There's STILL plenty of room at the bottom!

Andreas Olofsson

adapteva

Richard Feynman's Lecture (1959)

"There's Plenty of Room at the Bottom"
An Invitation to Enter a New Field of Physics

"Why cannot we write the entire 24 volumes of the Encyclopedia Britannica on the head of a pin?"

"The principles of physics, as far as I can see, do not speak against the possibility of maneuvering things atom by atom."

"I don't know how to do this on a small scale in a practical way, but I do know that computing machines are very large; they fill rooms. Why can't we make them very small, make them of little wires, little elements – and by little, I mean *little*. For instance, the wires should be 10 or 100 atoms in diameter, and the circuits should be a few thousand angstroms across."

Other Concepts:

- *Rearranging atoms, Micro-machines, Chemical synthesis, Micro-antenna arrays*

adapteva

Why the question is still relevant!

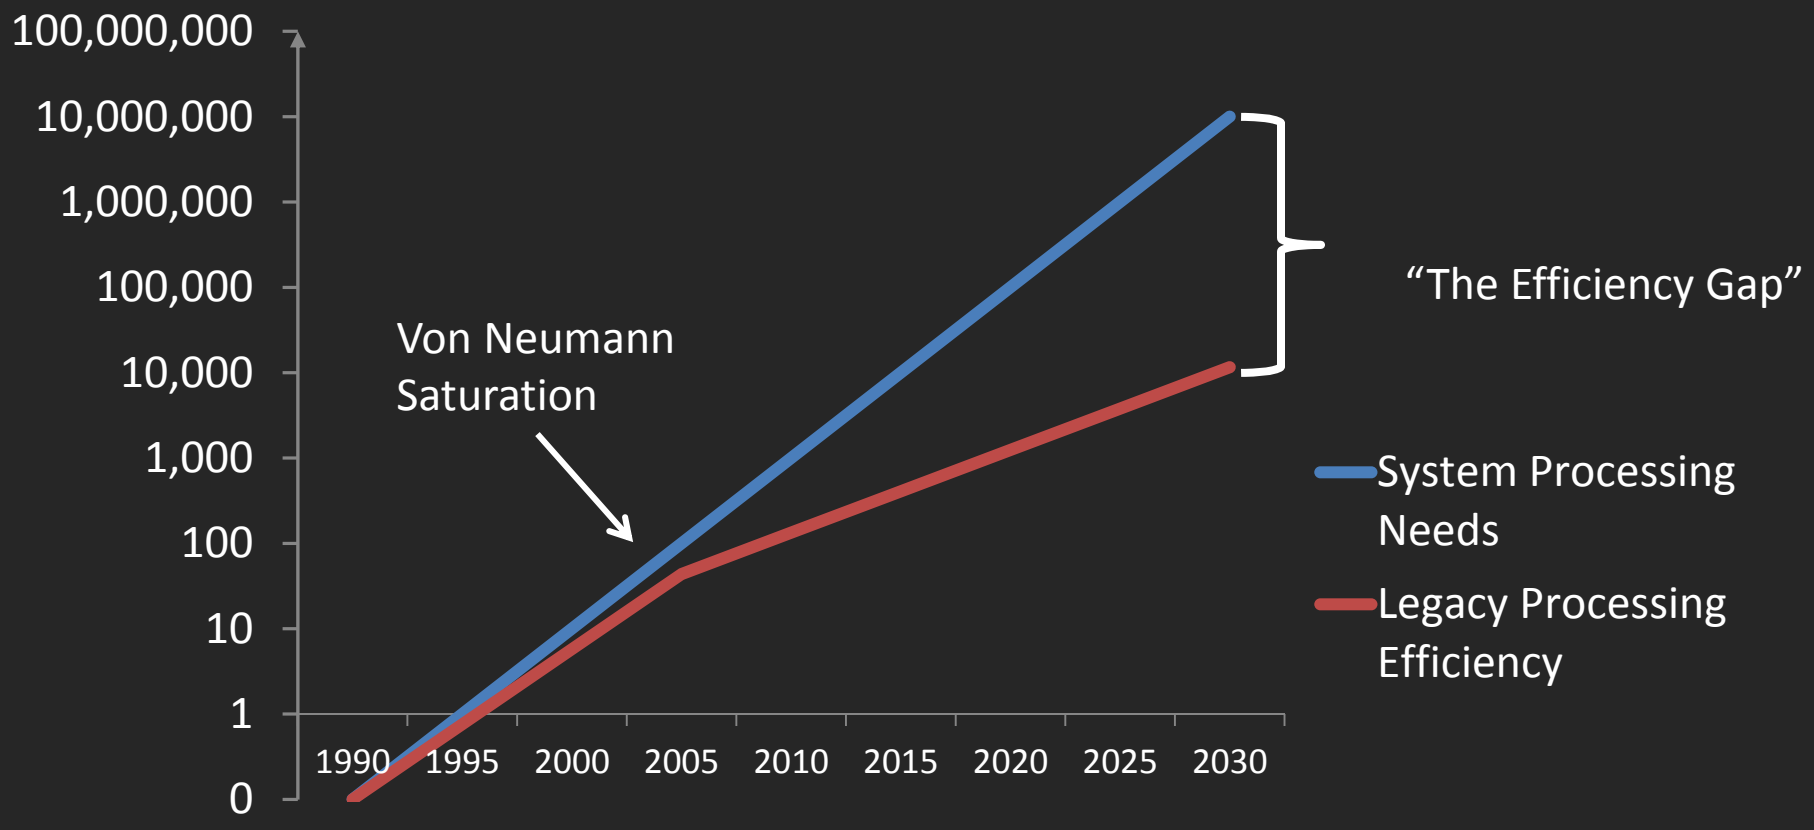

adapteva

Need for Efficient Processing is Universal

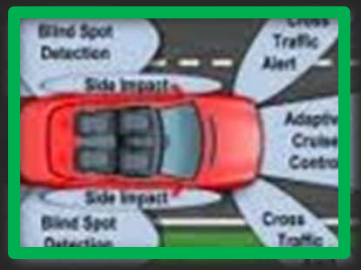

adapteva

Adapteva Company Introduction

Company History:

- Founded in 2008 by processor design team from Analog Devices
- [Coprocessor / manycore / DSP / multicore / accelerator / cpu] company
- Shipping 16-core 65nm product since May 2011
- Sampling 64-core 28nm product since July 2012
- Launched Parallella open computing platform in October 2012

Notable Achievements:

- 4 chips on < $2.5M in raised capital
- 50 GFLOPS/W demonstrated at chip level
- 28nm 64-core product only $10mm^2$
- Architecture scales to 1024 CPUs on-chip
- ~$2M in revenue so far (~break-even)
- ~5,000 customers

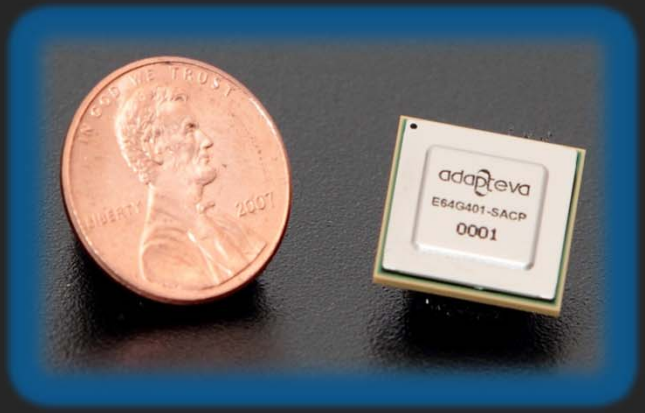

adapteva

Our Vision: True Heterogeneous Computing

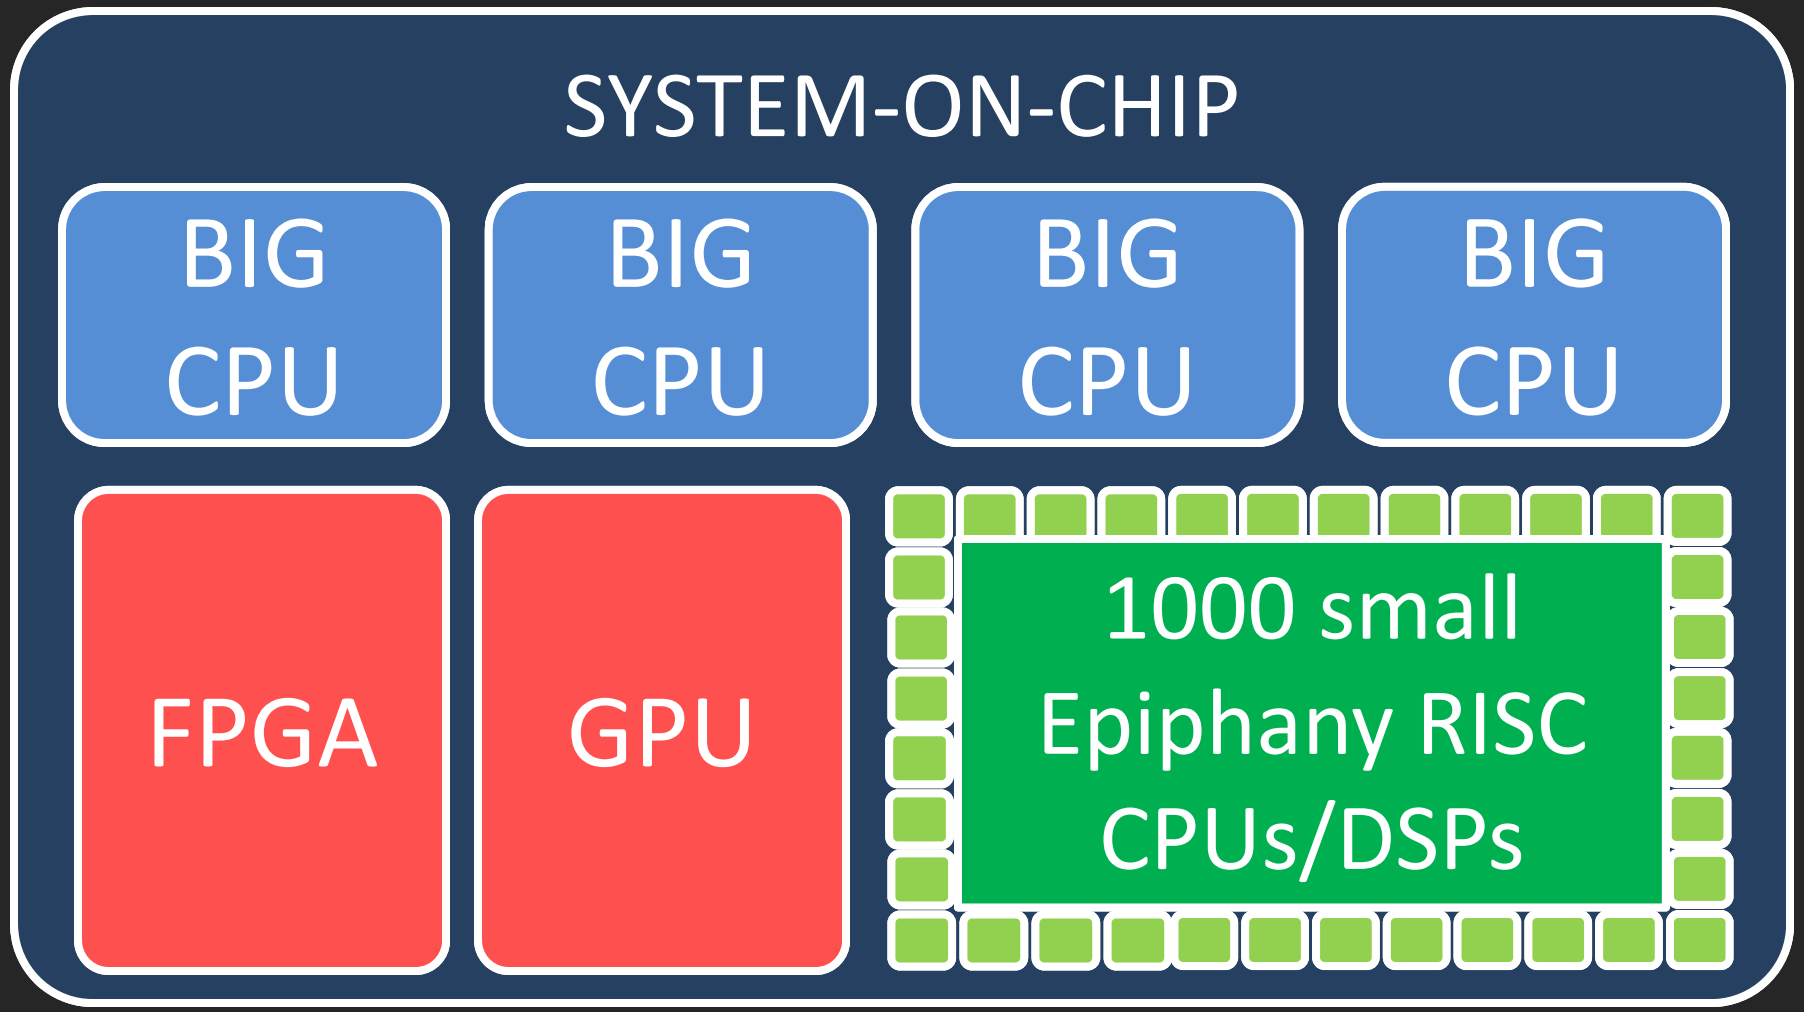

adapteva

HPC COMPARISON DATA

Technology	FPGA	DSP	GPU	CPU	Manycore
Process	28nm	40nm	28nm	32nm	28nm
Programming	VHDL	C++/Asmbly	CUDA	C/C++	C/C++
Area (mm^2)	590	108	294	**216**	**10**
Chip Power (W)	40	22	135	**130**	**2**
CPUs	n/a	8	16	**4**	**64**
Max GFLOPS	1500	160	3000	**115**	**102**
GHz * Cores	n/a	12	16	**14.4**	**51.2**
L1 Memory	6MB	512KB	2.5MB	**256KB**	**2MB**

Efficiency is everything

Peak performance means very little

No magic bullet!

adapteva

Epiphany: Massive Task-Parallelism

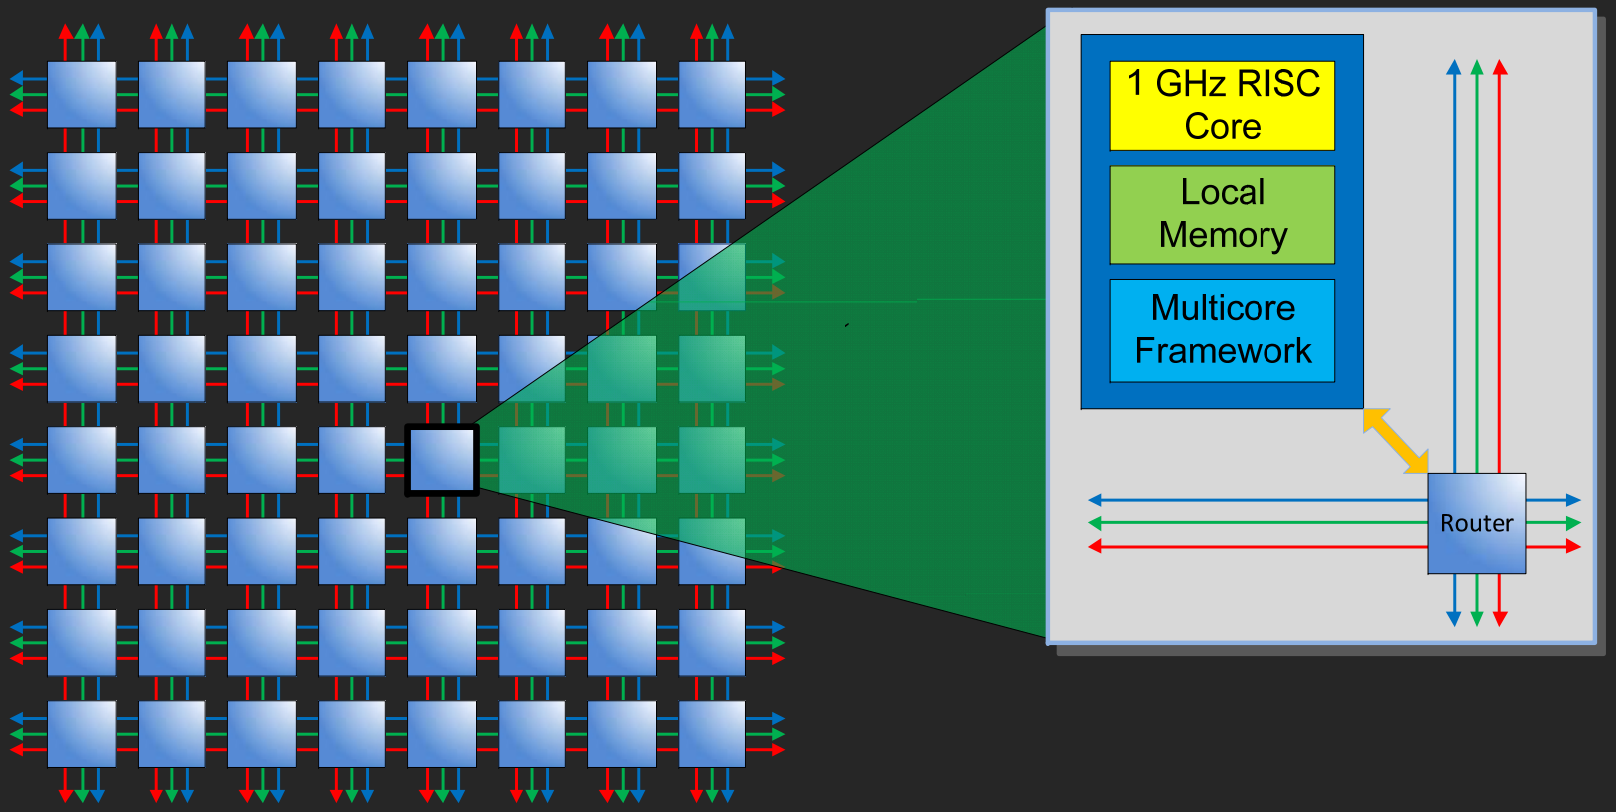

Coprocessor to ARM/Intel CPU

25mW per core

C/C++ programmable

adapteva

Programming Models

MODEL#1
TASK QUEUE MODEL
- Great for up to 2GFLOPS
- Supports standard C/C++
- "Cloud on a chip"

MODEL #2
DATA PARALLEL MODEL
- openCL programmable
- Easy integration of C/C++
- openMP roadmap

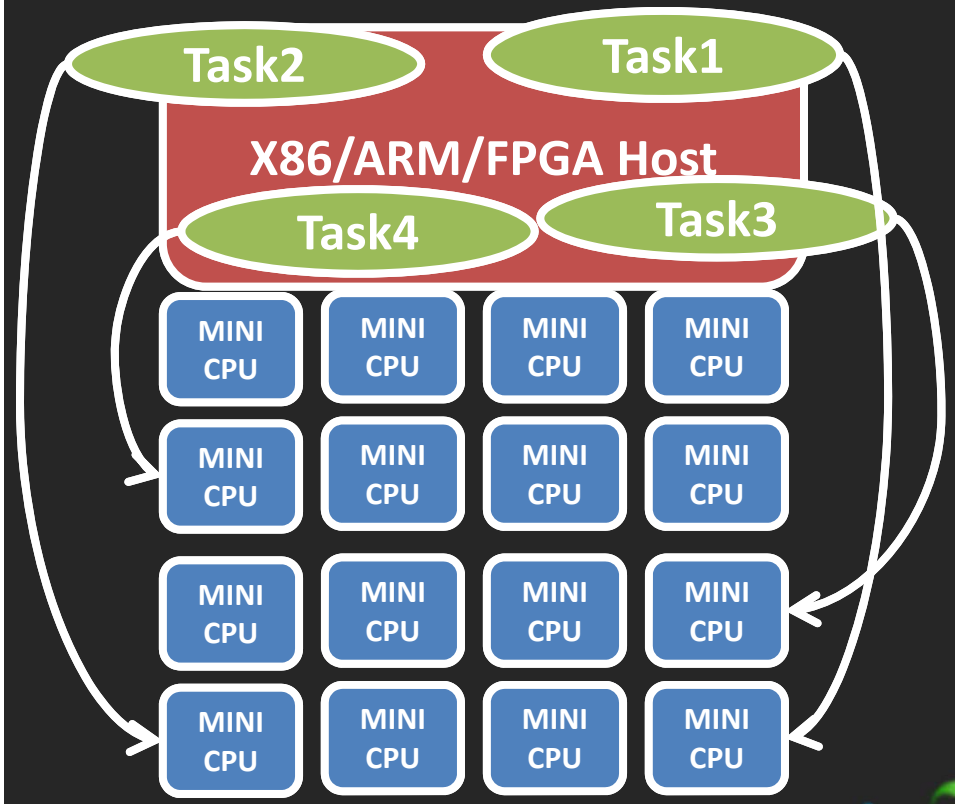

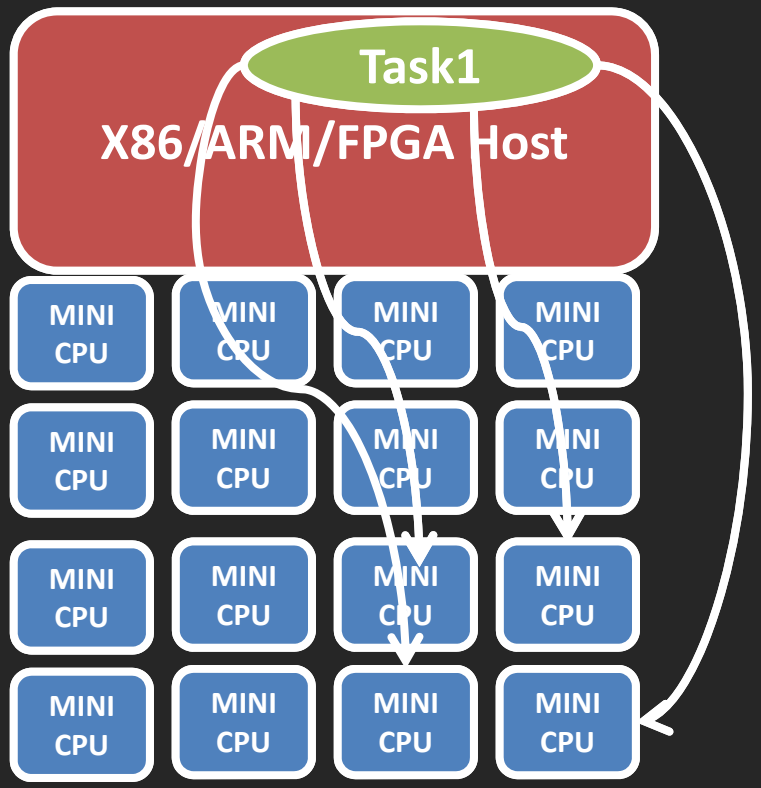

adapteva

The Future is Open HW Platforms!

Y2000 Model

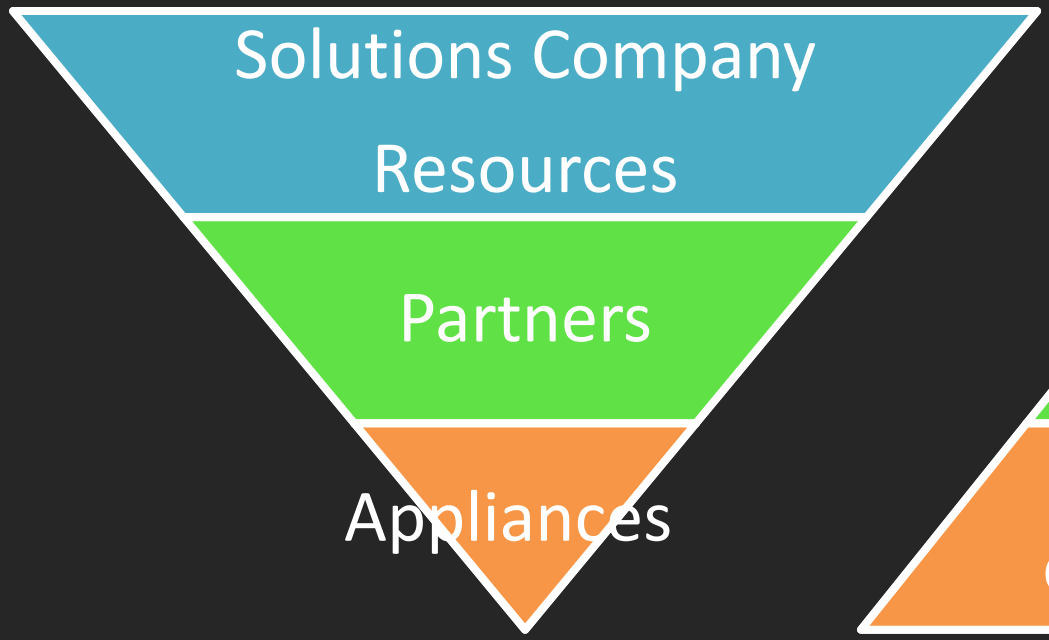

CLOSED

Y2012 Model

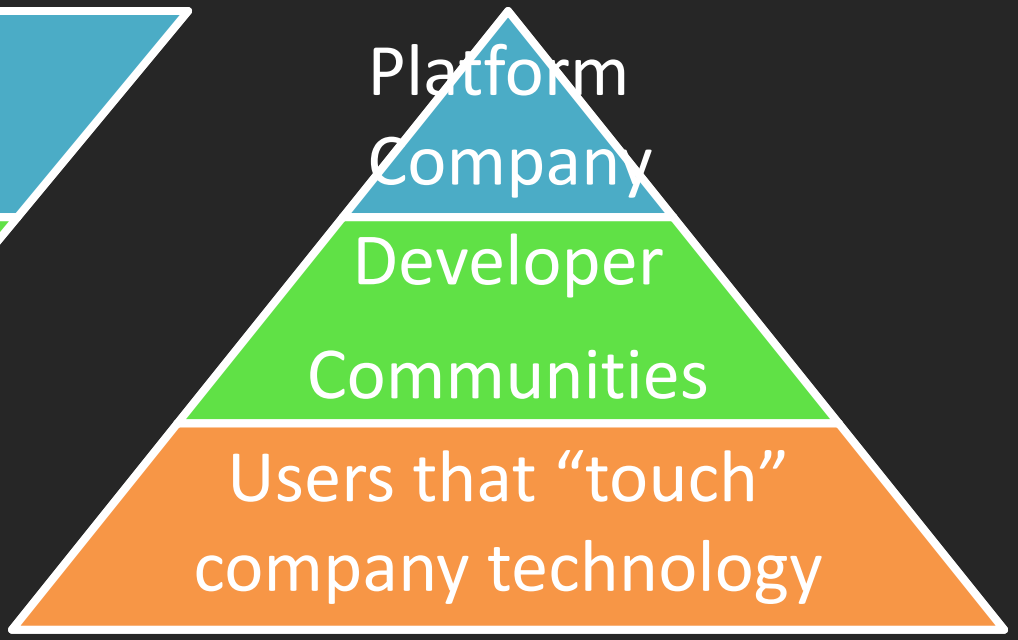

OPEN

adapteva

Parallella Open Computing

- Open (and ”free”):
  - Documentation
  - Board design files
  - Drivers
  - Software Tools
- Accessible (NO NDAs!)
- $100 entry point
- ~4000 devs signed up in 4 weeks

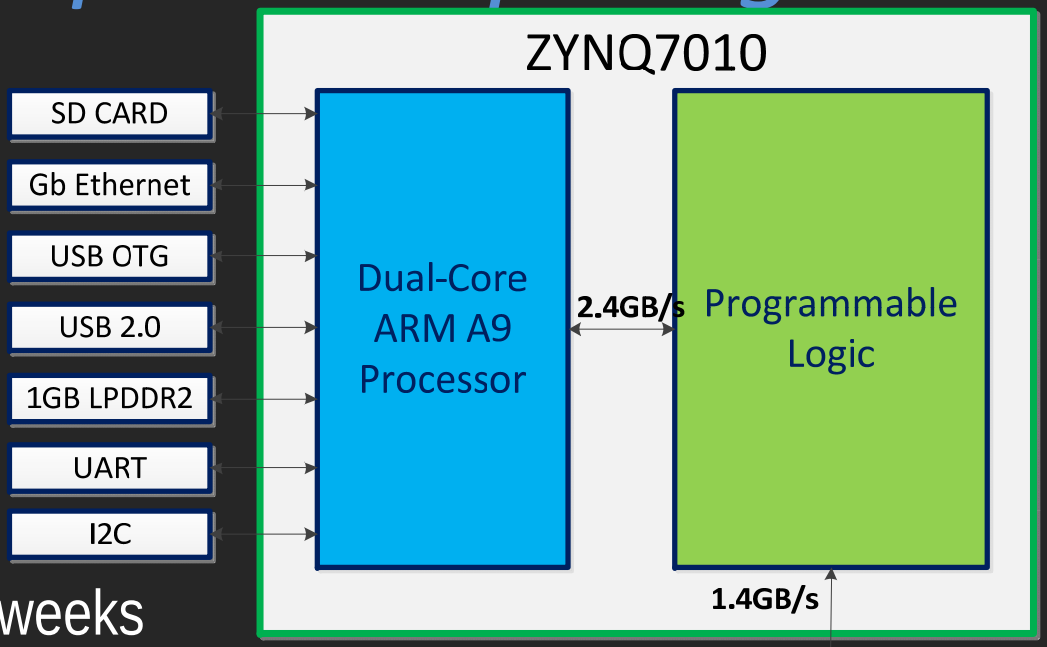

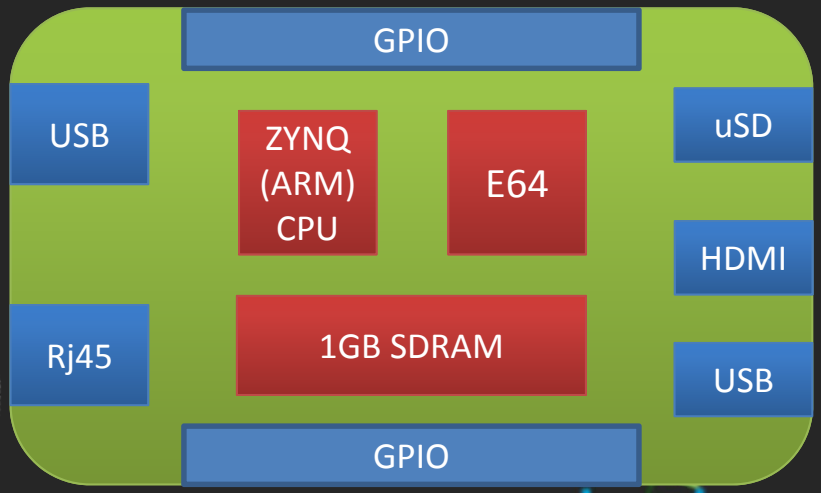

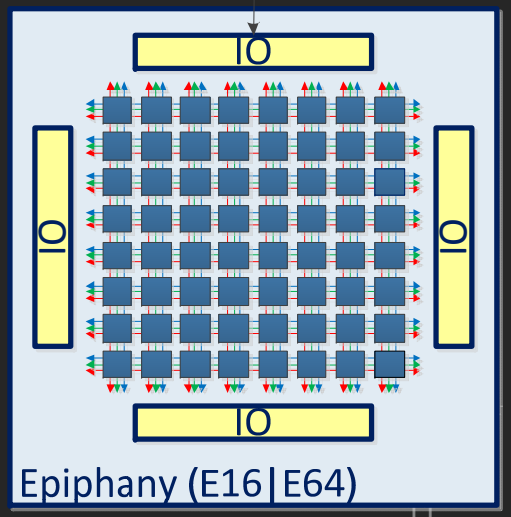

adapteva

Mind Boggling Progress

(1992)
Connection Machine 5

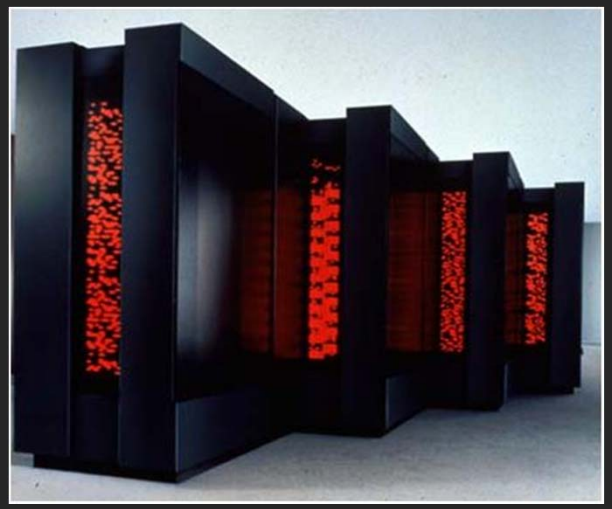

100 GFLOPS
100 KW
$10M

(2012/2013)
Parallella Board

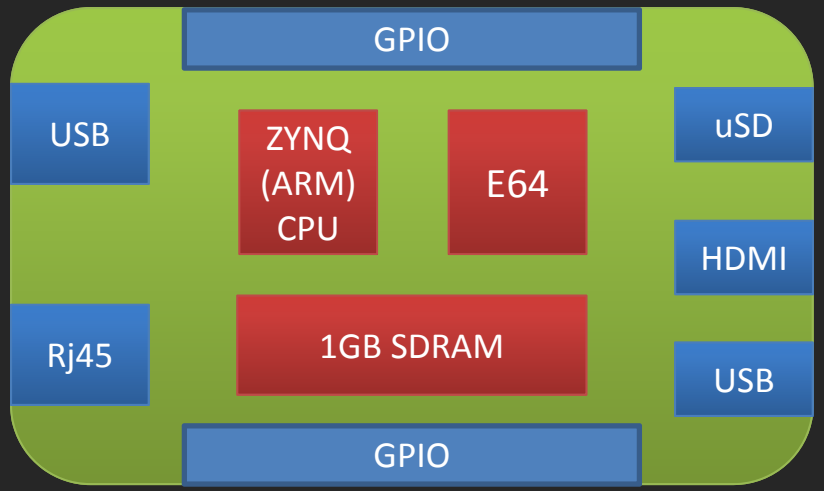

100 GFLOPS
5 W (20k X)
$200 (50k X)

adapteva

Smartphone: So much room…

iPhone4s
~58mm

What if your smartphone disappears?

A5X Chip
~16mm

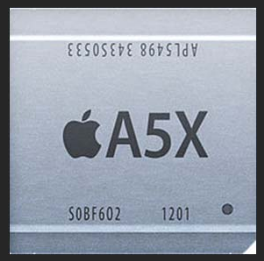

A5X-die
~13mm

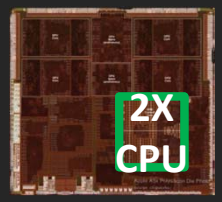

ARM
A9
~2mm

FPU
~0.15mm

62 cm^3

>1M X Volume Reduction

0.00003 cm^3

adapteva

HPC: So much room…

Exascale Costs Impractical?

What if we start over??

1 Exaflop Processor Only Costs (65nm)

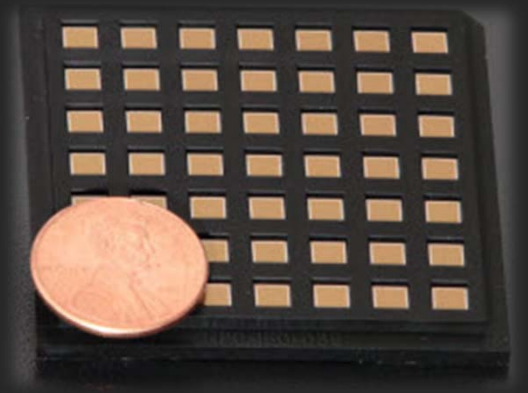

$1-10B → 100 - 1000x ratio → $10M?

adapteva

Capital Efficiency: So much room…

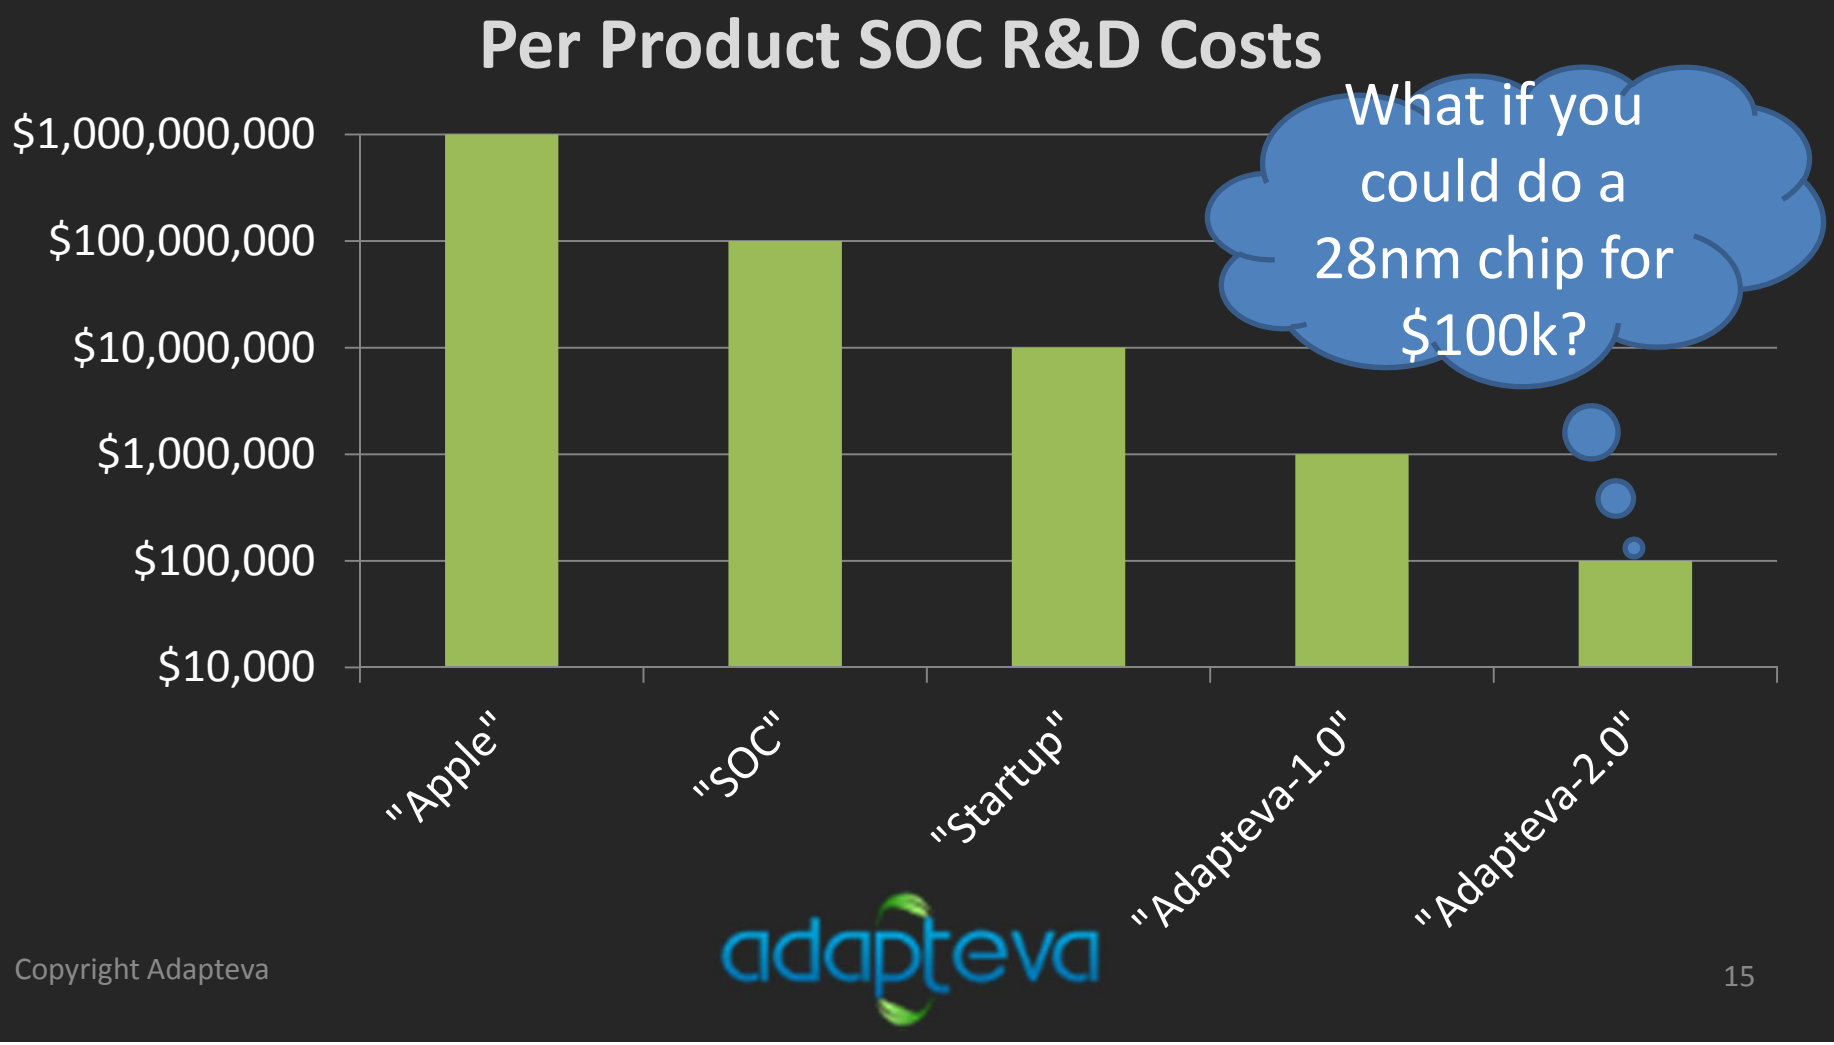

TRENDS THAT WILL SHAPE THE FUTURE OF COMPUTING

- **Efficiency**
- **Locality, Integration**
- **New Memory Hierarchy**
- **Redundancy**
- **Parallelism**
- **Programmability**
- **Distributed Architectures**
- **Ease of Use**
- **Thermal Balancing**
- **New Algorithms**

adapteva

TRENDS THAT WILL SHAPE THE FUTURE OF COMPUTING

- **Efficiency**
- **Locality, Integration**
- **New Memory Hierarchy**
- **Redundancy**
- **Parallelism**
- **Programmability**
- **Distributed Architectures**
- **Ease of Use**
- **Thermal Balancing**
- **New Algorithms**

adapteva

The Future of Computing:
Million Way Parallelism

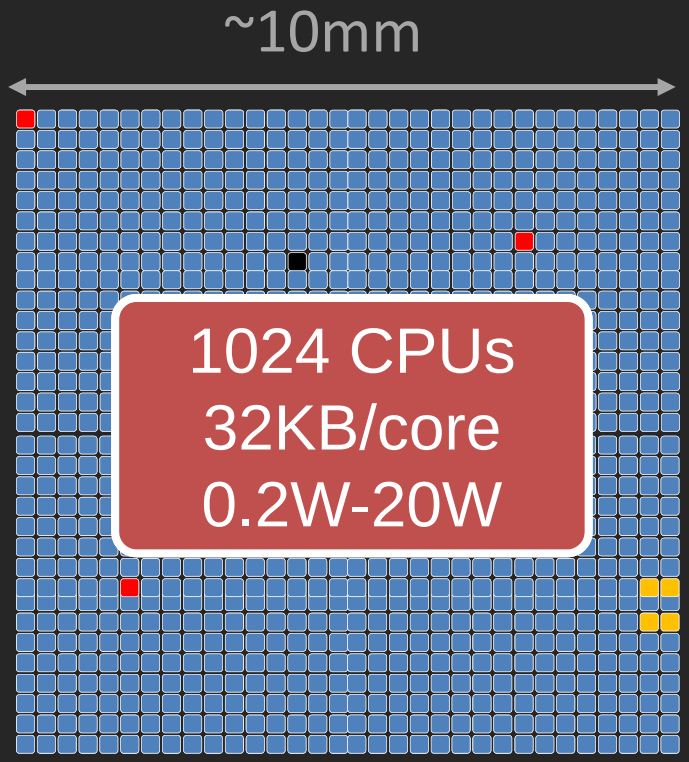

2012

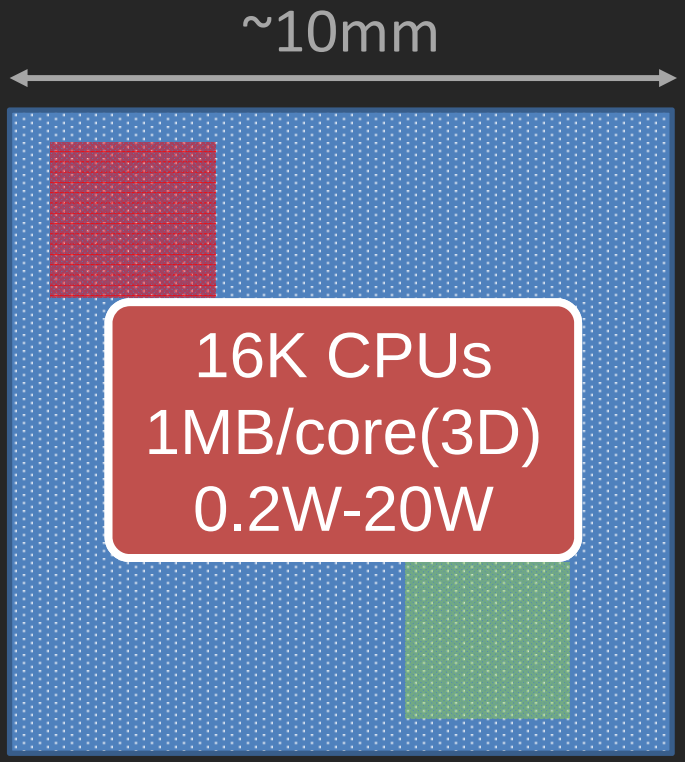

2022

adapteva

The Future is…

Open

Heterogeneous

Massively Task-Parallel

Efficient

Grande Challenges Ahead…

- **Rebuild the computer ecosystem**
- **Rewrite billions of lines of code**
- **Retrain millions of programmers**
- **Rewrite the education curriculum**

adapteva