



FPL Conference
Barcelona 09/09/2019



**Barcelona
Supercomputing
Center**
Centro Nacional de Supercomputación



EXCELENCIA
SEVERO
OCHOA

Barcelona Supercomputing Center

Prof. Mateo Valero,
BSC director



European Research Council

Supporting top researchers
from anywhere in the world




Barcelona Supercomputing Center Centro Nacional de Supercomputación

BSC-CNS objectives



Supercomputing services
to Spanish and EU researchers



R&D in Computer, Life, Earth and
Engineering Sciences



PhD programme, technology
transfer, public engagement

BSC-CNS is
a consortium
that includes

Spanish Government

60%



Catalan Government

30%

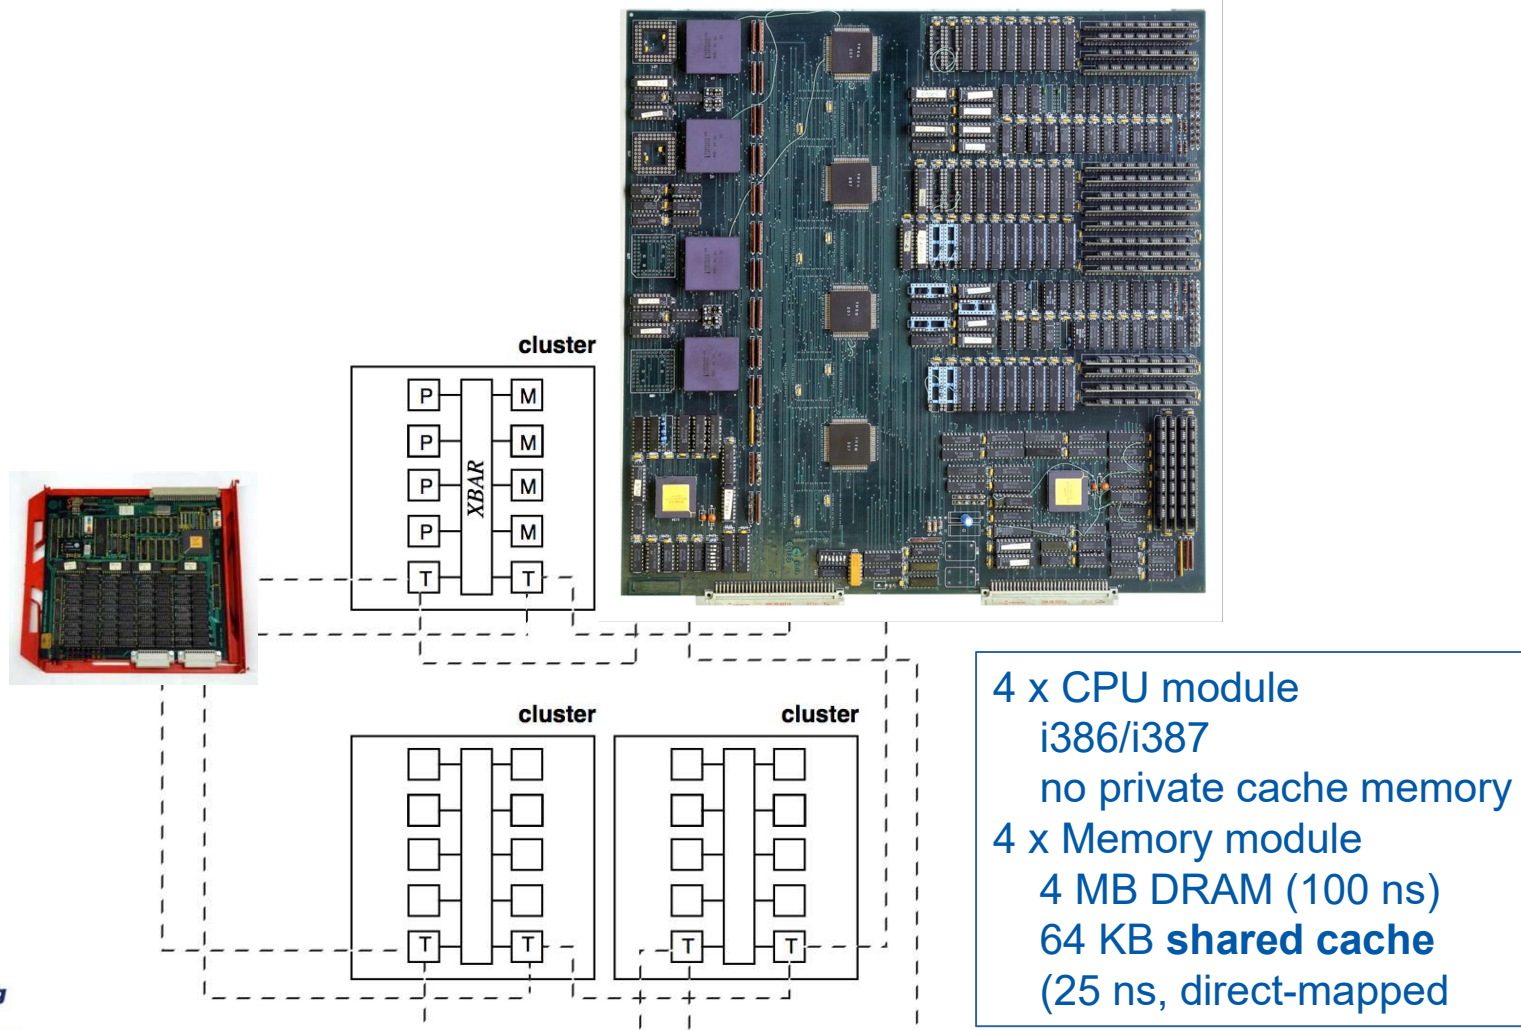


Univ. Politècnica de Catalunya (UPC)

10%



ICT386 multiprocessor architecture

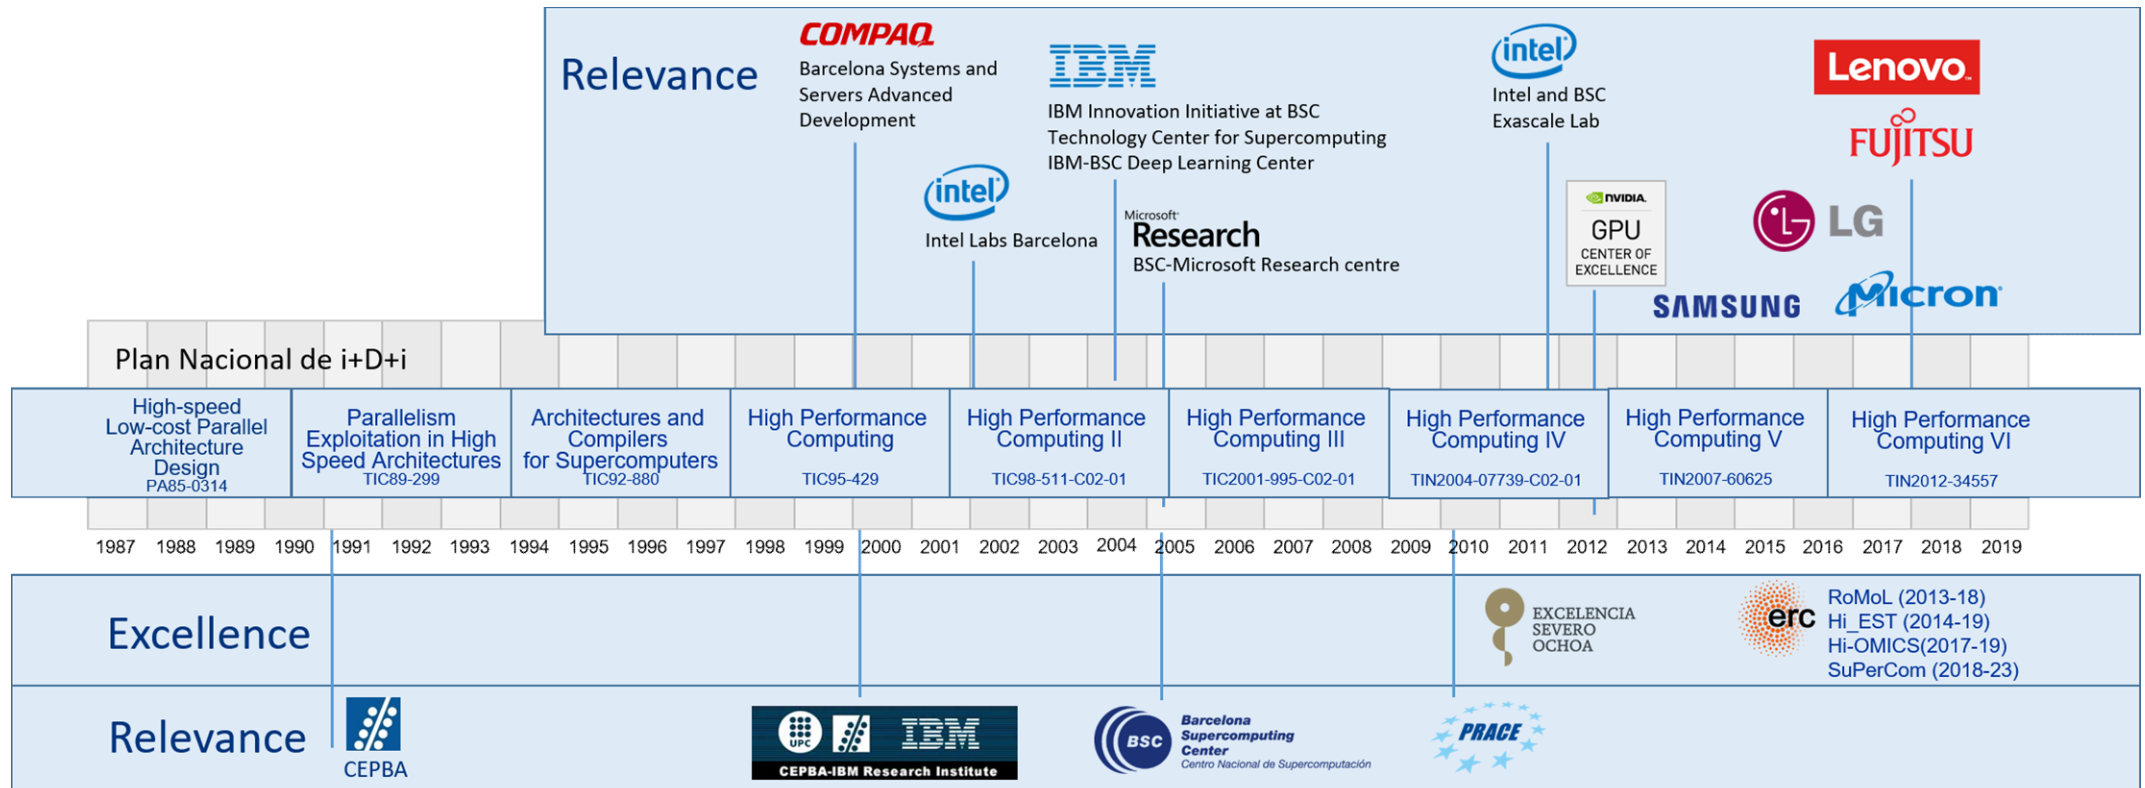


Professor Tomas Lang




Our Origins...

High-performance Computing group @ Computer Architecture Department (UPC)




Latency Has Been a Problem from the Beginning... ☹️



- Feeding the pipeline with the right instructions:
 - Software trace cache (ICS'99)
 - Prophet/Critic Hybrid Branch Prediction (ISCA'04)
- Locality/reuse
 - Cache Memory with Hybrid Mapping (IASTED87). Victim Cache ☺️
- A novel prefetching mechanism that boosts software prefetching (ICS'01)
- Virtual-Physical Registers (HPCA'98)
- Kilo Instruction Processors (ISHPC03,HPCA'06, ISCA'08)
 - Distant Parallelism (ICS99)

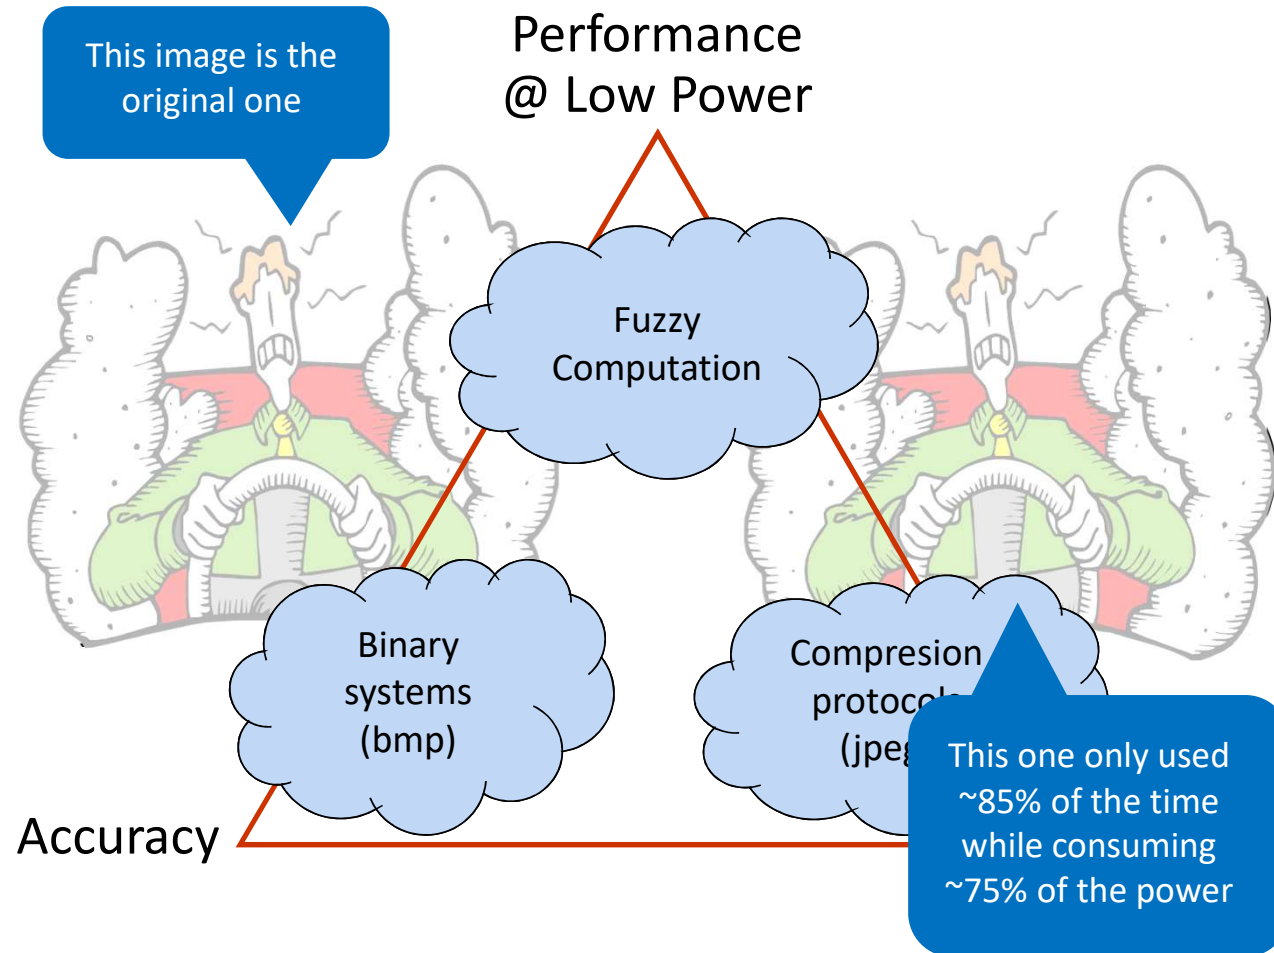
Memory Wall

... and the Power Wall Appeared Later ☹️☹️☹️




- Better Technologies
- Two-level organization (Locality Exploitation)
 - Register file for Superscalar (ISCA'00)
 - Instruction queues (ICCD'05)
 - Load/Store Queues (ISCA'04)
- Direct Wakeup, Pointer-based Instruction Queue Design (ICCD'04, ICCD'05)
- Content-aware register file (ISCA'09)
- Fuzzy computation (ICS'01, IEEE CAL'02, IEEE-TC'05). Currently known as Approximate Computing (HPC) and Reducing Precision (ML) 😊

Fuzzy computation




Vector Architectures... Memory Latency and Power ☺☺☺

- Out-of-Order Access to Vectors (ISCA 1992, ISCA 1995)
- Command Memory Vector (PACT 1998)
 - In-memory computation
- Decoupling Vector Architectures (HPCA 1996)
 - Cray X1
- Out-of-order Vector Architectures (Micro 1996)
- Multithreaded Vector Architectures (HPCA 1995)
- SMT Vector Architectures (HICS 1997, IEEE Micro 1997)
- Vector register-file organization (PACT 1997)
- Vector Microprocessors (Micro 1999, SPAA 2001)
- Architectures for Short Vectors (PACT 1997, ICS 1998)
 - Knights Corner
- Vector Architectures for Multimedia (HPCA 2001, Micro 2002)
- High-Speed Buffers Routers (Micro 2003, IEEE TC 2006)
- Vector Architectures for Data-Base (Micro 2012, HPCA2015, ISCA2016)




Awards in Computer Architecture



Eckert-Mauchly: IEEE Computer Society and ACM: (...) *“For extraordinary leadership in building a world class computer architecture research center, for seminal contributions in the areas of vector computing and multithreading, and for pioneering basic new approaches to instruction-level parallelism.”* **June 2007**




Seymour Cray: IEEE Computer Society: (...) *“In recognition of seminal contributions to vector, out-of-order, multithreaded, and VLIW architectures.”* **November 2015**



Charles Babbage: IEEE Computer Society: (...) *“For contributions to parallel computation through brilliant technical work, mentoring PhD students, and building an incredibly productive European research environment.”* **April, 2017**


OmpSs: data-flow execution of sequential programs



```
void Cholesky( float *A ) {
    int i, j, k;
    for (k=0; k<NT; k++) {
        spotrf (A[k*NT+k]) ;
        for (i=k+1; i<NT; i++)
            strsm (A[k*NT+k], A[k*NT+i]);
        // update trailing submatrix
        for (i=k+1; i<NT; i++) {
            for (j=k+1; j<i; j++)
                sgemm (A[k*NT+i], A[k*NT+j], A[j*NT+i]);
            ssyrk (A[k*NT+i], A[i*NT+i]);
        }
    }
}
```

```
#pragma omp task inout ([TS][TS]A)
void spotrf (float *A);
#pragma omp task input ([TS][TS]A) inout ([TS][TS]C)
void ssyrk (float *A, float *C);
#pragma omp task input ([TS][TS]A,[TS][TS]B) inout ([TS][TS]C)
void sgemm (float *A, float *B, float *C);
#pragma omp task input ([TS][TS]T) inout ([TS][TS]B)
void strsm (float *T, float *B);
```

Decouple how we write applications from how they are executed


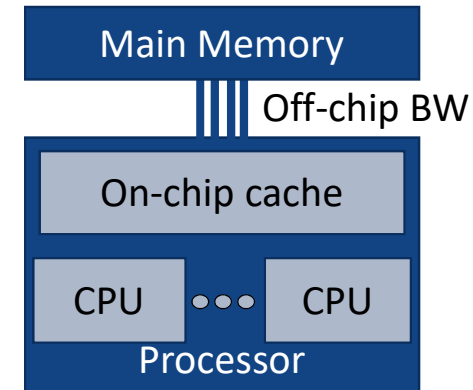


Clean offloading to hide architectural complexities

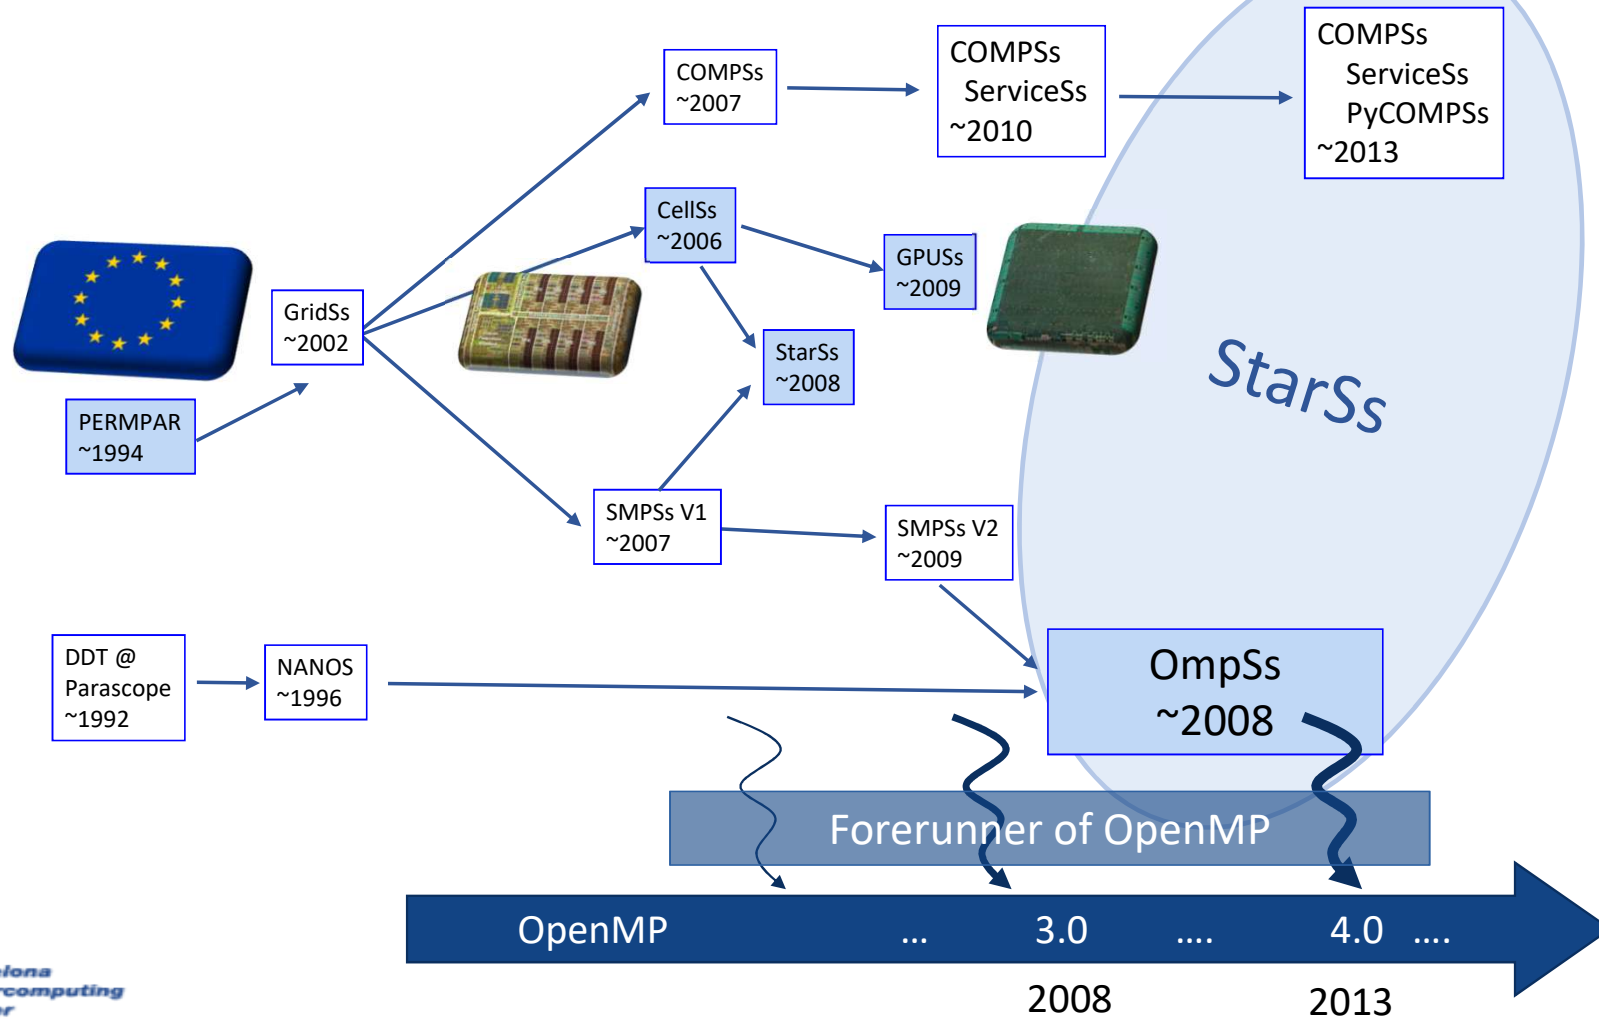


OmpSs: Potential of Data Access Info

- Flat global address space seen by programmer
- Flexibility to dynamically traverse dataflow graph “optimizing”
 - Concurrency. Critical path
 - Memory access: data transfers performed by run time
- Opportunities for automatic
 - Prefetch
 - Reuse
 - Eliminate antidependences (rename)
 - Replication management
 - Coherency/consistency handled by the runtime
 - Layout changes

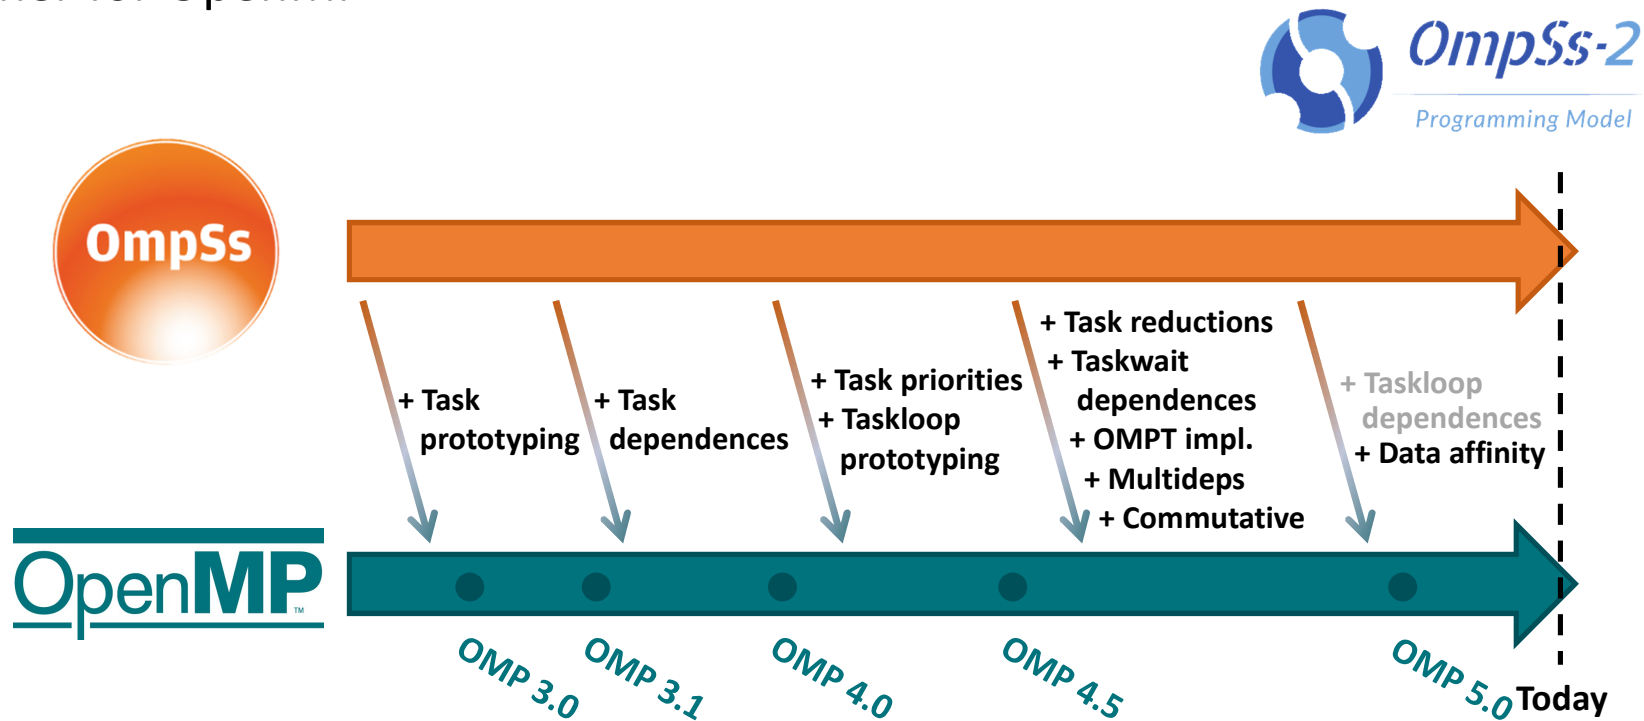


History / Strategy



OmpSs

- A forerunner for OpenMP



1st European researcher to receive Ken Kennedy Award



2017 Ken Kennedy Award Recipient

"For his contributions to programming models and performance analysis tools for High Performance Computing."

CSRankings: Computer Science Rankings

- CSRankings is a metrics-based ranking of top computer science institutions around the world:
 - <http://csranks.org/#/index?all>
- This ranking is designed to identify institutions and faculty actively engaged in research across a number of areas of computer science, based on the number of publications by faculty that have appeared at the most selective conferences in each area of computer science.
- All publication data is from [DBLP](#) (updated monthly; last update May 14, 2019).

Topic: High-Performance Computing in Europe from 1980

• #	Institution	Count	Faculty
• 1	▶ Polytechnic University of Catalonia	27.7	18
• 2	▶ ETH Zurich	17.9	5
• 3	▶ TU Munich	14.9	8
• 4	▶ VU Amsterdam	10.4	8
• 5	▶ Technion	7.7	9
• 6	▶ Ecole Normale Supérieure de Lyon	7.6	9
• 7	▶ University of Edinburgh	5.1	5



Topic: High-Performance Computing in the **world** from 1980

• #	Institution	Count	Faculty
• 1	▶ Ohio State University	36.3	11
• 2	▶ Univ. of Illinois at Urbana-Champaign	33.7	19
• 3	▶ Polytechnic University of Catalonia	27.7	18
• 4	▶ Georgia Institute of Technology	26.8	22
• 5	▶ University of Minnesota	26.5	10
• 6	▶ University of Chicago	25.9	7
• 7	▶ Purdue University	22.5	15
• 8	▶ Indiana University	22.3	10
• 9	▶ ETH Zurich	17.9	5
• 10	▶ University of California - Berkeley	17.7	11

Topic: Computer Architecture in Europe from 1980

• #	Institution	Count	Faculty
• 1	▶ Polytechnic University of Catalonia	29.5	14
• 2	▶ EPFL	22.1	8
• 3	▶ ETH Zurich	19.5	6
• 4	▶ University of Edinburgh	10.9	7
• 5	▶ Technion	10.1	10
• 6	▶ Tel Aviv University	4.8	4
• 7	▶ University of Cambridge	4.8	10



Topic: Computer Architecture in the world from 1980

• #	Institution	Count	Faculty
• 1	▶ University of Michigan	77.6	22
• 2	▶ Univ. of Illinois at Urbana-Champaign	61.3	22
• 3	▶ University of Wisconsin - Madison	61.1	17
• 4	▶ Stanford University	50.6	19
• 5	▶ Georgia Institute of Technology	33.2	17
• 6	▶ Princeton University	32.5	12
• 7	▶ Polytechnic University of Catalonia	29.5	14
• 8	▶ University of Washington	29.0	16
• 9	▶ Pennsylvania State University	27.7	13
• 10	▶ University of California - San Diego	26.5	14

Topics: “High-Performance Computing + Computer Architecture”
in the **world** from 1980

• #	Institution	Count	Faculty
• 1	▶ Univ. of Illinois at Urbana-Champaign	45.5	31
• 2	▶ University of Wisconsin - Madison	30.6	20
• 3	▶ Georgia Institute of Technology	29.8	30
• 4	▶ Polytechnic University of Catalonia	28.6	22
• 5	▶ Stanford University	25.5	20
• 6	▶ University of Michigan	23.1	25
• 7	▶ University of Chicago	22.1	14
• 8	▶ Purdue University	20.5	21
• 9	▶ ETH Zurich	18.6	8
• 10	▶ University of California - Berkeley	18.3	22

Our origins...



Barcelona Supercomputing Center
Centro Nacional de Supercomputación



Parsys Multiprocessor



Parsytec CCI-8D
4.45 Gflop/s



Compaq GS-140
12.5 Gflop/s




Compaq GS-160
23.4 Gflop/s



BULL NovaScale 5160
48 Gflop/s



Maricel
14.4 Tflops, 20 KW




Research prototypes



Convex C3800



SGI Origin 2000
32 Gflop/s



SGI Altix 4700
819.2 Gflops



SL8500
6 Petabytes



Transputer cluster



Connection Machine CM-200
0.64 Gflop/s



IBM RS-6000 SP & IBM p630
192+144 Gflop/s




IBM PP970 / Myrinet
MareNostrum
42.35, 94.21 Tflop/s

1985 1986 1987 1988 1989 1990 1991 1992 1993 1994 1995 1996 1997 1998 1999 2000 2001 2002 2003 2004 2005 2006 2007 2008 2009 2010



Barcelona Supercomputing Center
Centro Nacional de Supercomputación

The Killer Mobile processors™



« Microprocessors killed the Vector supercomputers

« They were not faster ...

« ... but they were significantly cheaper and greener

« History may be about to repeat itself ...

« Mobile processor are not faster ...

« ... but they are significantly cheaper and greener

HiPEAC Objectives

- ❑ to help companies identify and select the best architecture solutions for scaling up high-performance embedded processors in the coming years
- ❑ to unify and focus academic research efforts through a processor architecture and compiler research roadmap
- ❑ to address the increasingly slow progression of *sustained* processor performance by jointly developing processor architecture and compiler optimizations
- ❑ to explore novel approaches for achieving regular and smooth scaling up of processor performance with technology, and to explore the impact of a wide range of post-Moore's law technologies on processor architecture and programming paradigms.

Steering Committee





Network of c.2,000 European R+D experts in advanced computing: **high-performance** and **embedded** architecture and compilation

720 members, 449 affiliated members and 871 affiliated PhD students from **430** institutions in **46** countries.



**Barcelona
Supercomputing
Center**
Centro Nacional de Supercomputación

hipeac.net



HiPEAC


HiPEAC has received funding from the European Union's Horizon2020 research and innovation programme under grant agreement number 779656.

Barcelona Supercomputing Center Centro Nacional de Supercomputación

BSC-CNS objectives



Supercomputing services
to Spanish and EU researchers



R&D in Computer, Life, Earth and
Engineering Sciences



PhD programme, technology
transfer, public engagement

BSC-CNS is
a consortium
that includes

Spanish Government

60%



Catalan Government

30%

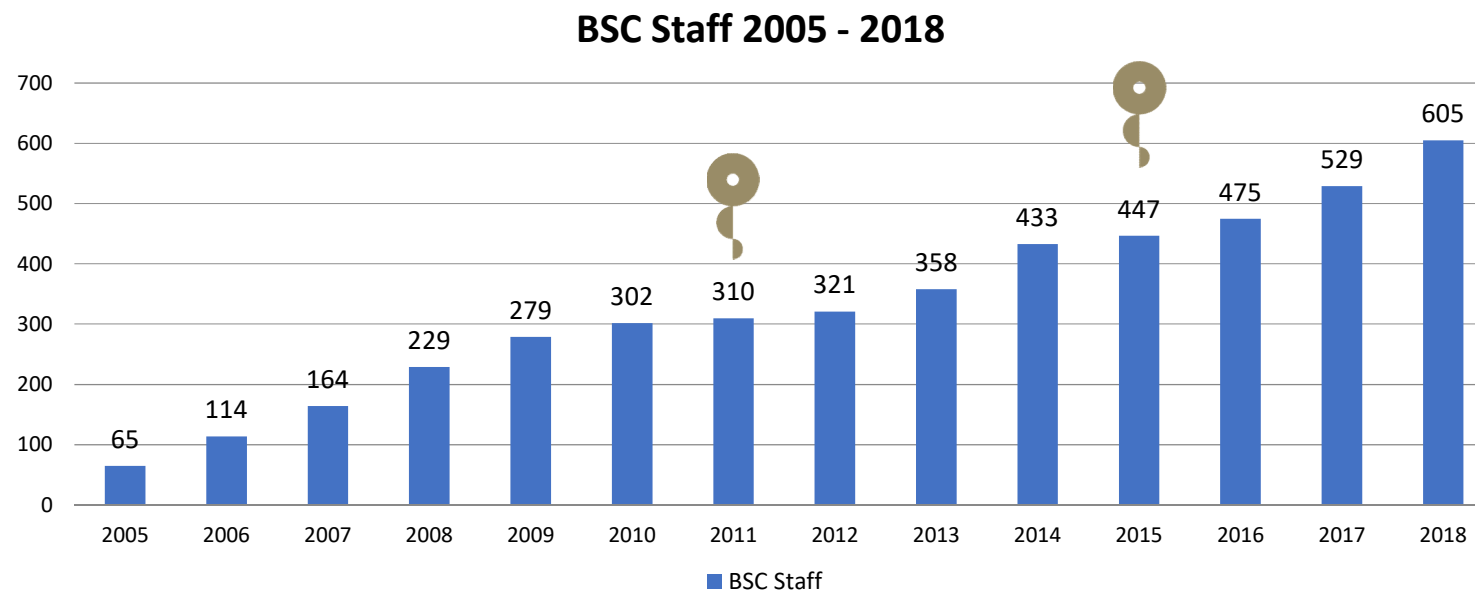


Univ. Politècnica de Catalunya (UPC)

10%



People evolution

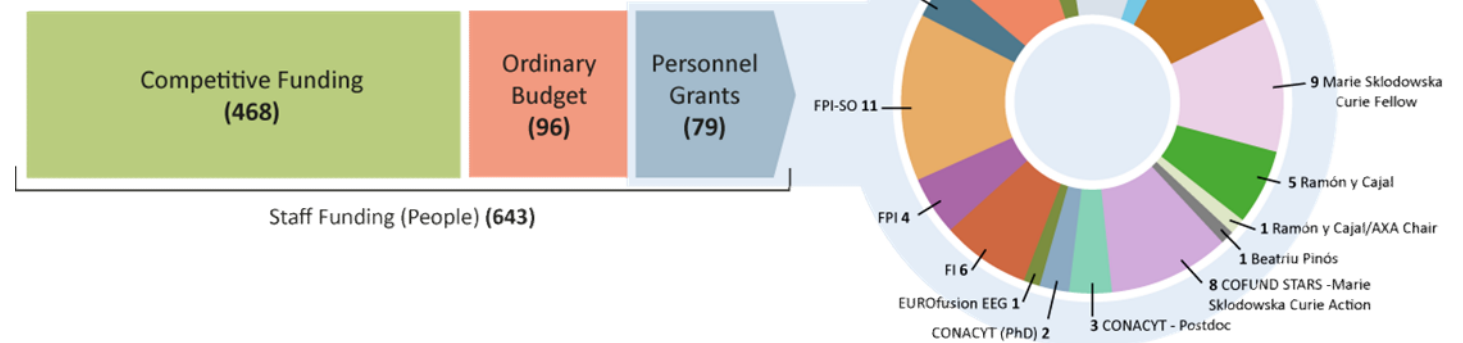




People



HR EXCELLENCE IN RESEARCH

Data as June 30, 2019



MareNostrum4

Total peak performance: **13,7 Pflops**

General Purpose Cluster:	11.15 Pflops	(1.07.2017)
CTE1-P9+Volta:	1.57 Pflops	(1.03.2018)
CTE2-Arm V8:	0.5 Pflops	(????)
CTE3-KNH?:	0.5 Pflops	(????)

MareNostrum 1

2004 – 42,3 Tflops
1st Europe / 4th World
New technologies

MareNostrum 2

2006 – 94,2 Tflops
1st Europe / 5th World
New technologies

MareNostrum 3

2012 – 1,1 Pflops
12th Europe / 36th World

MareNostrum 4

2017 – 11,1 Pflops
2nd Europe / 13th World
New technologies

Collaborations with Industry




Research into advanced technologies for the exploration of hydrocarbons, subterranean and subsea reserve modelling and fluid flows




Research on wind farms optimization and wind energy production forecasts




Collaboration agreement for the development of advanced systems of deep learning with applications to banking services



Simulations to improve the understanding of the rotating wheels flow physics and its impact over the aerodynamic performance



Advanced statistical methods to the optimization of maintenance, energy usage, and control of the city's water treatment and supply processes.



Research on efficient data sensing, algorithms for analysis of industrial processes and visualization of large datasets of industrial data




Artificial Intelligence and Big Data techniques to improve the quality of care and personalized diagnosis



BSC's dust storm forecast system licensed to be used to improve the safety of business flights.

Collaborations with Global IT industry 2019



BSC's spin-offs



**NOSTRUM
BIODISCOVERY, S.L.**

Applies supercomputing
**TO SPEED UP DRUG
DISCOVERY**

For the:

- **PHARMA
INDUSTRY**
- **BIOTECH
COMPANIES**



**MITIGA
SOLUTIONS, S.L.**

Provides operational
solutions **TO MINIMIZE
THE IMPACT OF
VOLCANIC ASH
HAZARDS**

For the:

- **AVIATION INDUSTRY**
- **ENGINE
MANUFACTURES**
- **CONSULTING SECTORS**



ELEM BIO, S.L.

Provides **BIOMECHANICS
SIMULATIONS**, offering
software-as-a-service
simulation tool focused
on cardiovascular and
respiratory systems

For the:

- **PHARMA INDUSTRY**
- **MEDTECH COMPANIES**
- **PUBLIC HEALTH**
- **EDUCATION**



NEARBYCOMP, S.L.

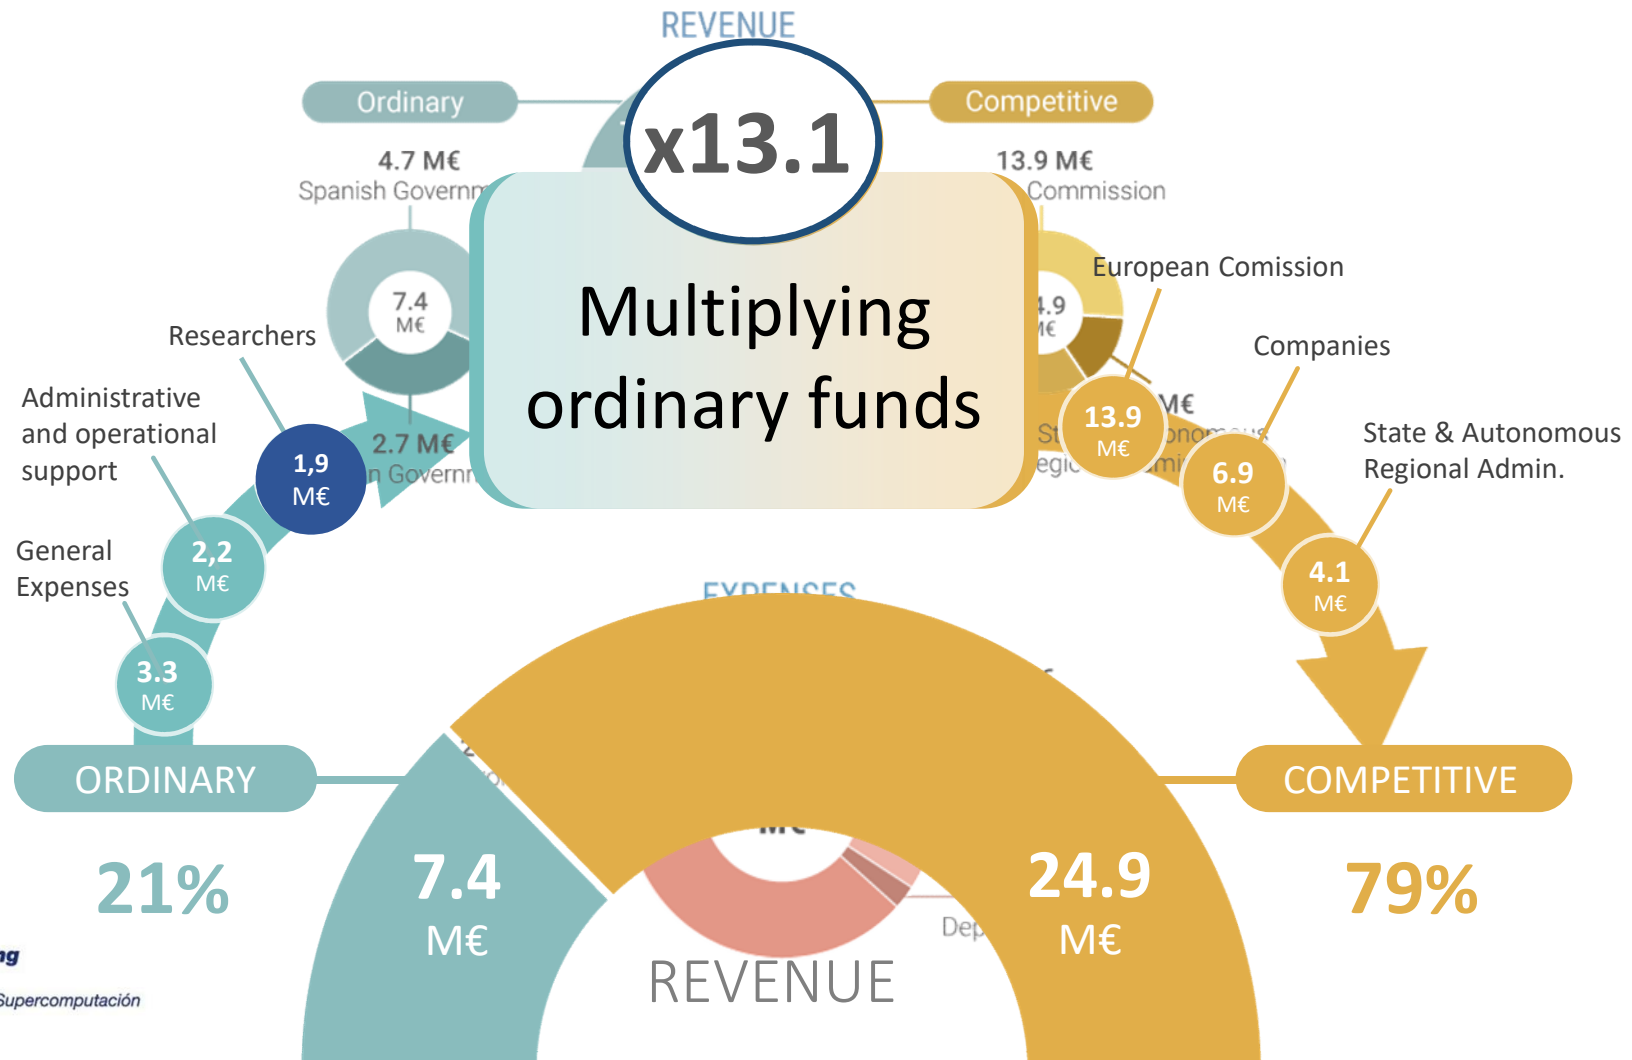
Provides **FOG
COMPUTING FOR
IOT**, delivering
customization services
for different scenarios
of FOG computing

For the:

- **5G**
- **IOT**
- **SMART CITIES**

BSC Resources

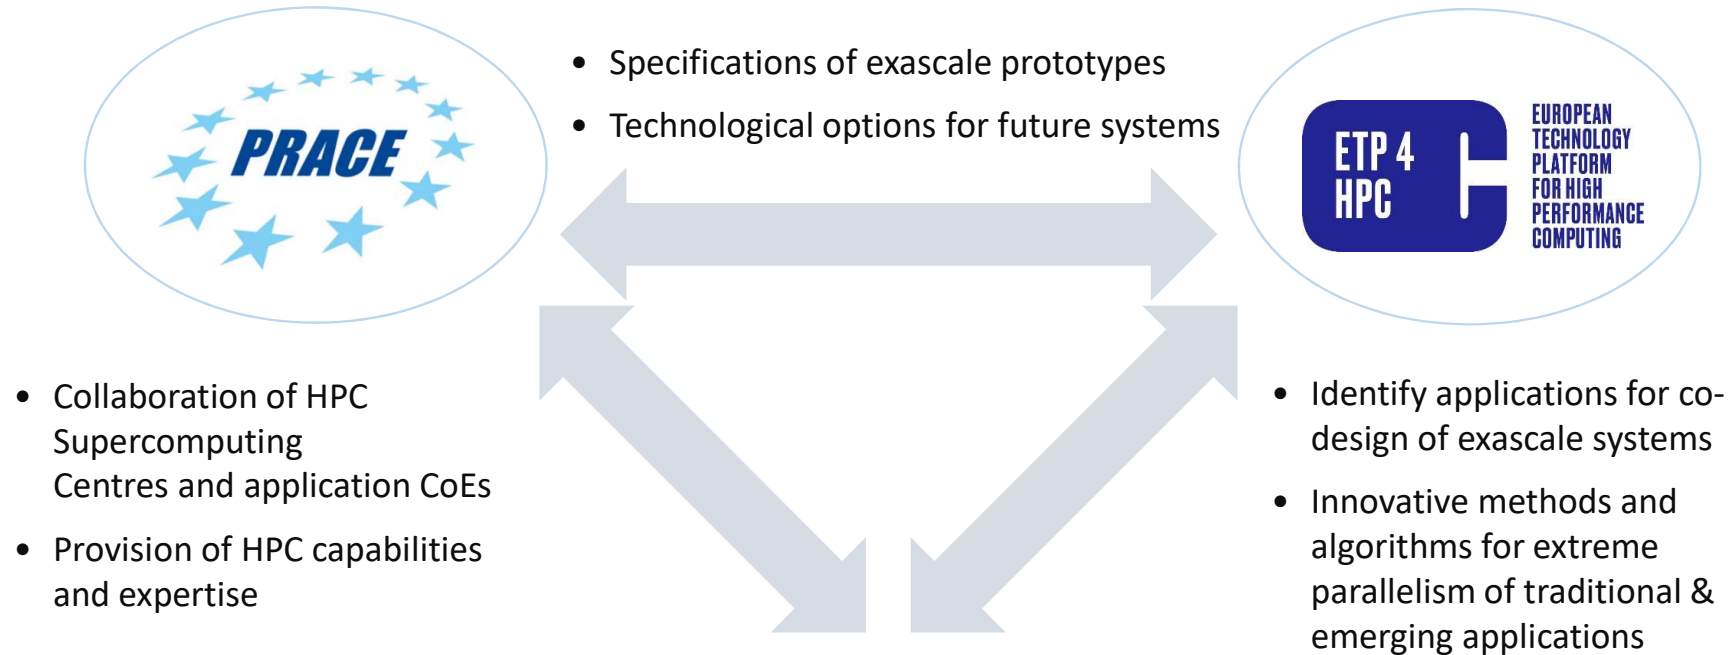
2018 executed budget



TOP-10 Spanish Organizations in Horizon 2020

Legal name	EU Contribution (€)	Project Participations
CSIC	230,434,008 €	536
Tecnalia	106,426,784 €	239
Barcelona Supercomputing Center	76,524.698 €	132
Universitat Politècnica de Catalunya	59,475,312 €	158
Universitat Pompeu Fabra	56,816,732 €	109
ICFO	56,517,896 €	78
Universitat Autònoma de Barcelona	56,322,646 €	117
Universidad Politécnica de Madrid	55,004.745 €	155
Universitat Politècnica de València	53.806.967	139
ATOS Spain	52,902,517 €	148

EU HPC Ecosystem



Centers of Excellence in HPC applications





Distributed Supercomputing Infrastructure

26 members, including

5 Hosting Members

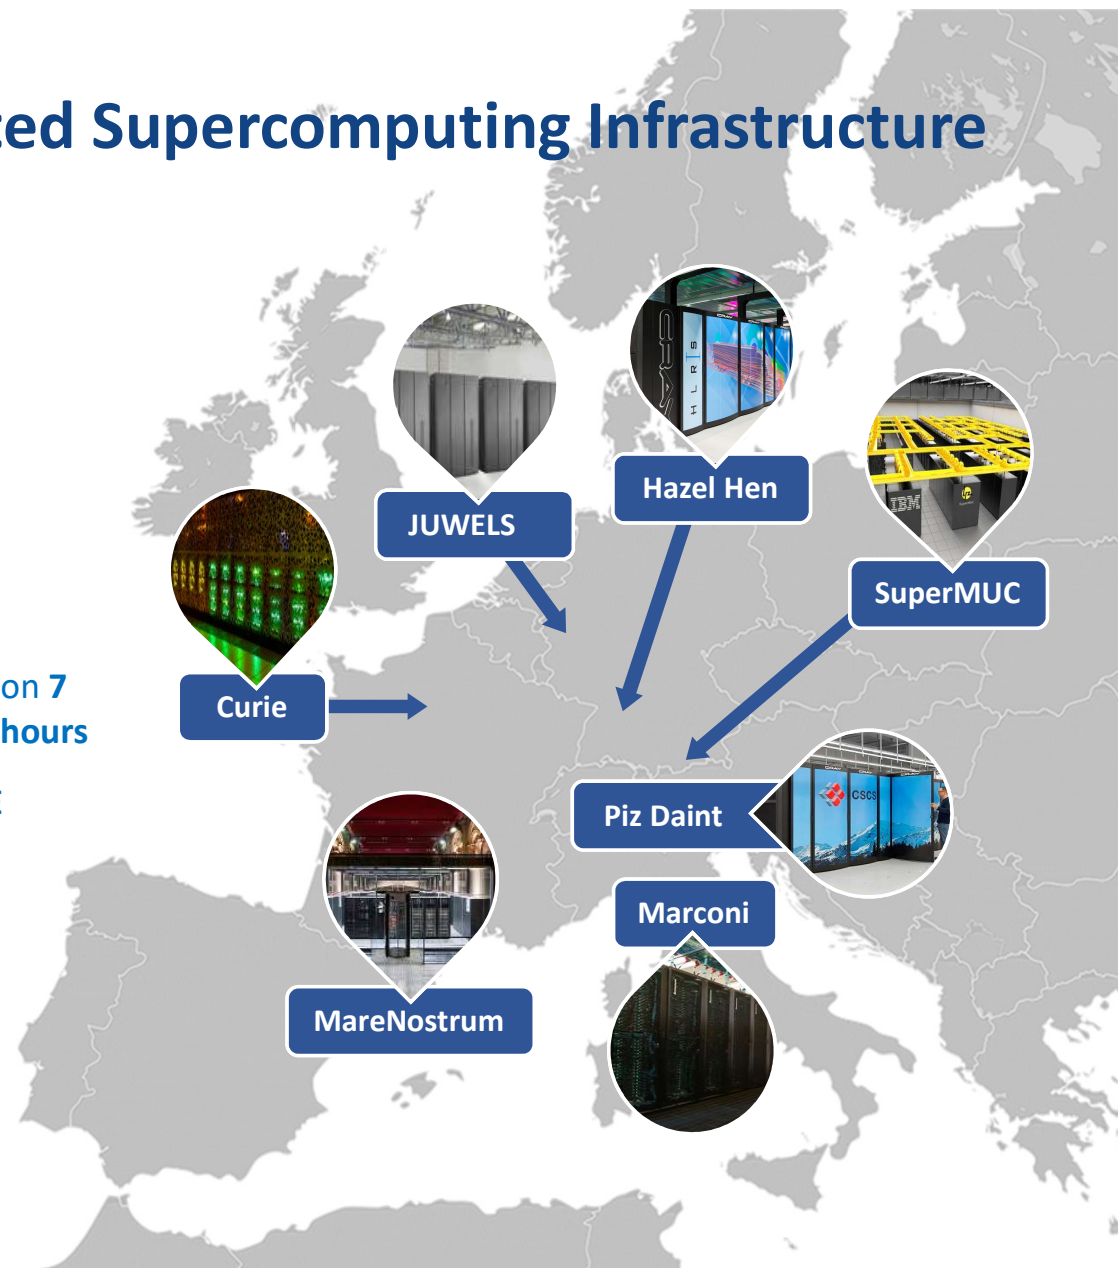
(Switzerland, France, Germany,
Italy and Spain)

688 scientific projects
enabled

110 PFlops/s of peak performance on **7**
world-class systems, **21 billions** of hours

>12.000 people trained by **6 PRACE**
Advanced Training Centers and
others events


Access prace-ri.eu/hpc-acces



The Codesign Challenge



participates




BioExcel

Centre of Excellence for
Biomolecular Research
(Led by KTH)



leads



ChEESE

Centre of Excellence in Solid
Earth
(Led by BSC)



participates



CompBioMed

Centre of Excellence on
Computational Biomedicine
(Led by Univ.
College of London)



participates



EoCoE

Energy oriented Centre
of Excellence
(Led by CEA)



participates




ESiWACE

Excellence in Simulation of
Weather and Climate in
Europe (Led by DKRZ)



participates




EXCELLERAT

Center of Excellence for
Engineering Applications
(Led by HPC HLRS)



participates



MAX

Materials design at the
eXascale (Led by CNR)



leads



PoP

Performance Optimization
and Productivity (Led by
BSC)



**Barcelona
Supercomputing
Center**

Centro Nacional de Supercomputación

RES: HPC Services for Spain

Red Española de Supercomputación is made up of
11 institutions and 12 supercomputers



Access

www.res.es



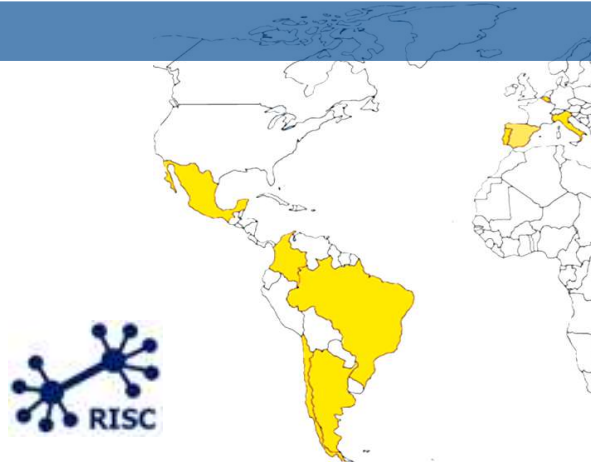
RISC Project

1. Identified research clusters for targeted research collaboration
2. Produced a Green Paper on HPC Drivers and Needs in Latin America
3. Produced a Roadmap for HPC strategic R&D in Latin America
4. Enhanced HPC R&D policy dialogue between policymakers and stakeholders



LATIN AMERICA

- Universidad Veracruzana
- Universidad de Chile
- Universidad de Buenos Aires
- Universidad Autónoma de Manizales
- Coppetec Fundação do Río de Janeiro



EUROPE

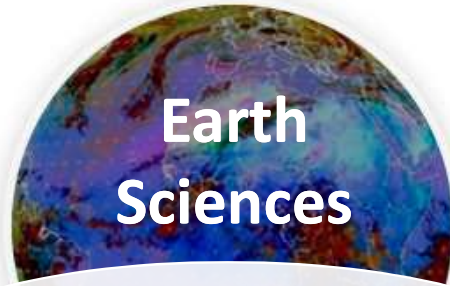
- BSC
- CINECA
- UPM
- Menon
- Uni Coimbra

Mission of BSC Scientific Departments



Computer Sciences

To influence the way machines are built, programmed and used: programming models, performance tools, Big Data and Artificial Intelligence, computer architecture, energy efficiency



Earth Sciences

To develop and implement global and regional state-of-the-art models for short-term air quality forecast and long-term climate applications



Life Sciences

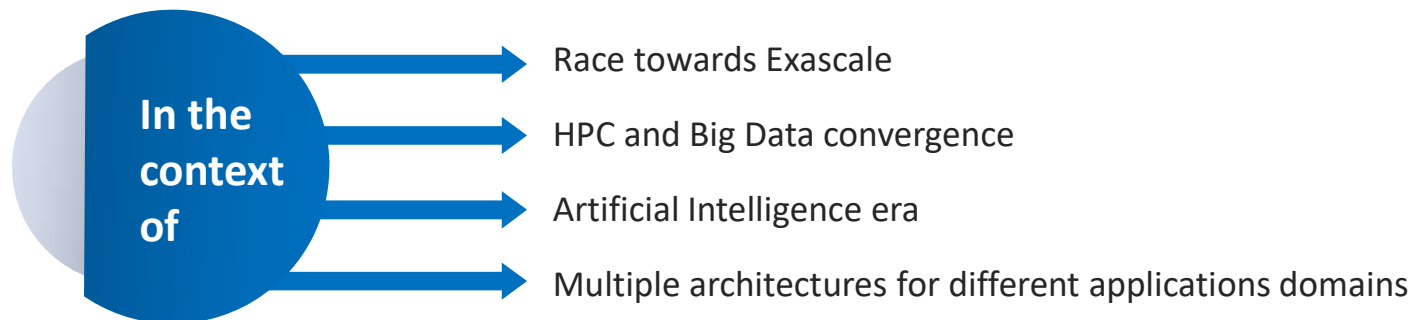
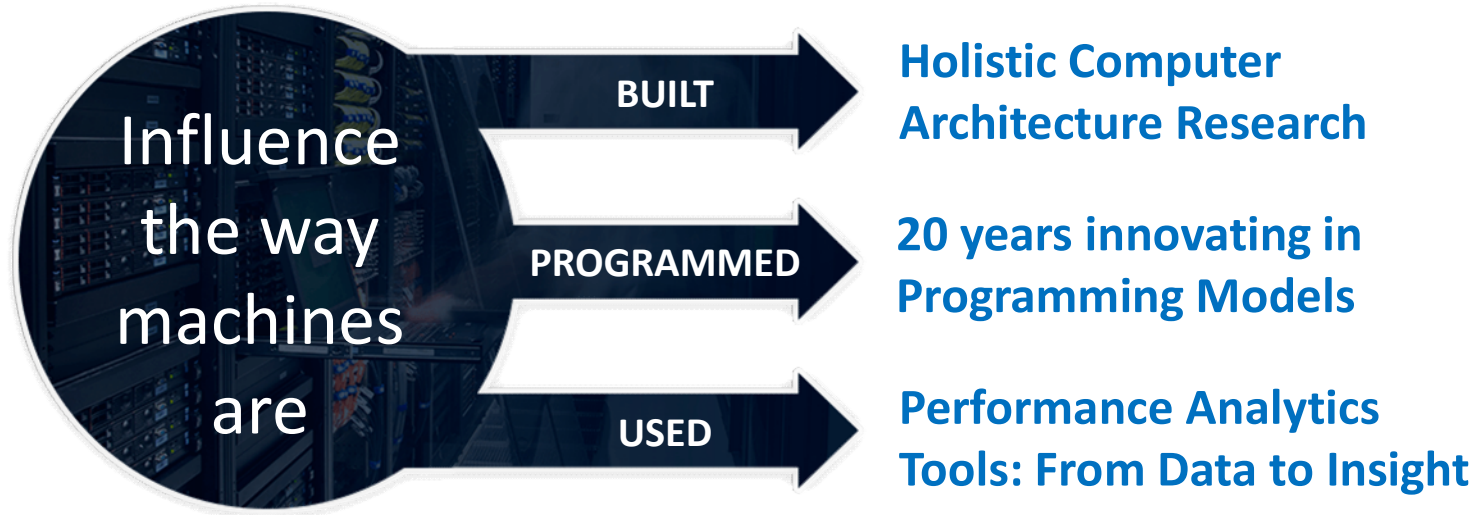
To understand living organisms by means of theoretical and computational methods (molecular modeling, genomics, proteomics)



CASE

To develop scientific and engineering software to efficiently exploit super-computing capabilities (biomedical, geophysics, atmospheric, energy, social and economic simulations)

Computer Sciences



- Transactional memory on FPGAs
 - ❑ EU FP7 VELOX (Integrated Approach to Transactional Memory on Multi- & Many-core Computers) – **ended**
 - ❑ Design and Profiling of Hybrid Transactional Memory on FPGAs (FCCM-2011, FCCM-2012)

- Database acceleration using FPGAs
 - ❑ EU FP7 AXLE (Advanced Analytics for Extremely Large European Databases)- **ended**
 - ❑ Hardware acceleration for SQL primitives, smart disks (FPL-2014, MICPRO-2017)

- FPGA-based DRAM access accelerators
 - ❑ Accelerators for applications with complex data access (FPT-2011, FPL-2012, FPL-2014, FPT-2014)

- Heterogeneous CPU/GPU/FPGA-based computing
 - ❑ EU FP7 ParaDIME (Parallel Distributed Infrastructure for Minimization of Energy) – **ended**
 - ❑ Trigeneous (CPU/GPU/FPGA) Low-energy platforms (FCCM-2015, HiPC-2015, Coolchips-2015, ICCD-2015)

- Aggressive supply voltage undervolting of commercial FPGAs
 - ❑ EU H2020 LEGaTO (Low-Energy Toolset for Heterogeneous Computing)- **ongoing**
 - ❑ Energy/performance/resilience trade-off study for FPGA-based DNNs (MICRO-2018, FPL-2018, PDP-2019)

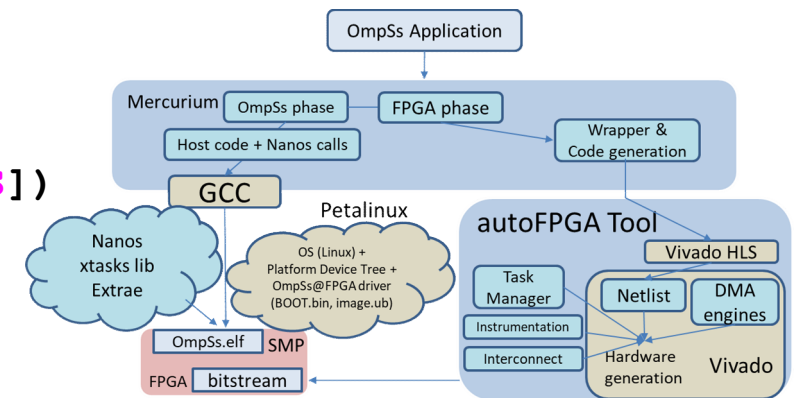
OmpSs@FPGA Ecosystem

- Improving programmers productivity
 - Programming model based on directives
 - Automatic code offloading to FPGA vendor tools
 - HLS and/or OpenCL compilers

```
#pragma omp target device(fpga) copy_deps  
#pragma omp task in([BS]a,[BS]b) inout([BS]c)  
void matrix_multiply(float a[BS][BS], float  
                      b[BS][BS],float c[BS][BS])
```


Application Acceleration on FPGAs with OmpSs@FPGA. FPT 2018: 70-77

OmpSs@Zynq all-programmable SoC ecosystem. FPGA 2014: 137-146



OmpSs@FPGA analysis

- Hardware instrumentation of FPGA IP cores
 - Generating trace information from inside the FPGA



Paraver trace with internal information

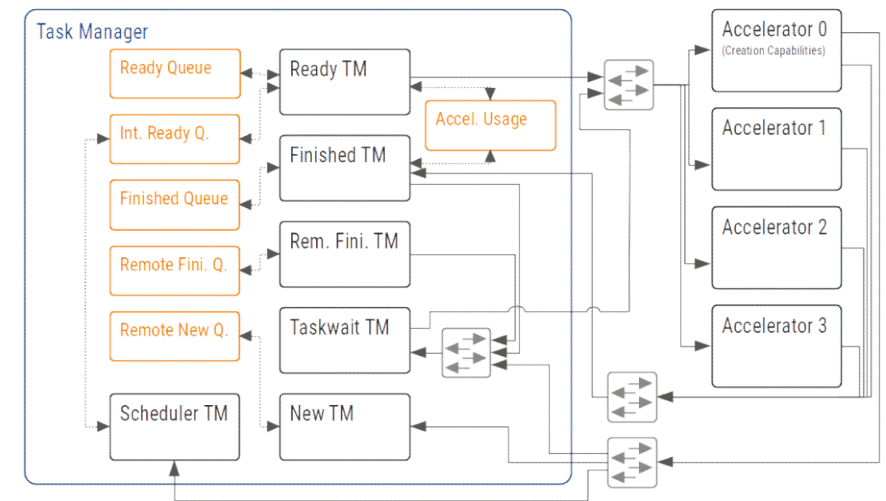
The Secrets of the Accelerators Unveiled: Tracing Heterogeneous Executions Through OMPT. IWOMP 2016: 217-236

OmpSs@FPGA Task Manager

- Ability to implement task creation and management inside the FPGA
 - Overcomes the overheads of fine granularity tasks management
- Picos HW: fast dependence management

Adding Tightly-Integrated Task Scheduling Acceleration to a RISC-V Multi-core Processor. To appear: MICRO 52 (2019)

A Hardware Runtime for Task-Based Programming Models. IEEE Trans. Parallel Distrib. Syst. 30(9): 1932-1946 (2019)



OmpSs@FPGA development

- European projects
 - AXIOM, EuroEXA, LEGaTO, EPEEC
 - Support for various boards
 - Application porting
 - User's support, tutorials...
- Collaboration with industry
 - IBM, Intel, Ikerگونه, Continental

BSC strategy for Artificial Intelligence

Projects with public/private institutions and companies

Precision medicine

Other domains

Genomic Analytics

Text Analytics

Medical Imaging

Organ simulation

Social & Personal
Data

Industrial CASE
apps

Earth Sciences

Data models and algorithms
(approximate computing -- reduced precision, adaptive layers, DL/Graph Analytics, ...)


Programming models and runtimes
(PyCOMPSs, TIRAMISU, interoperability current approaches)

Hw acceleration of DL workloads
(novel architectures for NN, FPGA acceleration)

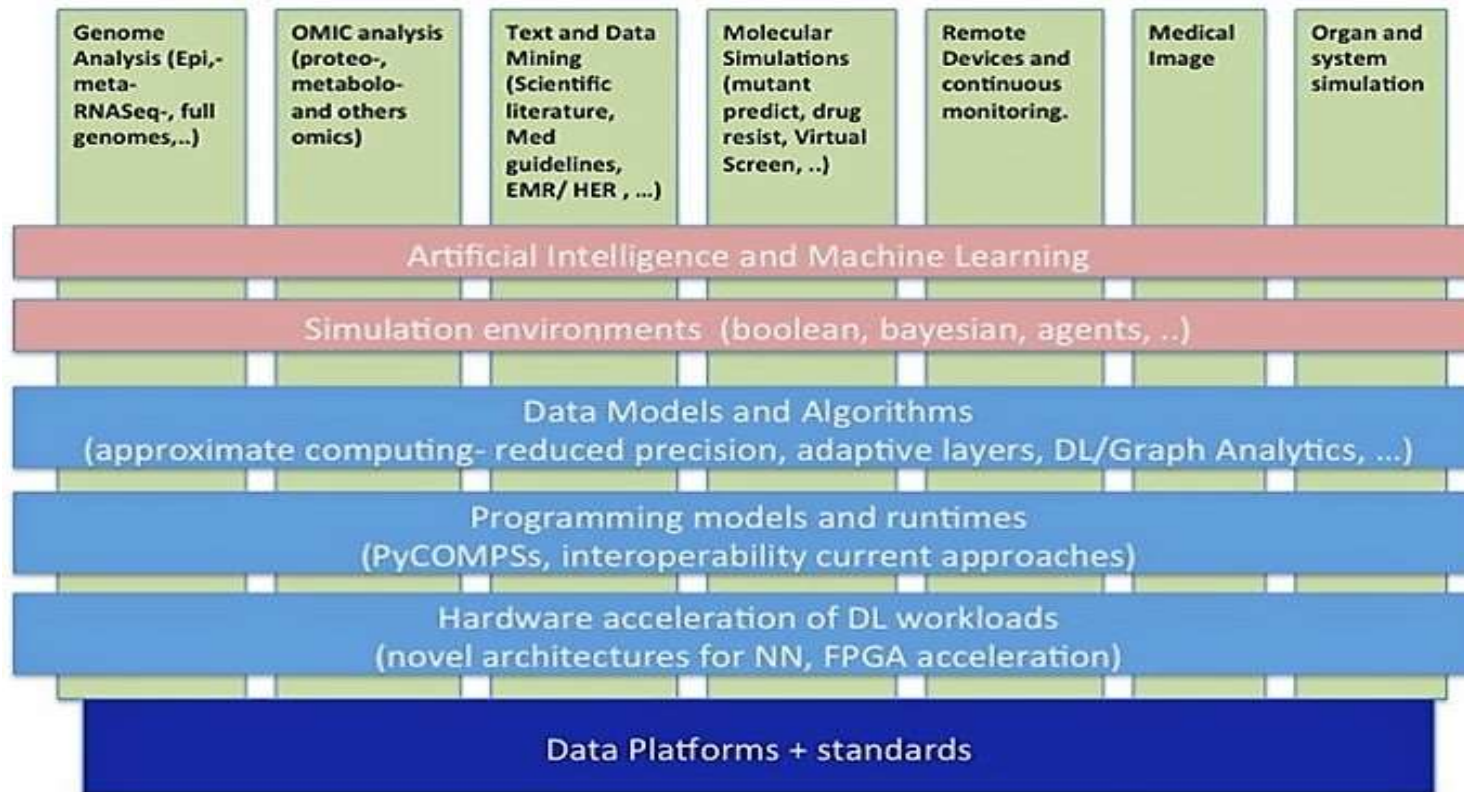
Data platforms + standards

Life Sciences

Understanding living organisms by theoretical and computational methods

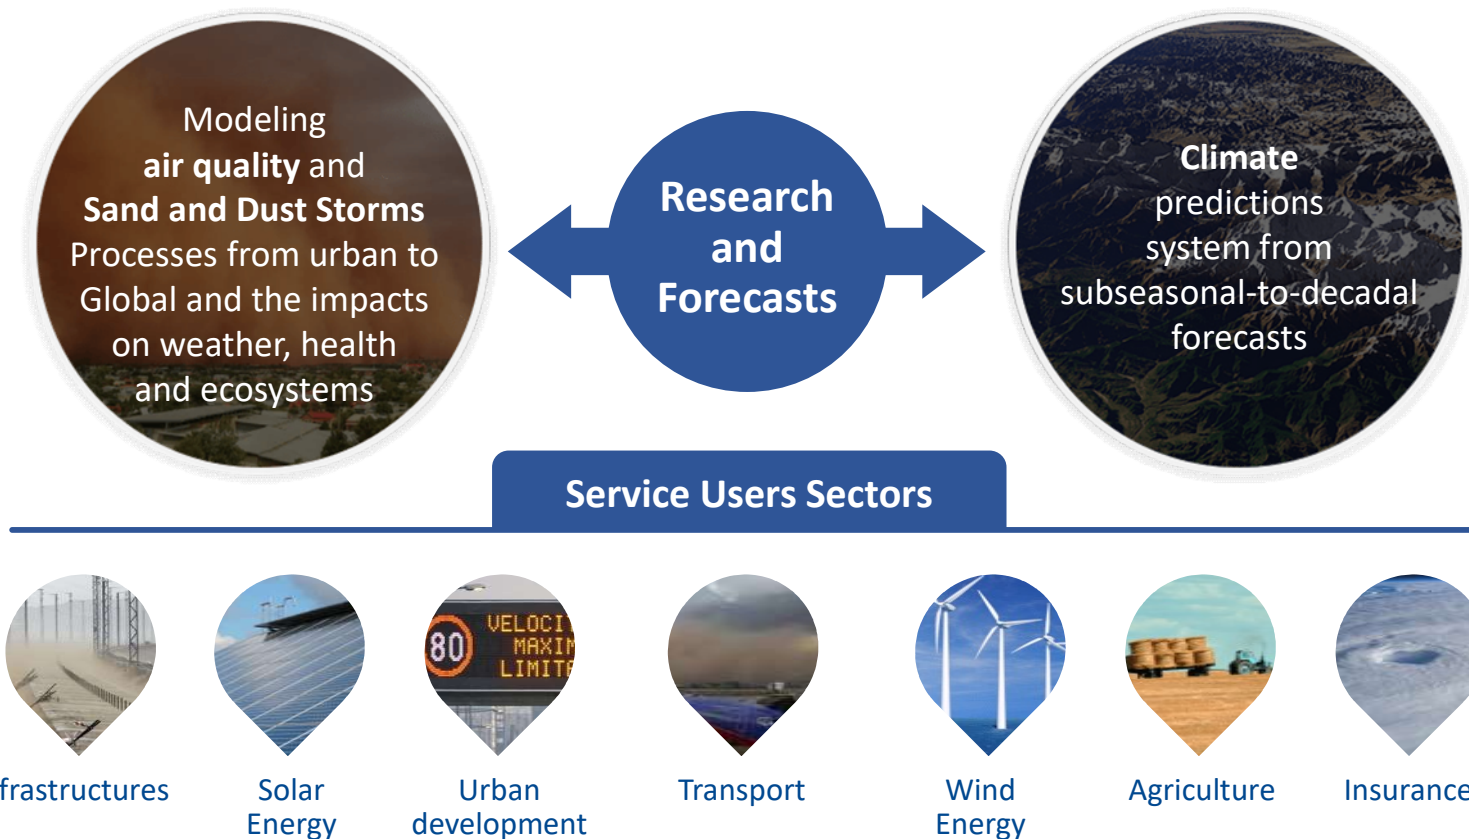


BSC technical strategy for Personalized Medicine



Earth Sciences

Environmental modelling and forecasting, with a particular focus on weather, climate and air quality



Computational Applications for Science and Engineering

Energy



Biomechanics




**INDUSTRY
ORIENTED
DEPARTMENT**

Smart & resilient cities



Data analytics & visualization



Continue international collaborations: Joint Laboratory on Extreme Scale Computing




In June 2014, the University of Illinois at Urbana-Champaign, INRIA, Argonne National Laboratory, Barcelona Supercomputing Center and Jülich Supercomputing Centre formed the Joint Laboratory on Extreme Scale Computing.

The Joint Laboratory focuses on software challenges found in extreme scale high-performance computers.

Researchers from the different centres regularly meet for workshops and at the last one in November 2014, researchers from Riken AICS also took part.

Education: Engagement with Universities



MSc in Innovation and Research in Informatics - HPC at FIB (UPC) Msc in Artificial Intelligence at FIB (UPC)

- Curriculum design
- Providing access to the BSC HPC facilities through internships
- BSC experts lecturing
- Advising Master Thesis
- Intensification on **HPC for AI** (forthcoming 2019-2020)

Double Diploma Agreement CIC(IPN-MEX)-FIB(UPC)

- Providing Master Thesis co-advising.

Doctoral Program Affiliation (main programs)

- Applied Math Program (UPC)
- Artificial Intelligence (UPC)
 - Double diploma (IPN-UPC) (forthcoming 2019-2020)
- **Computer Architecture Doctoral Program** (UPC)
 - Double diploma (IPN-UPC) (forthcoming 2019-2020)
- Environmental Engineering Doctoral Program (UPC)
- Biomedicine Doctoral Program (UB)
- Chemistry Doctoral Program(UB)

Post-Doctoral Programme CONACyT-BSC (2012-2020)

Bachelor Degree in Bioinformatics (UPF+UPC)

- Curriculum Design
- BSC experts lecturing

Fighting gender gap



“MareNostrum” song by a popular girls music group

Most women's presence in educational material



Girls' focused visits for primary school students

(6,000 the first year)



Enlivened by our superheroine



Guided by female educators trained in computing

8 fun activities introducing supercomputers



Top 10, June 2019

Computer	Cores	Accelerators	Rmax [PFlop/s]	Rpeak [PFlop/s]	Power (MW)	Effeciency [GFlops/Watts]
IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband	2.414.592	2.211.840	148,6	200,8	10,1	14,7
IBM Power System S922LC, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband	1.572.480	1.382.400	94,6	125,7	7,4	12,7
Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway	10.649.600		93,0	125,4	15,4	6,1
TH-IVB-FEP Cluster, Intel Xeon E5-2692v2 12C 2.2GHz, TH Express-2, Matrix-2000	4.981.760	4.554.752	61,4	100,7	18,5	3,3
Dell C6420, Xeon Platinum 8280 28C 2.7GHz, Mellanox InfiniBand HDR	448.448		23,5	38,7		
Cray XC50, Xeon E5-2690v3 12C 2.6GHz, Aries interconnect , NVIDIA Tesla P100	387.872	319.424	21,2	27,2	23,8	8,9
Cray XC40, Xeon E5-2698v3 16C 2.3GHz, Intel Xeon Phi 7250 68C 1.4GHz, Aries interconnect	979.072		20,2	41,5	7,6	2,7
PRIMERGY CX2570 M4, Xeon Gold 6148 20C 2.4GHz, NVIDIA Tesla V100 SXM2, Infiniband EDR	391.680	348.160	19,9	32,6	1,6	12,1
ThinkSystem SD650, Xeon Platinum 8174 24C 3.1GHz, Intel Omni-Path	305.856		19,5	26,9		
IBM Power System S922LC, IBM POWER9 22C 3.1GHz, Dual-rail Mellanox EDR Infiniband, NVIDIA Tesla V100	288.288	253.440	18,2	23,0		



System Overview

System Performance

- Peak performance of 200 petaflops for modeling & simulation
- Peak of 3.3 ExaOps for data analytics and artificial intelligence

Each node has

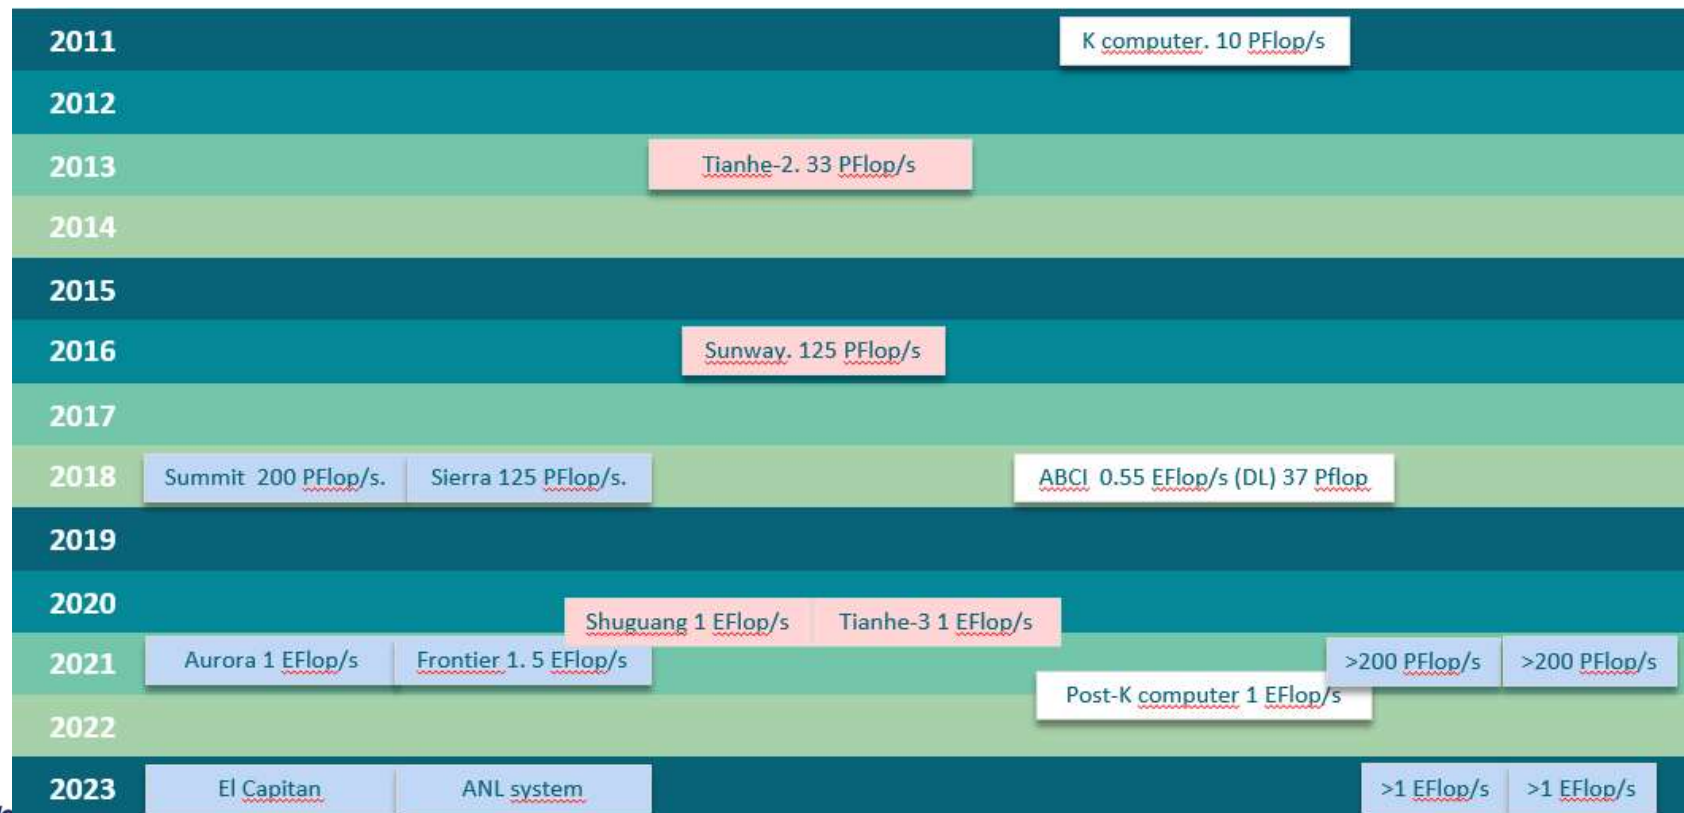
- 2 IBM POWER9 processors
- 6 NVIDIA Tesla V100 GPUs
- 608 GB of fast memory
- 1.6 TB of NVMe memory

The system includes

- 4608 nodes
- Dual-rail Mellanox EDR InfiniBand network
- 250 PB IBM Spectrum Scale file system transferring data at 2.5 TB/s



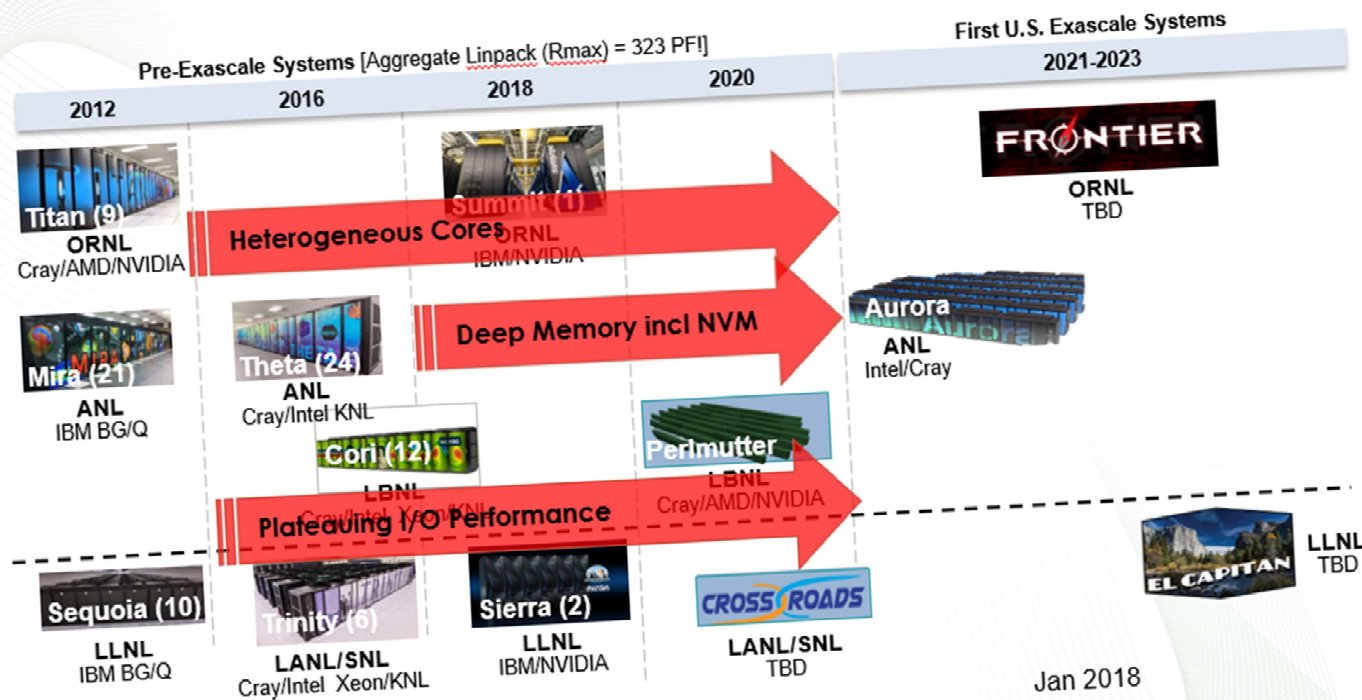
The Global Race Towards Exascale



The Exascale Race – The US example

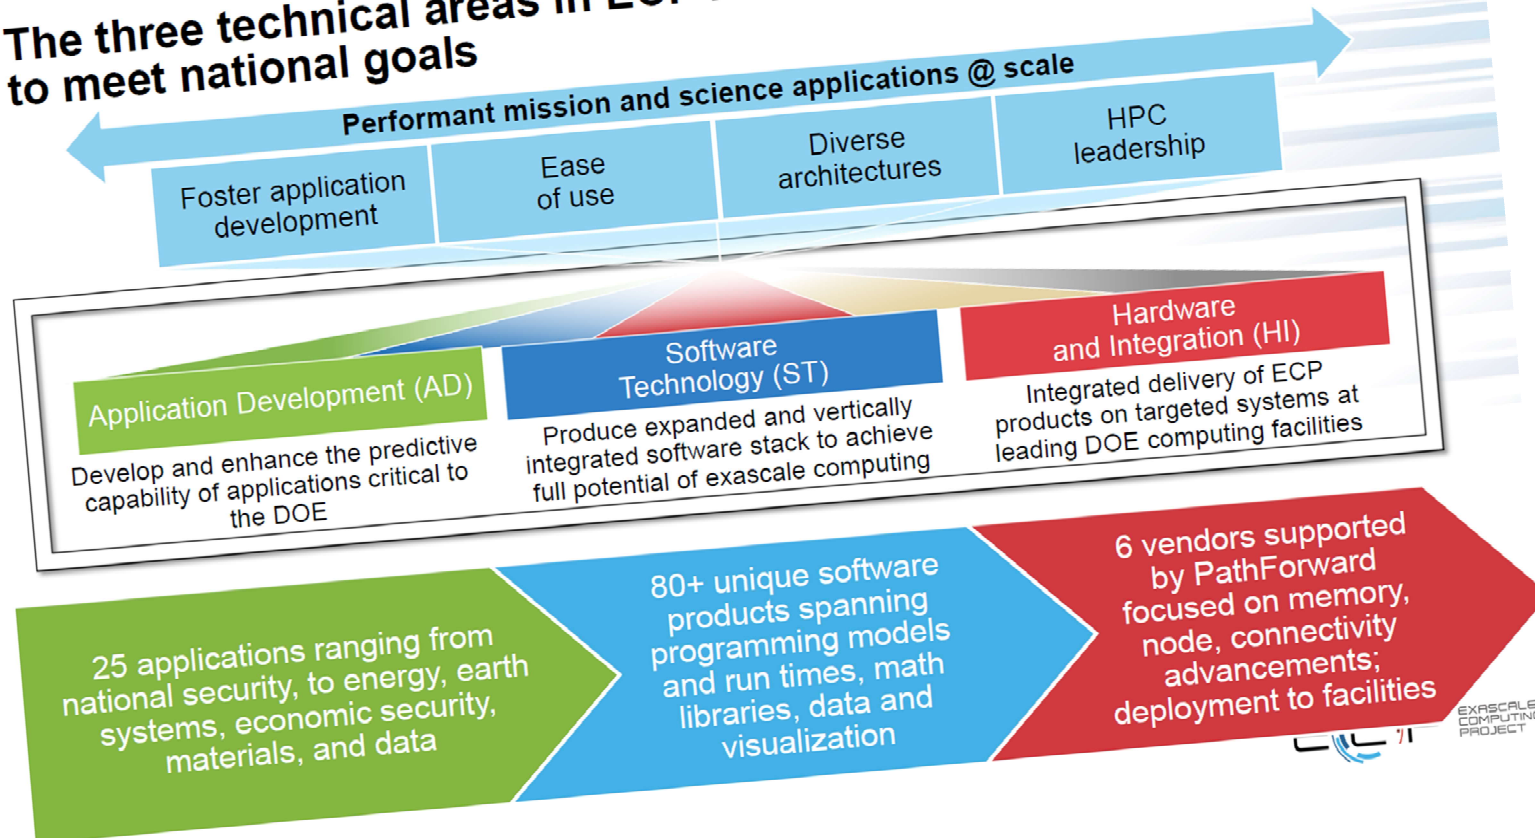
US Department of Energy (DOE) Roadmap to Exascale Systems

An impressive, productive lineup of *accelerated node* systems supporting DOE's mission



The Exascale Race – The US example

The three technical areas in ECP have the necessary components to meet national goals



EuroHPC: Unifying European HPC technologies



EuroHPC-JU members:

Austria, Belgium, Bulgaria, Croatia, Czech Republic, Denmark, Estonia, Finland, France, Germany, Greece, Hungary, Ireland, Italy, Latvia, Lithuania, Luxembourg, the Netherlands, Norway, Poland, Portugal, Romania, Slovakia, Slovenia, Spain, Sweden, Switzerland and Turkey



“A new legal and funding structure – the EuroHPC Joint Undertaking – shall acquire, build and deploy across Europe a world-class High-Performance Computing (HPC) infrastructure.

It will also support a research and innovation programme to develop the technologies and machines (hardware) as well as the applications (software) that would run on these supercomputers.”



**Barcelona
Supercomputing
Center**

Centro Nacional de Supercomputación

MareNostrum 5

A European pre-exascale supercomputer

- **200 Petaflops** peak performance (200×10^{15})
- **Experimental platform** to create supercomputing technologies “made in Europe”

Hosting Consortium:

Spain Portugal Turkey Croatia Ireland




|



Where Europe needs to be stronger

- ➔ Only 1 of the 10 most powerful HPC systems is in the EU
- ➔ HPC codes must be upgraded
- ➔ Vital HPC hardware elements are missing: general purpose processor and accelerators
- ➔ EU needs its own source of as many of the system elements as possible



BSC and the EC



Final plenary panel at ICT Innovate, Connect, Transform conference, 22 October 2015 Lisbon, Portugal.

"The transformational impact of excellent science in research and innovation"

*"Europe needs to develop an entire domestic exascale stack from the processor all the way to the system and application software",
Mateo Valero,
Director of
Barcelona
Supercomputing
Center*

Director of Barcelona Supercomputing Center, Mateo Valero, makes a pledge for developing a strong HPC ecosystem.

Published on 12/04/2016

Europe has the competence and skills to engage in the global competition towards Exascale Supercomputing. To fully benefit from the opportunities of the digital single market, Europe must strengthen the fundamental research on which digital transformation is based and build a stronger European High Performance Computing (HPC) ecosystem.

In a [quest blog post](#) on Commissioner Günther Oettinger's [website](#) Mateo Valero stresses the need for Europe to join the race towards Exascale supercomputing. According to him, there is an open window of opportunity for the High Performance Computing (HPC) development that would stimulate scientific breakthroughs and have tremendous impact on society and industry.



[Share](#)

ARM-based prototypes at BSC



2011
Tibidabo

ARM multicore

2012
KAYLA

ARM + GPU
CUDA on ARM

2013
Pedraforca


ARM + GPU
Infiniband
RDMA

2014
Mont-Blanc

Single chip ARM+GPU
OpenCL on ARM GPU



Mont-Blanc HPC Stack for ARM



Industrial applications



Applications



System software



Hardware



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

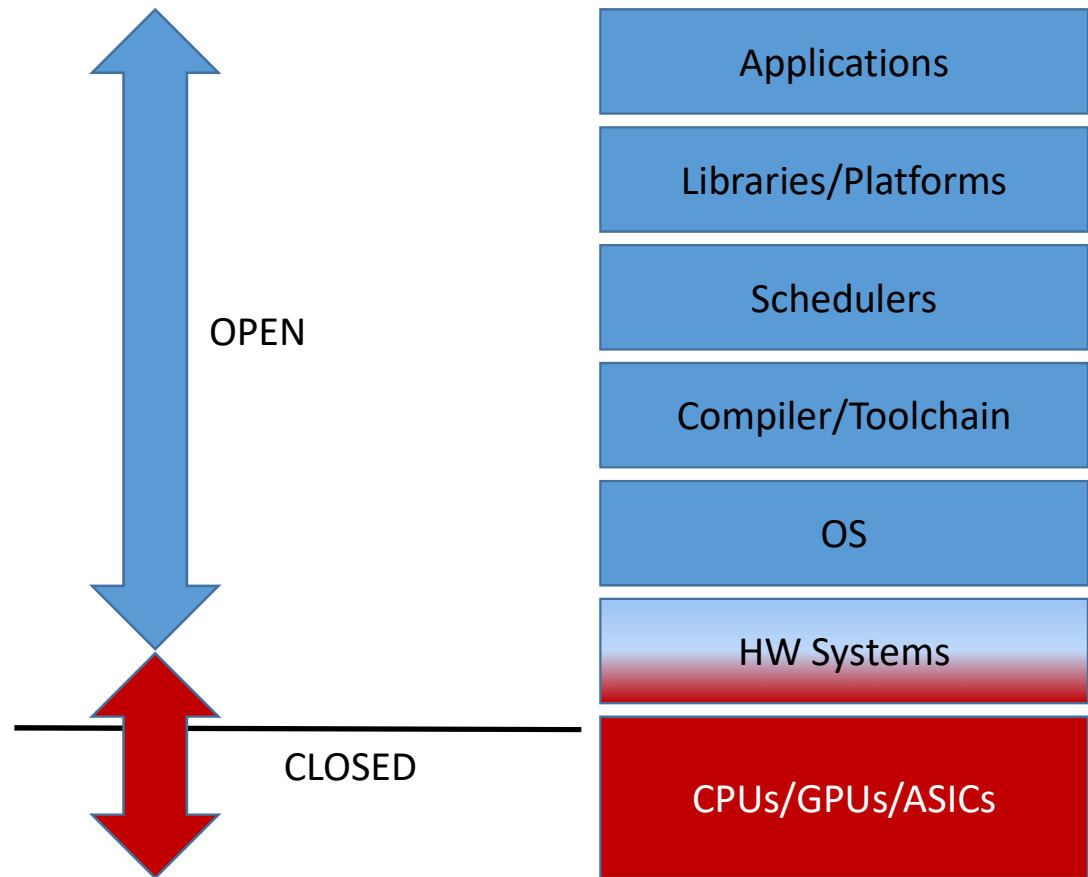


The US Cloud Act v The EU's GDPR - Data Privacy & Security




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?**




RISC-V is democratising chip-design



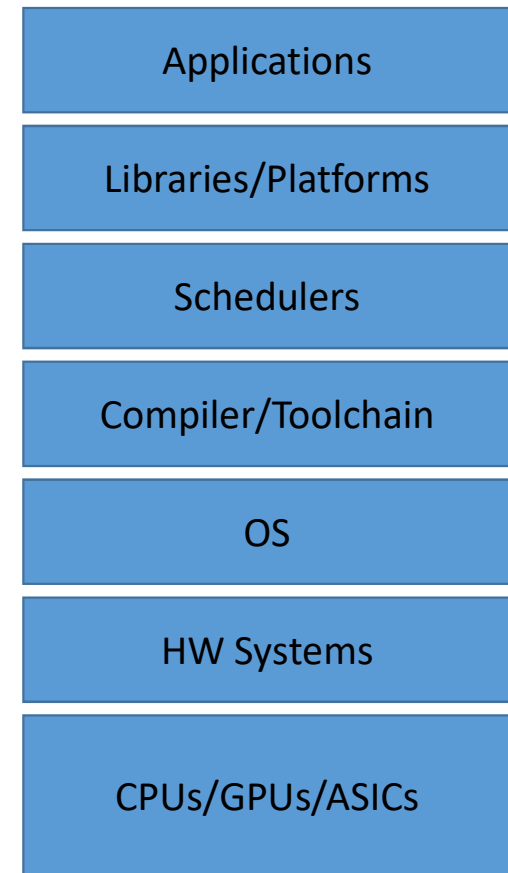
- ➔ More and more global IT actors are adopting RISC-V architectures to be vendor independent
 - ➔ Google
 - ➔ Amazon
 - ➔ Western Digital
 - ➔ Alibaba
- ➔ And of course the entire IoT ecosystem for lower performance, lower energy applications.
- ➔ Major opportunity for ICT industry also in Spain

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



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 centres

EPI Partners

BMW
GROUP



Rolls-Royce
Motor Cars Limited

Atos



Barcelona
Supercomputing
Center
Centro Nacional de Supercomputación



ETH zürich



Exascale supercomputing initiative at BSC

- ➔ Ground floor opportunity to design and build a European supercomputer at the best supercomputing center in Europe!
- ➔ The open-source hardware opportunity
- ➔ RISC-V HPC accelerator: from concept to implementation
- ➔ Latest silicon technologies: 7nm, 5nm and 3 nm
- ➔ Working with industrial and academic partners
- ➔ HPC, automotive, bio, meteorological and other workloads

Barcelona desarrollará el chip de los superordenadores europeos

La CE financia una tecnología clave para la soberanía informática del continente

El proyecto del chip europeo estará liderado por Barcelona

LA VANGUARDIA

Barcelona desarrolla el chip de los futuros superordenadores europeos

El superordenador MareNostrum 5 se lanzará a la conquista de procesadores y chips 'made in Europe'

El MareNostrum 5 incluirá una plataforma para crear chips europeos

El próximo superordenador contribuirá al desarrollo de tecnologías íntegramente desarrolladas en Europa

El súperordenador presentará batalla en la fabricación de chips y procesadores europeos

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 aim: 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**



How to implement this “Open Future World”?

- The BSC launches **LOCA**, the new **European Laboratory for Open Computer Architecture**, a joint long-term initiative to promote a vibrant RISC-V ecosystem, HQ in Southern Europe, supported by:
- The European Commission
- The BSC trustees
- The European Academics
- The main IT worldwide companies
- The digital technology industry in Spain?

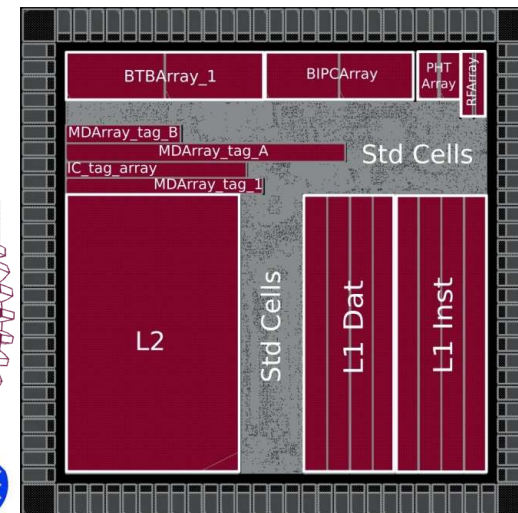



UPC, Cantabria, Chalmers, Rome Sapienza, Zagreb, Forth, ETH, EPFL.....



First Lagarto Tapeout

- Target design:
 - Simple in-order core with 5 stages, single issue
 - 16KB L1 caches, 64KB L2 cache, TLB
 - Memory controller on the FPGA side
 - FPGA – ASIC communication via packetizer
 - Debug ring via JTAG
 - Target technology: TSMC 65nm
 - Design fits in the total area budget of 2.5mm²
 - Submitted for fabrication in May 2019
- Collaborative project with different teams:
 - RTL Design: Lagarto (BSC + CIC-IPN)
 - Verification (BSC)
 - Logic Synthesis (UPC + BSC)
 - Physical design (IMB-CNM + BSC)
 - Tapeout and bringup (IMB-CNM + BSC)






The HPC Future is Wide Open!

- Can open source hardware play a big role like open source software?
- How do we build flexible accelerators?
- Can we leverage commodity components and merge them with HPC systems?
- Can we jumpstart HPC hardware development in Europe?

An Open Path to the Future

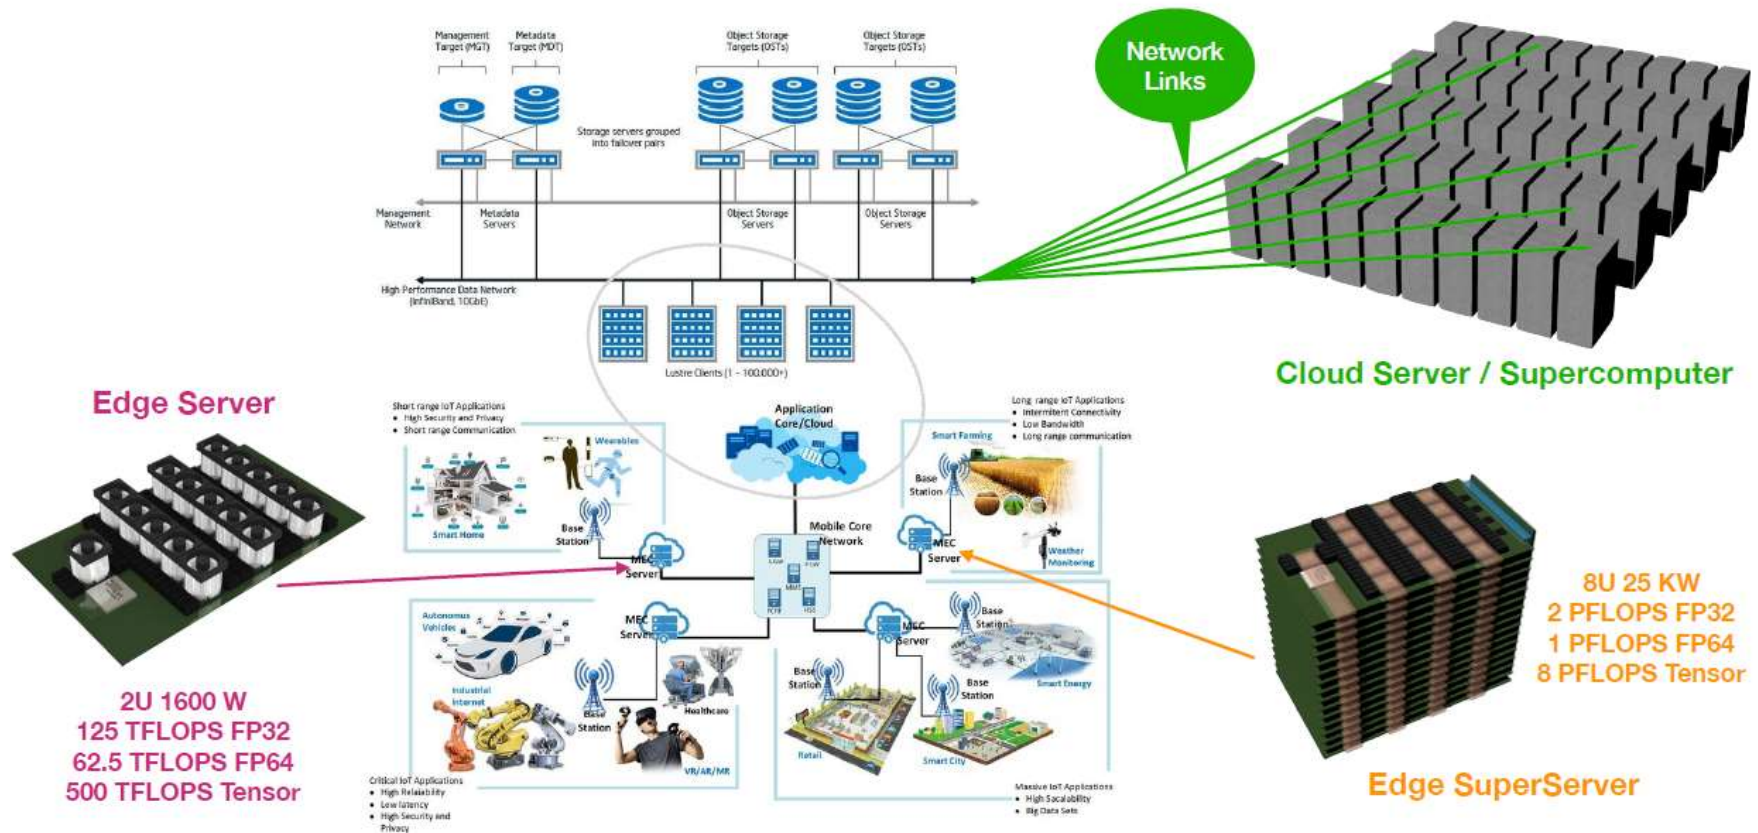
- We can change the balance of Host CPUs to Accelerators
 - OCP Accelerator Module (OAM): 1-8 or more accelerators per CPU
 - Partial to All-to-All OAM communication topologies:



enabling a variety of applications, including AI, ML, DL, etc.

- FPGAs: Accelerators, prototypes and emulators
- MareNostrum Experimental Exascale Platform

From IoT, Edge Computing, Clouds to Supercomputers





What does a 30 MW ExaFLOP SC look like?...

We have some ideas, come join the fun!

- ➔ 64 cabinets: 1.0 Exaflops
- ➔ Cabinet: 16 Petaflops, 400 KW (water cooled)
- ➔ 256 nodes, 24,576 cores
- ➔ 128 to 512 Terabytes DRAM
- ➔ 0.1 Byte/flop bandwidth ratio
- ➔ 40 Gflops/W efficiency
- ➔ 7nm initial, 5 and 3 nm follow-on designs

THE
NEW
NEW



BSC is hiring... Creating high value job opportunities in Spain

BSC is looking for talented and motivated professionals with expertise in the design and verification of IPs to be integrated into a European HPC accelerator. The design is based on a RISC-V architecture. This is a NEW project to build an energy efficient Exascale system.

Experienced professionals (Engineers and/or PhD holders) are wanted for:

- ➔ RTL / Microarchitecture
- ➔ Verification
- ➔ FPGA design
- ➔ Simulation
- ➔ Software: compilers/OS/RT

RISC-V has the opportunity to be like Linux. It would be global and go beyond Airbus and Galileo!




EuroHPC
Joint Undertaking



MareNostrum RISC-V inauguration 2021



MN6-RISC-V
2025???





**Barcelona
Supercomputing
Center**
Centro Nacional de Supercomputación



EXCELENCIA
SEVERO
OCHOA

Thank you

mateo.valero@bsc.es
martorell@bsc.es