



**Agile, eXtensible, fast I/O Module for the cyber-physical era**

IWES 2016 -- 1st Italian Workshop on Embedded Systems

Pisa, Italy 19-20 September 2016

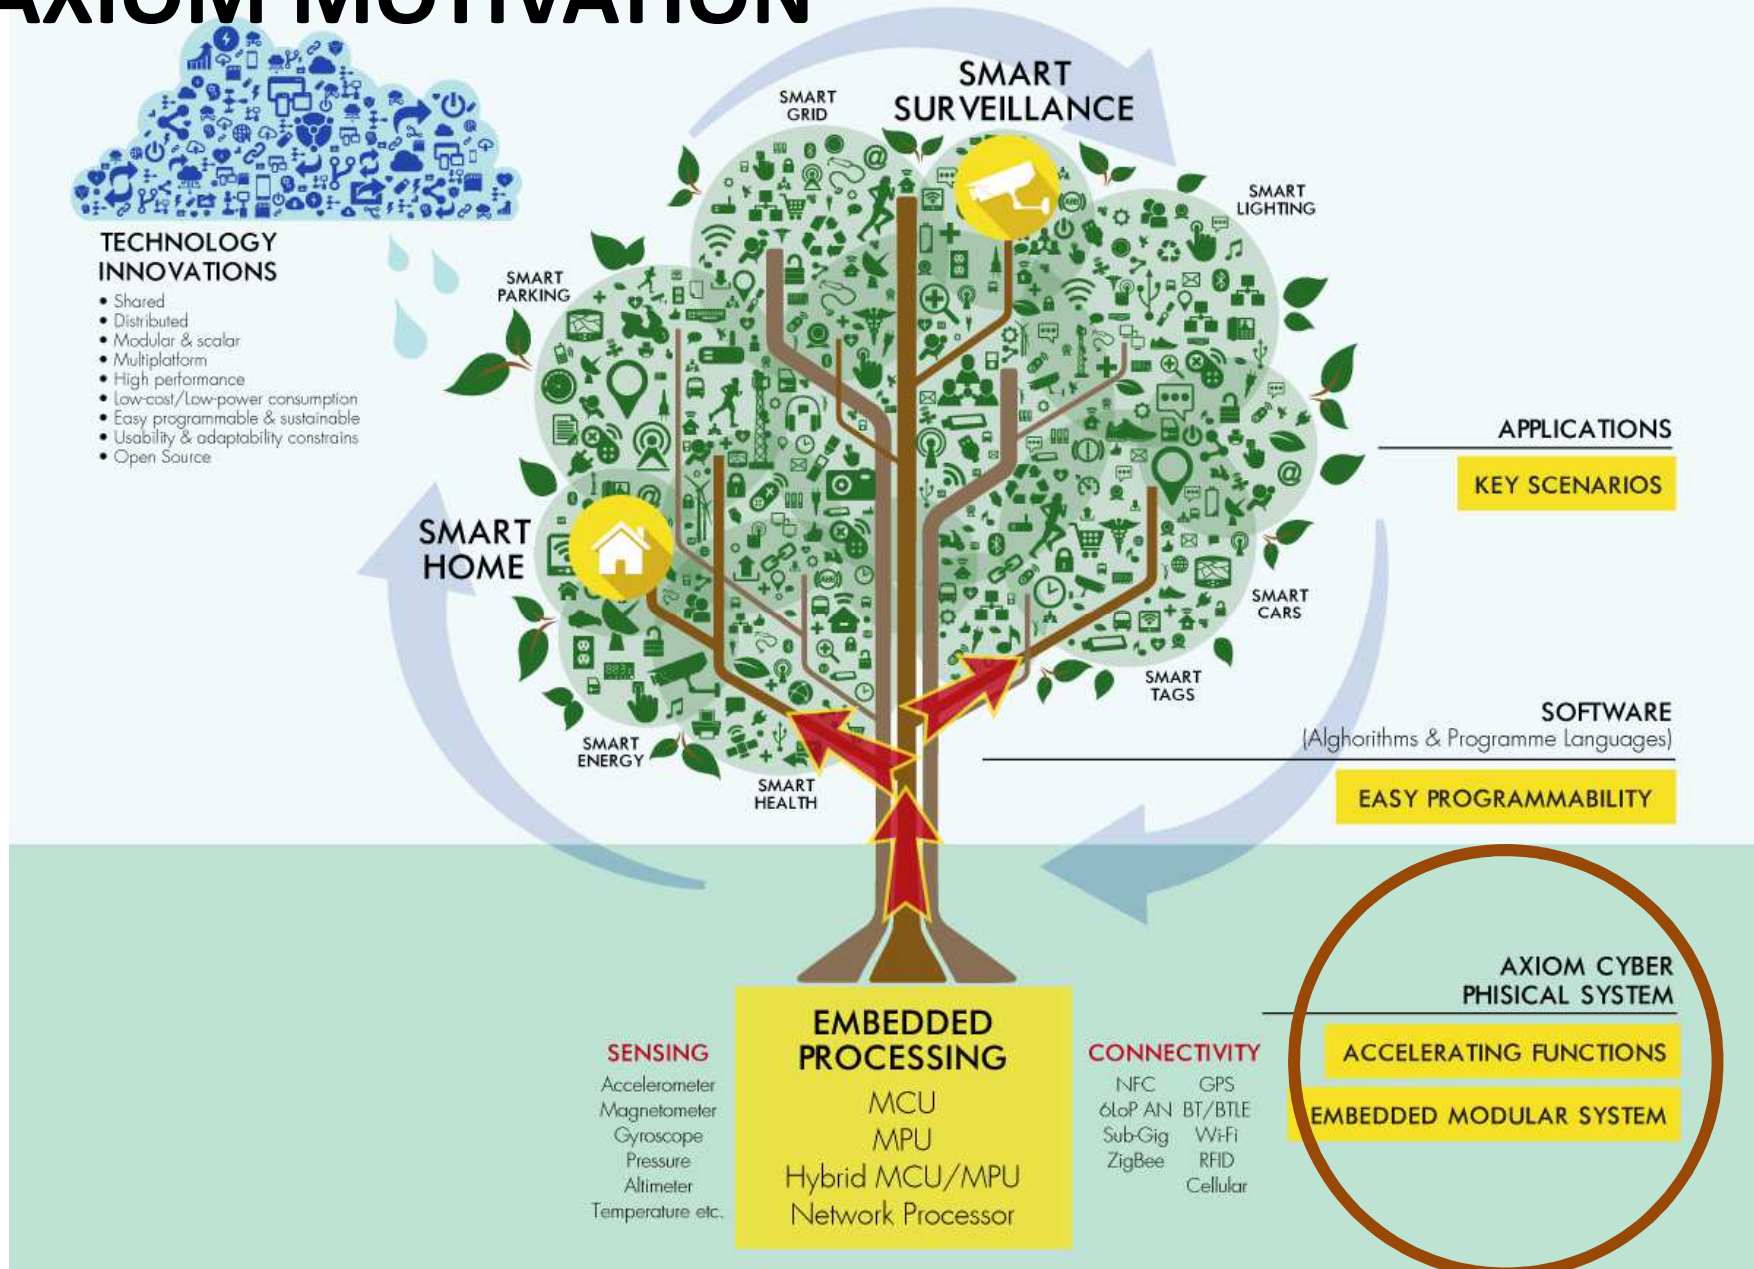
# **AXIOM: A 64-bit scalable embedded system including Arduino socket and on-chip FPGA**

Roberto Giorgi

University of Siena, Italy (Project Coordinator)



# AXIOM MOTIVATION



# Highlights of this talk

- 1) Exploring the concept of “scalable embedded system”
- 2) Indicating a way to achieve such scalability by supporting special threads called Data-Flow Threads (DF-Threads)
- 3) Illustrating how this concepts are integrated in the AXIOM project, which is focused to build a scalable Single Board Computer

# SECO (through AXIOM) is finalist for the 2016 Innovation Radar Prize (H2020)



# AXIOM OBJECTIVES

- **OBJ1) Producing a small board that is flexible, energy efficient and modularly scalable**
  - A as AGILITY, i.e. flexibility: FPGA, fast-and-cheap interconnects based on existing connectors like SATA
  - Energy efficiency: low-power ARM, FPGA
  - Modularity: fast-interconnects, distributed shared memory across boards
- **OBJ2) Easy programmability of multi-core, multi-board, FPGA**
  - Programming model: Improved OmpSs → X as EXTENSIBILITY
  - Runtime & OS: improved thread management
- **OBJ3) Leveraging Open-Source software to manage the board**
  - Compiler: BSC Mercurium
  - OS: Linux
  - Drivers: provided as open-source software by partners
- **OBJ4) Easy Interfacing with the Cyber-Physical Worlds**
  - Platform: integrating also Arduino support for a plenty of pluggable board (so-called “shields”) → “IO” as I/O
  - Platform: building on the UDOO experience from SECO
- **OBJ5) Enabling real time movement of threads**
  - Runtime: will leverage the EVIDENCE’s SCHED\_DEADLINE scheduler (i.e. EDF) included Linux 3.14, UNISI low-level thread management techniques
- **OBJ6) Contribution to Standards**
  - Hardware: SECO is founding member of the Standardization Group for Embedded Systems (SGET)
  - Software: BSC is member of the OpenMP consortium


# AXIOM – THE MODULE-v2

- KEY ELEMENTS
  - K1: ZYNQ FPGA (INCLUDES 6 ARM CORES)
  - K2: ARM GP CORE(S)
  - K3: HIGH-SPEED & INEXPENSIVE INTERCONNECTS
  - K4: SW STACK – OMPSS+LINUX BASED
  - K5: OTHER I/F (ARDUINO, USB, ETH, WIFI, ...)



# TOWARDS HPC + EMBEDDED CONVERGENCE


The HiPEAC vision for Advanced Computing in Horizon 2020



# SW+FPGA: is it useful?

## Accelerating Large-Scale Services – Bing Search

Currently deployed in 1600+ servers in production datacenters



J. Larus, Keynote2, HiPEAC Conf. , Jan 2015



# WHY OMPSS

```
1 #pragma omp target device(fpga_smp) copy_in
2 #pragma omp task in(a[0:64*64-1], b[0:64*64-1]) \
3         out(c[0:64*64-1])
4 void matrix_multiply(float a[64][64],
5         float b[64][64],
6         float out[64][64]) {
7     for (int ia = 0; ia < 64; ++ia)
8         for (int ib = 0; ib < 64; ++ib) {
9             float sum = 0;
10            for (int id = 0; id < 64; ++id)
11                sum += a[ia][id] * b[id][ib];
12            out[ia][ib] = sum;
13        }
14 }
15 ...
16 int main( void ){
17 ...
18 matrix_multiply(A,B,C1);
19 matrix_multiply(A,B,C2);
20 matrix_multiply(C1,B,D);
21 ...
22 #pragma omp taskwait
23 }
```

Application	Seq - DMA version	pthread version	OmpSs version
Cholesky	71	26	3
Covariance	94	29	3
64x64	95	39	3
32x32	95	39	3

# CAN WE DO THAT ?

UDOO



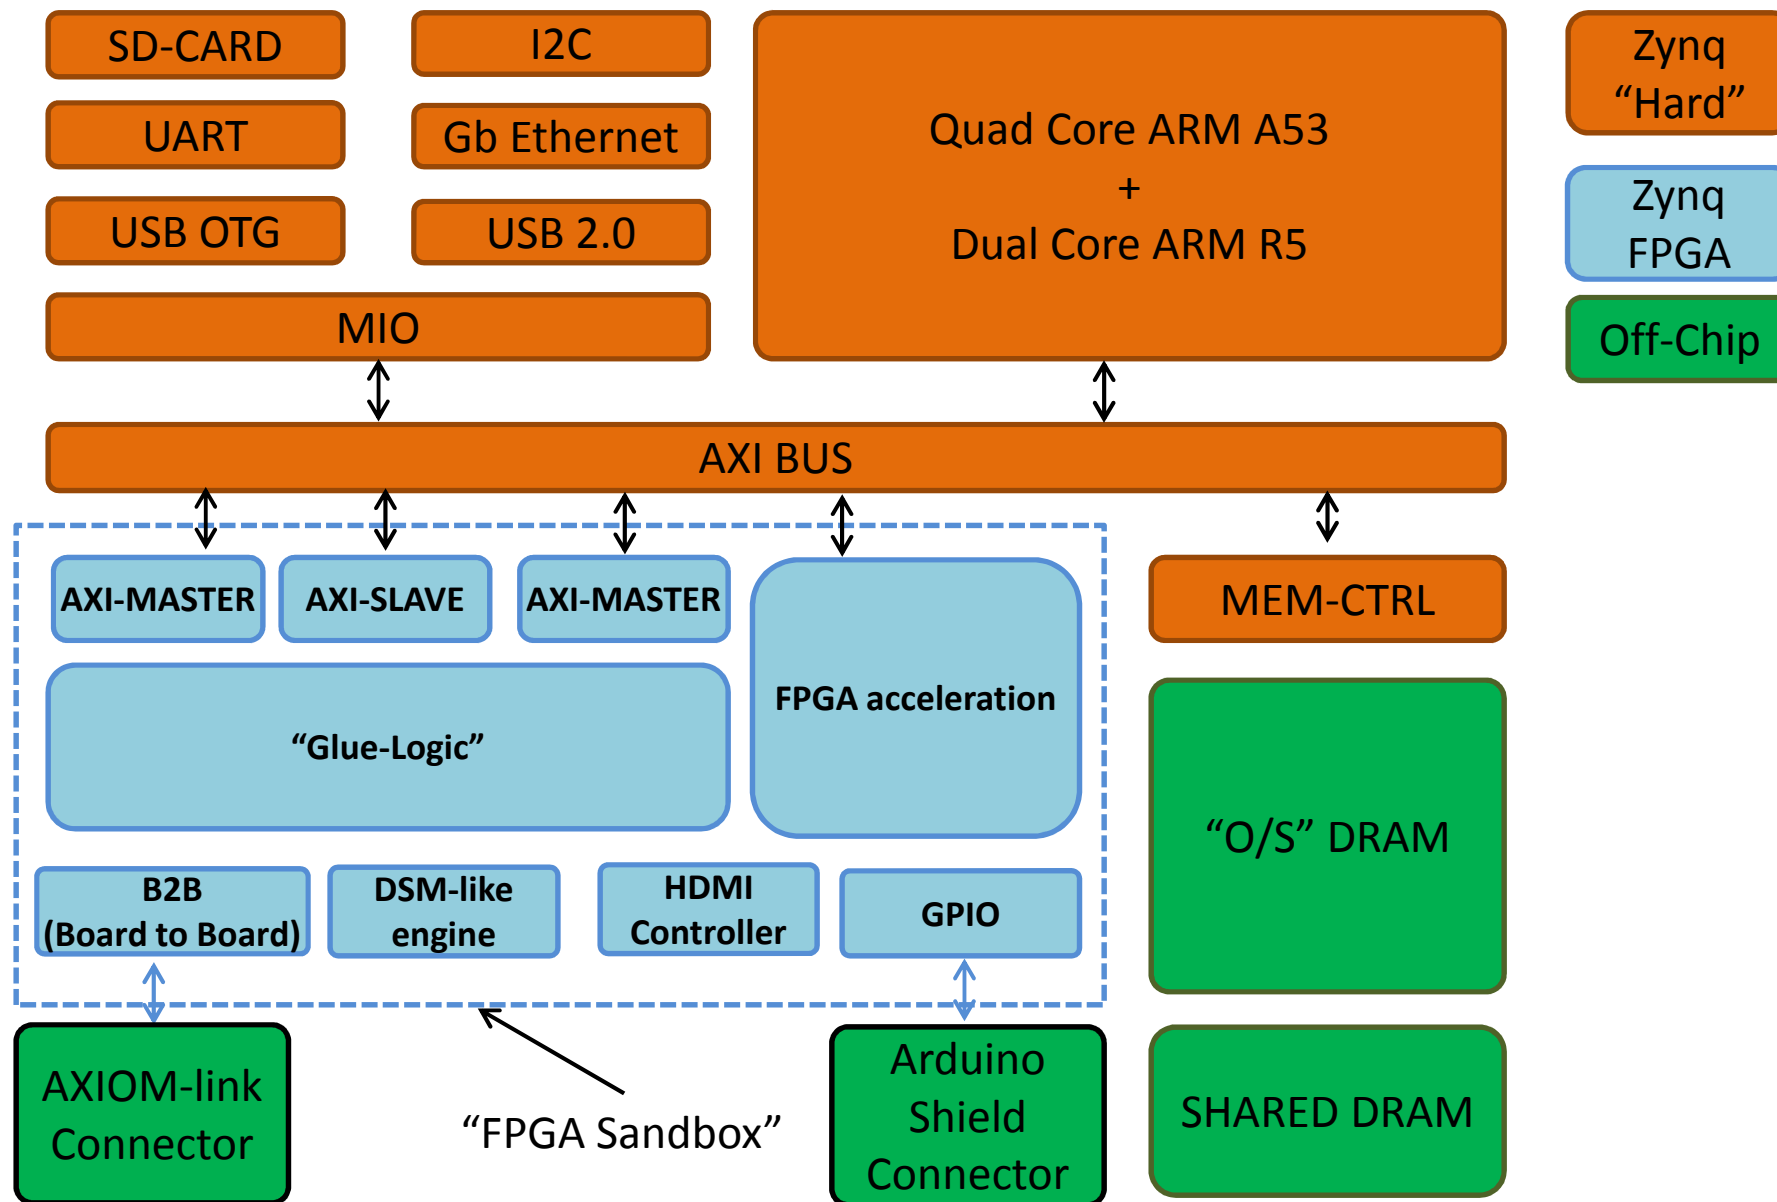
UDOO-NEO




- UDOO set-up and working in less than 2 years
  - Crowd-funding raised 600k \$ in 2 months by 4000+ backers
  - + additional 250k\$ for the UDOO-NEO + 800k\$ for the UDOO-X86



# AXIOM-v2 Architectural Template




# Bench setup @ UNISI






# Layout design (about 10x15 cm)



Disclaimer: subject to changes without notice.


# Testing Environment

- Problem to analyze





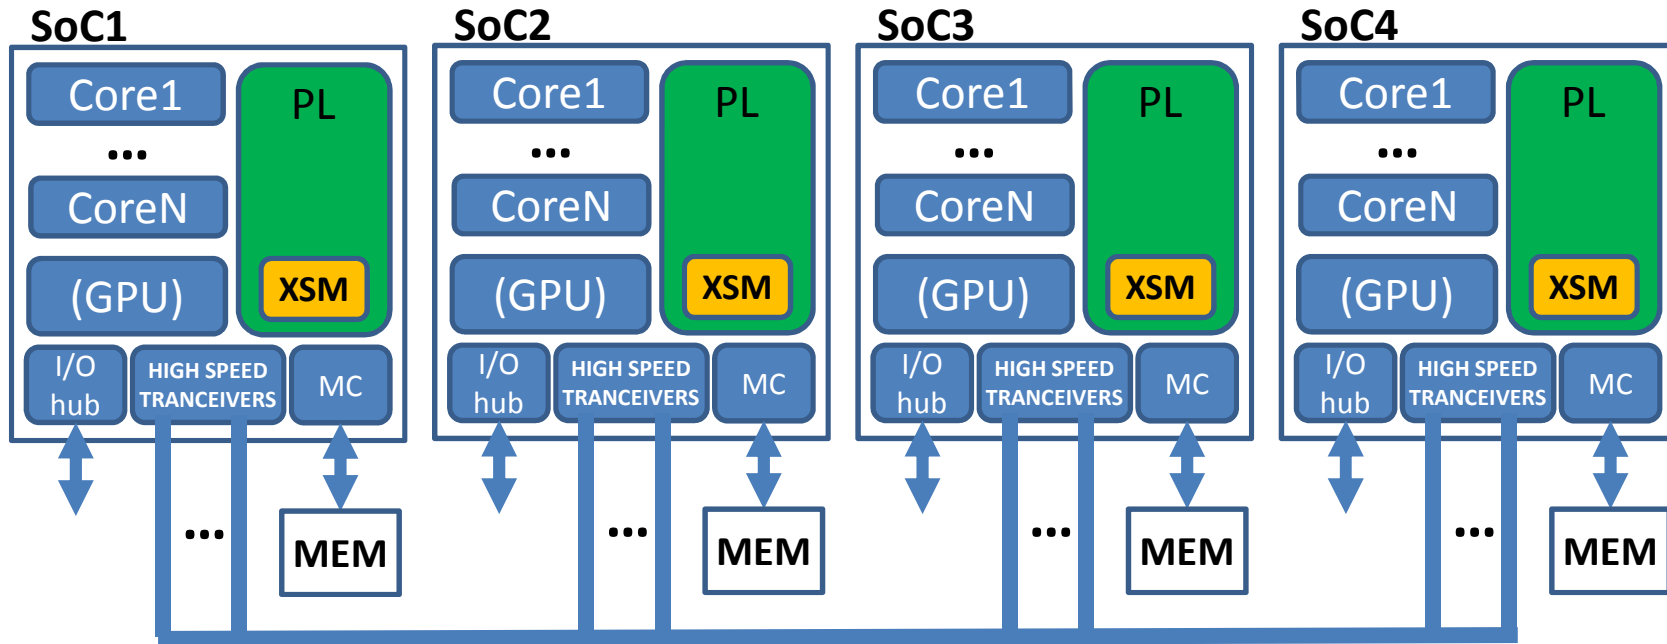
# XSM Low Level



- X-thread (new incarnation of DF-thread)
  - A function that expects no parameters and returns no parameters.
    - The body of this function can refer to any memory location for which it has got the pointer through XSM function calls (e.g., xpreload, xpoststor, xsubscribe, ...). An X-thread is identified by an object of type `xtid_t` (X-thread identifier). In other words:

```
typedef void (*xthread_t)(void)
```
- INPUT\_FRAME, OUTPUT\_FRAME
  - INPUT\_FRAME: A buffer which is allocated in the local memory and contains the input values for the current X-thread.
  - OUTPUT\_FRAME: A buffer which is allocated in the local memory and contains values to be used by other X-threads (consumer X-threads)
- SYNCHRONIZATION\_COUNT
  - A number which is initially set to the number of input values (or events) needed by an X-thread. The SYNCHRONIZATION\_COUNT has to be decremented each time the expected data is written in an OUTPUT\_FRAME.

# 4-board AXIOM System



# Modeled SoC

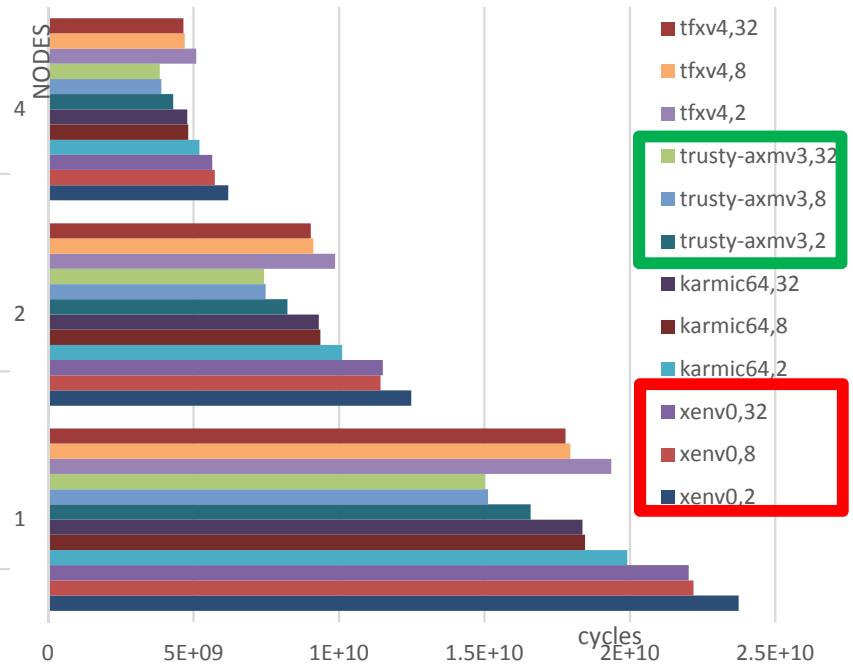
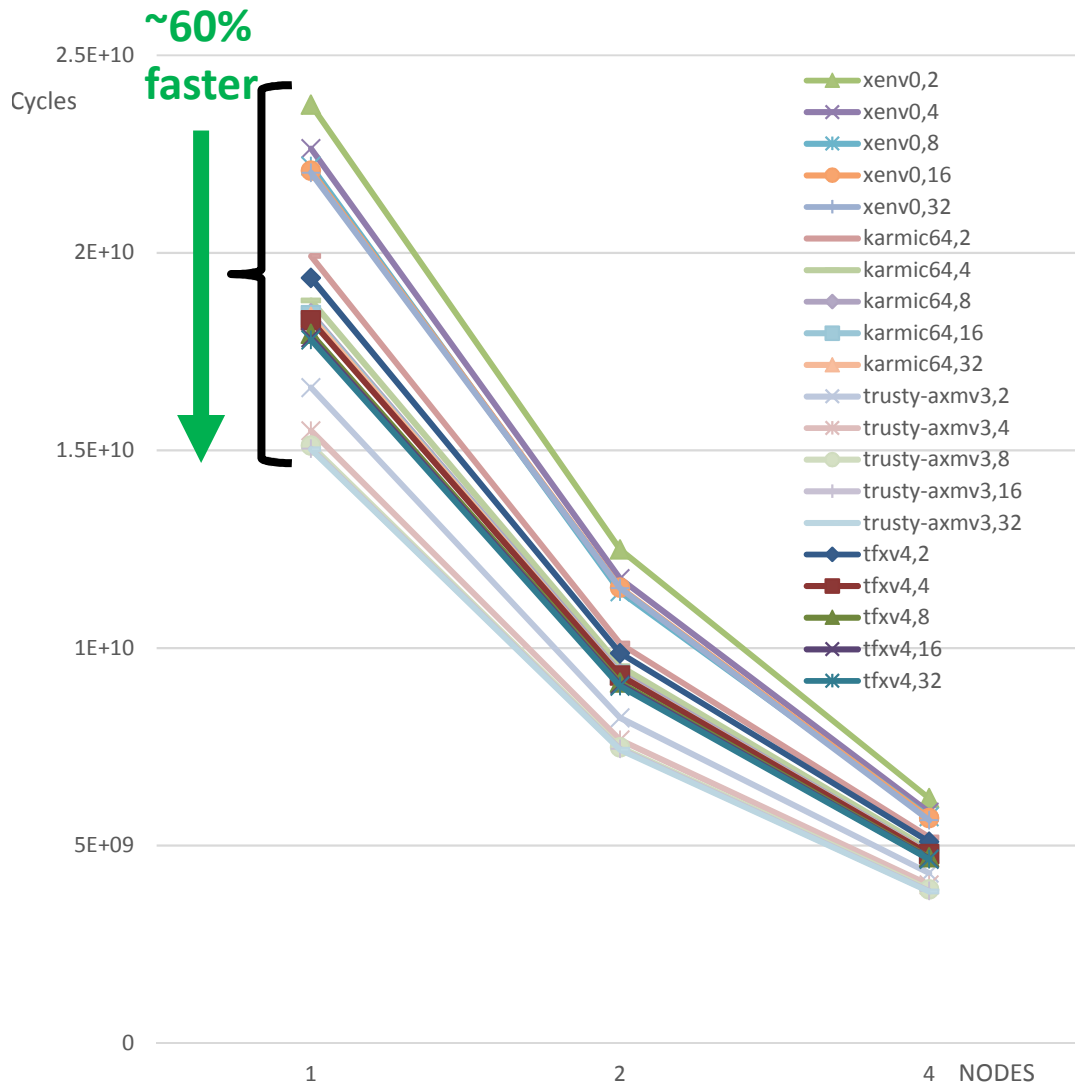
Parameter	Description
SoC	4-cores connected by a shared-bus, IO-hub, MC, high-speed transceivers
Core	1GHz, in-order superscalar
Branch Predictor	two-level (history length=14bits, pattern-history table=16kB, 8-cycle missprediction penalty)
L1 Cache	Private I-cache 32 KB, private D-cache 32 KB, 2 ways, 3-cycle latency
L2 Cache	Private 512 KB, 4 ways, 5-cycle latency
L3 Cache	Shared 4GB, 4 ways, 20-cycle latency
Coherence protocol	MOESI
Main Memory	1 GB, 100 cycles latency
I-L1-TLB, D-L1-TLB	64 entries, full-associative, 1-cycle latency
L2-TLB	512 entries, direct access, 1-cycle latency
Write/Read queues	200 Bytes each, 1-cycle latency

# Matrix-Multiply on COTSon/XSM

<http://cotson.sourceforge.net>

- Some experiments have been performed on the COTSon/XSMML with the following parameters
  - Square Matrix size  $n$  and block size  $b$ :
    - $n=160,200,250,320,400,500,640,800,1000,1280,1600,2000$   $b=5,10,25,50$
    - $n=128,256,512$   $b=8$
  - Different programming models
    - OpenMPI, Cilk
  - Different execution models
    - XSMML, Standard
  - Different Linux Distributions
    - Ubuntu 9.10 (karmic64), 10.10 (tfxv4), 14.04 (trusty-axmv3), 16.04 (xenv0)


# Total Cycles



~60% faster

# XSMML vs OpenMPI vs Cilk

Cycles Comparison



	1 NODE (1CORE)	4 NODES/CORES	XSMML SPEEDUP	
OPENMPI	54281301097	13223633943	3.63	3.49
CILK	45645234077	16738179585	3.05	4.41
XSMML	14941500251	3792215176		

\* For CILK we are using 4 cores instead of 4 nodes





Agile, eXtensible, fast I/O Module for the cyber-physical era

PROJECT ID: 645496

Roberto Giorgi — AXIOM project --- <http://www.axiom-project.eu>

Scalable Embedded Systems: Towards the Convergence of High-Performance and Embedded Computing

