



**Barcelona  
Supercomputing  
Center**  
*Centro Nacional de Supercomputación*



# **LOCA: Embracing Open Source Hardware to create an Open Source Ecosystem**

**John D. Davis, Ph.D.**

October 22, 2020

# An Open HPC Ecosystem?

- Technology Landscape
- HPC History
- LOCA
- Open Source Hardware Projects

# Today's technology trends




Massive penetration of Open Source Software

- IoT (Arduino),
- Mobile (Android),
- Enterprise (Linux),
- HPC (Linux, OpenMP, etc.)



Moore's Law + Power =  
**Specialization**

- More cost effective
- More performant
- Less Power



**SOFTWARE/  
HARDWARE  
CO-DESIGN**

# Today's technology trends




Massive penetration of Open Source Software

- IoT (Arduino),
- Mobile (Android),
- Enterprise (Linux),
- HPC (Linux, OpenMP, etc.)



Moore's Law + Power =  
**Specialization**

- More cost effective
- More performant
- Less Power



New Open Source Hardware  
Momentum from IoT and the  
Edge to HPC


- RISC-V
- OpenPOWER

# Linux History

- 1991: Started development, release
- 1992: X Windows released
- 1998: Adopted by many major companies
- 2004: BSC Supercomputer OS
- 2009: Basis for many new business systems and the cloud
- 2013: Android in 75% of the world smart phones
- 2015: De facto OS for IoT, mobile, cloud, and supercomputers

# HPC Today

- Europe has led the way in defining a common open HPC software ecosystem
- **Linux** is the de facto standard OS despite proprietary alternatives
- Software landscape from Cloud to IoT already enjoys the benefit of open source
- Open source provides:
  - A common platform, specification and interface
  - Accelerates building new functionality by leveraging existing components
  - Lowers the entry barrier for others to contribute new components
  - Crowd-sources solutions for small and larger problems
- **What about Hardware and in particular, the CPU?**
- Inhibits opportunities in holistic co-design
  - Facing barrier to innovation
  - Being able to have a conversation or not




# Mont-Blanc HPC Stack for ARM



## Industrial applications



## Applications



## System software



## Hardware



# The Exascale Race – The Japanese example



## Co-design from Apps to Architecture

- **Architectural Parameters to be determined**
  - #SIMD, SIMD length, #core, #NUMA node, O3 resources, specialized hardware
  - cache (size and bandwidth), memory technologies
  - Chip die-size, power consumption
  - Interconnect
- **We have selected a set of target applications**
- **Performance estimation tool**
  - Performance projection using Fujitsu FX100 execution profile to a set of arch. parameters.
- **Co-design Methodology (at early design phase)**
  1. Setting set of system parameters
  2. Tuning target applications under the system parameters
  3. Evaluating execution time using prediction tools
  4. Identifying hardware bottlenecks and changing the set of system parameters

Target applications representatives of almost all our applications in terms of computational methods and communication patterns in order to design architectural features.

Target Application	
Program	Brief description
① GENESIS	MD for proteins
② Genomon	Genome processing (Genome alignment)
③ GAMERA	Earthquake simulator (FEM in unstructured & structured grid)
④ NICAM+LETK	Weather prediction system using Big data (structured grid stencil & ensemble Kalman filter)
⑤ NTChem	molecular electronic (structure calculation)
⑥ FFB	Large Eddy Simulation (unstructured grid)
⑦ RSDFT	an ab-initio program (density functional theory)
⑧ Adventure	Computational Mechanics System for Large Scale Analysis and Design (unstructured grid)
⑨ CCS-QCD	Lattice QCD simulation (structured grid Monte Carlo)

# The Exascale Race – The Japanese example

**"Post-K" Arm64fx Processor is...**


- an Many-Core ARM CPU...
  - 48 compute cores + 2 or 4 assistant (OS) cores
  - Brand new core design by Fujitsu
  - Near Xeon-Class Integer performance core
  - ARM V8.2 --- 64bit ARM ecosystem
- ...but also a GPU-like processor
  - SVE 512 bit vector extensions (ARM & Fujitsu)
  - Integer (1, 2, 4, 8 bytes) + float (16, 32, 64 bytes)
  - Cache + access localization (sector cache) – similar to scratchpad
  - IBM2 OPI – Massive Mem BW (1TByte/s, Bytes/DPF ~0.4 same as K)
  - Streaming memory access, strided access, scatter/gather etc.
  - Intra-chip barrier synch. and other memory enhancing features
  - 40GByte/s Tofu-.D interconnect + PCIe 3

**#1 TOP500**

**GPU-like High performance in HPC, AI/Big Data, Auto Driving...**

RIKEN logo, RCCS logo, and a date stamp 20018/3/13 are visible on the slide.

# The European Open System Stack



# Why Europe needs its own processor

Processors now control almost every aspect of our lives

**Security** (back doors, etc.)

Possible **future restrictions on exports to EU** due to increasing protectionism

**A competitive EU supply chain** for HPC technologies will create jobs and growth in Europe

A group of researchers showed how a Tesla Model S can be hacked and stolen in seconds using only \$600 worth of equipment



NSA May Have Backdoors Built Into Intel And AMD Processors



A jet sale to Egypt is being blocked by a US regulation, and France is over it



Google 'suspends some business with Huawei' after US blacklist



Amazon and Super Micro urge Bloomberg to retract 'unsupported' spy chip report




Car Hacking Remains a Very Real Threat as Autos Become Ever More Loaded With Tech



The US Cloud Act v The EU's GDPR - Data Privacy & Security



# RISC-V is democratizing chip-design




- ➔ More and more global IT actors are adopting RISC-V architectures to be vendor independent
  - ➔ Google
  - ➔ Amazon
  - ➔ Western Digital
  - ➔ Alibaba
- ➔ And of course the entire IoT ecosystem for lower performance, lower energy applications.
- ➔ Major opportunity for ICT industry also in Spain

# RISC-V Worldwide Growth Projection

CAGR for RISC-V CPU cores between 2018 and 2025

	Computer	Consumer	Communications	Transportation	Industrial	Other RISC-V	Total
<b>Adv. Perf Multicore SoC</b>	60%	73%	224%	153%	106%	171%	144%
<b>Value Multicore SoC</b>	60%	92%	222%	159%	122%	201%	177%
<b>Basic SoC</b>	63%	115%	217%	166%	127%	215%	190%
<b>FPGA</b>	62%	72%	190%	163%	102%	176%	149%
<b>Total</b>	61%	81%	209%	160%	110%	185%	158%



President of Semico Research, Jim Feldhan, commented: “Based on the already sizeable adoption of RISC-V, we forecast that the market will consume a total of 62.4 billion RISC-V cores by 2025, signaling that RISC-V is on the fast track to mainstream adoption.”


# RISC-V History

- 2010: Started development and initial proposal
- 2015: RISC-V Foundation formed
- 2019: Adopted by many major companies
  - Starting in the embedded market with already over 1 Billion CPUs
- 2020
  - RISC-V Foundation moves to Switzerland

The time is now to embrace and support RISC-V from IoT to HPC

# HPC Tomorrow

- Europe can lead the way to a completely **open SW/HW stack for the world**
- RISC-V provides the open source hardware alternative to dominating proprietary non-EU solutions
- Europe can achieve complete technology independence with these foundational building blocks
- Currently at the same early stage in HW as we were with SW when Linux was adopted many years ago
- **RISC-V can unify, focus, and build a new microelectronics industry in Europe.**



# Open Source Beyond 2020



**“Open Source has become mainstream across all software industry during the past 10 years. To a large extent, open software re-use has proven economically efficient. The level of maturity of Open Source Hardware (OSH) remains far lower than that of Open Source Software (OSS). However, business ecosystems for OSH are developing fast so that OSH could constitute a cornerstone of the future Internet of Things (IoT) and the future of computing.”**

*- DG Connect & DG IT Workshop, Brussels, Nov. 14-15, 2019*

Source: <https://ec.europa.eu/digital-single-market/en/news/workshop-about-future-open-source-software-and-open-source-hardware>



# LOCA @ BSC



LOCA

European Laboratory for Open  
Computer Architecture

# BSC full stack



HPC Applications

Specialization using  
HW/SW Co-Design

HPC Hardware



European Laboratory for Open  
Computer Architecture



# LOCA Goals


- Mechanism to extend open source ecosystem to include H/W
  - Add H/W expertise to BSC and European partners, leverage S/W expertise
  - Provide proven/usable Open Source H/W
  - Intersection of academia and industry
  - Open European IP repository → rapid implementation
  - Catalyst to reinvigorate European ICT industry
  - **Global** collaboration and training center
  - Incubator for European IP



# Open



- Participation from Academia across Europe
  - Student training
  - Affiliate Labs at other institutions
  - Research collaboration
  - Repository for research work



- Participation from European and global industry
  - Engagement with European Research and Development
  - Access to European talent
  - European visibility and branding
  - Innovation incubator

# European Collaboration & Education



## Casteller

(human tower)



**Barcelona  
Supercomputing  
Center**  
Centro Nacional de Supercomputación

Traditional chip design is done in a Master/Apprentice environment

LOCA recreates this environment by bringing in Masters from industry to collaborate with a variety of people, pushing beyond RTL

Professors, students, and industry veterans all together


Ideal sandbox for creative and innovative work

Research and Design to chip fabrication




# From IoT, Edge Computing, Clouds to Supercomputers

## Accelerator Use Model



# Current BSC RISC-V European Accelerator & CPU Landscape




# Rebuilding the European CPU Industry

Production

Closed + Open

Open

Research



Design

Exascale

eProcessor

EPI RISC-V  
PILOT

MEEP

EPI

IP/Infra

IP/Infra

IP/Infra

IP/Infra

# The European Processor Initiative

- In the same way BSC led the development of ARM processors for HPC in the various MontBlanc projects, now it leads the RISC-V HPC accelerator development in EPI
- EPI is a 100% funded EuroHPC project (120 M€ ) to develop European processor technology by 2022
- BSC was the original initiator of EPI and most active proponent in the scientific and technical community
- EPI is led by Atos/Bull with 28 partners from leading HPC industrial and academic centers

# EPI Partners



Rolls-Royce  
Motor Cars Limited



**Barcelona  
Supercomputing  
Center**  
Centro Nacional de Supercomputación



ALMA MATER STUDIORUM  
UNIVERSITÀ DI BOLOGNA



CHALMERS



UNIVERSITÀ DI PISA



UNIVERSITY OF ZAGREB  
FACULTY OF  
ELECTRICAL  
ENGINEERING  
AND COMPUTING



PROVE & RUN



**Barcelona  
Supercomputing  
Center**  
Centro Nacional de Supercomputación



# EPAC architecture

- **Objective:** Develop & demonstrate **fully European** processor IPs based on the **RISC-V ISA**. Build on existing EU IP, leverage EU background and vision
- Provide a very **low power** and high computing throughput **accelerator for HPC & Emerging -> Automotive**

## Main Contributors:

### ■ RISC-V in the Tile:

- 8 Vector cores, 8 STX, 1 VRP

### ■ Vector Core :


- Spain + Italy + Croatia
- C: RISC-V Scalar core
- V: RISC-V Vector core: 8 lanes

### ■ STX

- Switzerland + Germany
- RISC-V NTX (AI/ML/DL) + Stencil

### ■ VRP

- France
- RISC-V + Extended precision FPU




# Rebuilding the European CPU Industry

Production

Closed + Open

Open

Research



Design

Exascale

eProcessor

EPI RISC-V  
PILOT

**MEEP**

EPI

IP/Infra

IP/Infra

IP/Infra

IP/Infra

# MareNostrum 5

## A European pre-exascale supercomputer

- 200 Petaflops peak performance ( $200 \times 10^{15}$ )
- Experimental platform (MEEP) to create supercomputing technologies “made in Europe”

Hosting Consortium:

Spain Portugal Turkey Croatia Ireland



|



# MEEP: MareNostrum Experiment Exascale Platform

MEEP is a flexible FPGA-based emulation platform that will explore hardware/software co-designs for Exascale Supercomputers and other hardware targets, based on European-developed IP. The project provides two functions:



An evaluation platform of pre-silicon IP and ideas, at speed and scale.



A software development and experimentation platform to enable software readiness for new hardware.



## MEEP

MareNostrum Experimental  
Exascale Platform



The MEEP project has received funding from the European High-Performance Computing Joint Undertaking Joint Undertaking (JU) under grant agreement No 946002. The JU receives support from the European Union's Horizon 2020 research and innovation programme and Spain, Croatia, Turkey.

MEEP will enable software development, accelerating software maturity, compared to the limitations of software simulation. IP can be tested and validated before moving to silicon, saving time and money.

# Accelerated Compute and Memory Engine (ACME) Architecture






# Rebuilding the European CPU Industry

Production

Closed + Open

Open

Research



Design

Exascale

EPI RISC-V  
PILOT

**eProcessor**

MEEP

EPI


IP/Infra

IP/Infra


IP/Infra

IP/Infra

# eProcessor Stack



# eProcessor System diagram




# Rebuilding the European CPU Industry

Production

Closed + Open

Open

Research



Design

Exascale

**EPI RISC-V  
PILOT**

eProcessor

MEEP

EPI


IP/Infra

IP/Infra



IP/Infra

IP/Infra

# The European PILOT



# The European PILOT Accelerator Architecture




# Rebuilding the European CPU Industry

Production

Closed + Open

Open

Research



Design

**Exascale**

EPI RISC-V  
PILOT

eProcessor

MEEP


EPI

IP/Infra


IP/Infra

IP/Infra

IP/Infra



# The Future: Flagship RISC-V Exascale Accelerator



# A Successful EC Roadmap for the Future

- Embrace Open Source Hardware
  - Build infrastructure to support open source
  - Create an environment that links research to the ICT industry
  - Supported by DG Connect and other funding instruments
- Teach, train, and collaborate
- Leverage the Global Technology ecosystem
- Requires **SIGNIFICANT** funding for programs to build made in Europe IP
  - Many **focused** projects: Horizon Europe, Digital Europe, ...
  - Many **focused** teams: Collaborative IP and TRL development
  - Total and integrated vision for the future
  - Large Accelerator and CPU investment (> 1B€)

# The Future is Wide Open!

- There is an urgent need, from mobile phones to supercomputers: more compute at lower power
- The RISC-V ecosystem is in the nascent period where it can become the de facto open hardware platform of the future
  - An opportunity for Europe to lead the charge to creating a full stack solution for everything, from supercomputers down to IoT devices
- Our main goal: Create European chips that meet the needs of future European and global markets across HPC, cloud, automotive, mobile to IoT
- **This is the framework for the Exascale Supercomputing Initiative at BSC**



# Focus, Freedom, and Forward

- BSC is embracing a Open Source Hardware to complement OSS
  - RISC-V has the momentum to succeed as the Open Source ISA, much like Linux
- BSC is building the infrastructure to support the future of computing
  - Combining applications, system software, hardware design and hardware
- BSC can organize and coordinate full stack efforts
  - More resources = faster time to success
  - More partners (Academia/Industry)
  - Research/Resource multiplier
- Europe and the world can unite around the BSC Vision to move forward faster, together



**Barcelona  
Supercomputing  
Center**  
*Centro Nacional de Supercomputación*

**BSC is Made in Europe**



# LOCA is Bigger than Barcelona!



**Casteller**  
(human tower)



European Laboratory for Open  
Computer Architecture



European  
Open  
Hardware,  
Architecture  
Initiative

We must stand on the shoulders of giants to build great things.  
**We are assembling many giants and hope you can join us.**



**Barcelona  
Supercomputing  
Center**  
Centro Nacional de Supercomputación



# Thank you

[john.davis@bsc.es](mailto:john.davis@bsc.es)