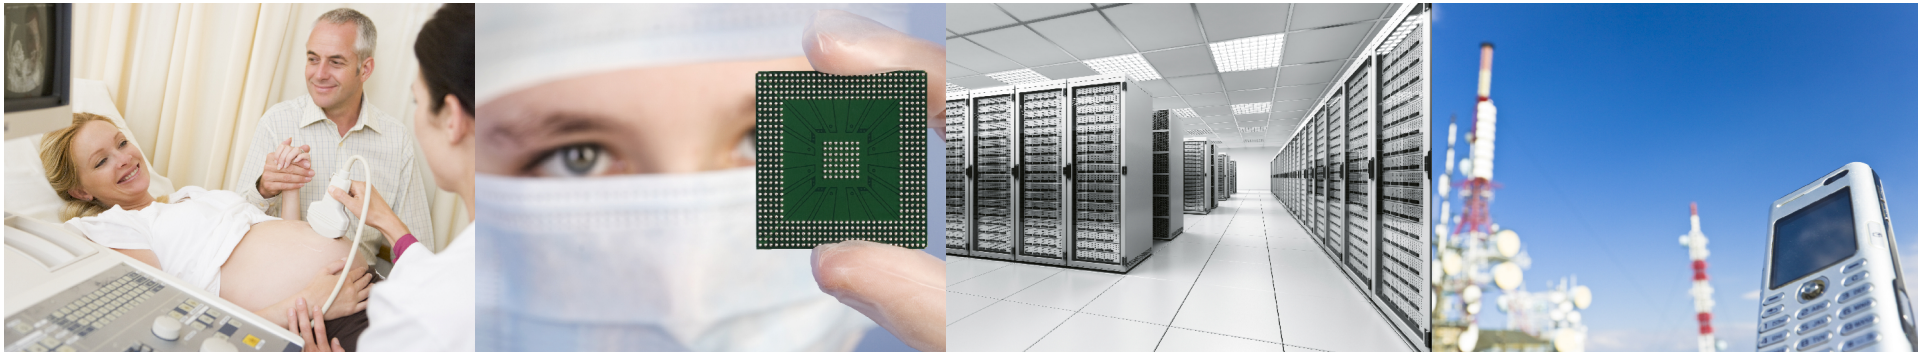

A Scalable Processor Architecture for the Next Generation of Low Power Supercomputer

PRACE Workshop, October 2010

Andreas Olofsson

adapteva

Company Introduction

- Company founded in 2008 with mission to produce programmable processors with 10x the energy efficient of existing products while using 1/10th of the development budget
- Veteran processor design team with 10 successful low power products, 100M units shipped, and over $200M in product related revenue
- Adapteva has completed development of a ground breaking programmable soft accelerator with energy efficiency of up to 50 GFLOPS/Watt
- Now sampling first product to lead customers and strategic partners

Why efficient supercomputing is a critical need!

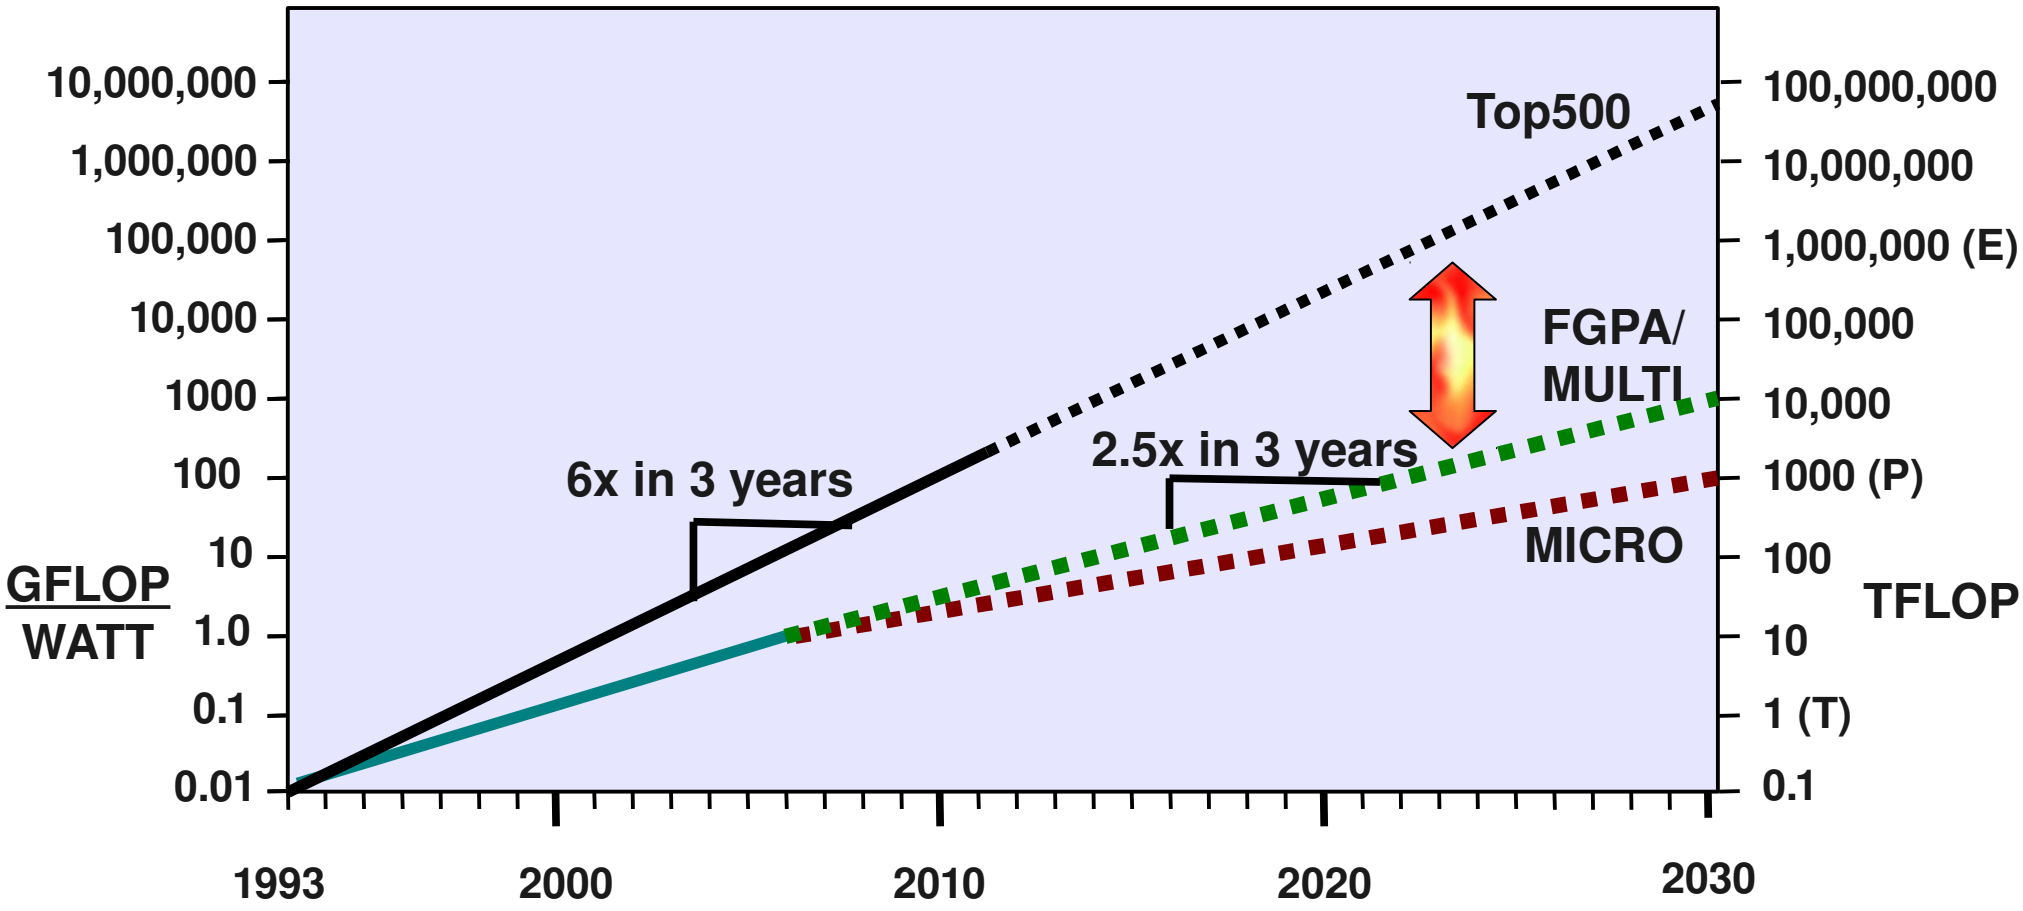

Data from James Anderson at Lincoln Labs and Top500

adapteva

DARPA UHPC Program

System Element	Goals
Cabinet	
Form Factor	Cabinet: width < 24 inches; height < 78 inches; and depth < 40 inches
Energy Efficiency	50 GFLOPS/W LINPACK (HPL)[14] benchmark
Peak Performance	1 PFLOPS (HPL)
Maximum Cabinet Power	57 kW including: UHPC System, storage system, fans, self contained cooling, high bandwidth I/O, etc.
Cooling	Self contained within cabinet. All approaches not requiring external resources are allowable.
Module	
I/O Capability	Support of massive streaming sensor data
Numeric Format – Floating Point	IEEE754 single and double precision, sufficient to support HPC compatible implementations of IEEE754
Numeric Format – Fixed Point	16, 32, and 64-bit supporting all arithmetic and

Big Winners:

INTEL, NVIDIA, MIT, SANDIA

adapteva

Bad News for Semi. Companies

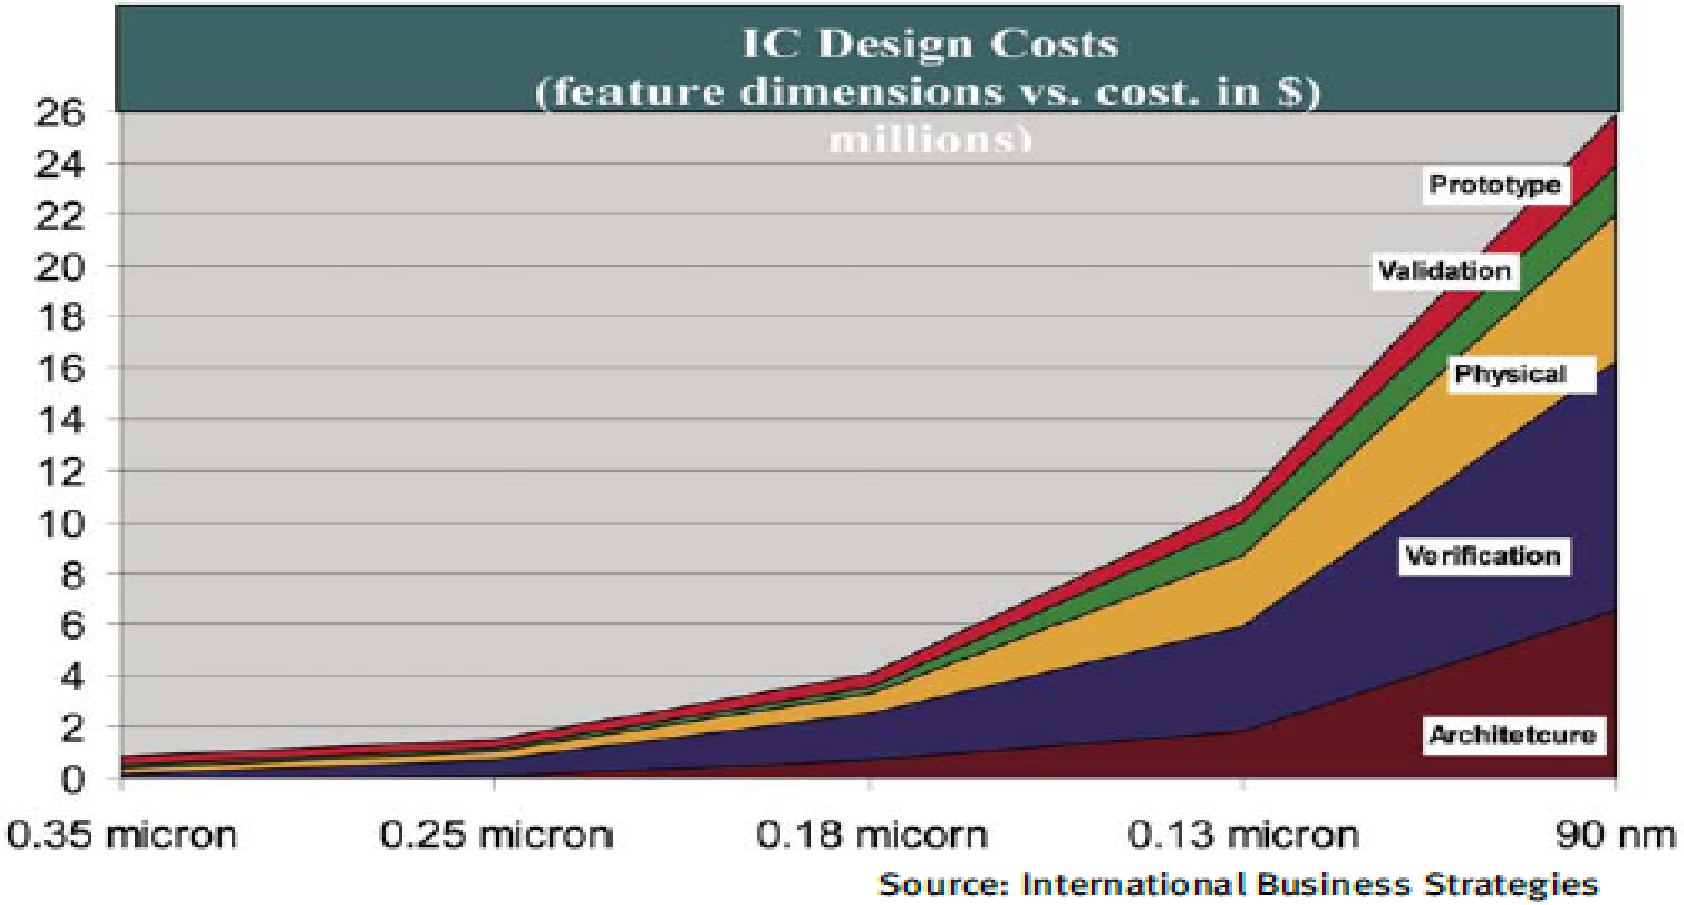

More Bad News..

$2.36 in 10ku

- $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

VS.

$4

- Coffee beans
- Water..

adapteva

Area vs Max Performance (65nm)

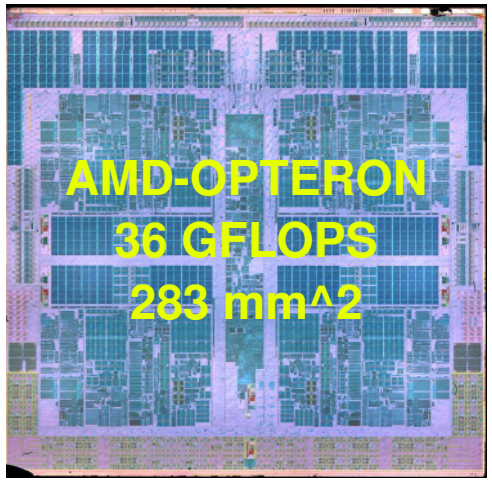

Ease of Use

Efficient CPU

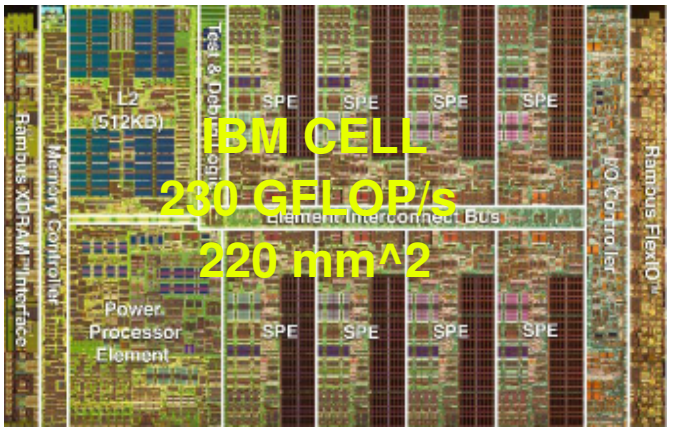

What if we could borrow only the best features?

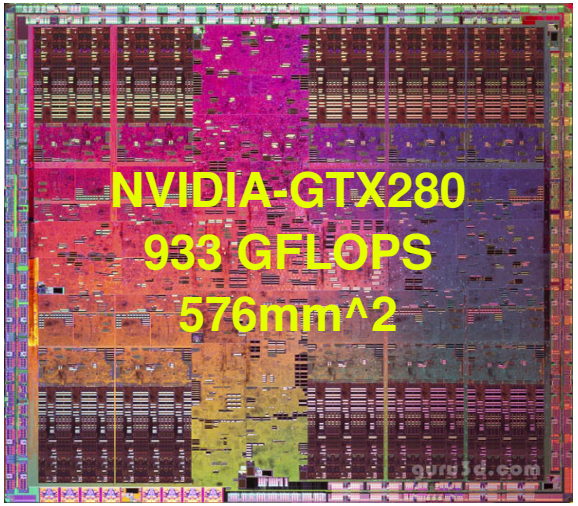

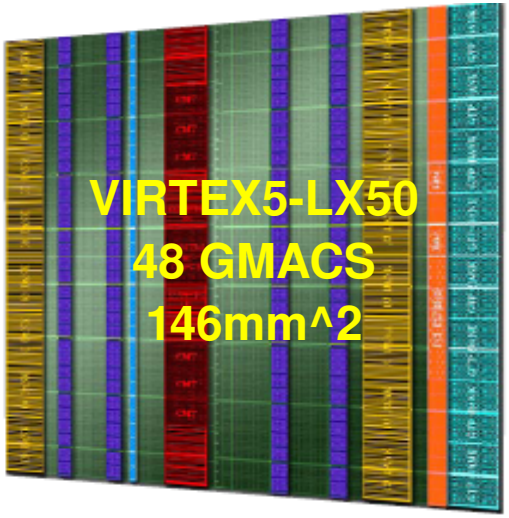

Accelerator

Scale

adapteva

A History of Embedded 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 embedded processing is programmable massively parallel multicore chips!

adapteva

Epiphany™ Signal Processing Fabric

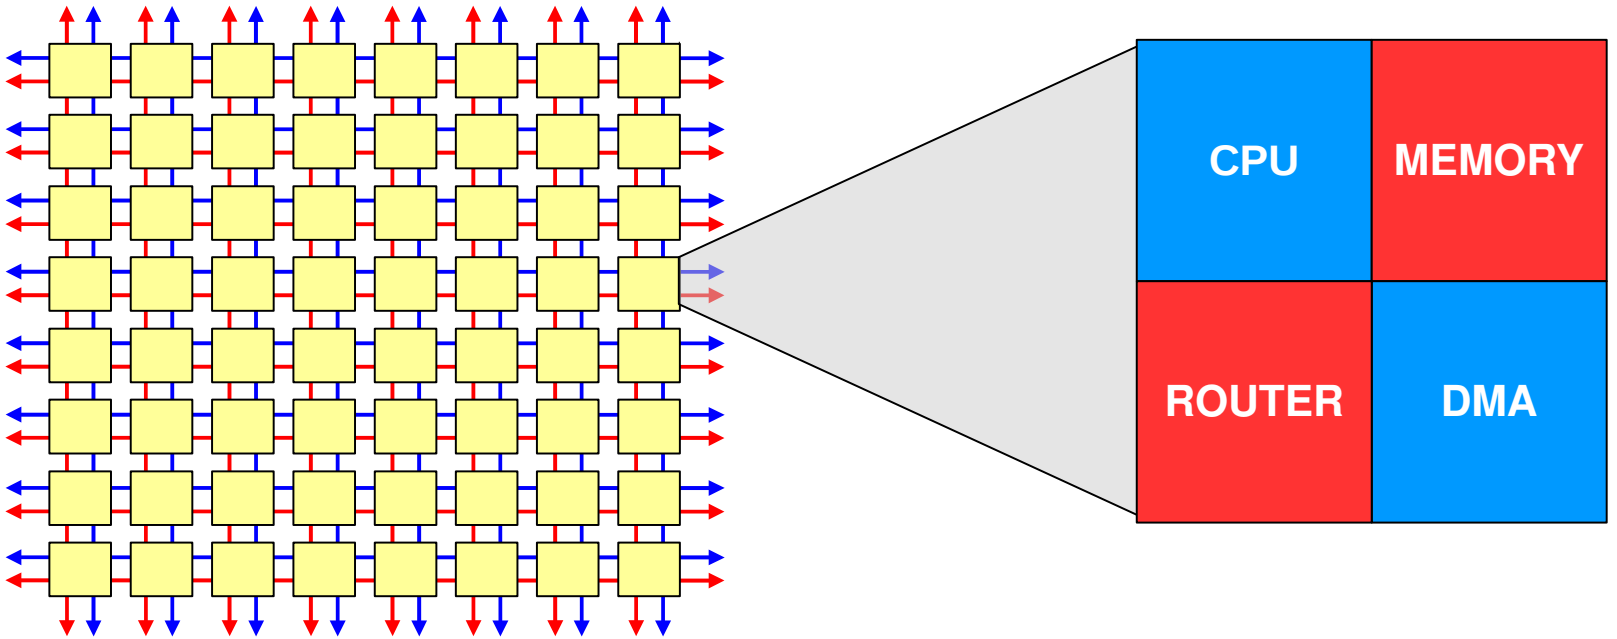

- Easy to use programming model
- Scalable to 1000's of cores
- IEEE Floating Point Compliant

- 25-50 GFLOP/W at 65nm
- 0.7mm per processing tile
- 100 GFLOP/W at 28nm

adapteva

Architecture Summary

Memory Model:

- **HW driven caches kill energy efficiency!**
- **Cache misses very expensive!**
- So..NO HW caches.
- Distributed Flat Shared Memory
- All memory accessible by all cores
- 32KB multi bank SRAM per core

Network- On-Chip:

- **Wires are cheap!**
- 104 bit atomic packets
- One packet per clock cycle
- 1 ns per hop

CPU:

- **Stability is expensive!**
- Optimized for math
- "Good enough" for control code
- 1 GHz operation
- 2 FLOPS/cycle
- 64 entry register file

A Very Simple Flat Memory Model

- Completely unprotected
- Examples of core to core communication:
  - STR R0, [R1,#42]
  - LDR R0, [R1,#42]

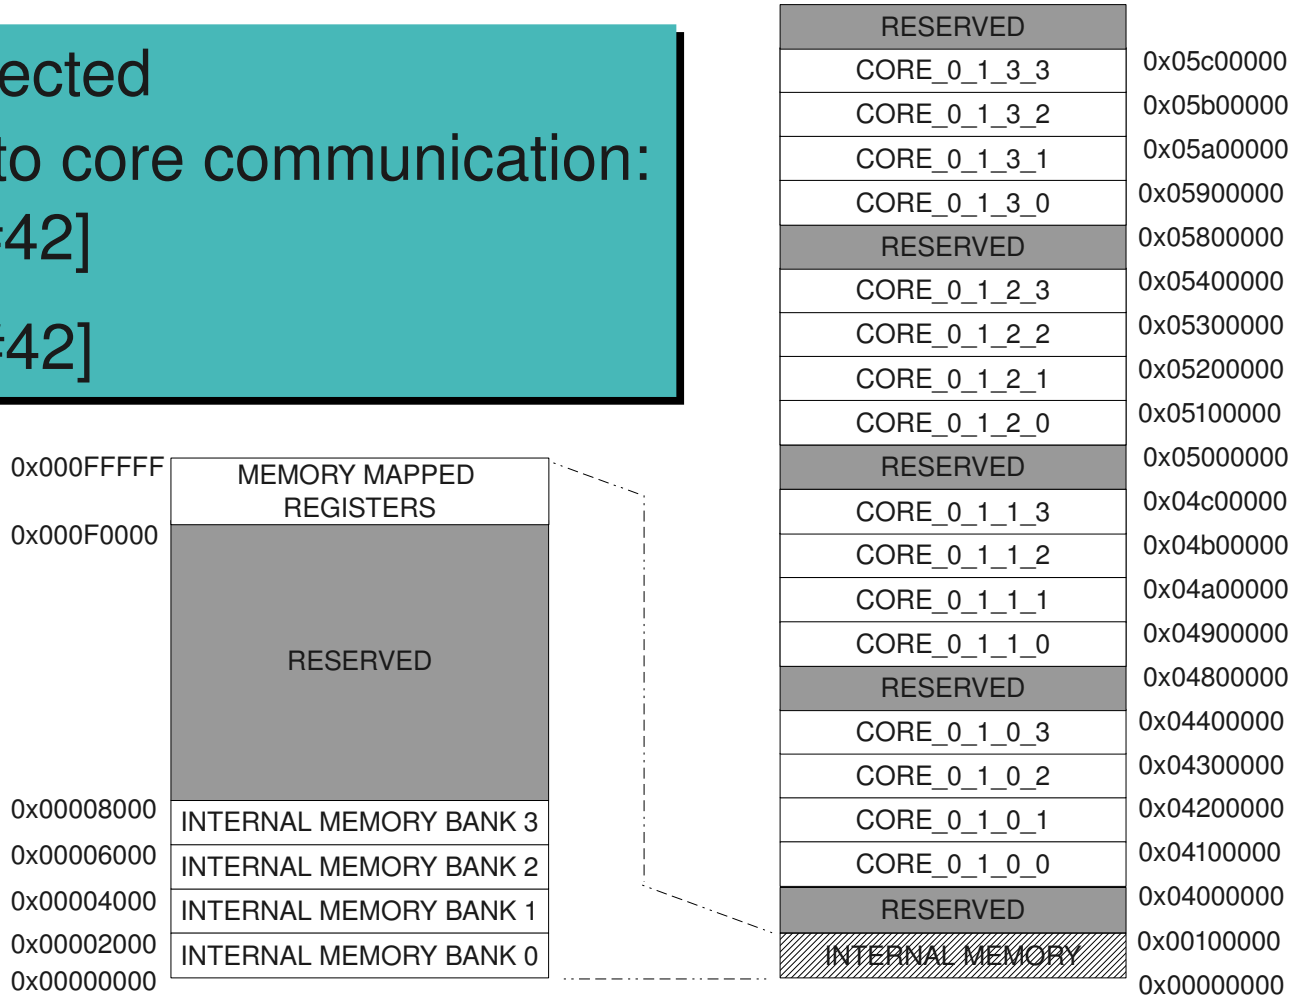

Network On Chip Routing Algorithm

- Each processor node has x,y coordinate
- Destination address compared to mesh node address to determine direction
- Reads are “write requests”
- X-first, then Y

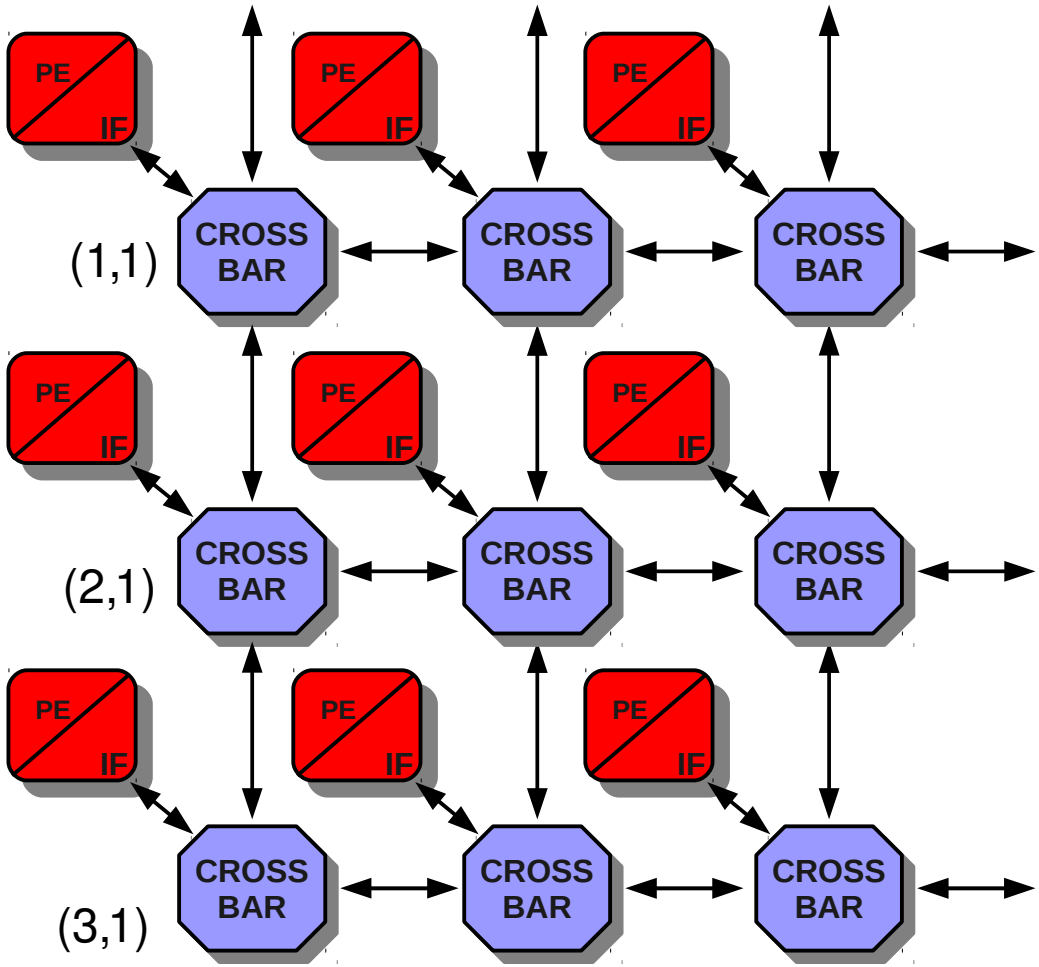

adapteva

Architecture Comparison

	Intel80[2]	Tile64[3]	Intel48[5]	Renesas[4]	(This work)
Process	65nm	90nm	45nm	45nm	65nm
Frequency	3.13GHz	850MHz	2GHz	648MHz	500MHz
Cores	80	64	48	8	16
Area (mm^2)	275	n/a	567	16	12
Transistors (Million)	100	615	1300	n/a	40
Performance (GLFOPS)	1000	No Native Floating Point	n/a	36	16
Power (W)	200	10	125	0.85	0.35
Area/Core(mm)	3.43	n/a	11.81	2	0.72
Watt/ (GHZ*Cores)	0.8	0.18	1.3	0.16	0.04

adapteva

A “Programmer's Architecture”

- “Task-Channel” Programming Model
- ANSI-C Programmable using efficient GNU C-compiler
- Runs floating point C-programs out of the box!
- No special program constructs needed!
- Native single cycle floating point instruction support!
- Shared memory architecture
- Orthogonal register and instruction set

Robust Programming Environment

MULTICORE PROGRAMMING MODEL/LIBS
(IN DEVELOPMENT)

ECLIPSE IDE
(AVAILABLE)

GNU COMPILER
(AVAILABLE)

WYSIWYG SIMULATOR
(AVAILABLE)

GNU DEBUGGER
(AVAILABLE)

GNU BINUTILS
(AVAILABLE)

FFT 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

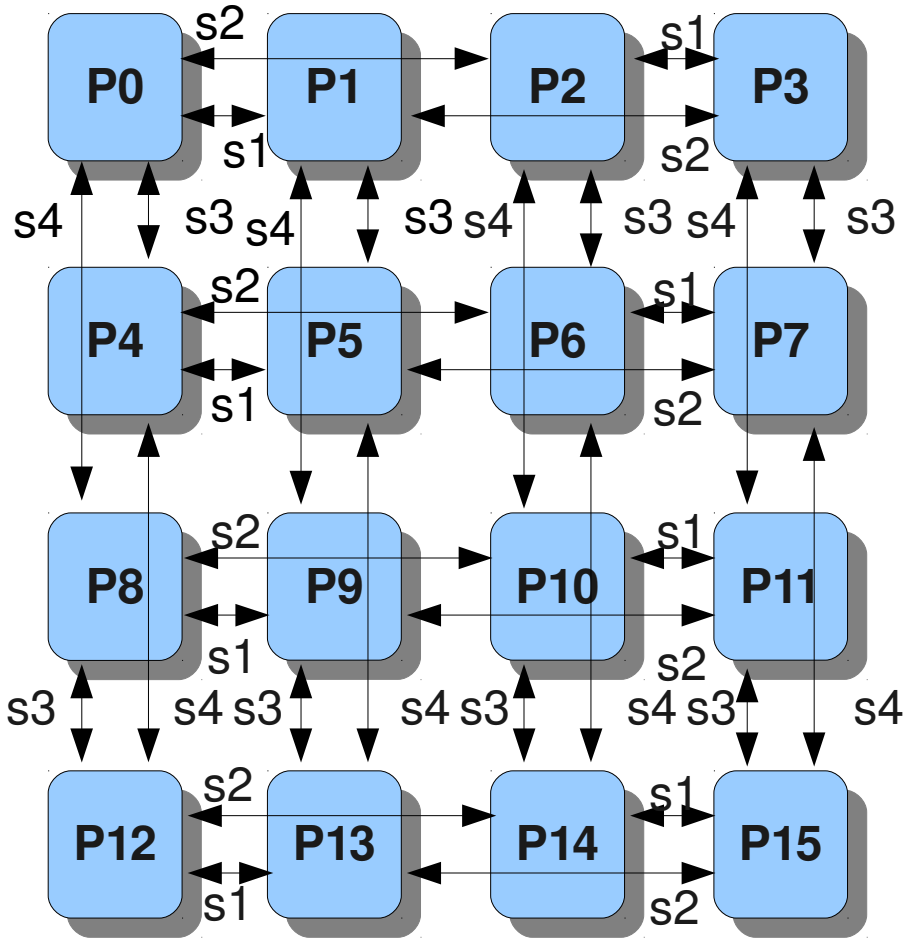

Compromised Computing

What users want:
An easy to use C programmable device with infinite energy, zero power consumption, and all interconnect standards

Soft Floating Point Accelerator

The compromise:
- Microprocessor takes care of the really complex stuff
- FPGA driven connectivity
- A soft accelerator takes care of the “10% code”

FPGA

adapteva

A Radar Signal Chain Example

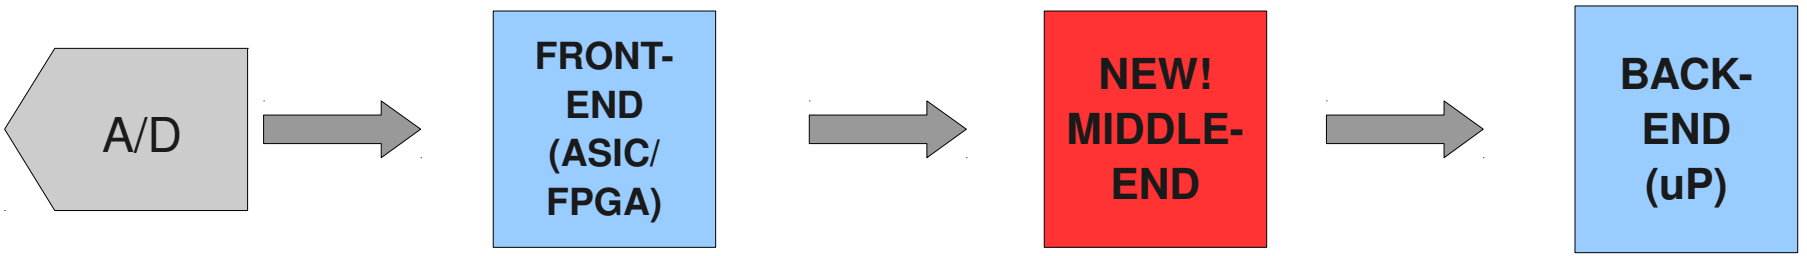

Mult-channel High Data Rate	10,000+ Giga-ops Filterering Fixed Point Low Complexity	1000-10,000 Gigaflop Adaptive Beamforming Pulse Compression Floating Point Medium Complexity	10-100 Gigaflop Clutter Target Detection Large Code High Complexity

- Implementing complex floating point algorithms in FPGA too complex
- Microprocessor based middle end processing far too power hungry
- **The answer is a new device optimized for this type middle-end floating point signal processing!**

adapteva

Arcitecture Application Sweet-spot

What our customers want!

- Floating point
- Massive performance
- Lower power (<3 W) per chip
- C-programmability and ”ease of use”
- Scalability

GREAT FIT APPLICATIONS BASED ON CUSTOMER INTERVIEWS	
Portable Ultrasound	Antenna Pre-distorition
Airborne Radar	Data Center/HPC
Automotive Radar	Soft real time compression
Antenna Beam-forming	

adapteva

E16xx Product Intro (samples available!)

Features:

- 16 Independent Processor Cores
- 4 full duplex FPGA LVDS links
- 4Mb On-chip Shared Memory
- 65nm IBM Process Platform

Performance

- 350mW core power at 500MHz!
- 32 GFLOPS/sec max performance
- 8 GB/sec Off-chip Bandwidth
- 512 GB/sec On-chip Bandwidth

Availability:

- Silicon in lab confirms power advantage
- Currently sampling to lead customers

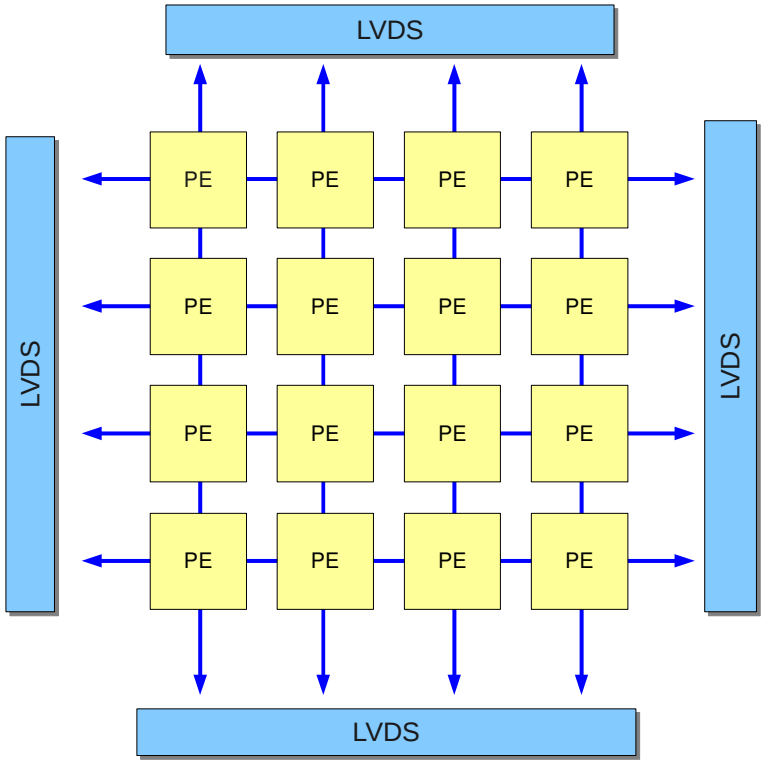

Short Term Product Roadmap

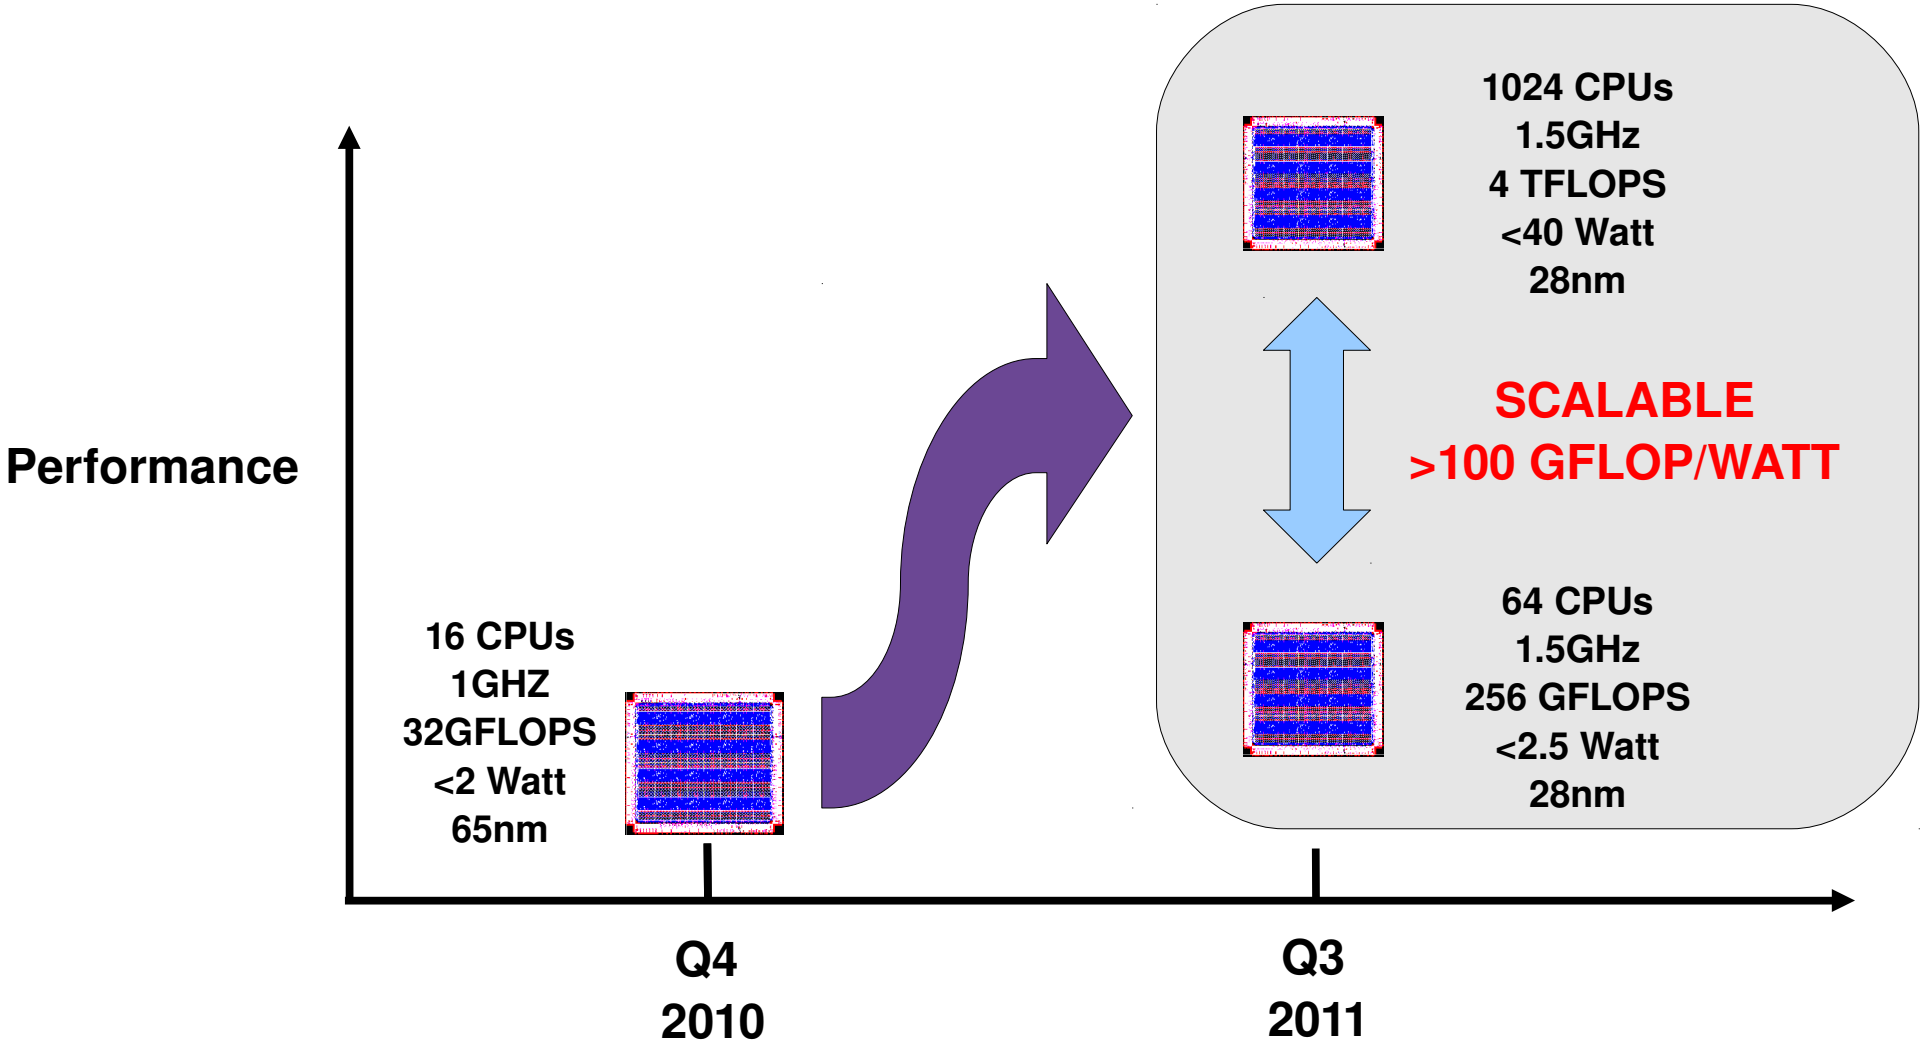

Development ECO System

- **State of the art Tool Chain: (available now!)**
  - Optimized GNU-C compiler
  - Multi-core debugger
  - Standard Eclipse IDE
  - Standard C and Math library
  - Unique guaranteed 100% accurate WYSIWYG chip simulator
- **Evaluation Boards: (available now!)**
  - Plugin boards for standard Altera FPGA development kits
  - FPGA emulation board available for SW/HW co-development
- **Software Libraries: (available November)**
  - Multi-core FFT implementation
  - Multi-core programming libraries

Adapteva HPC Capabilities

- Adapteva, Bittware, and Petapath brings together key expertise in chip, board, and software design enabling fast execution of prototype HPC systems.
- 64-128 GFLOP FMC (VITA 57) accelerator card available 3/15/11
- 1 TFLOP 50-Watt accelerator card in the works for Q2 2011
- Will continue to grow library support base on customer feedback
- Can spin double precision floating point chip with very limited funding

Adapteva Board Offering

- Based on Altera Stratix IV Family
- Attaches to ATCA carriers or other cards equipped with AMC bays and used in MicroTCA system
- Advanced FPGA framework simplifies system design.
- <50 Watts with daughter card
- Base card available today
- Daughter card available 3/15

BittWare SIGNAL PROCESSING SYSTEMS

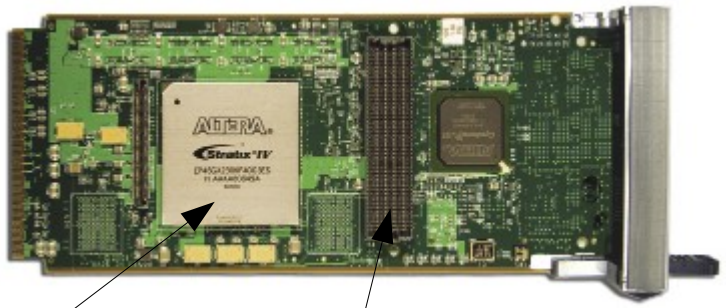

FPGA handles back plane connectivity

FMC Connector enables up to 128 GFLOPS of soft performance with Adapteva chips.

adapteva

A straw-man 1 TFLOP, 50 Watt AMC card

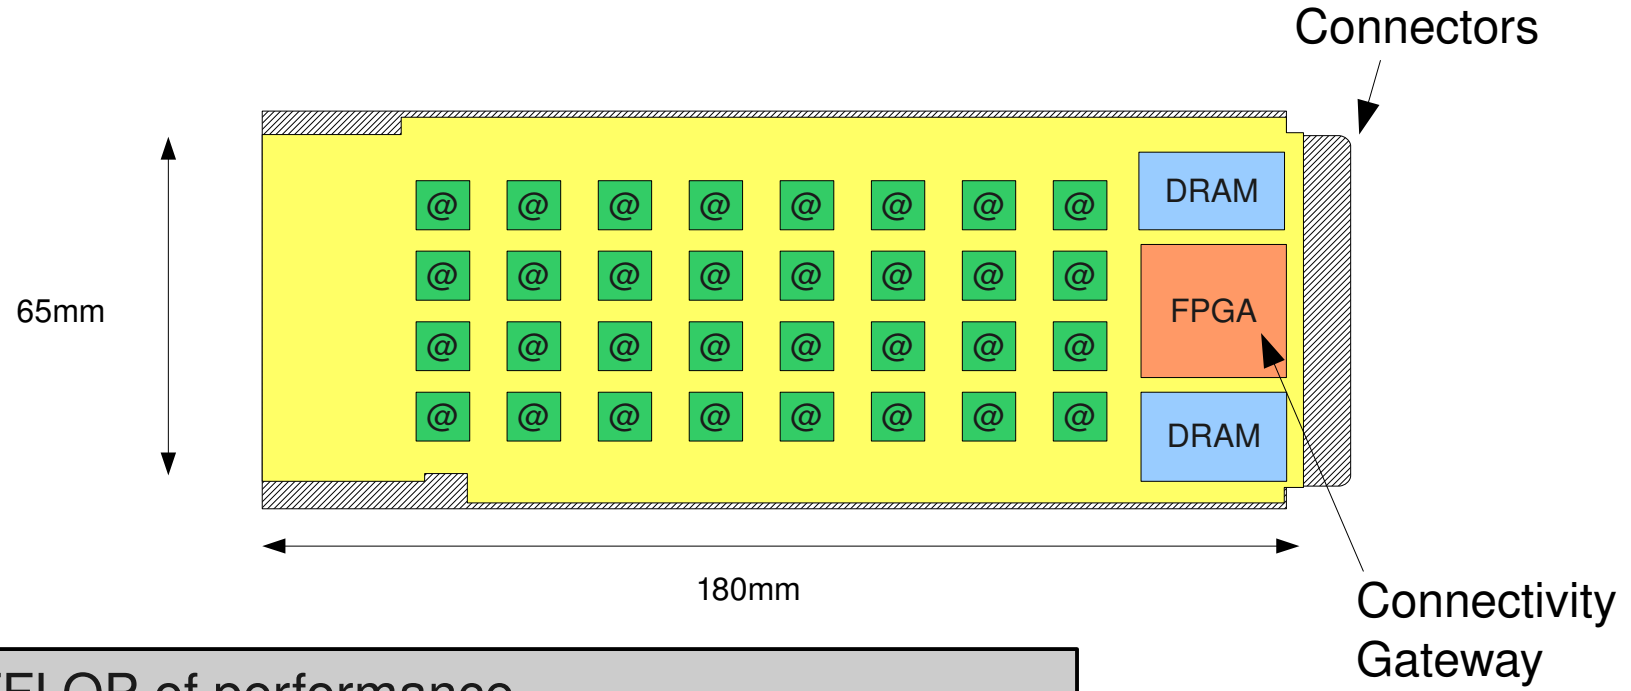

- 1 TFLOP of performance
- 16 MB of distributed memory
- X GB of on-board memory
- Shows off performance density achievable
- Probably not practical for any real applications

Conclusions

- The next generation supercomputers will not be practical without a radically different architecture
- The only way to improve energy efficiency significantly is through compromise
- Adapteva has developed a software programmable floating point accelerator demonstrating 25 GFLOPS/W in silicon and 100 GFLOPs/W on the road map for 2011
- The programming model needs to change..but not drastically

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