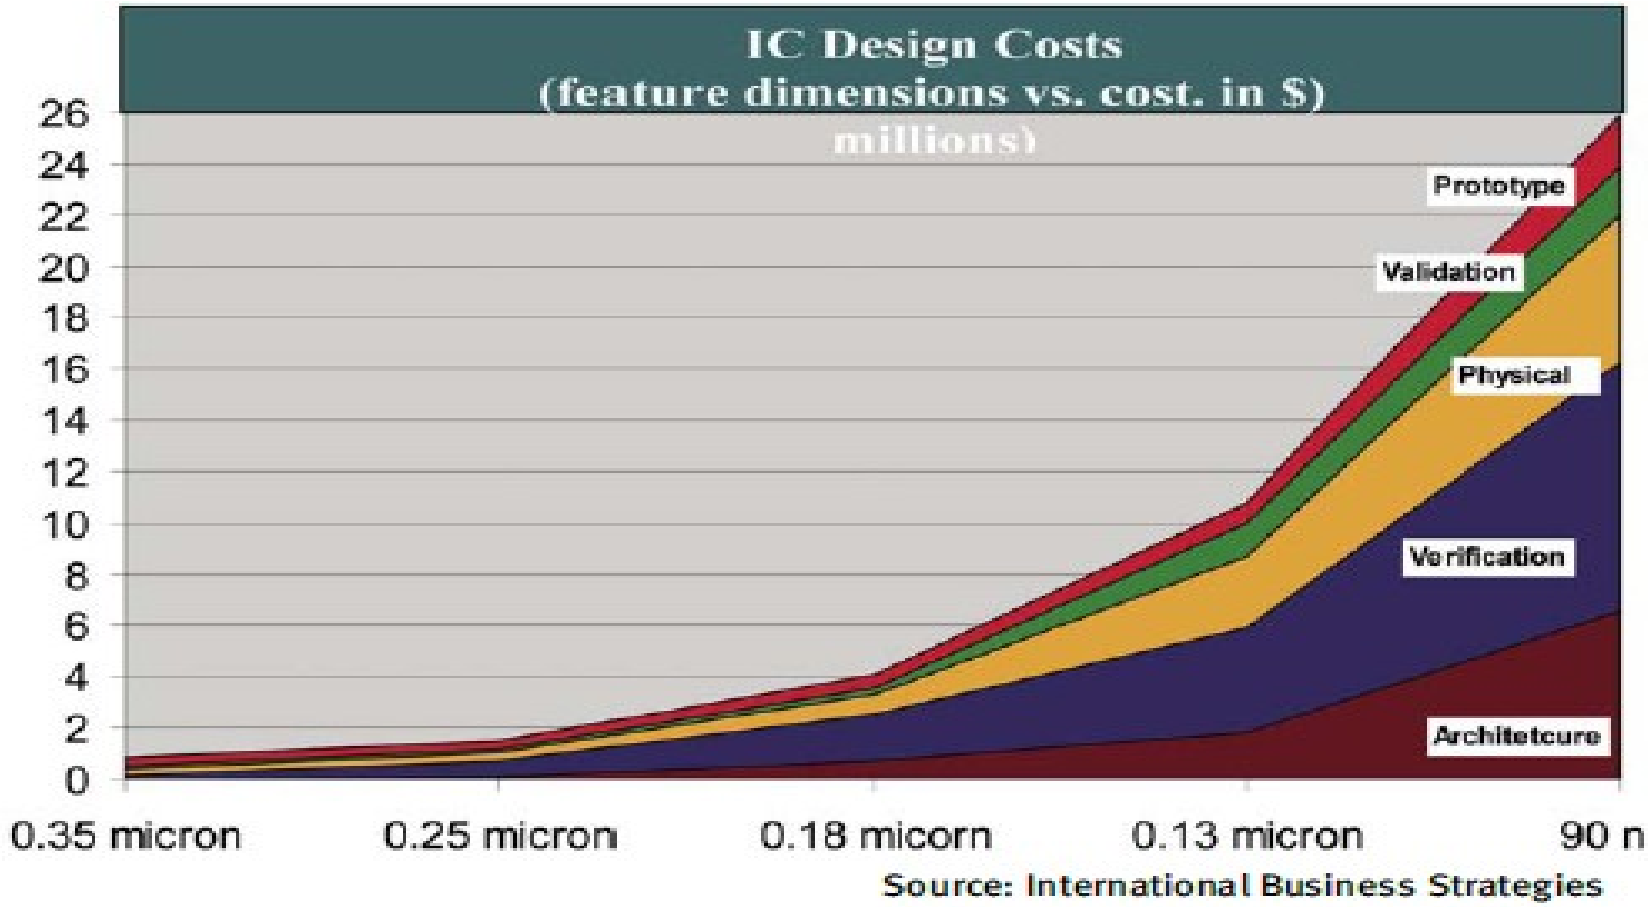

Ashl… …ailable dollars for companies like NVidia, Qualcomm and Apple. And thatis without the fabrication process to actually produce the chip … as a $… million… Collet… percent of this cost is labor and that the major part of the ove… is verification.

adapteva

Honest VC Response to Semi-Entrepeneur

- Assumption: SOCs cost $50-70M
- Observation: Chip companies are risky
- Observation: Chip companies take time
- Rule: VCs want 10x return
- Trend: P/E for HW companies are BAD!

- Result: Chip companies don't get funded

THIS IS
BAD NEWS FOR THE INDUSTRY!

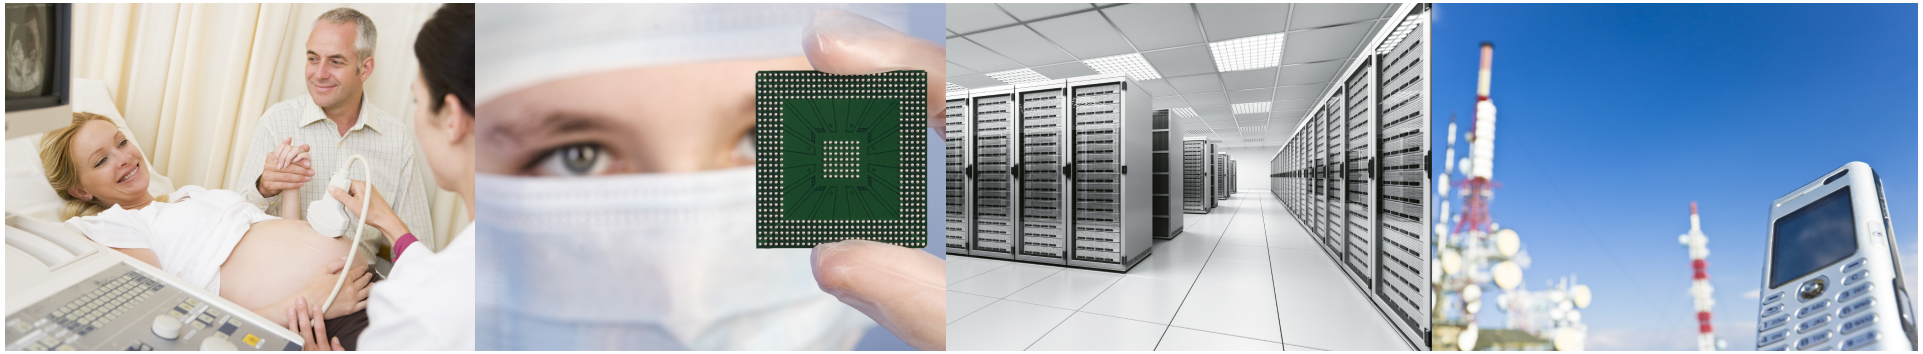

There has to be a better way...

adapteva

Let's Question The Assumptions

- It costs $50-70M to create a semiconductor company
  - The EDA tools are too expensive
  - The development tools are too expensive
  - Tapeouts cost millions of dollars
  - The reference applications cost too much
- A programmable architecture is 100x less efficient than an ASIC
- Floating Point Units are too expensive to include in massively parallel chips
- You can't design a processor architecture without a cache

A Brief History of Signal Processing

	"MAC-DSP" ERA	VLIW/SIMD ERA	FPGA ERA	**"MANYCORE" ERA**
PERIOD	1985-1995	1995-2005	2005-2015	**2015-**
STRENGTHS	ENERGY EFFICIENCY	ENERGY EFFICIENCY & EASE OF USE	ENERGY EFFICIENCY & FLEXIBILITY & SCALABILITY	**ENERGY EFFICIENCY & SCALABILITY & EASE OF USE**
ACHILLES HEEL(s)	PROGRAMMING MODEL	NOT SCALABLE	AREA INEFFICIENT	**??**
NORMALIZED PERFORMANCE (GFLOP/W)	1-2	0.5-2	0.5-5	**50**

The future of DSP is massively multicore and Adapteva have the silicon to prove it

adapteva

The Cost of Programmability

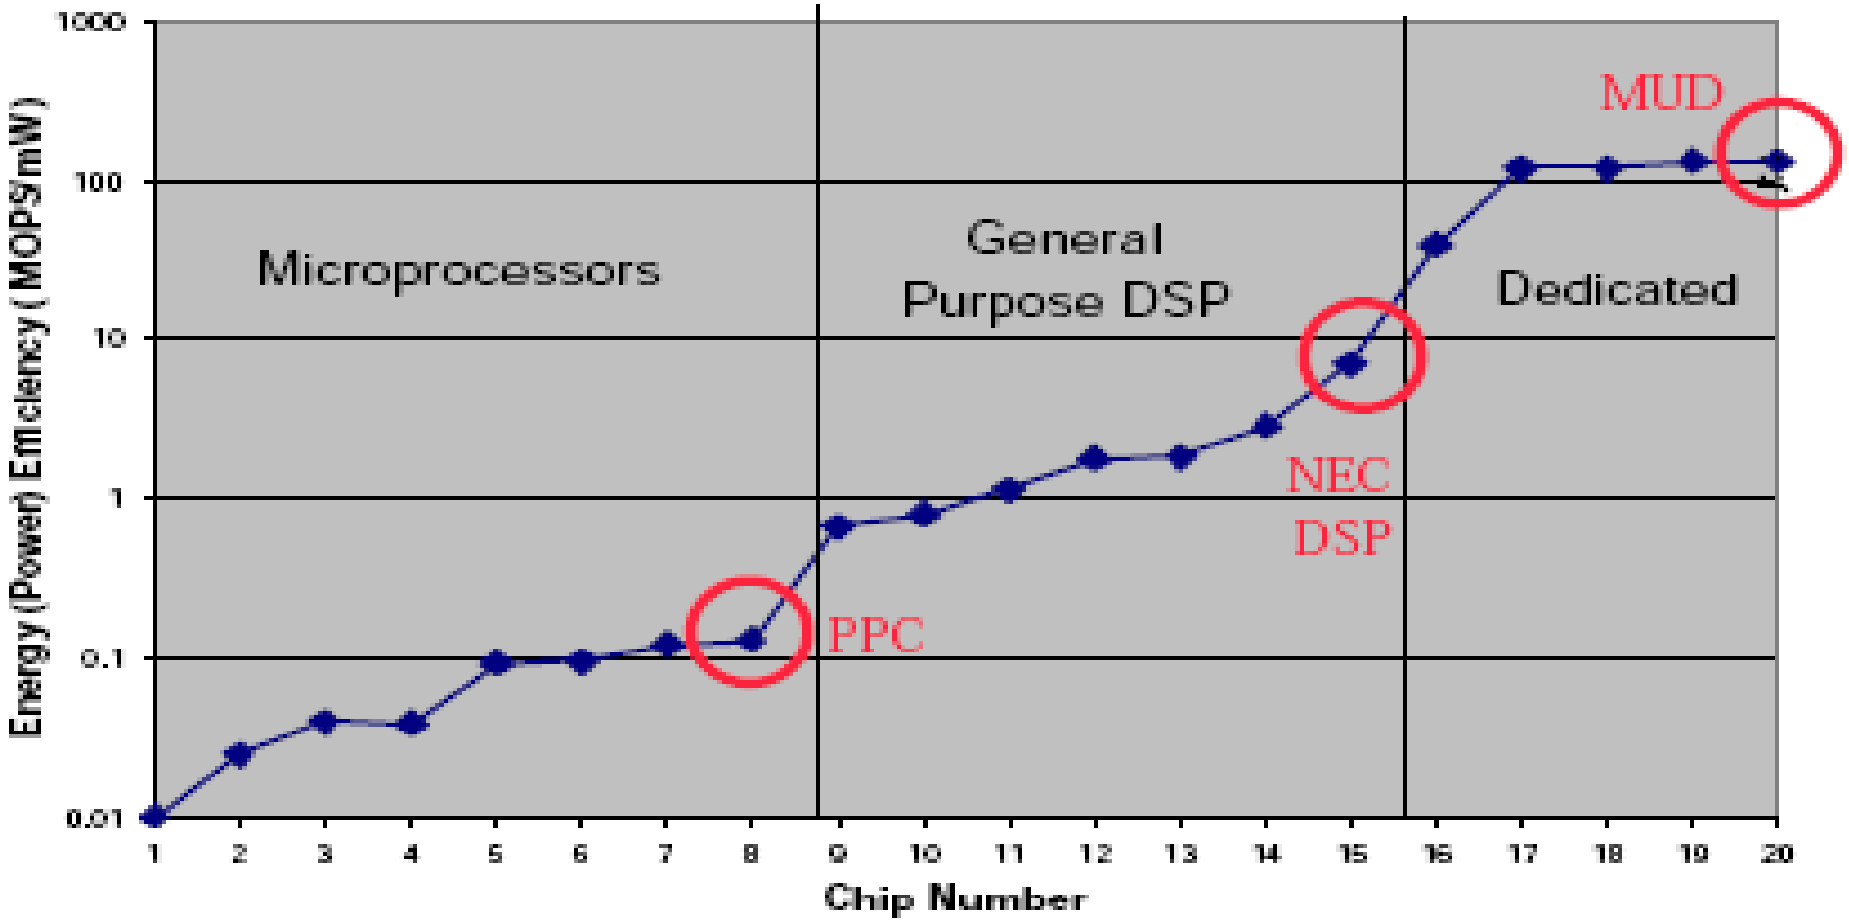

Borrowed from Bob Brodersen's EE225C SOC Design Lecture

adapteva

Area vs Max Performance (65nm)

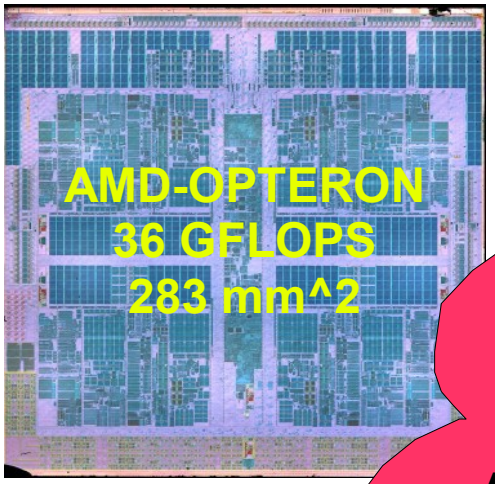

The quest for the holy grail :
An architecture that is flexible,
easy to use, and power efficient.

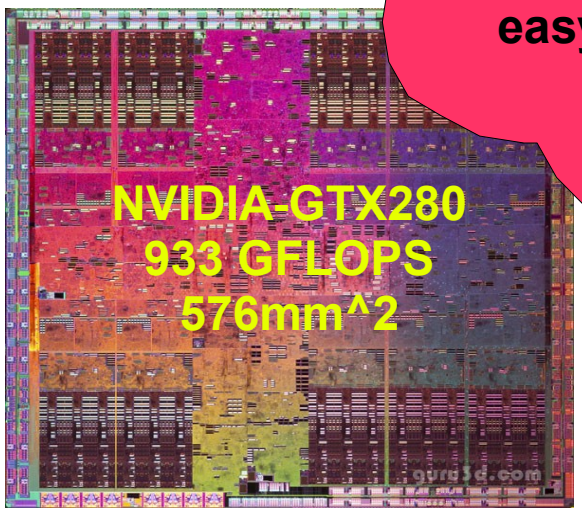

adapteva

Doomsday “Black Silicon” Effect

- FPGAs, GPUs have outgrown Moore's law in the last few years by increasing die sizes to maximum allowed by process node
- This trend has already stopped!
- FPGA's have a killer leakage problem due to configuration SRAM
- At more advanced nodes, leakage power will dominate and maximum die sizes will start to shrink
- This means the free ride is over and we HAVE to make more efficient use of every single transistor (ie reduce silicon area per operation)!

Adapteva derives its power efficiency advantage from its area efficiency. Thus our power advantage will become even more pronounced at lower geometries.

The Low Power Need #2 (HPC/SuperComputers)

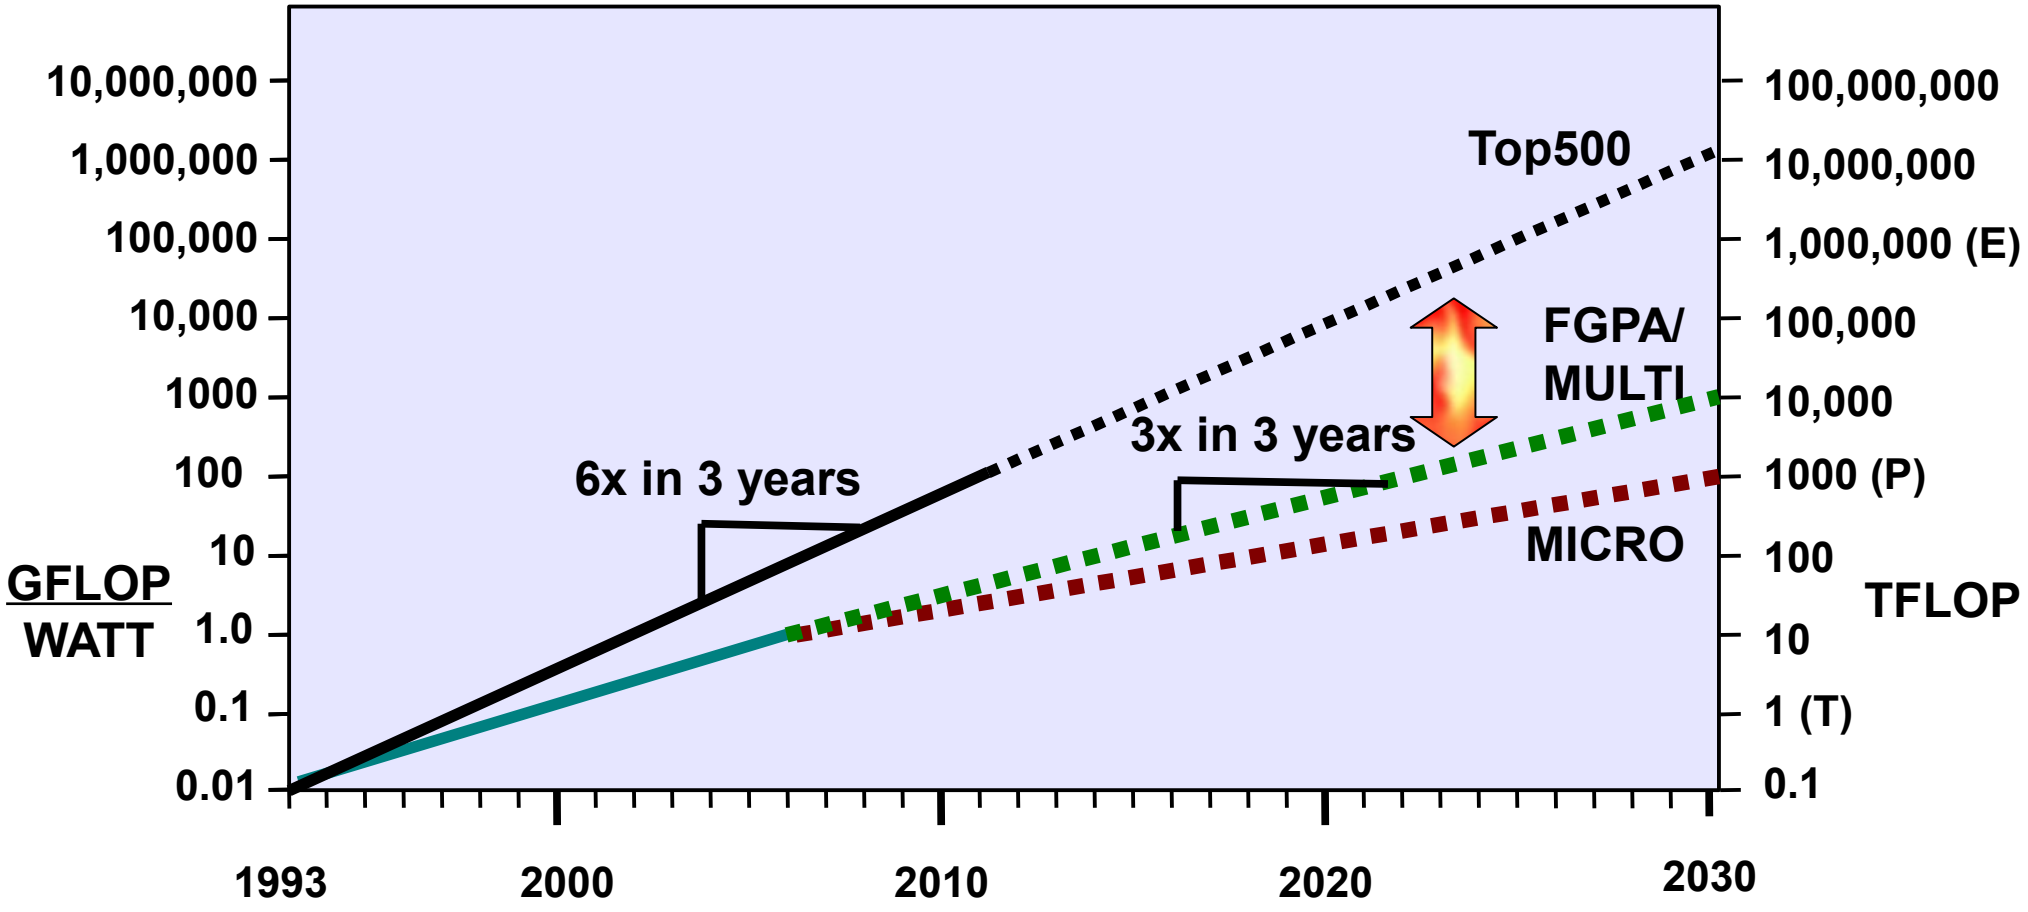

Data from James Anderson at Lincoln Labs and Top500

adapteva

The Epiphany™ Scalable Processor Fabric

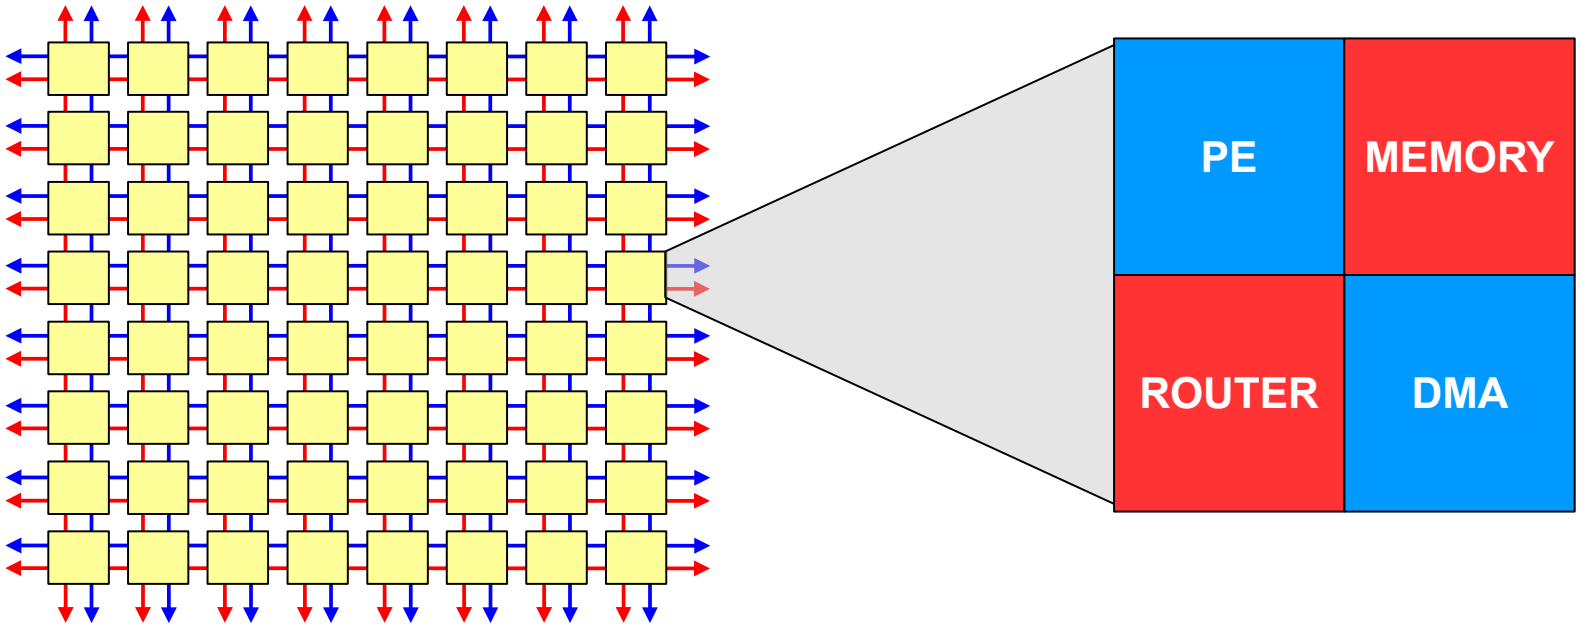

- Shared Memory Architecture
- ANSI C-Programmable
- IEEE Floating Point Processing
- "Infinite" Scalability

- 50 GFLOP/W at 65nm
- 2 GFLOP per Processor Element
- 8GB/sec Inter-core BW
- 100 GFLOP/W at 28nm

adapteva

The Initial Programming Model...

- ANSI-C Programmable with efficient C-compiler
- Runs your existing floating point programs out of the box!
- No special program constructs needed!
- Native single cycle throughput floating point instruction support!
- No SIMD, VLIW, or pragma constructs
- Shared memory architecture
- Orthogonal register and instruction set
- Data-flow programming model

Best suited for "die-hard" embedded system developers!

FFT Multicore Mapping Example

- **Approach:**
  - 1024 point FFT is spread over 16 processors
  - s1,s2,s3,s4 refer to the four FFT stages for combining data with 64 point complex data movements
  - Lower # procs transfer W0 to higher # procs.
  - Lower # proc calculates Wj0+Wj1 x Cj, higher # proc calculates Wj0-Wj1 x Cj
- **Results:**
  - <3us execution time
  - High efficiency
  - Work in progress, still room for improvement

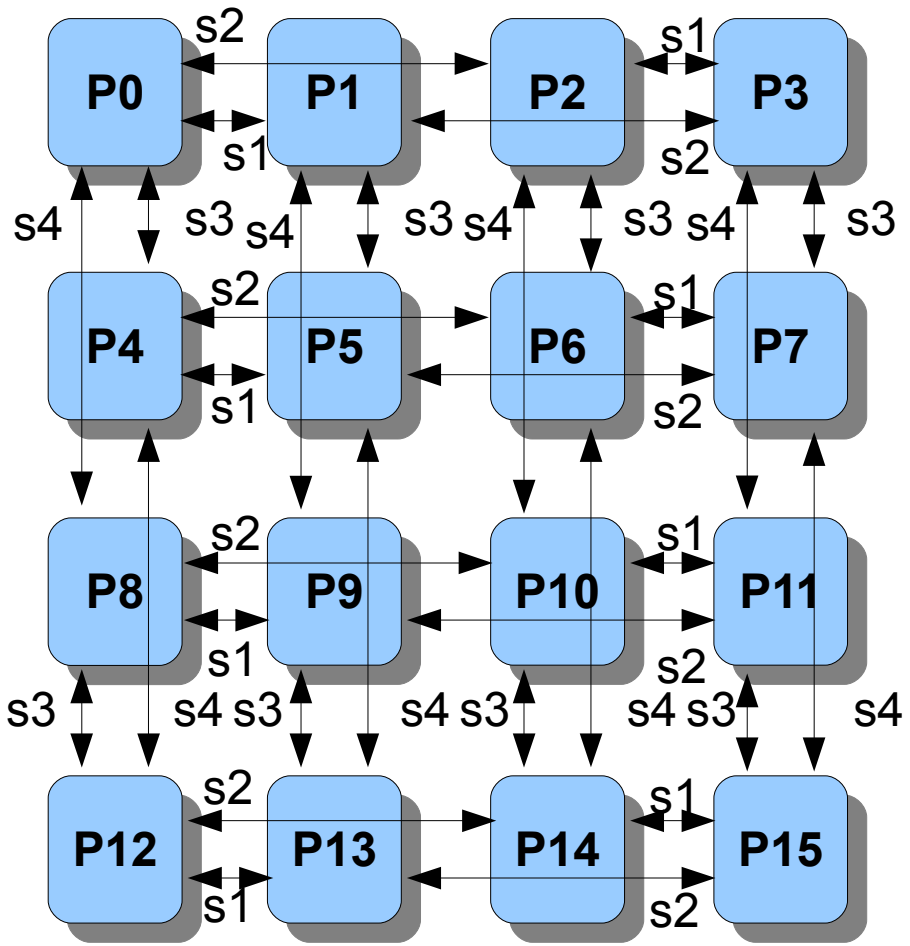

adapteva

Exciting Possibilities in Parallel Programming

- Distribute time critical kernels into sub steps and communicate between cores using high speed on chip network
- Program “sub-kernels” in C and use DMA to move data to next core.
- Allows us to set records in absolute latency computation!!

	On Chip Mesh	On Board Mesh
Total BW	4 Tb / sec	512 Gb / sec
Latency	1 ns / hop	10-1000 ns /hop
Power Efficiency	40 Tb / W	60 Gb / W

What if...

- User wants a single threaded program model..
- The user runs out of local memory..
- There are 1000 cores on a chip..
- The user has to port his legacy C-code..
- The user has to port his Verilog/VHDL code...
- The user doesn't want to learn a new architecture...
- The user has x86 binaries that he wants to port...
- The user wants to run an O/S...

Any suggestions welcome!

Contact Information

General Contact Information

Email: info@adapteva.com
Address: Adapteva Inc.
1666 Massachusetts Ave, Suite 14
Lexington MA 02420
Mobile: +1 (781) 325-6688 (Andreas)
Office: +1 (781) 328-2646