Inventing the Future of Computing

An Alternative to GPU Acceleration For Mobile Platforms

Andreas Olofsson
andreas@adapteva.com

50th DAC
June 5th, Austin, TX

adapteva

Adapteva Achieves 3 "World Firsts"

1. First commercial chip to reach 50 GFLOPS/W

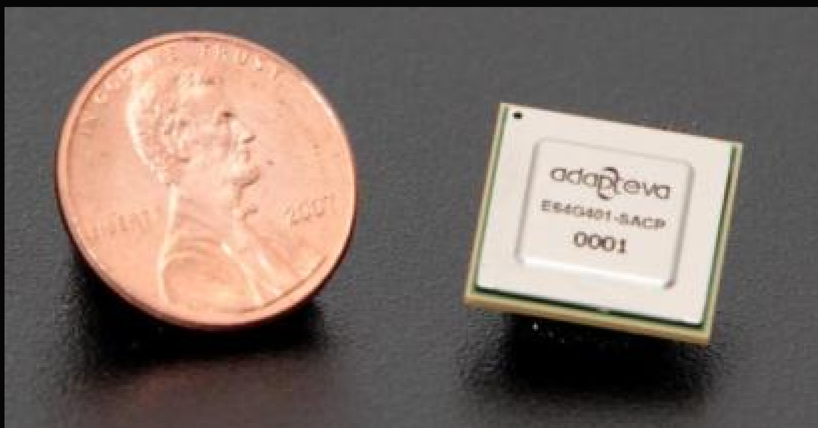

2. First mobile processor with an open source OpenCL™ SDK

3. First semiconductor company to successfully crowd-source project

KICKSTARTER

adapteva

What is Adapteva

Company History:

- Fabless semiconductor company founded in 2008
- 16-core 65nm Epiphany-III chip product sampling since May 2011
- 64-core 28nm Epiphany-IV chip product sampling since July 2012
- Parallella open computing platform launched in October 2012

Notable Achievements:

- #1 in microprocessor energy efficiency
- 4 chips on $2.5M in raised capital
- $2M in total revenue to date
- 5K customers, 6,300 boards pre-sold
- 18 Patents pending

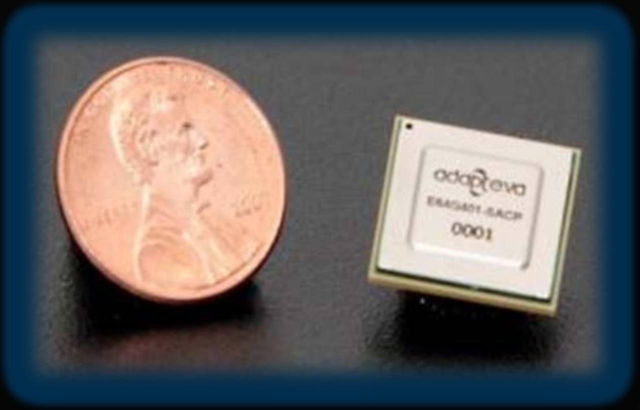

Our guiding light

Parallel

Efficient

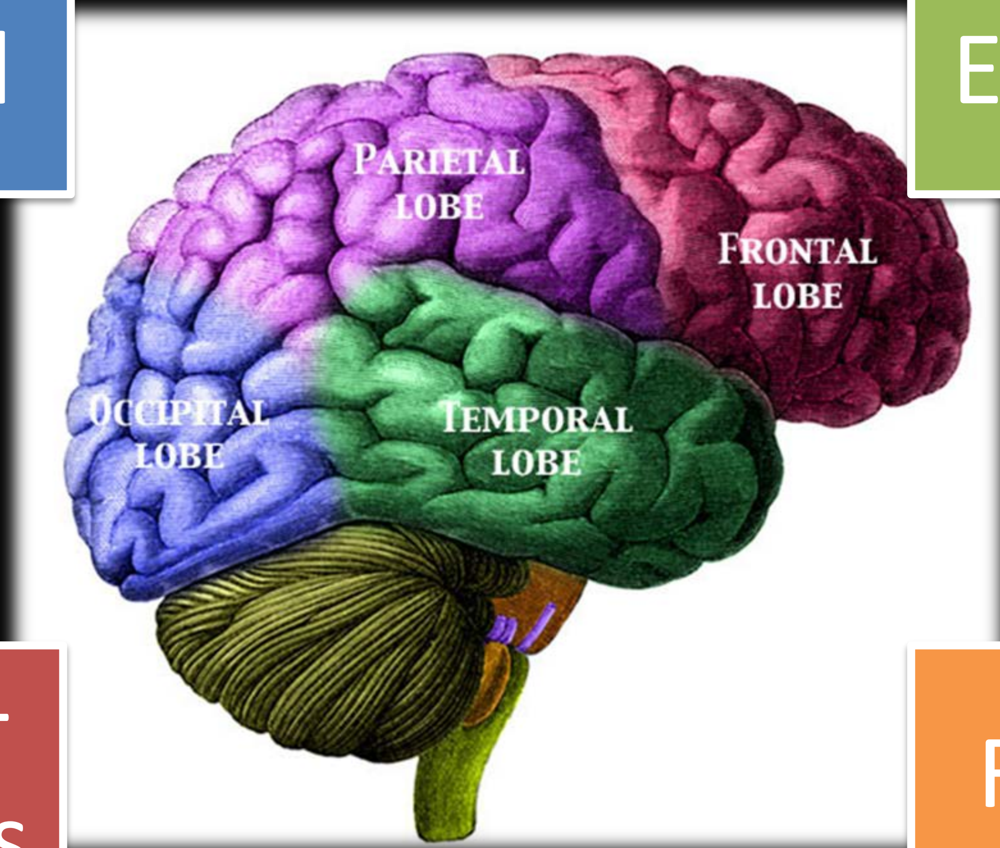

Hetero-
geneous

Robust

Any Reason to Think the Future of Computing is NOT Parallel?

No Computing | Parallel Computing

No Electronic Computing
-1943

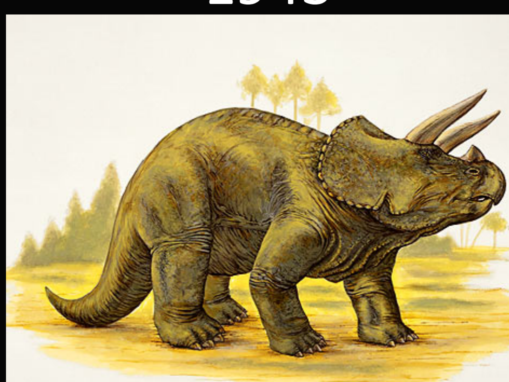

"Von Neumann Age"
Serial Computing
1943-2013?

Parallel Computing
2013-??

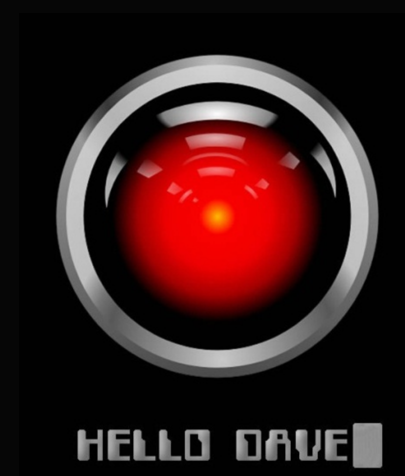

A Practical Start:

True Heterogeneous Computing

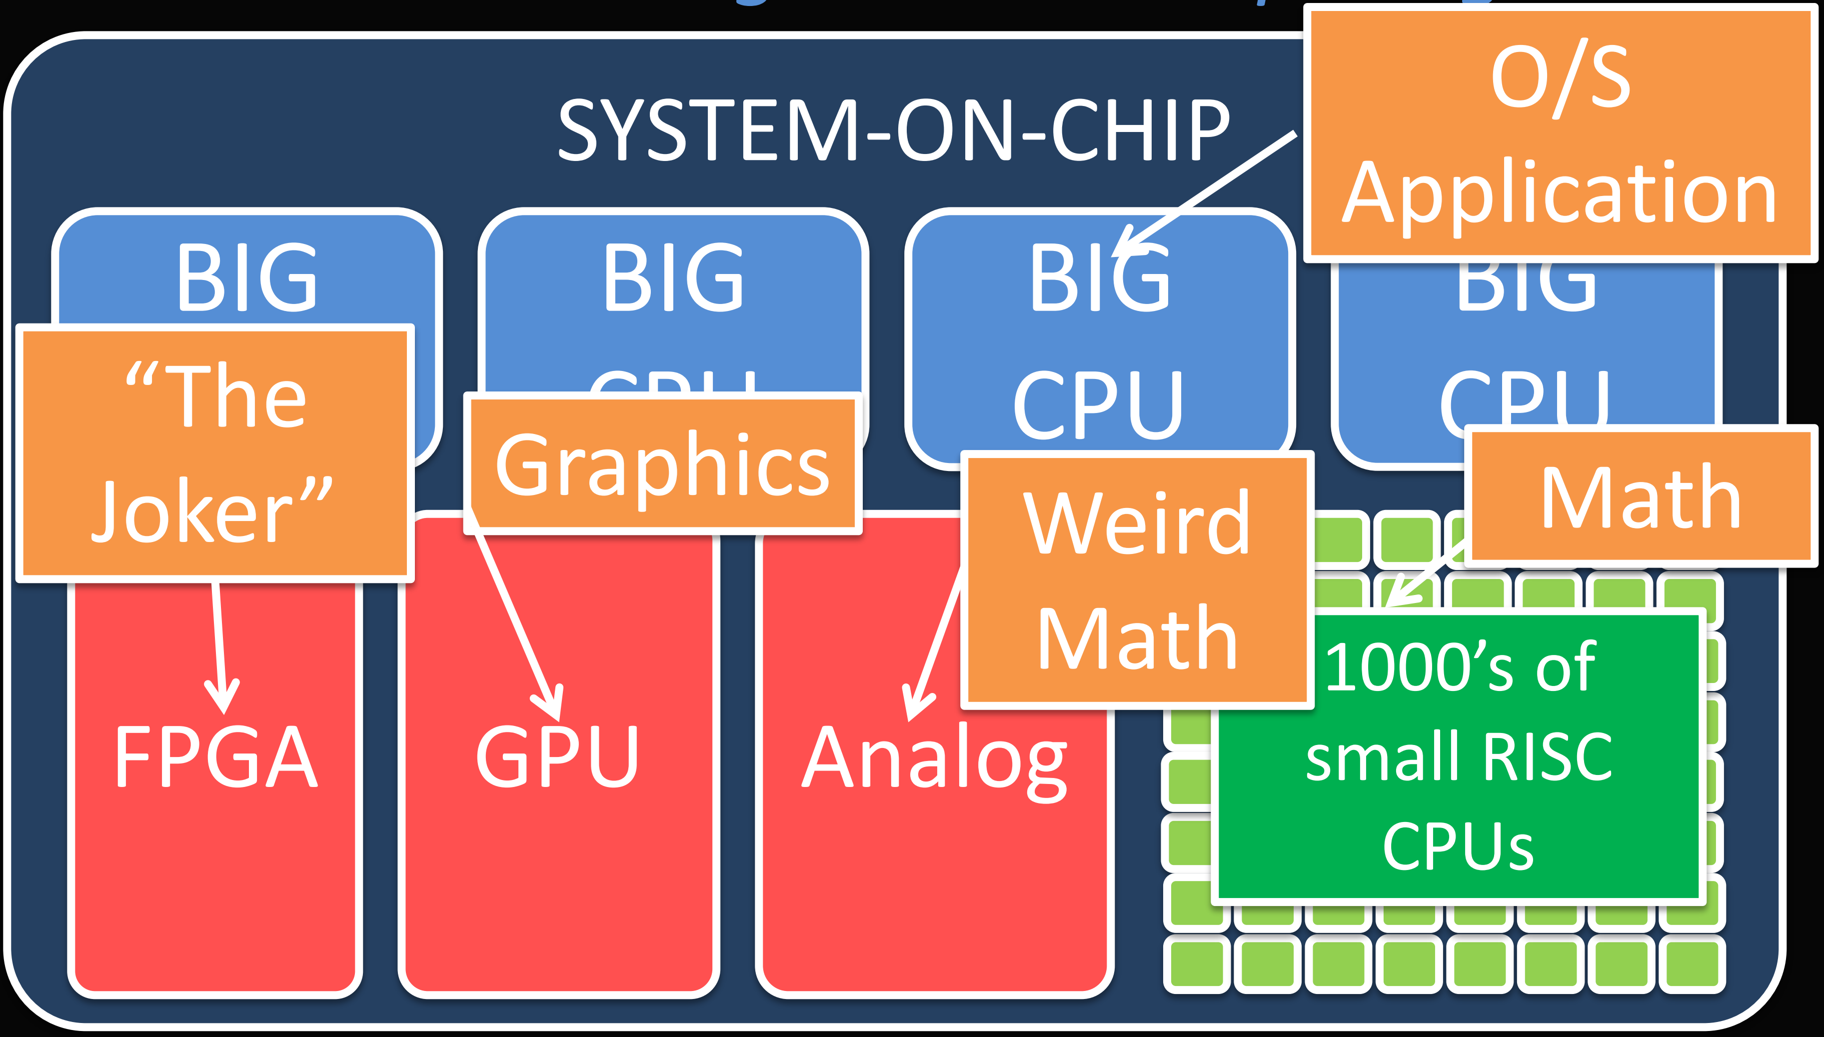

The Accelerator Challenge

Status Quo Approach (~1.3X speedup)

Smart Coprocessor (>10X speedup?)

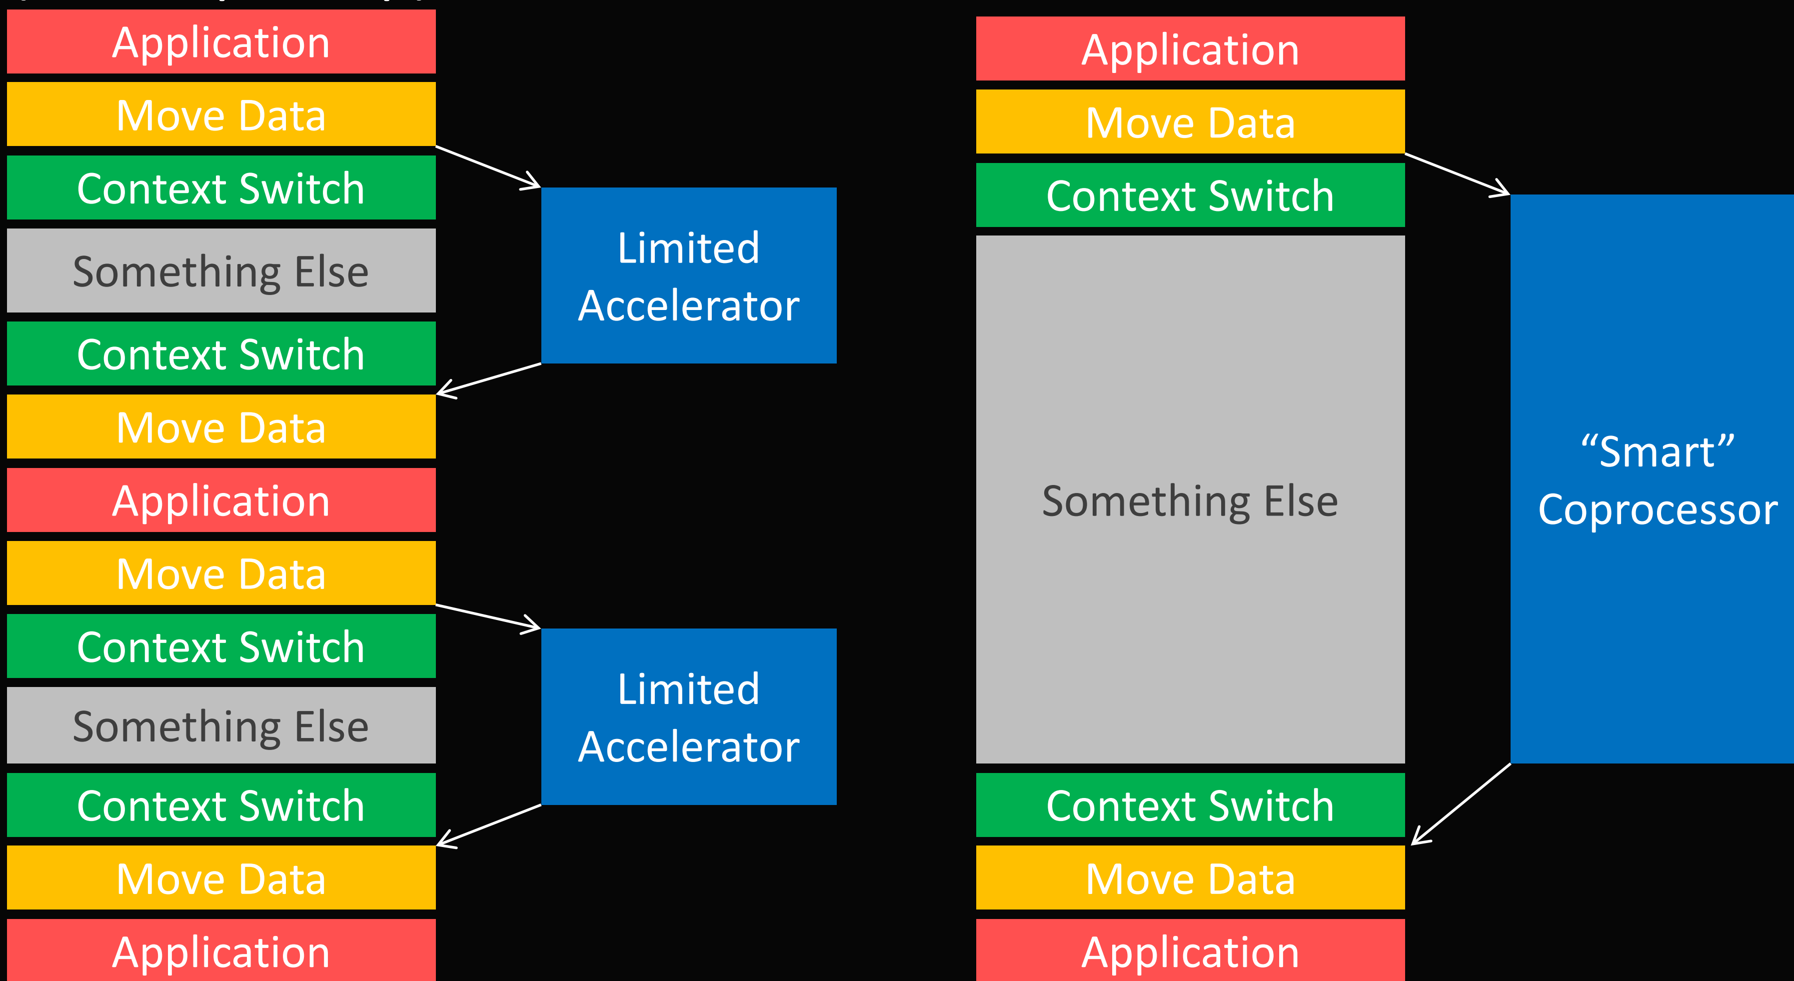

The Epiphany Coprocessor

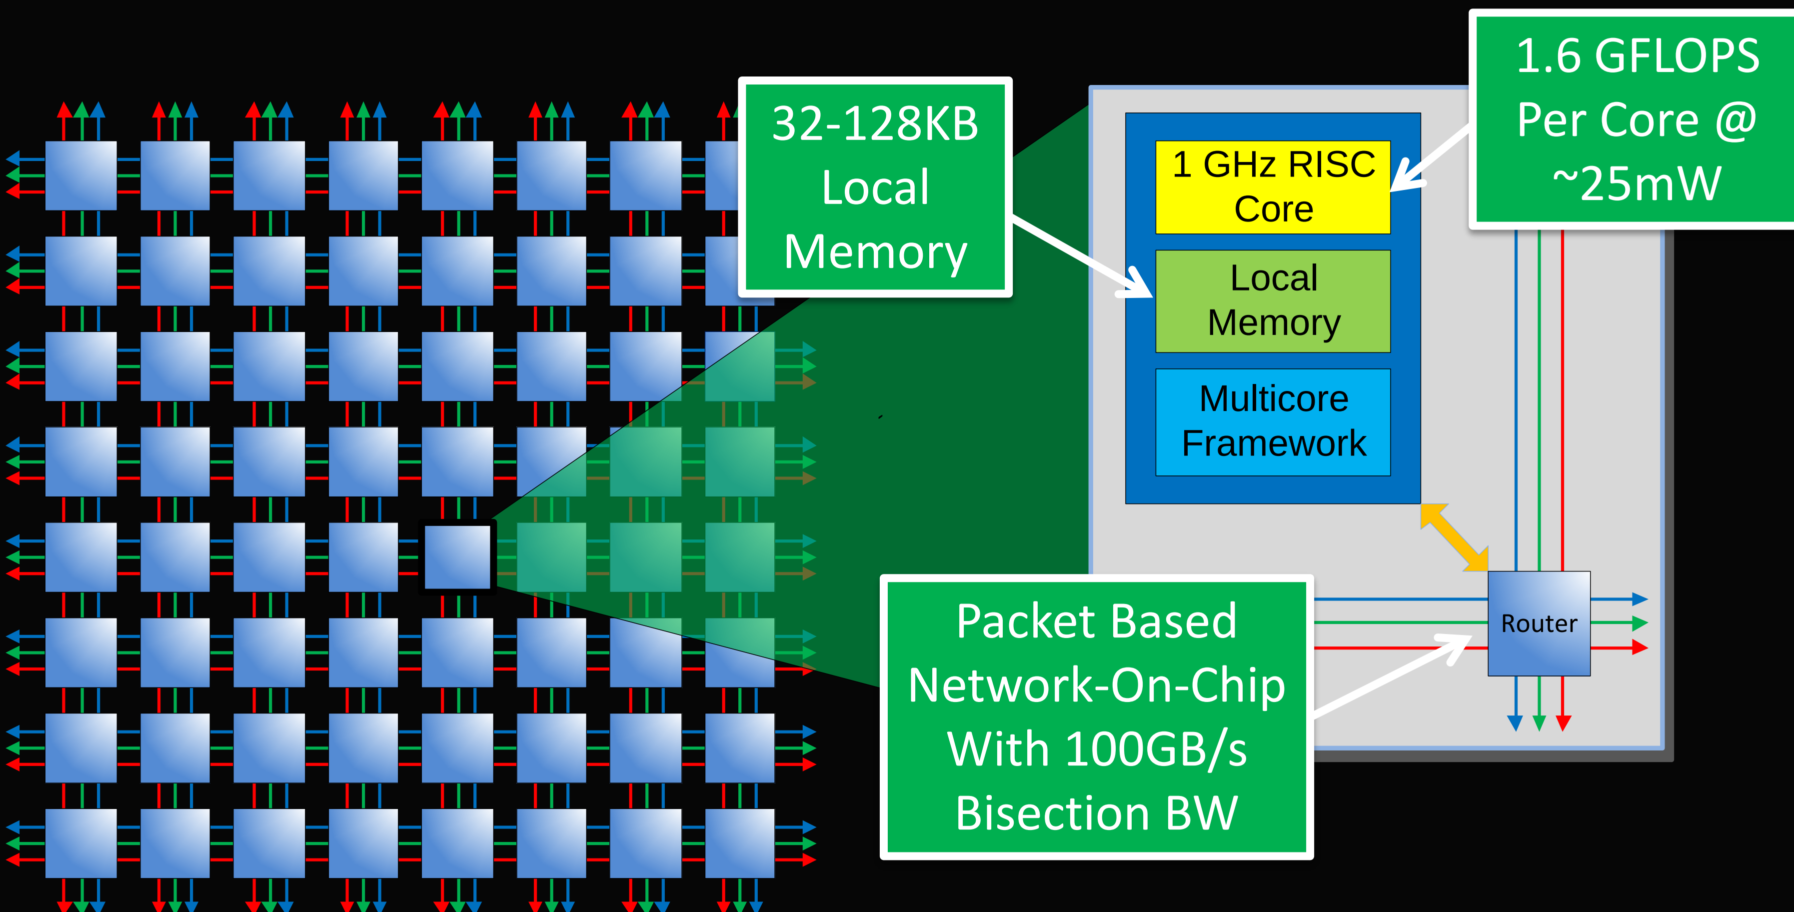

Coprocessor for ARM/x86 Host

<20pJ / FLOP !

MIMD/Task-Parallel Accelerator

Epiphany-IV -- GLOBALFOUNDRIES 28SLP IP

- 64 CPUs
- IEEE Floating Point (SP)
- 800 MHz Max Frequency
- 100 GFLOPS Performance
- 6.4 GB/s IO BW
- 200 GB/s peak NOC BW
- 1.6 TB/sec on chip memory BW
- 25 Billion Messages/sec
- 2MB on chip memory
- 10 mm^2 total silicon area in 28nm
- 2 Watt total chip power
- 324 ball 15x15mm BGA
- Sampling since July, 2012

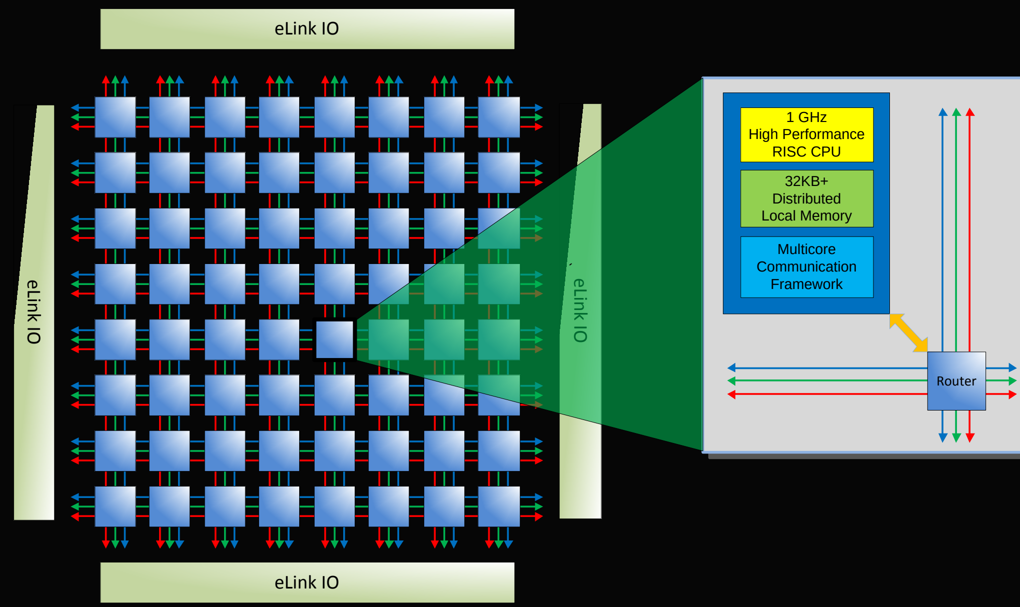

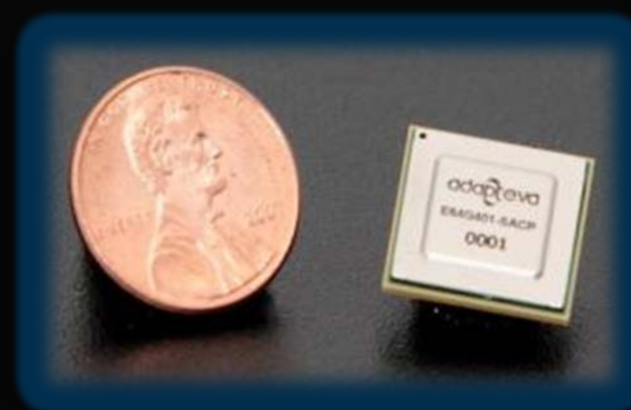

adapteva

Epiphany ANSI-C Benchmarks

(Cycles)	Naïve C	Optimal C	Theoretical	C-Efficiency
8x8 Matrix Multiplication	2852	773	512	66%
16 Tap FIR Filter (32 points)	1562	620	512	82%
Bi-quad IRR Filter (32 points)	n/a	991	768	77%
Dot-product (256 point)	800	557	256	49%

1 day per benchmark

(compare to GPUs?)

	Adapteva E64 800 MHz	Tilera GX36 1.4GHz	Intel Xeon L5640 2.2GHz	Nvidia Tegra-2 1GHz
CoreMark TM Score	**77,912**	165,276	118,571	5,866
# Cores	**64**	36	8	2
Power	**2W**	**~30-50W**	**~50-100W**	**~1-2W**
1024-Core Chip	**2,493,184**	n/a	n/a	n/a

Server Level Performance at 2Watts!!

adapteva

Architecture Comparison

Technology	FPGA	DSP	GPU	CPU	Epiphany
Process	28nm	40nm	28nm	32nm	28nm
Programming	VHDL	OCL/C++/C	CUDA/OCL	OCL/C/C++	OCL/C/C++
Area (mm^2)	**590**	**108**	**294**	**216**	**10**
Chip Power (W)	40	22	135	130	2
"CPUs"	n/a	8	32	4	64
Max GFLOPS	**1500**	**160**	**3000**	**115**	**102**
GHz * Cores	n/a	12	32	14.4	51.2
Compile Time	Hours	Minutes	Minutes	Minutes	Minutes
L1 Memory	**6MB**	**512KB**	**2.5MB**	**256KB**	**2MB**

Efficiency is everything

Peak performance means very little

No magic bullet!

adapteva

Epiphany: A Truly Scalable Architecture

A Single Unified Instruction Set Architecture!

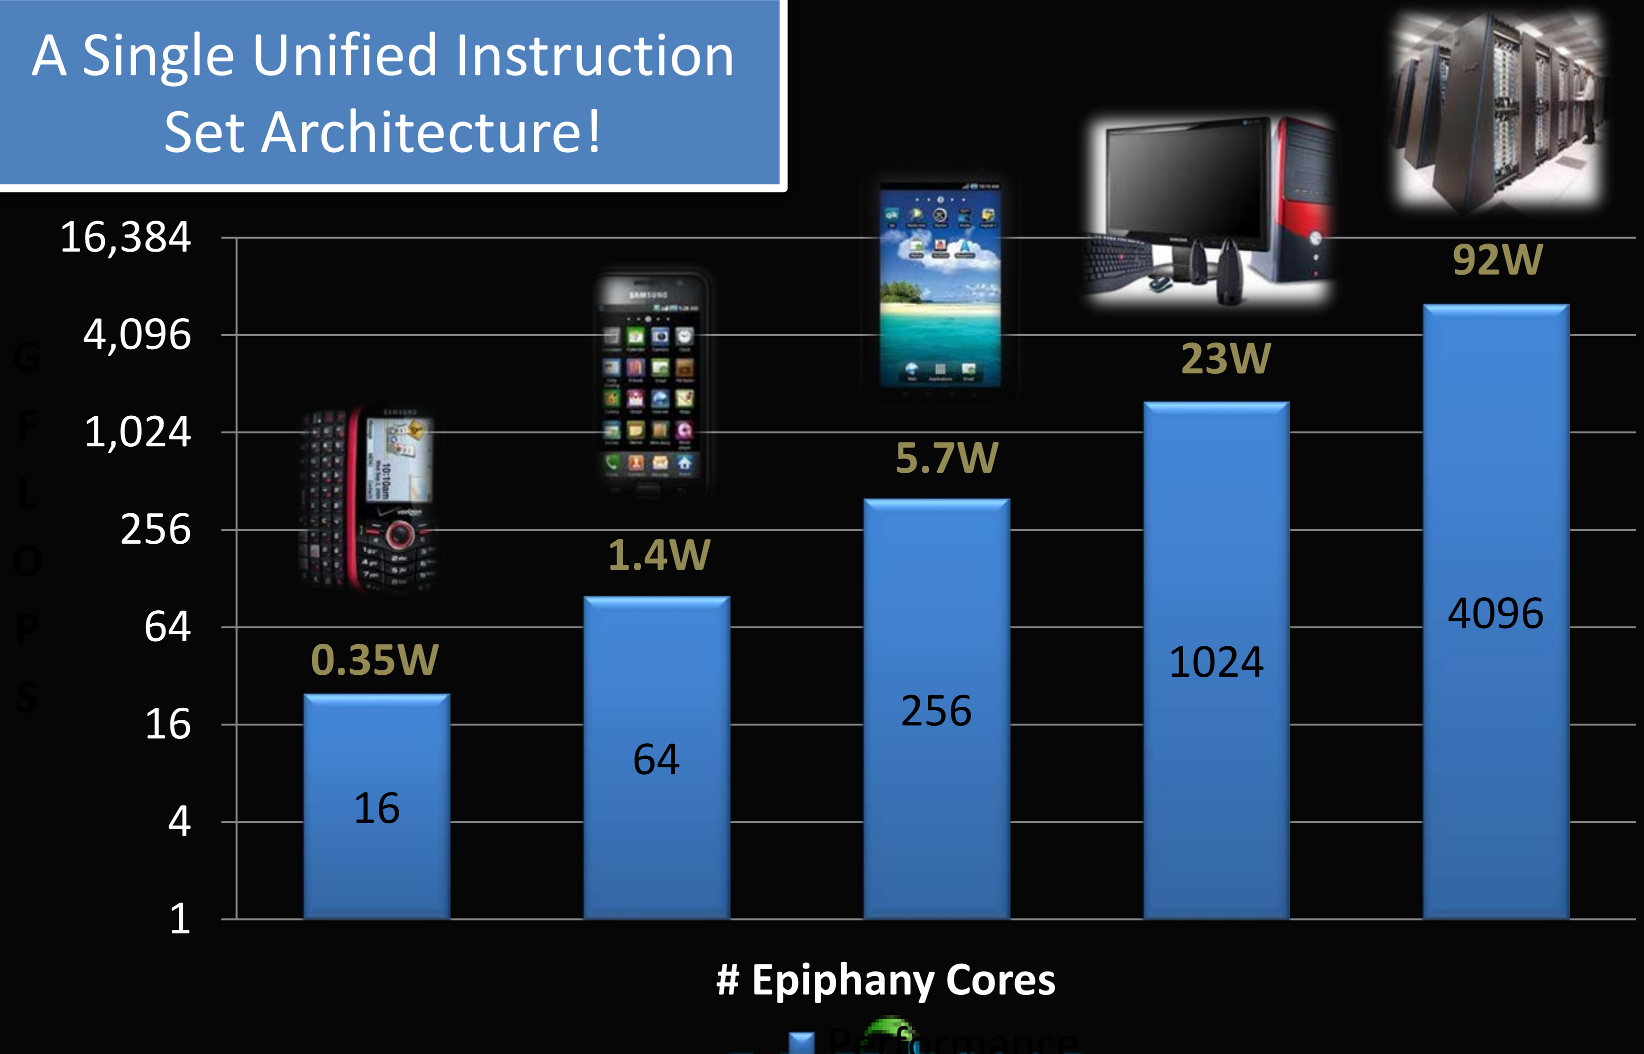

How the $#@% Do We Program This Thing?

Epiphany Programming Models

MODEL #1
DATA PARALLEL MODEL
- openCL programmable
- Easy integration with C/C++
- openMP/MPI roadmap

MODEL #2
WORKER BEE MODEL
- Great for up to 2GFLOPS
- Supports standard C/C++
- "Cloud on a chip"

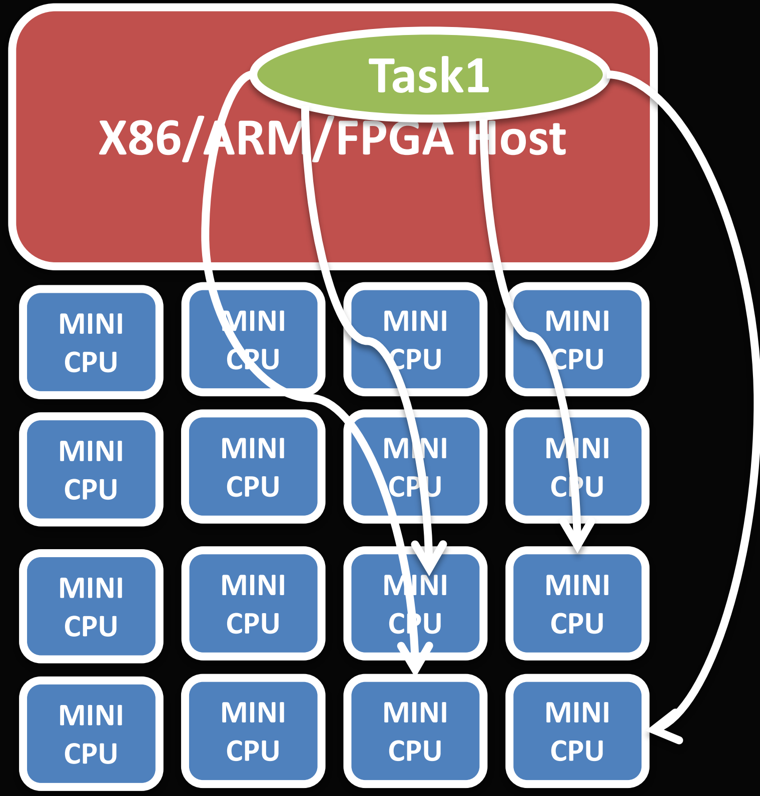

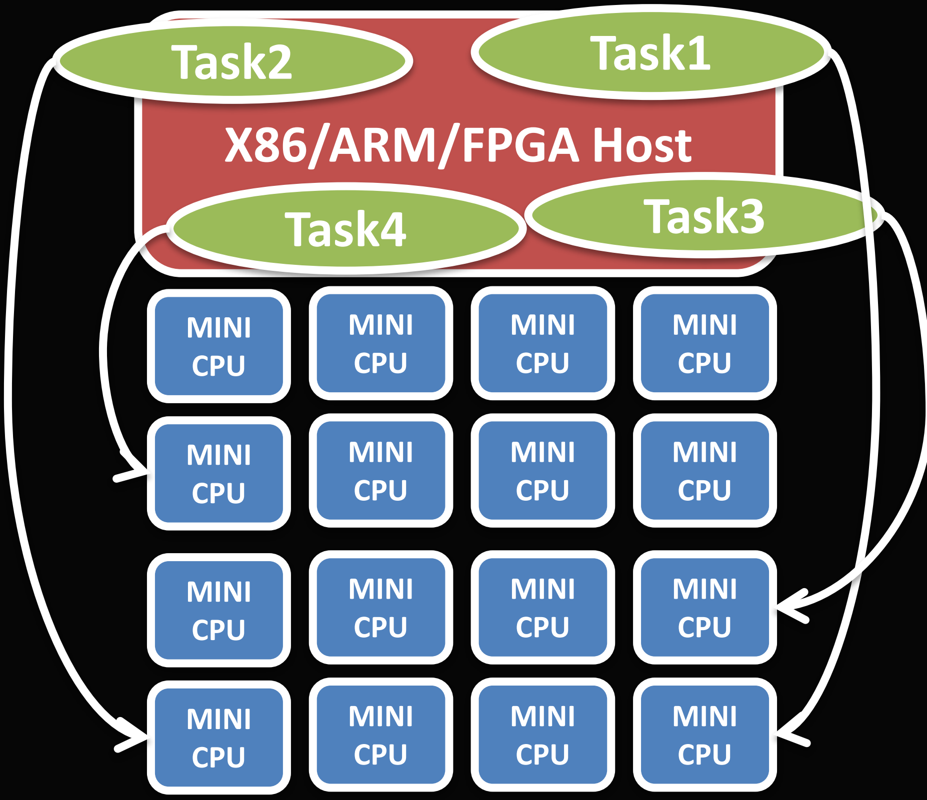

Parallel Programming Frameworks

Erlang	SystemC	Intel TBB	Co-Fortran	Lisp	Janus
Scala	Haskell	Pragmas	Fortress	Hadoop	Linda
Smalltalk	CUDA	Clojure	UPC	PVM	Alef
Julia	OpenCL	Go	X10	Posix	XC
Occam	OpenHMPP	ParaSail	APL	Simulink	Charm++
Occam-pi	OpenMP	Ada	Labview	Ptolemy	StreamIt
Verilog	OpenACC	C++ Amp	Rust	Sisal	Star-P
VHDL	Cilk	Chapel	MPI	MCAPI	?????????

adapteva

Stupid Hurdles That Hinder Collaboration

- Proprietary SDKs and programming frameworks
- Lack of datasheets/documents
- Closed source drivers
- Expensive lock-in hardware
- NDA requirements
- Exlcusive access

Open HW is now following the same successful path as open SW!

adapteva

Parallella: Our "Secret Weapon"

- A $99 single board ”parallel” computer that runs Linux
- Open source (SDK, board files, drivers) (github.com/parallella)
- Open documentation (adapteva.com/all-documents)
- Open to all (forums.parallella.org)

The Parallella Board

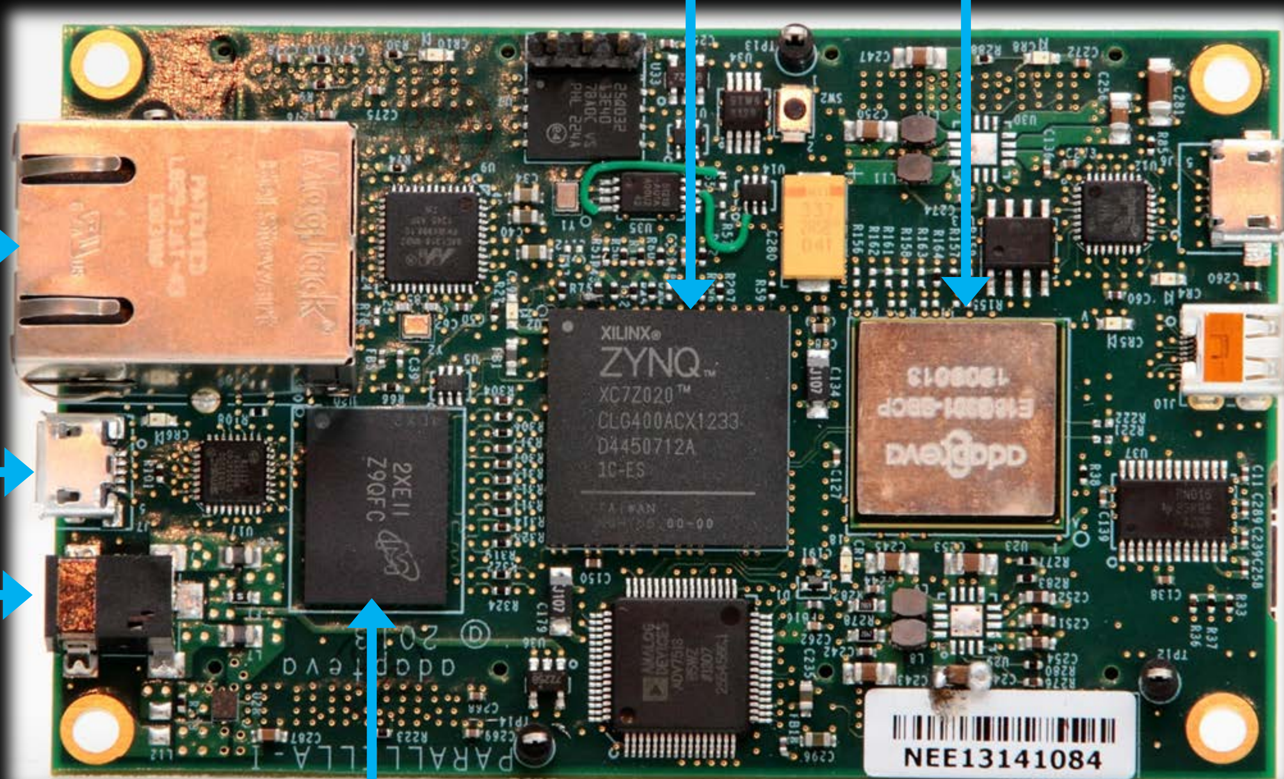

adapteva

Parallella Kickstarter Campaign

- 5,000 customers
- 6,300 boards ”pre-sold” in 4 weeks
- 67 countries, all 50 US states
- 50-75% of backers are developers
- 12,000 more signups since Jan 1st
- Backer Application Interest:
  - Software Defined Radio
  - Ray tracing/rendering
  - Image processing
  - Robotics
  - Gaming

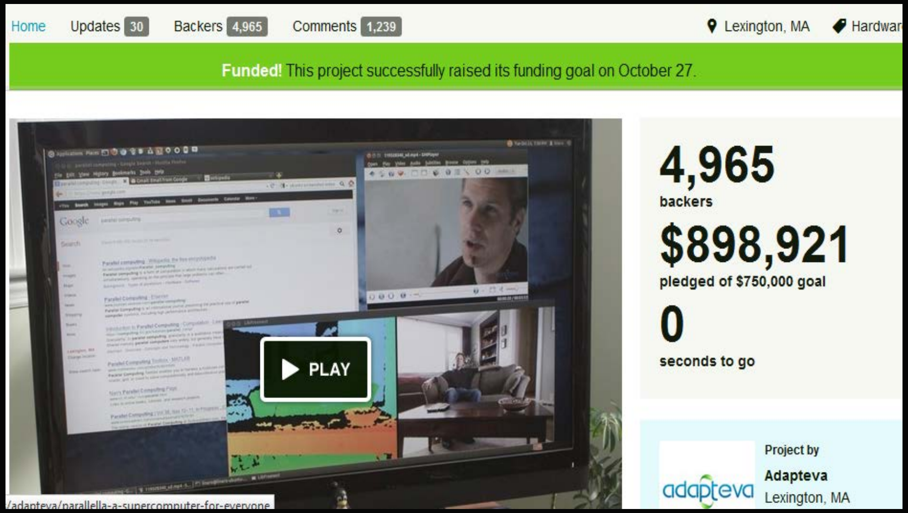

- Cryptography
- Parallel computing research
- Distributed Computing
- Machine Learning
- HPC

Epiphany IP Conclusions

- #1 in processor energy efficiency at 70 GFLOPS/Watt (core)
- Silicon proven in GLOBALFOUNDRIES 28SLP node
- Only multicore IP that is scalable to 1000’s of cores on chip
- Easier to use than GPGPUs