

# Parallella: A Love Story

Heterogeneous..

Parallel..

Efficient..

Open..


*Andreas Olofsson*

*MIT, Jan 7, 2013*



# Adapteva Achieves 3 “World Firsts”

1. First processor company to reach 50 GFLOPS/W



2. First open source OpenCL™ SDK in the mobile market

3. First semiconductor company to successfully crowd-source project

**KICKSTARTER**


*“OpenCL and the OpenCL logo are trademarks of Apple Inc. used by permission by Khronos.”*



# Prologue



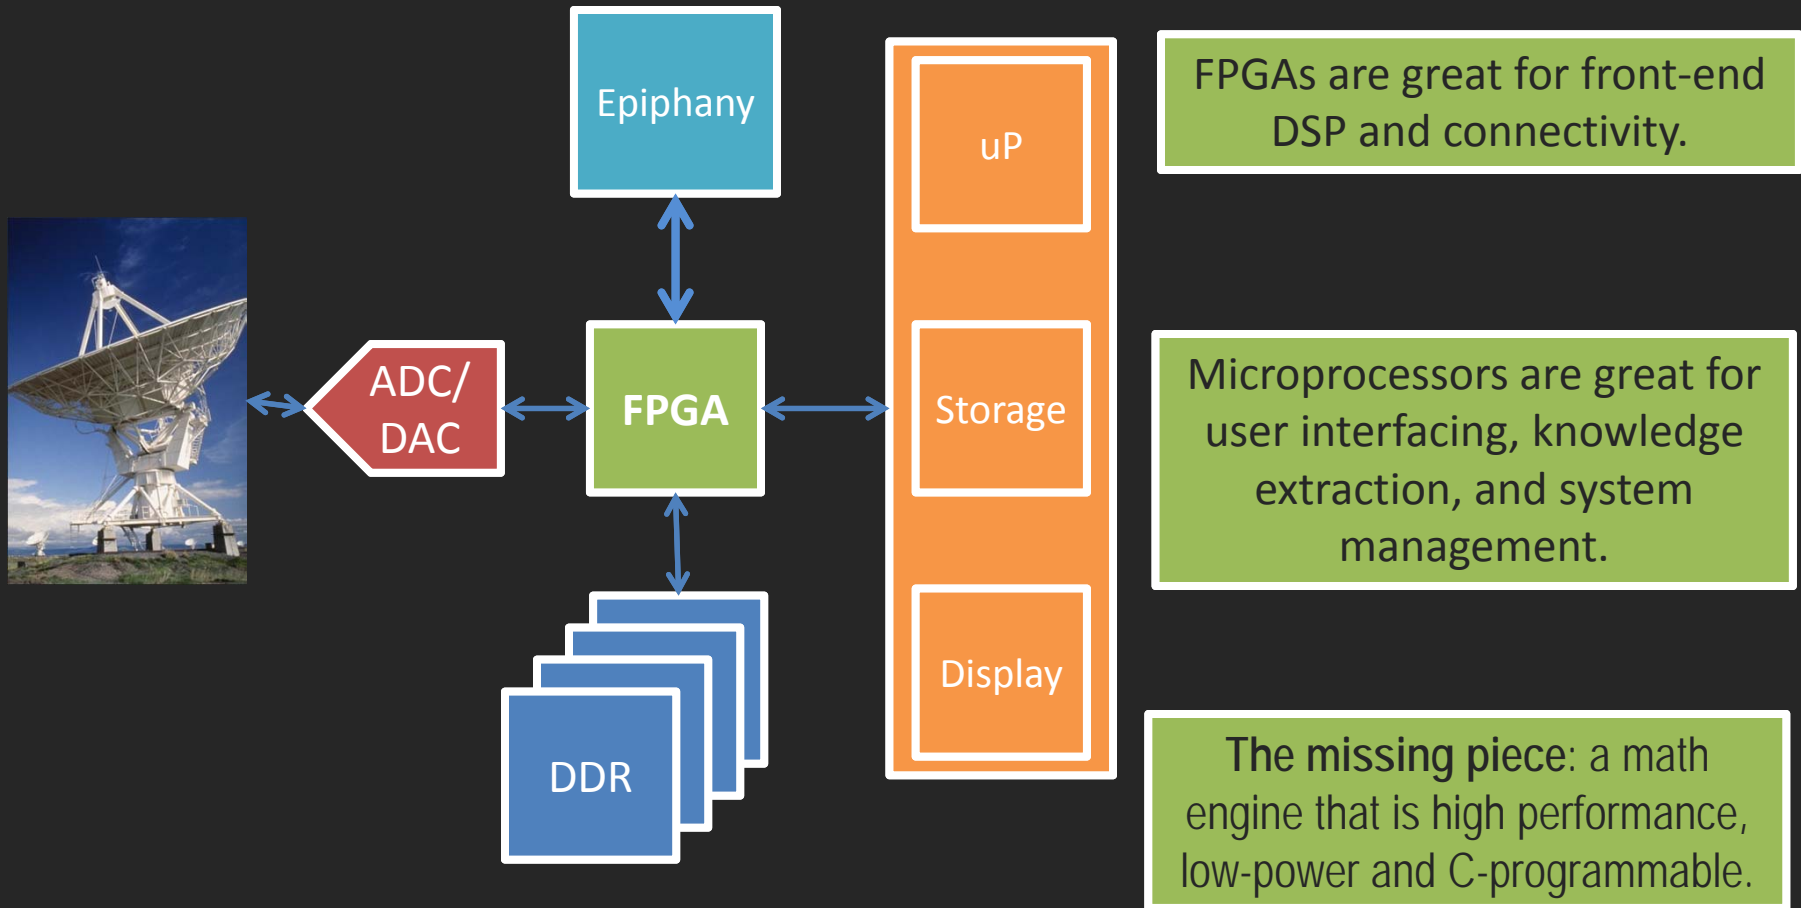
# Why we need heterogeneous and parallel platforms



# ASIC, FPGA, DSP, CPU?

	ASIC	FPGA	DSP	CPU
Flexibility	Poor	Great	Good	Good
Efficiency	Great	Good	Good	Fair
Development Cost/Risk	High	Medium	Medium	Low
Leverage	Minimal	Modest	High	Huge



# A Practical Radar System Example



# Why SOC integration is so disruptive



62 cm<sup>3</sup>



What if your smartphone disappears?


>1M X Volume Reduction

0.00003 cm<sup>3</sup>



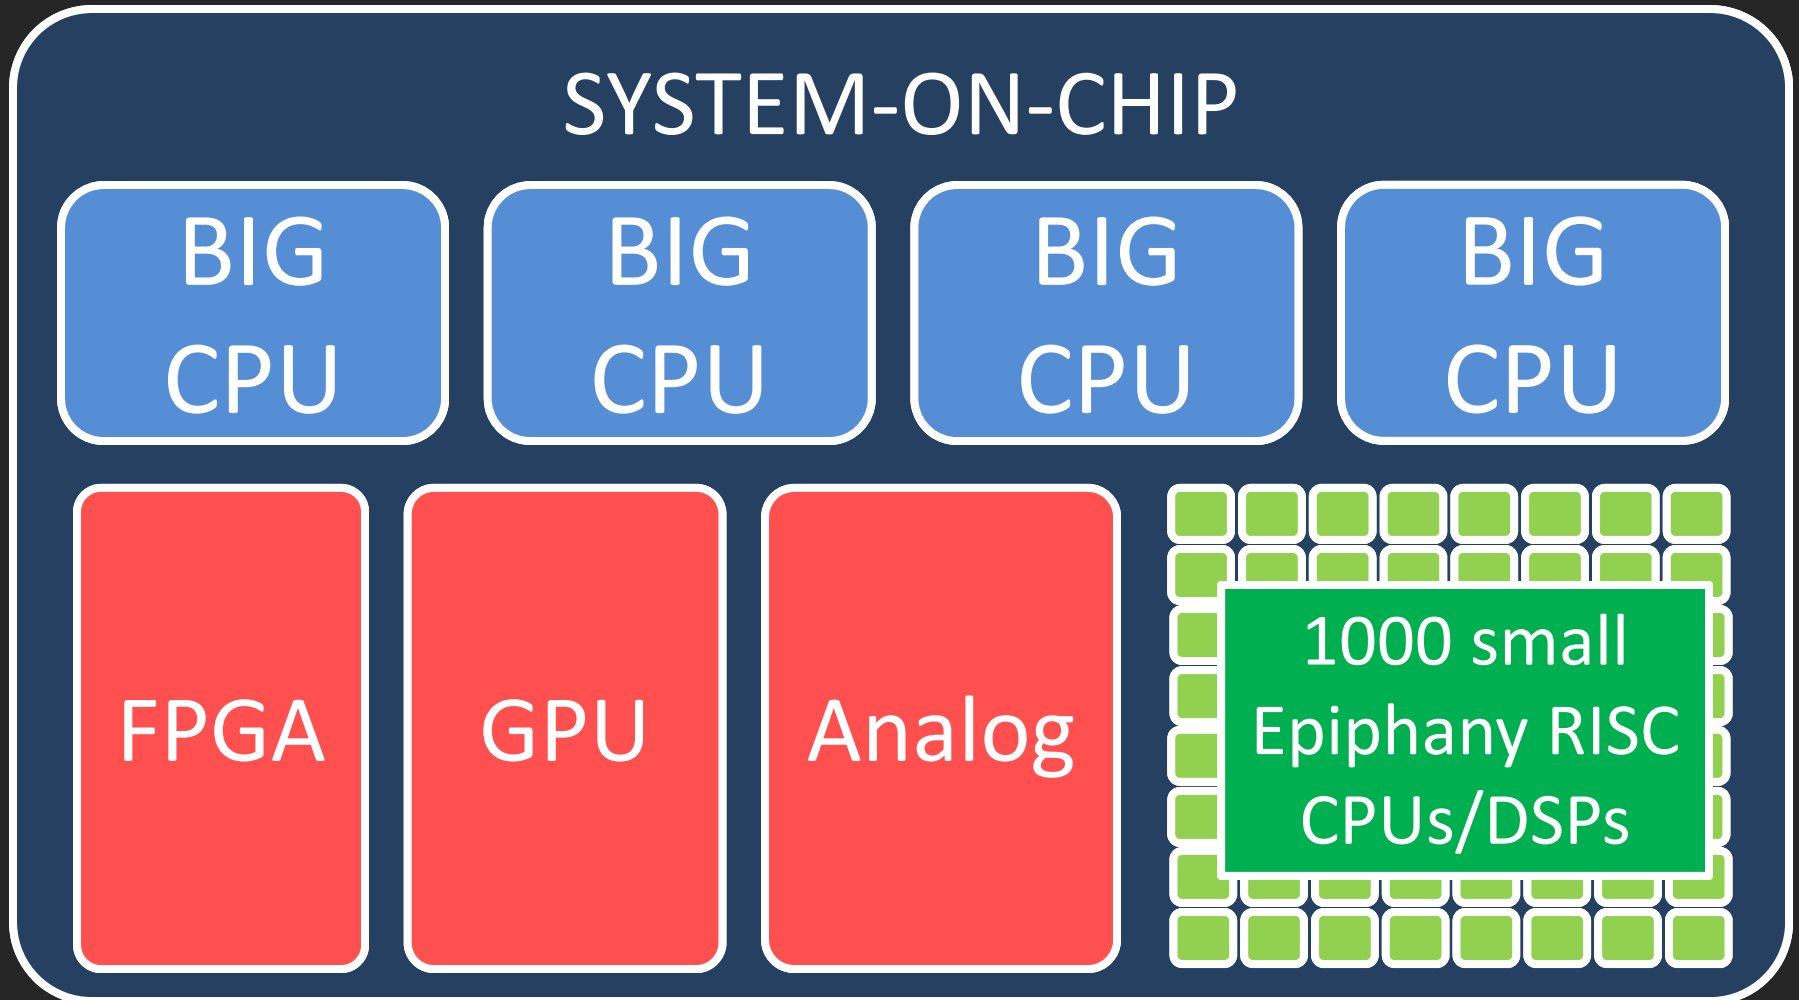
# The Problem: SOCs are complex!

## Per Product SOC R&D Costs

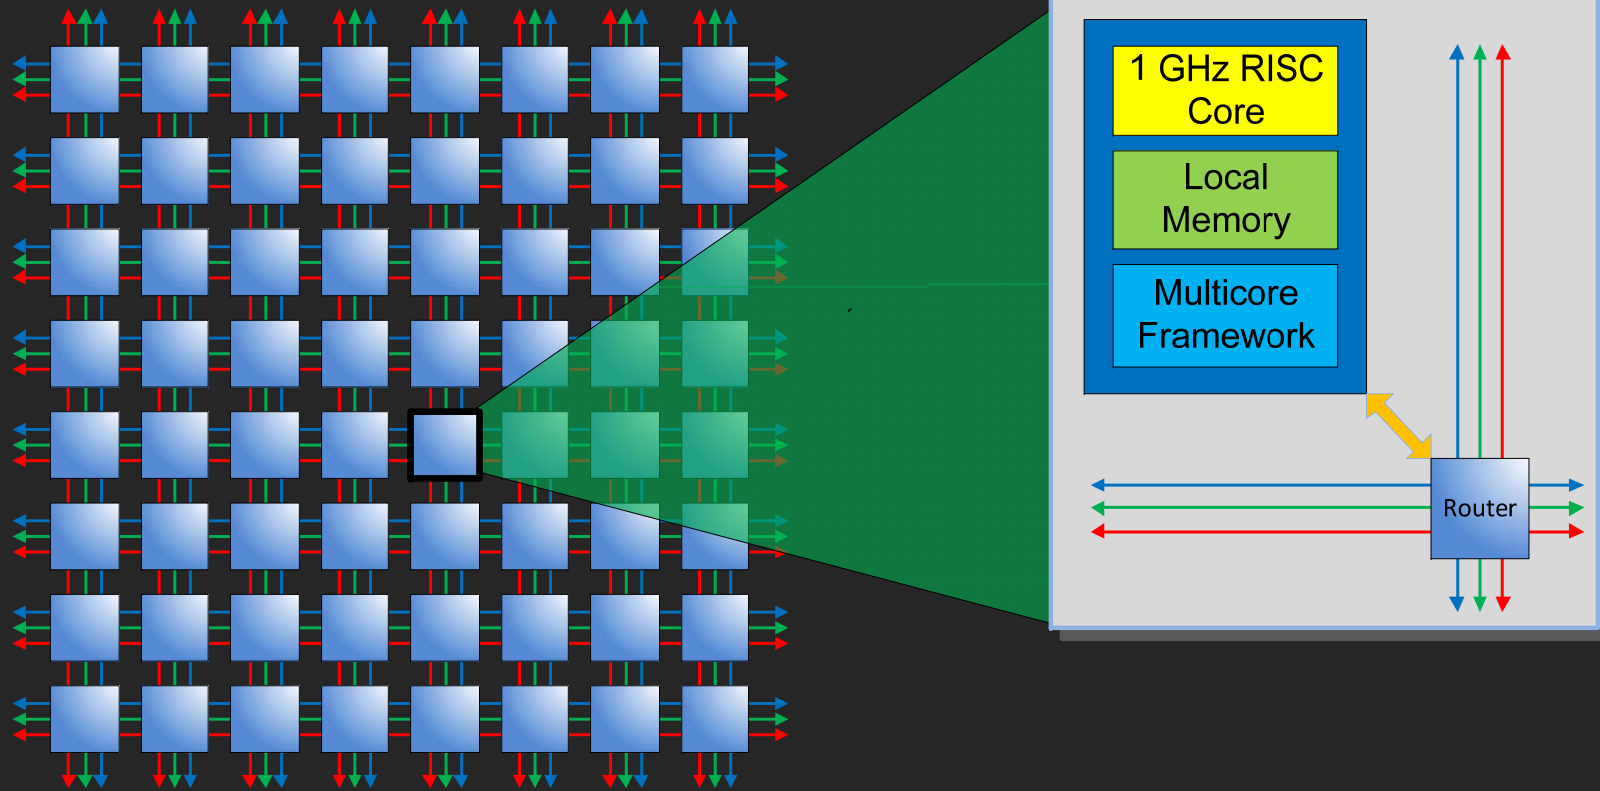




# *Our Vision: True Heterogeneous Computing*



# Epiphany: Massive Task-Parallelism



Coprocessor to  
ARM/Intel CPU

25mW per core

C/C++ programmable




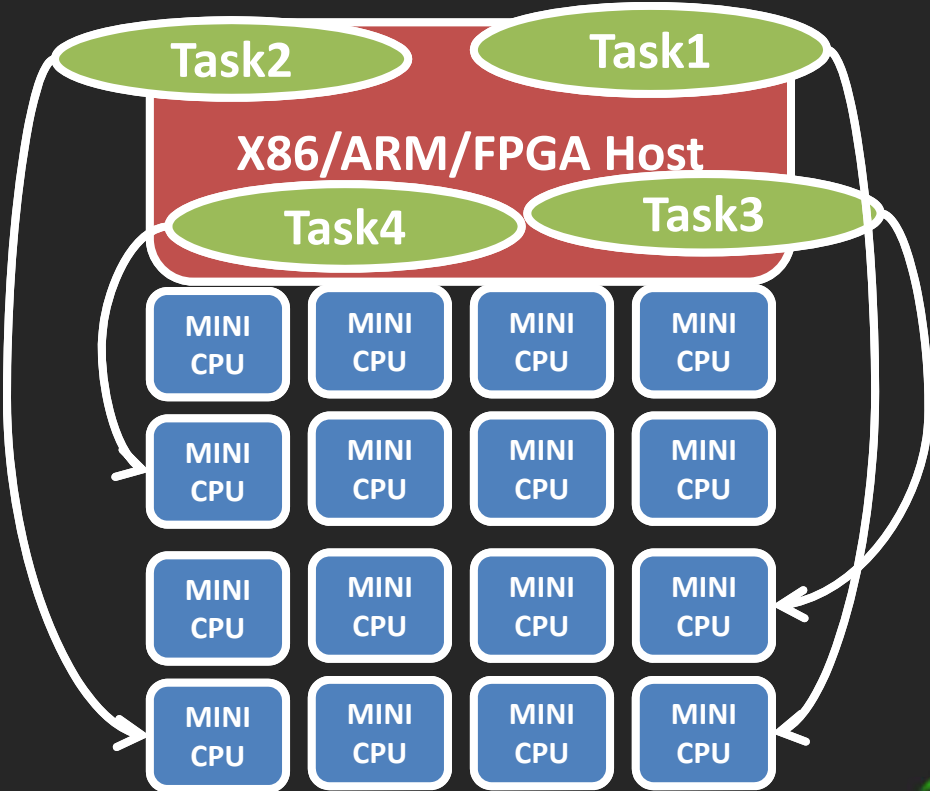
# Programming Models

**MODEL #1  
TASK QUEUE MODEL**



- Up to 2 GFLOPS/core
- Supports standard C/C++
- "Cloud on a chip"

**MODEL #2  
DATA PARALLEL MODEL**

- openCL programmable
- Easy integration of C/C++
- openMP/MPI roadmap



# Epiphany Silicon Devices




## Features:

- 16 RISC CPU cores
- 512KB distributed memory
- IEEE Floating Point
- 32 distributed DMA engines
- 4 off-chip serial links
- 65nm

## Specifications:

- 1 GHz
- 32 GFLOPS
- 2 Watt Max Chip Power
- 512 GB/sec memory bandwidth
- 8GB/sec off chip BW



## Features:

- 64 RISC CPU cores
- 2MB distributed memory
- IEEE Floating Point
- 128 distributed DMA engines
- 4 off-chip serial links
- 28nm

## Specifications:

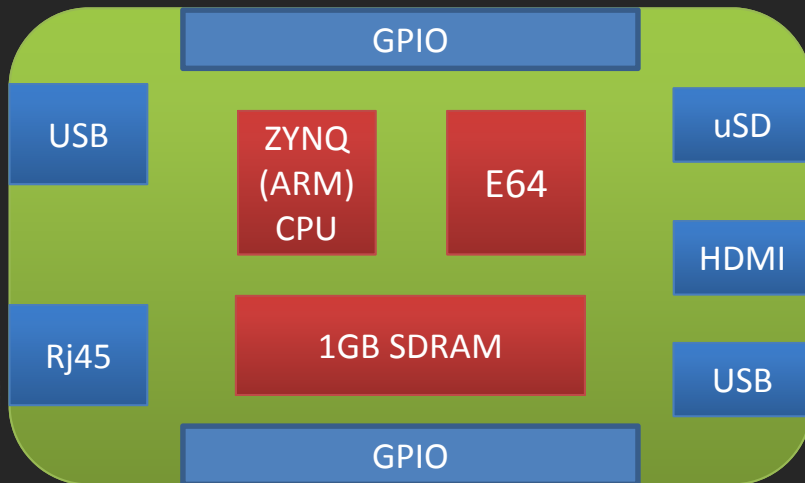

- 800 MHz
- 100 GFLOPS
- 2 Watt Max Chip Power
- 1.6 TB/sec memory bandwidth
- 8GB/sec off chip BW

# Parallella

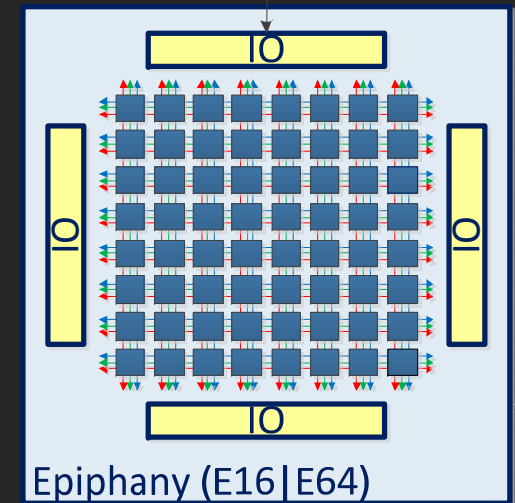


# Parallella Open Computing

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




adapteva



# How cool is this?

(1992)

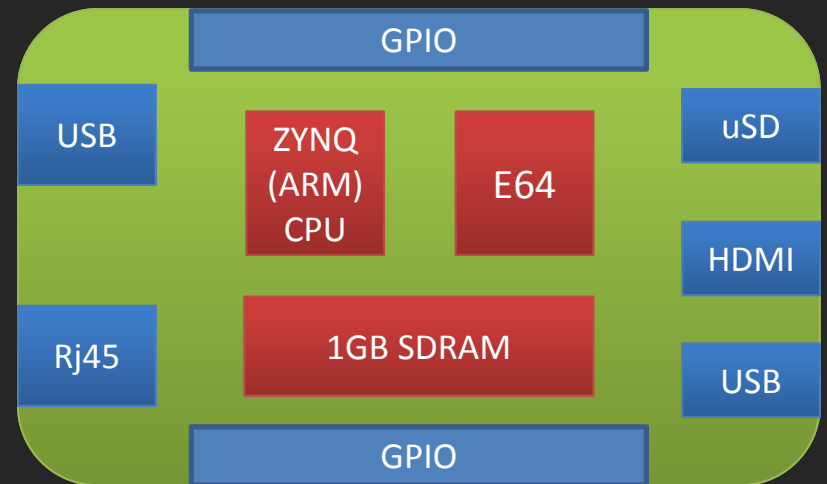
Connection Machine 5



100 GFLOPS  
100 KW  
\$10M

(2012/2013)

Parallella Board



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





# *Parallella Coprocessor Approach*

ARM runs Linux


Epiphany  
accelerates key  
tasks

Programmable  
logic “makes  
anything possible”

## **Program Flow**

1. ARM boots Linux. First stage boot loader from Flash, everything else from SD card.
2. “Main” application executes on ARM
3. Application sends critical tasks send to Epiphany using OpenCL or simple threads
4. ARM/Epiphany communication through shared DRAM buffer outside virtual memory of O/S.

# Zedboard Introduction



**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**

