

# CINECA 2019 - 2020

HPC annual report



LEONARDO





Cineca HPC Report 2019 - 2020  
website: [www.hpc.cineca.it](http://www.hpc.cineca.it)  
mail: [info-hpc@cinca.it](mailto:info-hpc@cinca.it)

Contents are © CINECA 2019 - 2020 and are released under  
the Attribution-NonCommercial-NoDerivs (CC BY-NC-ND)  
Creative Commons license, version 3.0.

Editorial Office  
Cineca Consorzio Interuniversitario  
Via Magnanelli 6/3  
40033 - Casalecchio di Reno (BO) - Italy

Editorial Board  
Sanzio Bassini  
Simona Caraceni  
Francesca Delli Ponti  
Elda Rossi

Design and Layout project  
Francesca Delli Ponti

Print  
nuova MONOGRAF snc  
Via E. Collamarini, 5/H-I  
40138 - Bologna - Italy

Thanks to:

Fabio Affinito; Paola Alberigo; Jorge Amaya; Claudio Arlandini; Isabella Baccarelli; Lorenzo Baraldi; Gabriele Brizzi; Francesco Califano; Carlo Cavazzoni; Lorenzo Celli; Daniele Cesarini; Mirko Cestari; Balasubramanian Chandramouli; Beatrice Chiavarini; Riccardo Ciolfi; Patrizia Coluccia; Massimiliano Corsini; Maurizio Cremonesi; Rita Cucchiara; Roberto Da Vià; Francesco Falciano; Luca Ferraro; Federico Ficarelli; Franz Fischnaller; Daniela Galetti; Bruno Giacomazzo; Luca Govoni; Massimiliano Guarrasi; Antonella Guidazzoli; Francesco Iannone; Silvano Imboden; Jay Vijay Kalinani; Wolfgang Kastaun; Giovanni Lapenta; Maria Chiara Liguori; Alessandro Marani; Fabio Massaccesi; Antonio Memmolo; Riccardo Mengoni; Sergio Orlandini; Roberto Orosei; Daniele Ottaviani; Enzo Papandrea; Eric Pascolo; Fabio Pitari; Lara Querciagrossa; Marco Rorro; Gabriella Scipione; Stefano Tagliaventi; Debora Testi; Giuseppe Trotta; Roberta Turra; Cinzia Zannoni.

# CINECA HPC REPORT

## 2019 - 2020



Dear colleagues,

in presenting the compendium of the results achieved and the activities carried out during the past year 2019, allow me first of all to remember that last year was the 50th anniversary of the founding date of Cineca.

Cineca was established in 1969 as a non-profit inter-university consortium between the Universities of Bologna, Padua, Florence and Venice, by decree of the President of the Republic on the proposal of the then Minister of Education, University and Research, with the aim of take over the management of the data centre intended to host the Control Data 6600 system, the first supercomputer installed in Italy, acquired in order to support public and private research of the national scientific community. Since then, Cineca has grown both institutionally, currently almost all Italian public universities and all the National Research Institutes are members of the Cineca Consortium, as well as the current Ministry of University and Research, both organizationally, supporting the national academic system also for the management and administrative aspects.

As part of this development process, the founding mission of acting as a national supercomputing centre has remained at the centre of institutional activity. The relevance of Cineca's institutional mission as a supercomputing centre has been re-confirmed by the national mandate to express its candidacy for the European competitive call to host one of the pre-exascale systems acquired and financed by the EuroHPC Joint Undertaking. The Leonardo project, conceived for this purpose, has been positively evaluated in the European competitive process and is currently being implemented.

We celebrated the anniversary of the founding of Cineca and the brilliant success of the Leonardo project in certain occasions during 2019, more widely presented in the following of this report. The dynamic speed of the evolution of our supercomputing architecture and of the inherent services imposed by technological innovation, as well as guidelines to keep our development persistent, is such that the occasion of our annual report is useful for presenting a look into the future. The elements of development are already mentioned in the sections that present the most relevant ongoing activities.

Finally, a section for future trends has been included in this report which describes the objectives we set for the near future. First of all, the completion of the Marconi project with the beginning of production of the MARCONI-100 partition, the description of the Leonardo project for the pre-exascale system, as well as the introduction of a service for access to quantum computing platforms, looking ahead about long terms trends.

Nevertheless, the year 2020 will be a year of intense work for the construction of the new exascale data centre at the Bologna Technopole, which will host the ECMWF data centre, the Cineca HPC data centre and the INFN/CNAF data centre in the same campus. And it will be the year to select the provider to whom to assign a 120 million Euro order for the supply of a supercomputing system with power in the warp of hundreds of PFlop peak. It will also be the year to consolidate the role of Cineca within the national system of excellence in science and industrial innovation, also integrating the national supercomputing system managed by Cineca in the framework of the national big data system of the National Agencies and IT system of the public administrations, to support the public and private research system to better face the challenges of our time, from climate change to personalized medicine, from the circular economy to safeguarding the territory and the patrimony of our cultural and environmental heritage.

With kindest regards,



Sanzio Bassini

Director of High Performance Computing Department of Cineca

## 2019 IN NUMBERS

- 2 50 years of HPC milestones
- 4 HPC Infrastructures
- 6 Resource allocation
  - 7 Prace
  - 9 Iscra
  - 14 Agreements
  - 15 Meteo/Industry
- 16 Usage report
- 18 Users' statistics
- 20 Training
- 21 Staff
- 22 Events

## SCIENTIFIC OVERVIEW

- 30 AI agents that can see, speak and interact
- 32 Searching for water on Mars
- 34 Binary Neutron Star Mergers
- 36 HPC powers large scale de novo drug design
- 38 3D Saint Peter's Cathedral in Bologna
- 40 Last Supper Interactive



## EUROPEAN PROJECTS AND COLLABORATIONS

PRACE 6IP is looking forward to the exascale	44
PRACE SHAPE: projects for the SME	45
HPC for EUROfusion	46
SUPER	47
AIDA	48
Highlander	50
Mistral	52
DARE	54
Cybele	55
Io Twins	56
AI4EU	57

## FUTURE TRENDS

Completion of the MARCONI project	60
EuroHPC: a European and Italian success story	62
Accelerated High Performance Computing	64
Supporting Eni HPC infrastructure	66
National ecosystem for Industry 4.0	67
Quantum Computing	68



2019  
IN NUMBERS

Cineca HPC is the main centre for scientific computing in Italy. We run a large computing infrastructure and make it available to Italian and European researchers, as well as to important Italian companies within a programme for supporting national industrial competitiveness. Here we present data about usage and users, projects and events and, last but not least, our educational activity in the field of High Performance Computing.

# 50 years of HPC milestones

Elda Rossi  
Cineca

Cineca was born in 1969, and this means that 2019 is a very important anniversary for us: the 50<sup>th</sup> year of important milestones reached in the field of Advanced Scientific computing.

The idea of an innovative centralised infrastructure for supporting computational research dates back to 1967 when four Italian Universities (Bologna, Florence, Padua and Venice) signed an agreement for creating an inter-university consortium for the management of a computing centre in the north-east of Italy. The original meaning of the acronym CINECA is in effect: Consorzio Interuniversitario del Nord Est per il Calcolo Automatico (Interuniversity Consortium in the North-East for Automatic Computing).

The statute papers were approved in late 1969 and few months later, in February 1970, they were published on the Italian Official Gazette, thus starting the incredible adventure that goes until today.

The first computer acquired by the Consortium was a CDC 6600, generally considered the first supercomputer, with performance of up to three MegaFlops.

Over the years Cineca continued to acquire very powerful computers, such as the Cray X-MP in 1985, that was the fastest computer at the time with a dual processor system and a performance of 800 MegaFlops.

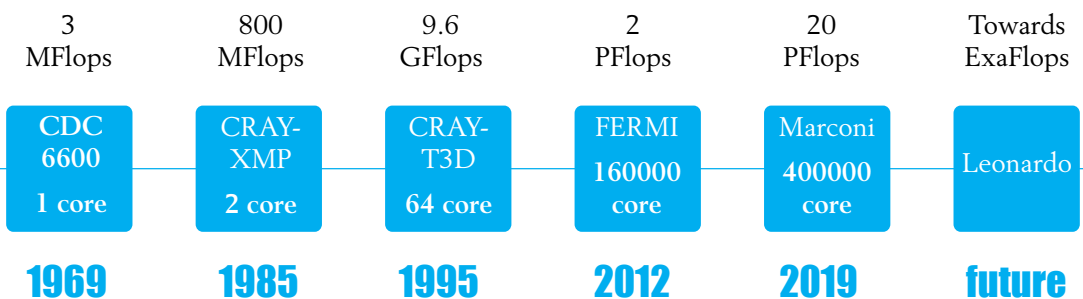
Another important milestone was in 1995 when the new massively parallel architecture took place with Cray-T3D. It consisted of 64 Processing Elements, connected by a three-

dimensional torus network. The performance increased by one order of magnitude, reaching 9.6 GigaFlops.

Less than 20 years after, in 2012, the top machine was FERMI, an IBM BlueGene/Q system that, thanks to its 160 thousand processing elements, was able to increase the performance of 200 thousand times. The system was the 7<sup>th</sup> more powerful system worldwide with a peak performance of 2 million GigaFlops (2 PetaFlops), near 1 billion times more powerful than the first CDC supercomputer.

Today Cineca hosts the Marconi system, a 20 PetaFlops cluster based of conventional and scale-out processors, as described later in this report. For the next near future, we expect a new large increase toward the exascale computing, with the EuroHPC initiative that is expected to bring here Leonardo, a new pre-exascale accelerated computer that will be part of the new pan-European supercomputing infrastructure.

In year 2013 Cineca merged the other Italian consortia for scientific computing (CILEA in Milan and CASPUR in Rome) to give rise to a large supercomputing centre for high performance computing with a national scope. The European dimension is guaranteed by the PRACE initiative (Partnership for Advanced Computing in Europe). Since 2009 Cineca is one of the main partners of PRACE by hosting on its supercomputing systems the numerical projects of the European scientists.



**1969:**

The first Italian supercomputer in Cineca, a CDC 6600, designed by the father of the supercomputing Seymour Cray



**1985:**  
Cray-XMP  
The first Cray supercomputer in Cineca



# HPC infrastructures

Elda Rossi  
Cineca

In year 2019, our main HPC system (Marconi) remained unchanged. It is our Tier-0 system that started its life in mid-2016, with the set-up of the first partition (A1) based on Broadwell chips and a computational power of 2 PFlops peak. After several upgrades and inclusions, at the end of 2019 it was configured in two partitions: A3, made of 3216 SkyLakes nodes; and A2, a scale-out partition made of 3.600 many-cores node (KNL), with a total peak performance of about 20 PFlops. This supercomputer takes advantage of the Intel® Omnipath architecture, which provides the high performance interconnectivity required to efficiently scale the system's thousands of servers. A high-performance Lenovo GSS storage subsystem, that integrates the IBM Spectrum Scale™ (GPFS) file system, is connected to the Intel Omni-Path Fabric and provides data storage capacity for about 10 PByte net. Marconi was ranked for the TOP500 list in Novembre 2018 and reached position no. 19 with a sustained performance (Rmax) of 10385 TFlops.


During the year, the main event affecting the HPC infrastructure has been the enhancement of Galileo, our Tier-1 "National" level system. Galileo, introduced first in 2015, then reconfigured at the beginning of 2018, was further enhanced in August 2019. Now the system is made of 1022 nodes (Intel Broadwell) connected with an Intel Omnipath, 100Gb/s internal network. Sixty nodes are equipped with NVIDIA K80 GPUs. This is the Tier-1 system, available for the Italian research community. The last HPC cluster of the infrastructure, still at the Tier-1 level, is D.A.V.I.D.E. (Development of an Added Value Infrastructure Designed in Europe), based on Power Architecture and coupled with NVIDIA GPUs with NVLink. Another big enhancement in year 2019 has been the total upgrade of the internal connectivity of the HPC environment. Now it is based on new equipment supporting a 100 Gbit Ethernet, thus allowing for a faster and most robust interconnection of the clusters and the storage facilities.

## Tier-0

	CPU	Total nodes/ cores	Memory/node
Marconi A2 - KNL	Intel Knights Landing 1x Intel Xeon Phi7250 @1.4GHz 68 cores each	3600 nodes 244800 cores	96 GB
Marconi A3 - SKL	Intel SkyLake 2x Intel Xeon 8160 @2.1GHz 24 cores each	3216 nodes 154368 cores	192 GB

## Tier-1

	CPU	Total nodes/ cores	Memory/node
Galileo Lenovo NextScale Omnipath 100Gb/s	Intel Broadwell 2 x Xeon E5-2697 @2.1GHz 36 cores each	1022 nodes 36792 cores	128 GB
D.A.V.I.D.E. OpenPower cluster Infiniband EDR (2Gb/s)	D.A.V.I.D.E. OpenPower cluster Infiniband EDR (2Gb/s) 2 x Power8+ 4 Tesla P100 SXM2 16 cores each	45 nodes 720 cores	256 GB

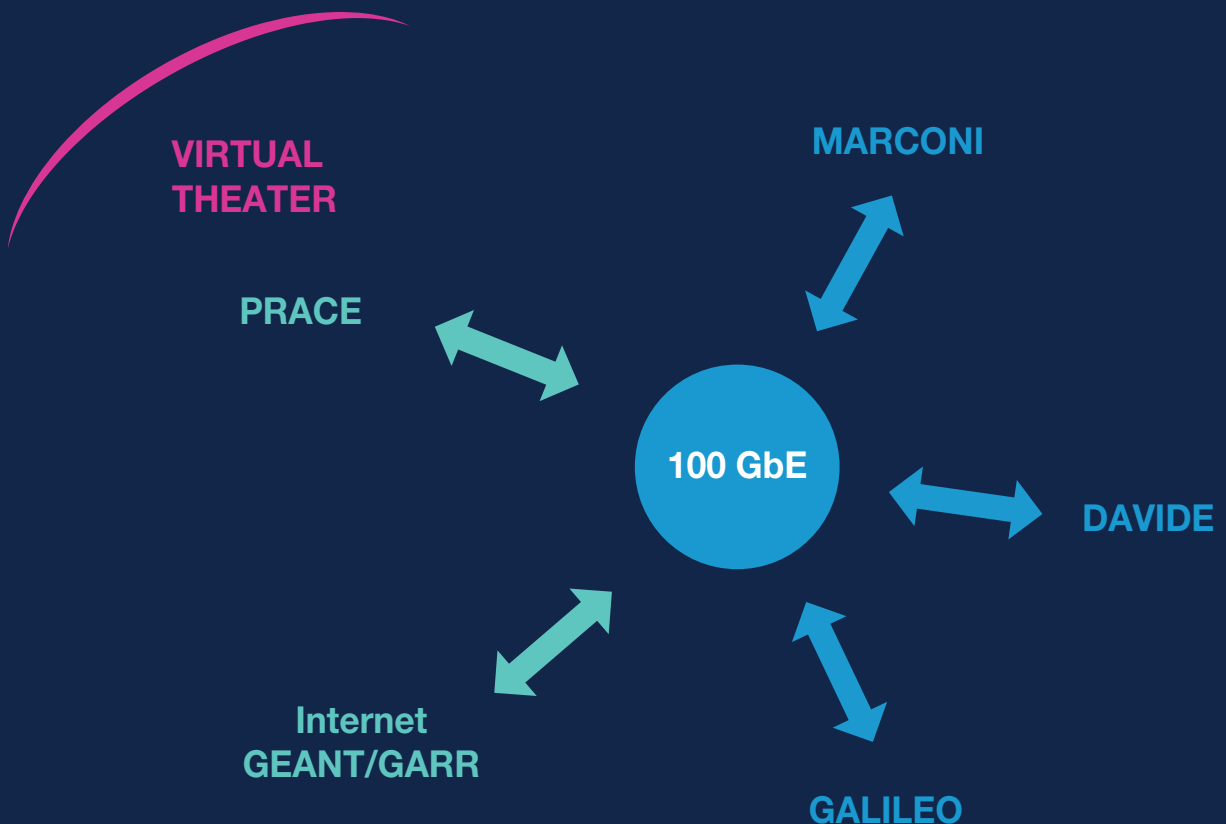


### Million core-h computational capacity

The computational capacity of these systems is usually expressed in terms of core-h, the number of computing hours that can be exploited for each single core of the system. In the picture above the total annual capacity is reported for the whole HPC environment, for a grand total of 3.7 billion core-h.

Galileo:	187 Mcore-h
Marconi A2:	2144 Mcore-h
Marconi A3:	1340 Mcore-h
D.A.V.I.D.E.:	6 Mcore-h

3677 Mcore-h



# Resource allocation

The HPC clusters in Cineca are used by Italian and European researchers for running their computational research. The access to the HPC facilities is based on “computational projects”. Some of them are based on peer-review mechanisms: PRACE is at European level; ISCRA is for national researchers. Other projects, defined under the category “Agreements”, refers to collaborations with the main Italian research Institution; “EUROfusion” reports for projects activated within the agreement with the European consortium for nuclear energy; “Industrial” represents activities for supporting the national industrial competitiveness.

Each project has also a starting and ending date, a budget (in terms of core-h), a PI (Principal Investigator), several collaborators, and a science domain.

The numbers below represent the allocated resources, in terms of million core-hours, for the different type of projects. The EUROfusion

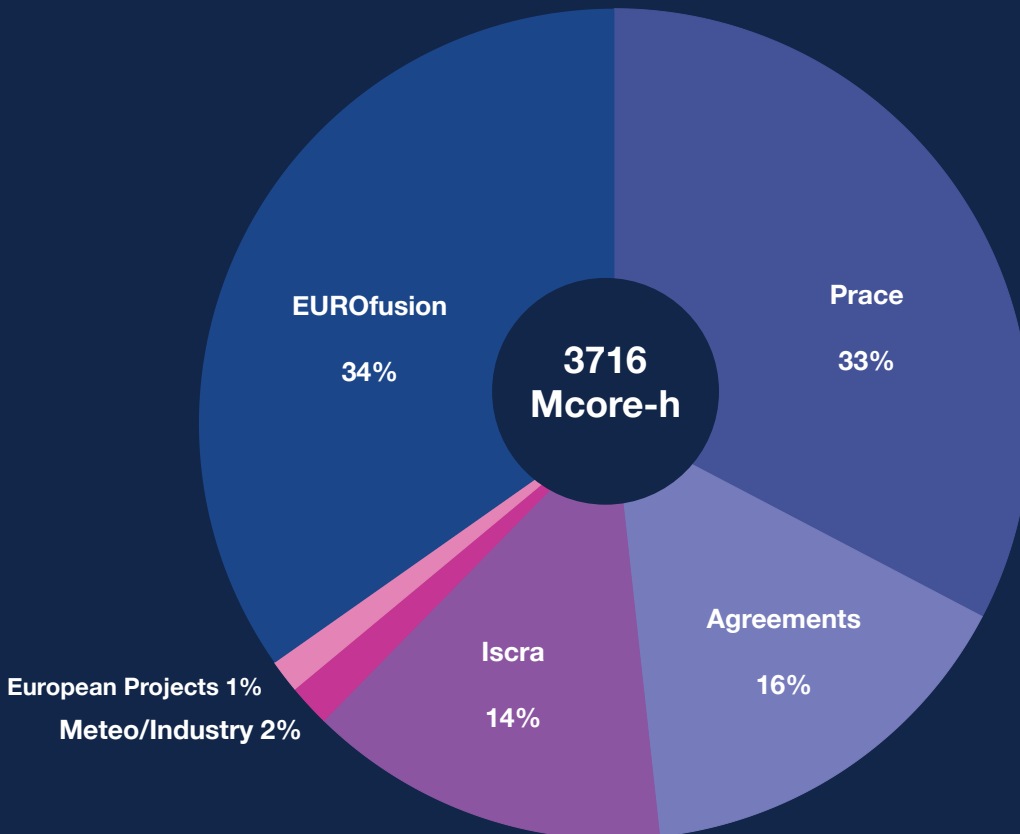
consortium has a dedicated partition on the Marconi-A2 cluster (449 nodes), on Marconi-A3 (2410 nodes) and on D.A.V.I.D.E. (40 nodes). On these partitions the allocation is made autonomously by the dedicated committee. The numbers reported correspond to the whole partition capacity.

On the Galileo cluster there is a dedicated partition of 124 nodes for weather forecast activities and 48 nodes for industrial projects. Also in this case, the numbers reported for dedicated partitions correspond to the whole capacity of the partition itself.

One third of the resources are distributed to PRACE, one third to EUROfusion and the last third equally distributed between agreements and ISCRA projects.

If we focus only on the shared partitions, in total, on all HPC clusters in Cineca, the allocations have covered all available resources: the rate of coverage has been 118%.

**Resource allocation for project categories: percentage of core-h**





# Prace

Massimiliano Guarrasi  
Cineca

PRACE (Partnership for Advanced Computing in Europe) gives computational resources and technical support to European researchers to enable new high-impact scientific discoveries. It provides HPC resources through Preparatory Access (for code scaling and optimization) and through Project Access (for production). Project Access calls (2 per year) are intended only for large-scale, computationally intensive projects.

In 2019 Cineca was one of the major contributors of PRACE, making available 1211 Mcore-h on the Marconi cluster.

Two new calls (18 and 19) were launched in year 2019, awarding 27 projects, with respect of 82 total projects all over Europe.

Over 1126 Mcore-h were allocated, with respect of 3320 M total.

			CINECA		Tot. PRACE	
CALL	Start	End	N. of awarded projects	Awarded resources (million core hours) Marconi KNL	N. of awarded projects	Awarded resources (million core hours)
18 (Project Acc.)	02/04/2019	01/04/2020	13	526	35	1420
19 (Project Acc.)	01/10/2019	30/09/2020	14	600	47	1900
<b>TOTAL</b>			<b>27</b>	<b>1126</b>	<b>82</b>	<b>3320</b>
Prep. Acc.	01/01/2019	31/12/2019	25	3		



# Iscra

Paola Alberigo  
Cineca

ISCRA (Italian SuperComputing Resource Allocation) gives computational resources and technical support to researchers affiliated to Italian institutions, in order to support the Italian research.

It provides HPC resources through Class C projects (for code testing and pre-production) and Class B projects (full production).

IscraC projects are evaluated and selected on a

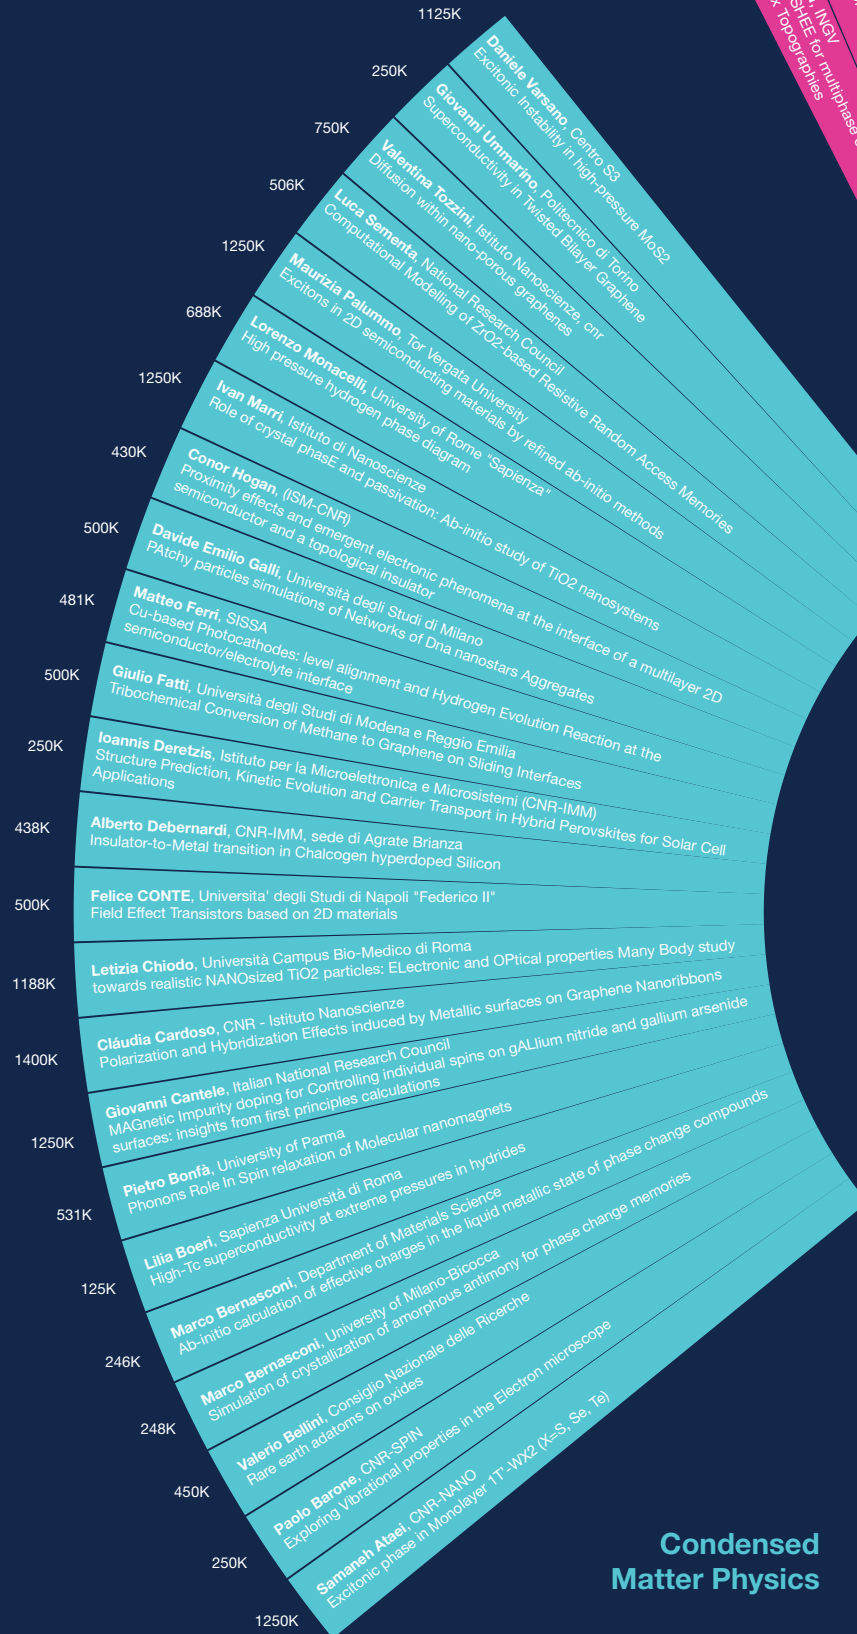
continuous basis, IscraB projects twice a year. In 2019 Cineca provided to the Italian research community 521 Mcore-h on Galileo and Marconi.

In particular two IscraB calls were launched as reported in the table below (19 and 20). We report also the last call of 2018 (call 18) which run mainly during 2019.

			CINECA		
CALL	Start	End	N. of awarded projects	Awarded resources (million core hours)	
				Galileo	Marconi KNL
19 B	10/05/2019	09/05/2020	32	18	122
20 B	09/10/2019	08/09/2020	35	9	217
<b>TOTAL 2019</b>			<b>67</b>	<b>27</b>	<b>339</b>
18 B	30/11/2018	31/11/2019	42	34	133

The awarded Iscra projects, grouped for science domain, for each one the core-h allocated

## Earth and Climate Science



# ISC

## Condensed Matter Physics

# RA

## Computational Chemistry

1250K	<b>Gianfranco Bocchinfuso</b> , Università di Roma "Tor Vergata" Activation Mechanism of Phosphatase SHP2 Investigated through REMD and Simulations
875K	<b>Carlo Camilloni</b> , University of Milano Protein dynamics and toxicity in light chain amyloidosis
492K	<b>Mauro Causà</b> , Università degli studi di Napoli Federico II Computational interpretation of spectroscopic features of defective Metal Organic Frameworks (MOF)
288K	<b>Giovanni Di Liberto</b> , Università degli studi di Napoli Federico II Understanding metal-support interactions in model catalysis operating in electrochemical and high photoactivity
813K	<b>Suzana Drita</b> , Università di Padova Modelling of Functionalized Gold Surfaces for Biosensing
481K	<b>Matteo Farnesi Camellone</b> , CNR-IOM c/o SISSA Substitutional atomic defects in monolayer-protected clusters for controlled release
250K	<b>Giuseppe Forte</b> , Catania Adsorption Behavior of Water Molecules at the Magnetite/Water Interface
400K	<b>Alessandro Fortunelli</b> , CNR Design of Chirality in Monolayer-protected Clusters (MPC)
492K	<b>Francesco Gentile</b> , Università degli studi di Torino Perovskite Interfacial Catalysis by Confinement in Energy and Environment
41K	<b>Hongsheng Liu</b> , University of Milano-Bicocca Light-driven chiral transitions in excited states of conjugated polymers
233K	<b>Matteo Maestri</b> , Dipartimento di Ferrara A systematic first-principles investigation of catalysis by confinement in ionic liquids
1238K	<b>Simone Meloni</b> , University of Ferrara PeRovskite Interfacial Engineering with Ionic Liquids
625K	<b>Adriana Pietropaolo</b> , Consiglio Nazionale delle Ricerche Injection Rates in OrgaNIC Electronic Devices
1090K	<b>Alessandro Pecchia</b> , University of Catanzaro Design of a protein mimicking a broadly neutralizing Dengue virus antigen
401K	<b>Walter Rocchia</b> , IIT-Istituto Italiano di Tecnologia Light-driven chiral transitions in excited states of conjugated polymers
444K	<b>Angelo Spinello</b> , CNR-IOM Design of a protein mimicking a broadly neutralizing Dengue virus antigen
369K	<b>Sergio Paolo Tosoni</b> , Università di Milano-Bicocca Unravelling the mechanism of action of novel small-molecule splicing modulators: A step forward towards precision medicine of hematologic cancers
1000K	<b>Piero Ugliengo</b> , Università degli Studi di Torino Ab-initio molecular dynamic simulations on astrochemical reactions occurring on interstellar icy mantles
500K	<b>Andrea Vittadini</b> , CNR Porphyrin Extraction of Titanium from Rutile and Anatase titanium dioxides

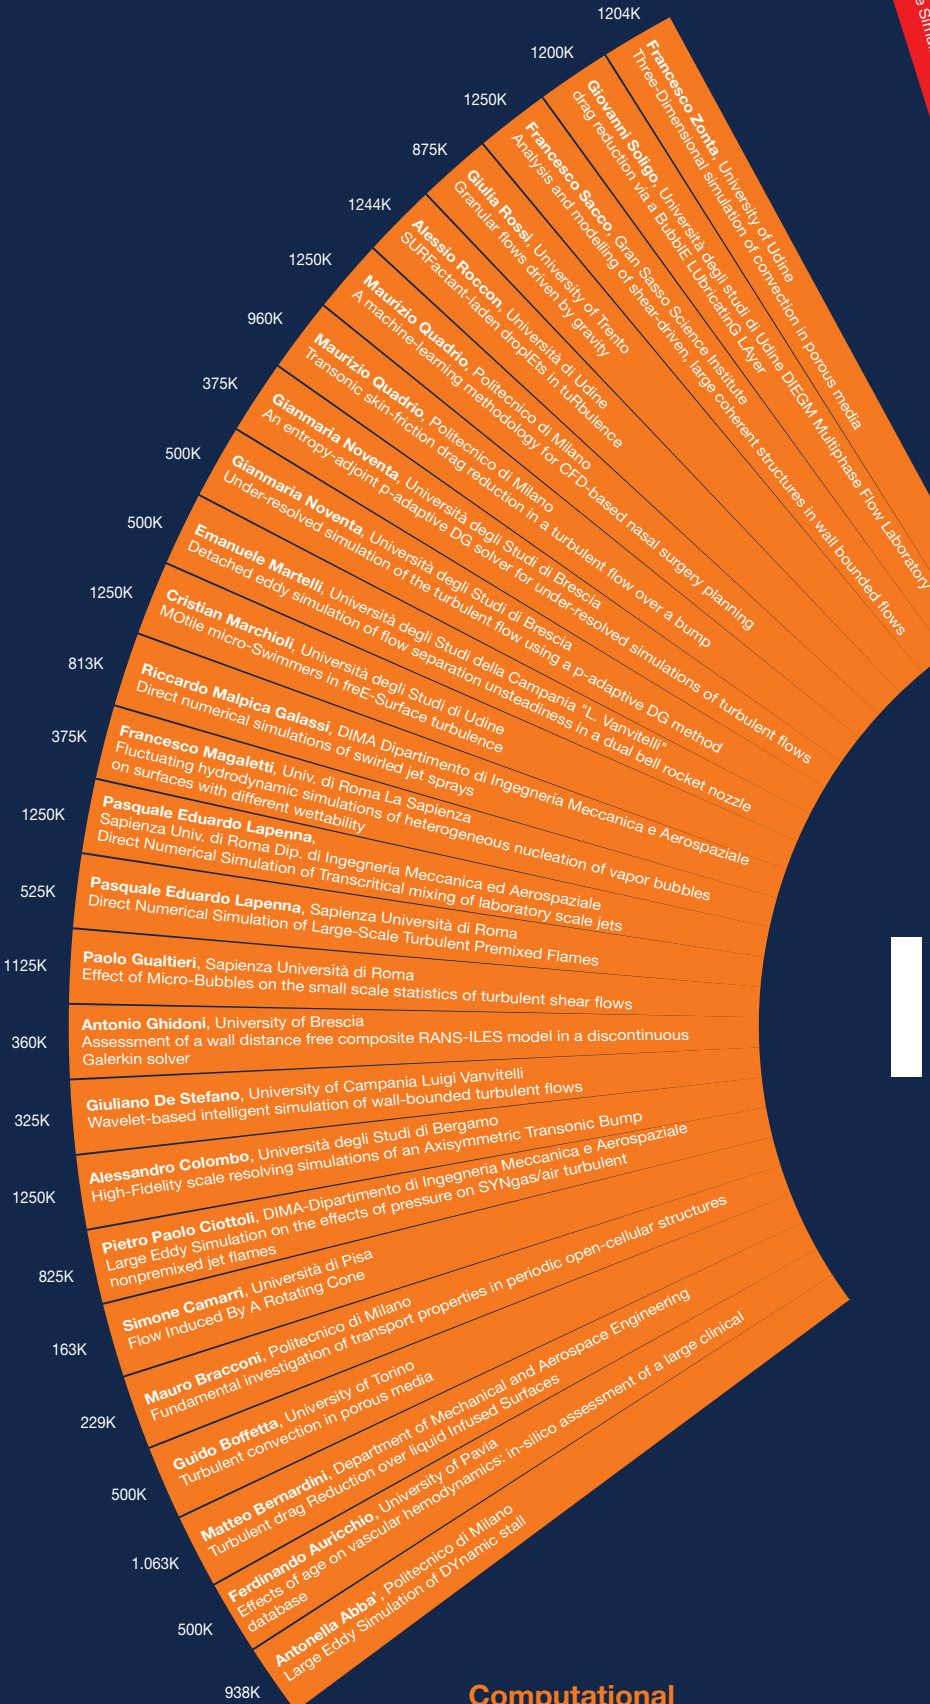
## Life Science Computational Biology

490K	<b>Laura Bonati</b> , University of Milano-Bicocca Quaternary Architecture of the AhR:ARNT Protein Dimer investigated by Multiscale Molecular Dynamics simulations
500K	<b>Grazia Cottone</b> , University of Palermo Studying the function of Botulinum Neuro Toxins through the analysis of their conformational dynamics and of their inter-molecular interactions
484K	<b>Alessandra Magistrato</b> , DEMOCRITOS Modeling Center for Research in aTOMISTIC Simulation Unraveling the structural/functional impact of the Estrogen Receptor alpha N terminal domain via Experimentally Driven Molecular Simulations
567K	<b>Luca Maragliano</b> , Istituto Italiano di Tecnologia Molecular Dynamics Simulations of Ion Permeation and Selectivity in Human Sodium Channels
333K	<b>Francesco Musiani</b> , University of Bologna Going down the routes of metal trafficking through metadynamics
1175K	<b>Francesca Spyraakis</b> , University of Turin Investigating Nuclear Receptors Dynamic cross-relationships
500K	<b>Egilio Ugo D'angelo</b> , University of Pavia Realistic cerebellar multi-scale network reconstruction

## Life Science Bioinformatics

The awarded Iscra projects, grouped for science domain, for each one the core-h allocated

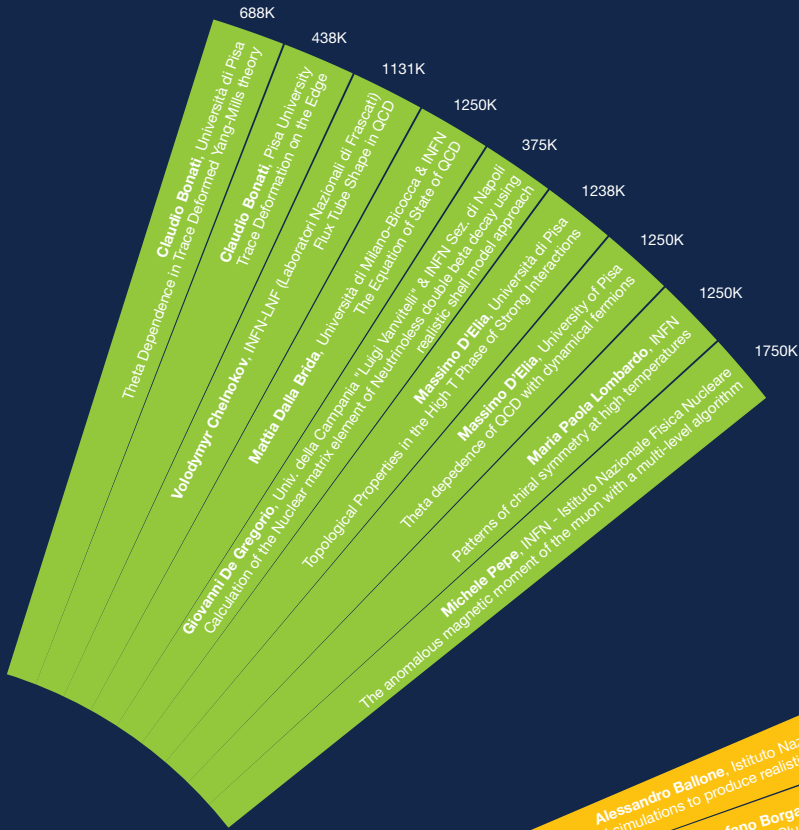
## Computational Engineering



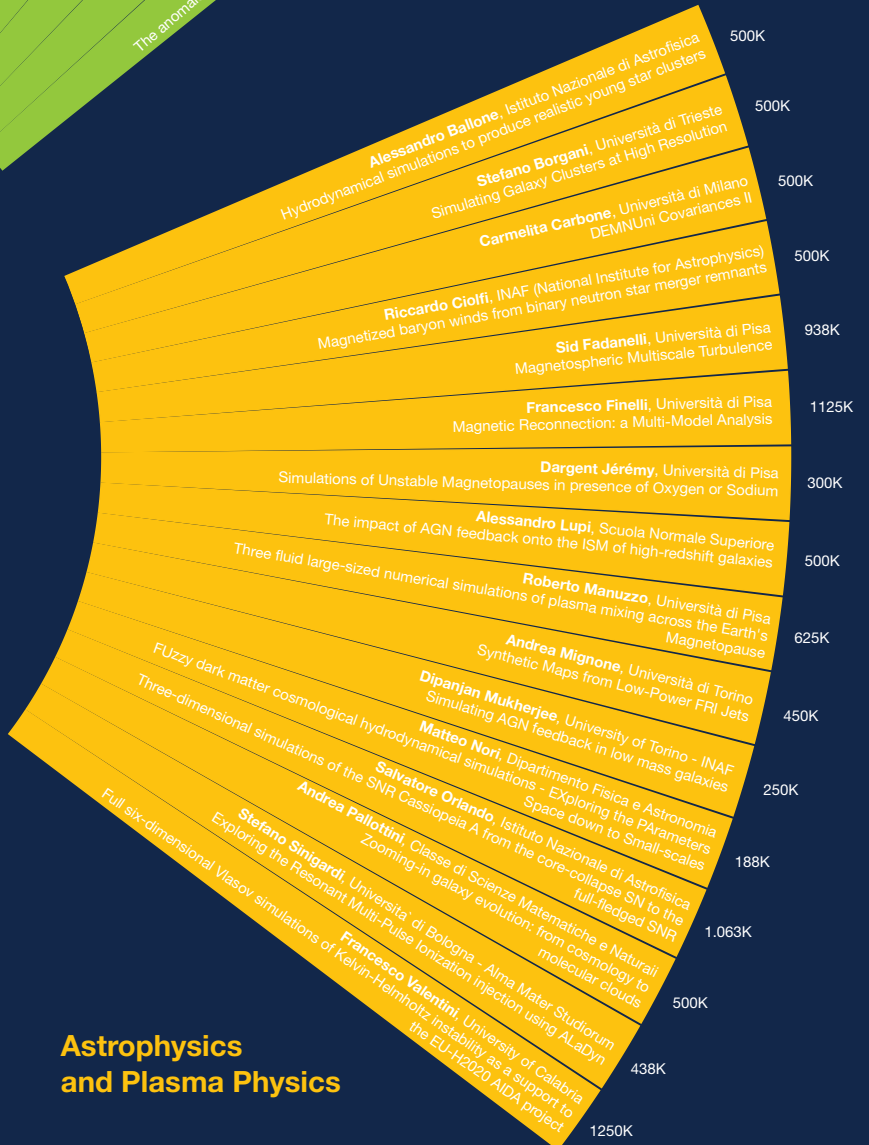
# ISC

## Computational Fluid Dynamics

## Particle Physics



# RA



## Astrophysics and Plasma Physics

# Agreements

Maurizio Cremonesi  
Cineca

Cineca computing platforms may be accessed also by agreements for collaborations.

Indeed several Universities and Research Institutions decided to support their researchers by defining agreements reserving a certain amount of computing and storage resources for their researchers and technicians.

These agreements grant access to resources from one to a few years, without the need of submitting a project proposal, so they are particularly suited for special projects and small experimental activities.

The research and academic institutions with active agreements in 2019 are listed below.

Among them, there are also well known public

institutions in the field of Genomics and Life Science research.

In total 577 Mcore-h were allocated to agreement projects. The largest availability corresponds to INFN that, together with SISSA, ICTP and University of Milano Bicocca, in 2018 signed a three years agreement for the exploitation of Marconi A3.

INAF holds an important agreement (64 Mcore-h in 2019) thanks to a three years activity on Marconi A3.

The institutions in the field of Life Science research all together sum up to a 6 Mcore-h grant. A similar amount corresponds to the other 8 Universities and Polytechnics.

Institution	Award resources (Mcore-h)
INFN (Istituto Nazionale Fisica Nucleare)	301
SISSA (Scuola Internazionale Superiore di Studi Avanzati)	108
ICTP (international Centre for Theoretical Physics)	49
UniBicocca (Università degli Studi di Milano Bicocca)	35
INAF (Istituto Nazionale di Astrofisica)	64
IIT (Istituto Italiano di Tecnologia)	7
OGS (Istituto Nazionale di Oceanografia e Geofisica Sperimentale)	
Telethon	6
Elixir (Distributed INfrastructure for Biological Data)	
OspPBG (Ospedale Pediatrico Bambin Gesù)	
INMI (Istituto Nazionale Malattie Infettive Spallanzani)	
SIGU (Società Italiana di Genetica Umana)	
Istituto Ronzoni	
Università di Trieste	7
Politecnico Milano	
Università Milano	
Università di Genova	
Università di Brescia	
Università di Bergamo	
Università Bocconi	
<b>TOTAL</b>	



# Meteo/Industry

Claudio Arlandini, Elda Rossi and Cinzia Zannoni  
Cineca

## Innovation projects

In the evolutionary context of technologies and paradigms linked to the fourth industrial revolution, better known as Industry 4.0, HPC and Data analytics have assumed an absolute centrality being the enabling factor of this revolution. Towards industries, and especially SMEs, Cineca acts as a Digital Innovation Hub, that is a one-stop-shop that help companies to become more competitive with regard to their business/production processes, products or services using digital technologies. The industries that are regular customers or are developing a project with Cineca include:

- Dompé, a major Italian biopharmaceutical company, with whom Cineca is developing the EXSCALATE project, the first computational platform dedicated to addressing the need to promptly respond to a pandemic crisis
  - Nolangroup, leader in the design and production of motorcycle helmets for leisure and sport
  - Axyon AI, a fintech start-up providing financial time series forecasting and predictive insights with its AI-powered platform
- In 2019 industrial partners were allocated significant computational resources on the Galileo cluster, typically in dedicated mode.

## Collaboration with ENI

The collaboration with Eni (the most important Italian Energy Company) dates back to over 50 year ago, consisting in the set-up of the company HPC ecosystem, development and implementation of advanced applications,

and management of the company HPC systems. Since 2013, when Eni built its own Computing Centre to host its proprietary HPC infrastructure, Cineca continues to assure the system management service. In addition to that, in 2019 Cineca reserved to Eni a dedicated partition on Galileo to help the development of the proprietary applications.

## Meteo/climate services

Over nearly 30 years, Cineca is the closest partner of Arpae (Agenzia regionale per la prevenzione, l'ambiente e l'energia dell'Emilia-Romagna) in activities related to the national weather forecast service. In 2019 Cineca continued to assure the service related to the last tender won in 2018, providing supercomputing resources for predictive modeling weather and sea organized by Hydro-Climate-Weather service (Arpae-SIMC) of Arpae.

Cineca is also collaborating with OGS (National Institute of Oceanography and Applied Geophysics) to produce data on bio-geo-chemical analysis of the Mediterranean sea through the project "Copernicus: marine environment monitoring service", a service that provides information at European level on the quality of the sea and on the environmental conditions of seas and oceans skirting the continent.

Other partnerships, with Arpa-Piemonte, CIMA (Centro Internazionale in Monitoraggio Ambientale) and CMCC (Centro Euro-Mediterraneo sui Cambiamenti Climatici), relate with simulations and predictive modelling both in meteorological and weather-marine in relation to the quality control air and seas.

In order to guarantee a suitable level of service, these collaborations rely on dedicated resources on the Galileo HPC cluster.

# Usage report

In year 2019, out of a theoretical capacity of 3677 million core-h, the usage was of 3386 core-h, the 92% of the whole capacity.


The usage data of the dedicated partitions on Marconi (for EUROfusion) and on Galileo (for meteo/industrial collaborations) was not measured, but set equal to the total capacity of the partitions.

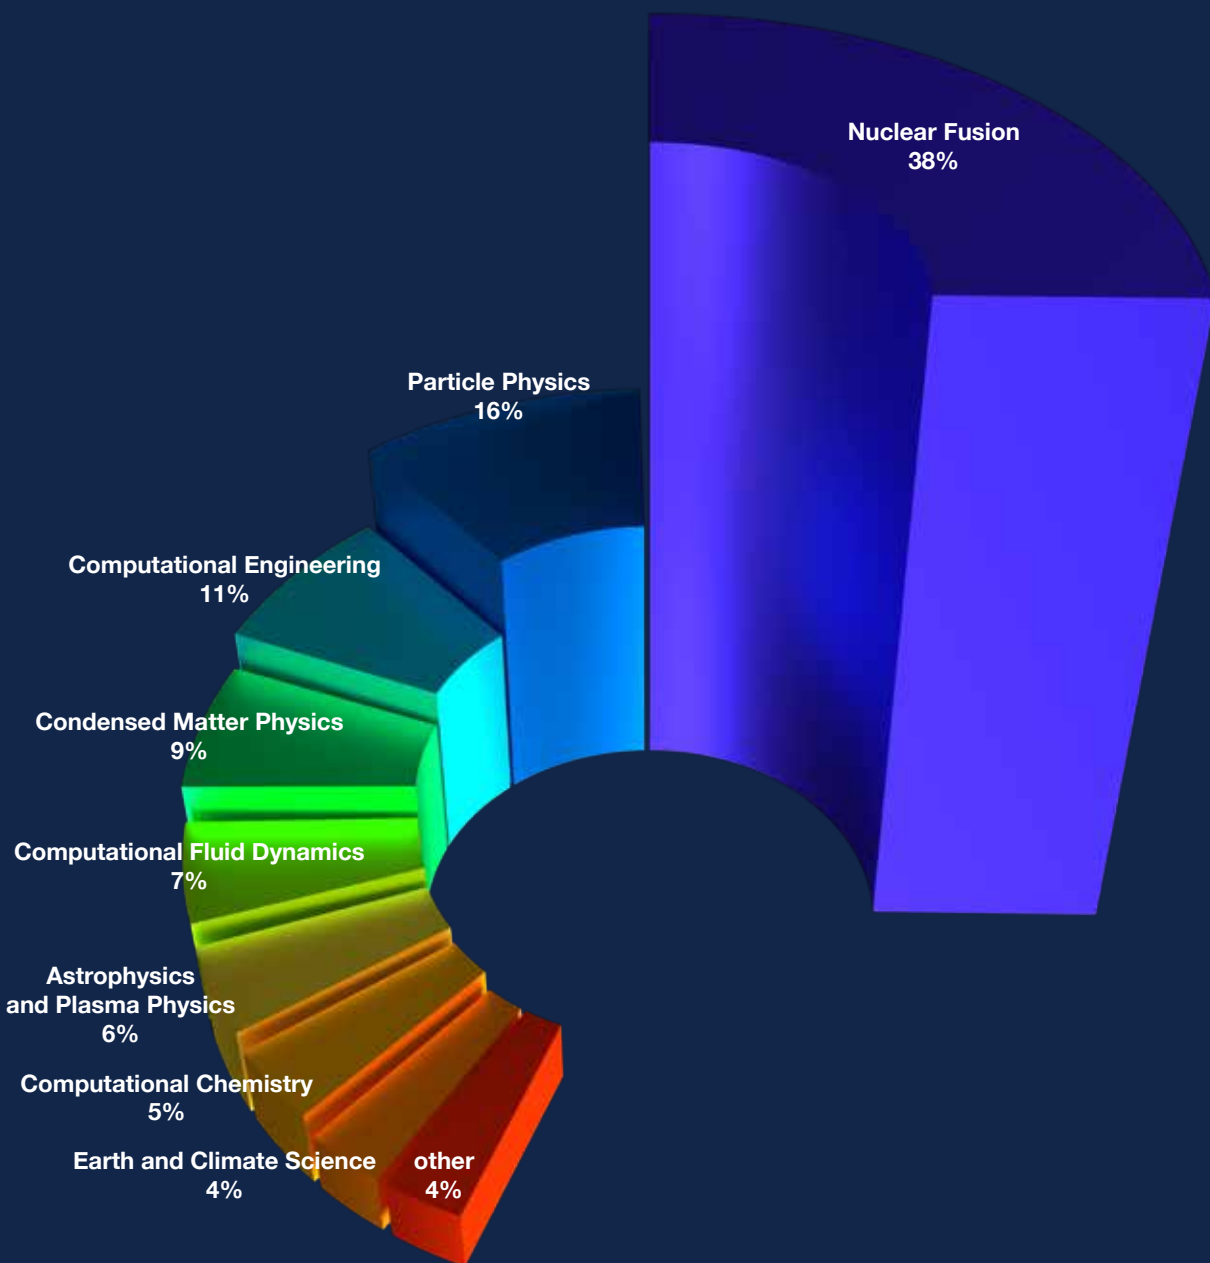
Nearly one third of resources was used by PRACE projects, 38% by EUROfusion, another third equally distributed among Iscra and

Agreements projects.

With respect of the different scientific domains: Nuclear Fusion is the first one (due to the presence of EUROfusion researchers), followed by Particle Physics, thanks to the collaboration with INFN (National Institute for Nuclear Physics) and Computational Engineering. Nevertheless, the most traditional disciplines of Chemistry, Condensed Matter Physics and Astrophysics are still well represented with 22% of usage.

Effective resources use for project categories: percentage of core-h





Core-h on HPC clusters used by scientific areas in 2019

	Capacity (Mcore-h)	Assigned (Mcore-h)	Used(*) (Mcore-h)
Galileo	187	202	129
Marconi A2	2144	2174	1938
Marconi A3	1340	1333	1312
D.A.V.I.D.E.	6	7	7
	<b>3677</b>	<b>3716</b>	<b>3386</b>

(\*) reserved partitions are considered fully allocated & used


# Users' statistics

At the end of year 2019 the active (personal) users on our HPC systems were 3602.

The great majority of the users are males (79%), aged between 31 and 45 (48%), and working for Italian institutions (68%), mainly Universities or other public research bodies in Italy (63%) or abroad (28%). The large Italian cities, clustering multiple institutions, are well represented: Milan (552 users), followed by Rome (363), Trieste (294), Bologna (270) and Turin (146). Among the more represented foreign nationalities: Germany and France (thanks to


the EUROfusion community), Spain, India and United Kingdom.

Our users represent scientists in all scientific disciplines that take advantage of computational resources, the most represented being: Chemistry and Material Science (23%), Computational Fluid Dynamics (13%), Astro and Plasma physics (11%), Life science (11%). In the picture in the opposite page the number of users belonging to best represented Italian institutions are reported.

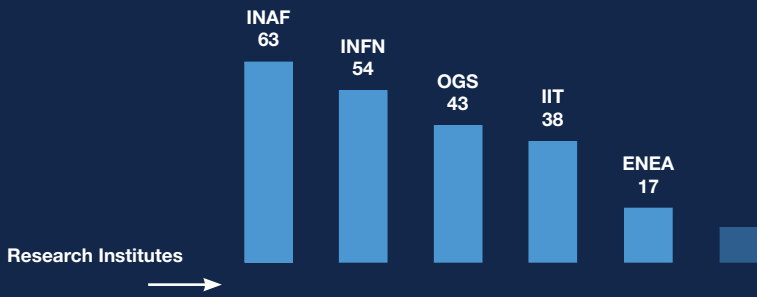
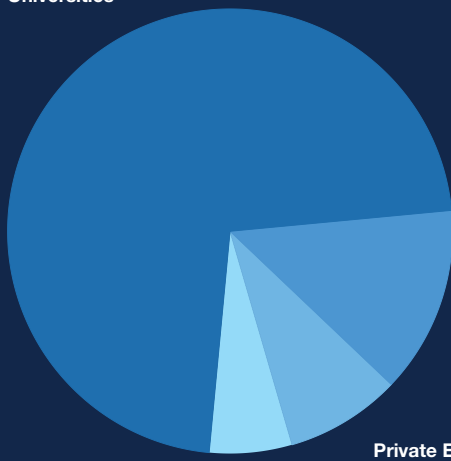


Courtesy of BVTk - ECMWF summer of weather code 2019  
Lorenzo Celli, Silvano Imboden, Gabriele Brizzi, Enzo Papandrea, Antonella Guidazzoli


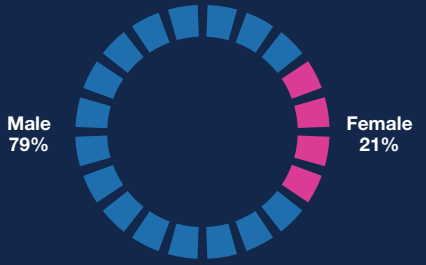

Temperature anomalies between 1981 and 2019.  
(Blue, less than 2.5° Celsius; red, higher than 2.5° Celsius)



↑ Universities



Private Entities  
CNR



# Training

Francesco Falciano  
Cineca

Training has always been a hallmark of the support activities carried on by Cineca for the Italian research community and users. Starting from 2012 it was extended to European researchers, Cineca being recognised as a PRACE Advance Training Centre (PATC) in HPC.

During 2019 we reached quite impressive numbers: 26 classes, distributed across the three sites of the consortium (Bologna, Rome and Milan) and 6 schools. More than 50 people in the department, for a total of 119 days of lessons and highly specialised training. In total, more than 530 Italian and European researchers took advantage of our training program, with high satisfaction levels (8.4/10), as shown by the surveys we regularly administer.

Our experts were involved in 3 academic courses (Masters and Doctorate schools) in several Italian Universities teaching about HPC

and Big data topics.

Furthermore in 2019 we performed Cineca guided tours for students belonging to 20 high school coming from all the country.

Other important initiatives: Summer of HPC (a PRACE initiative offering young students from all over Europe to work with HPC teams); European Researchers Night (Cineca is the coordinator of SOCIETY, a European project organising the Night in Bologna and surroundings in 2018-2019); HPC-Europa3 (we are the coordinator of the project offering transnational access to researchers visiting HPC centres).

The Training activity is regulated by two specific procedures that refers to the ISO standard: ISO 27001:2013 "Information Technology - Security techniques - Information security management system".

**530**  
Italian and European  
Students

**119**  
Days of lectures

**26**  
classes  
(8 PTC)

**20**  
Guided tours

**50**  
Teacher

**6**  
Schools  
(3 PTC)

**3**  
Academic Classes

# STAFF

In the HPC department of Cineca work about 88 people, distributed among the three premises: 65 in Bologna, 7 in Milan and 16 in Rome.

Several new colleagues were enrolled in 2019: we wish them a very fruitful and interesting job.

## **Daniele Cesarini**

I graduated in Computer Engineering at the University of Bologna, where I also earned a PhD degree in Electronics, Telecommunications, and Information Technologies Engineering. In Cineca I am currently working in the area of performance optimization on large-scale scientific applications for new generation of heterogeneous HPC architectures. My expertise include parallel programming model, shared and distributed memory systems, high-performance computer architectures and runtime systems.

## **Balasubramanian Chandramouli**

I hold a PhD and strong research experience with an expertise in Data Analytics. I work in the Data Management & Analytics division at Cineca, where I am involved in applying machine learning techniques for mining high dimensional data.

## **Beatrice Chiavarini**

I got a degree in Architecture at the University of Ferrara. After my graduation I specialized in the field of web developing and I collaborated with Cineca Visit Lab as a developer for Augmented and Virtual Reality applications and online interactive 3D environments. In Cineca I work as a web and WebGL developer for "Data Management & Analytics" group.

## **Roberto Da Vià**

I got my PhD degree in Mechanics and Advanced Engineering Sciences at the University of Bologna. During the PhD I gained expertise in the field of Computational Fluid Dynamics, in particular for the simulation of heat transfer in liquid metal turbulent flows. At Cineca I work in the HPC - POC solution group where we provide support for CFD and CAE HPC applications to industrial partners and user community.

## **Antonio Memmolo**

I got my Ph.D degree in Aeronautical and Space Engineering at Sapienza - Università di Roma, where I have also been a Postdoc for a year. I have gained expertise in Computational Fluid Dynamics, playing with massively-parallel in-house codes used to run Direct Numerical Simulations of Navier-Stokes equations. During both my Ph.D and Postdoc I've learnt how to fix both software and hardware problems, also causing some of them.

## **Riccardo Mengoni**

I obtained my master degree in theoretical physics at University of Camerino with a thesis on quantum computing. I then moved to the computer science department of University of Verona for my PhD, where I worked on quantum machine learning. In summer 2018 I was selected for the USRA Quantum Academy that gave me the opportunity to be a NASA intern working on quantum annealing applications. Here at Cineca I was granted a scholarship on "Quantum Computing and Quantum Annealers".

## **Lara Querciagrossa**

I achieved a PhD in Chemical Science at the University of Bologna with thesis title "Modelling and simulations of some anisotropic soft-matter systems: from biaxial to chiral nematic". I held various postdoc research positions at the same University. Using Molecular Dynamics and mainly coarse-grained models, I have studied various systems including biological systems, liquid crystals and organic semiconductors.



# Events

## Summer of HPC

**Massimiliano Guarrasi**  
Cineca

Summer of HPC (SoHPC) is an initiative funded by the European Union through the PRACE project, which aims to allow university students from all over Europe to have a study and work experience abroad at major European computer centres. The selected students work under the supervision of a mentor (generally a researcher) and a company tutor (coming from the computer centre that will host them), on projects related to High Performance Computing (HPC) and scientific visualization using calculation resources provided by PRACE (see page 46 for more detail on this European project).

Italy through Cineca is one of the 5 PRACE hosting members (the others are Germany, France, Spain and Switzerland), i.e. it provides computing resources on the basis of competitive calls to researchers from all over the world.

PRACE Summer of HPC aims to increase the interest of new generations in HPC and related technologies by disseminating the experiences of program participants via social media.

During 2019, Cineca hosted all the selected students for the first one week collective training, after that each couple of students reached the assigned super computing centre. The school was focused on parallel programming, visualisation and social media for the dissemination.



As the previous year two students worked at Cineca during the 2019 summer on two different projects:

Martin Molan (Exploring HPC faults),

Li Juan Chan (In Situ/Web Visualisation of CFD Data Using OpenFOAM).

**First week training school at Cineca**





# Events

## Leonardo

Patrizia Coluccia  
Cineca

In the framework of the pre-events to the European Researchers' Night, Cineca organized the "Leonardo pre-exascale" event to present the supercomputer Leonardo to the general public and celebrate 50 years of High Performance Computing in Italy. The event was the occasion to retrace the evolution of computational research over the last five decades and to consider the perspectives Leonardo could open up in different scientific fields. Directors and representatives of some of the most important Italian Research Institutions took part in the event on behalf of the wider Italian scientific community involved in high performance computing.

Sanzio Bassini, Head of Supercomputing at Cineca, commented: "With the Leonardo project, the European Commission recognizes the important path that Cineca has taken in the last fifty years, hosting HPC systems capable of enabling research and innovation. In addition to the classic research areas, Leonardo opens up further opportunities that can be accomplished today because we have the skills to combine computing power, big data and knowledge".

The topics addressed in the event covered the steps that led from the Sixties and Seventies, and the first applications in Italy of high performance computing in various disciplines, to the near exascale systems, and ranged from high energy physics to meteorology, astrophysics as well as cultural heritage.

Attention focused also on the importance of the teaching programs in the field of Scientific Computing, and the role of supercomputing in industrial research.

Roberto Viola Director General of the DG Connect - EU Commission, via conference call from Brussels, stated: "With the founding of Leonardo, the aim of the EU Commission is to promote collaboration among the European Research institutions, building a network of knowledge and collaboration between all centres that can profit by the use of supercomputing, but also to provide supercomputing resources to the SMEs and start-ups, and enhance the innovation process".



The event was moderated by Federico Taddia (scientific journalist), with the participation of:

- Roberto Viola (Director General of DG CONNECT at the European Commission)
- Antonino Rotolo (pro-rector for Research at the University of Bologna),
- Patrizio Bianchi (Emilia-Romagna regional Councillor),
- Sanzio Bassini (Director of the HPC department of Cineca),
- Stefano Tibaldi (CMCC Foundation),
- Tiziana Paccagnella (Director of the ARPAE-SIMC),
- Claudio Zannoni (University of Bologna),
- Roberto Orosei (INAF),
- Franz Fischnaller (Albertina Academy of Fine Arts of Turin),
- Antonio Zoccoli (President of INFN).

# Events

## The European Researchers' Night

Simona Caraceni  
Cineca

The European Researchers' Night (ERN) is an initiative promoted by the European Commission since 2005 that involves thousands of researchers and research institutions every year in all European countries.

It takes place every year across Europe and neighboring countries on the last Friday of September. The goal is to create opportunities for researchers and citizens to meet and spread scientific culture and knowledge of the research professions in an informal and stimulating context. In Bologna and in the cities of the University of Bologna Campuses the ERN is organized by Society, founded from European Community in the last four years. Society is a partnership of all the major local Research Institutes such as (in alphabetic order) Cineca, that coordinate the project, CNR, INAF, INFN, INGV, the University of Bologna and the scientific communication partner ComunicaMente.

Researchers involved in the Night event in Bologna and the nearby University Campus cities last 27 September 2019 welcomed more than 9000 visitors with a wide offer of activities, capturing different audiences. Many activities were designed for kids and youngsters, offering hands-on activities and entertaining science displays and games. The increase of the percentages of interviewees who attended ERN for the first time in 2019 was 64% in Bologna and 65% in the other cities. These numbers, which are the most significant indicators of the effectiveness of the event are even better than in 2018, where they were respectively 61% and 56%.

The success of SOCIETY is also testified by the increasing interest from Institutions/ Schools/ Companies, the so-called third-parties. The list of third-parties/collaborators has doubled from 43 in 2018 to a huge 86 in 2019. SOCIETY has become a reference to which an increasing number of public and private subjects ask for collaboration, in order to give greater value and larger diffusion to their own initiatives.

The rather ambitious program of about 30 pre-events distributed in the months preceding ERN has been widely successful, being attended by about 7300 people.

In Bologna most of the activities took place in via Zamboni, in the pedestrian academic district, right in the historical city centre. The inauguration ceremony took place in Palazzo Magnani, the stunning seat of Fondazione del Monte and Unicredit that host the important "Quadreria" and Carracci frescoes, where all the Night activities started at the presence of the most important personalities, students and researchers.

Almost 80 stands were distributed mostly in Scaravilli square and along the Zamboni street, where porticoes ensured protection in case of rain, that didn't come gifting the event of a warm late summer starry night. Researchers were available at their stand throughout the entire evening (18:00-24:00), in most cases performing their activity on demand. In other cases activities were organized in shifts and some required reservation. A limited number of activities were also held indoor, as the lucky one-to-one conversations ("A tu per tu"), held in Scuderie. Many high-school students were directly involved in outreach activities, introducing scientific contents to the general public, after long-term collaboration with scientists. Everybody could enjoy guided tours or simply wander through the many stands, exploring a broad spectrum of research disciplines, spanning from hard science to humanities, from outer space to human health, from applied technology to philosophy. For the entire duration of the Night, the stands were crowded with people making the 2019 edition a Night to remember.






Opening Ceremony of Bologna European Researchers' Night 2019

From the left:  
Francesco Ubertini, Rector of Bologna University,  
Marilena Pillati, Deputy Mayor of Bologna City Council,  
Sanzio Bassini, Director of HPC Cineca



HPC Stand for European Researchers' Night 2019 in via Zamboni, Bologna, from the left: Roberto Da Vià, Alessandro Marani and Fabio Pitari (Cineca)



Focus on experiments

# Events

## SC19 HPC is now: Denver, we're back!

Daniela Galetti  
Cineca

The SC conference is the most important event of the year for High Performance Computing, Networking, Storage and Analysis, where find and discuss the roadmaps of the manufacturers, the news of the vendors of HPC related hardware and software, the most used Operating systems, the most efficient Resource managers or the most used applications in the different fields.

The 2019 edition of Supercomputing took place in Denver. The town is a well known place for supercomputing community, in fact it hosted the conference 3 times in the last 6 years. Cineca participated to the SC conference running a booth to exploit its activities and applications in the HPC domain.

Like and more then last year, HPC infrastructures offered their power to machine learning algorithms, Artificial Intelligence solutions and big data analysis to improve the human environment, with work automation of uncomfortable job or weather and climate studies, reduction of pollution or discovering of new drugs and treatments.

A lot of these applications take advantage of accelerators, and especially the graphic process units, that are becoming more important than the CPUs inside the computing servers.

Big data instead focuses on more capable and flexible storage where to save and elaborate this big amount of information. Standard areas organized in huge filesystem will put aside object storage and the border between the different kind of memories is going to become thinner and thinner.

SC is also a very important occasion, beyond the technical program, to share experiences, solutions and good practices with colleagues and all people belonging to the High Performance Computing world during Panels, Workshops or simply visiting others booths or welcoming curious and interested persons at Cineca booth.

The Cineca staff participating in the SC conference has ensured regularly the presence at the kiosk desk, in order to provide information and clarifications to visitors and to discuss possible collaborations paving the way for new international projects and activities .

In 2020 the HPC circus is going to move on the right side of USA, in Atlanta, with the motto "more than HPC", so we expect that more new applications, algorithms and exploitation of HPC infrastructure will be shown for the science and human benefits.



**Sanzio Bassini, Carlo Cavazzoni and Daniela Galetti (Cineca) together with Jensen Huang (NVIDIA CEO)**



Above the render of the booth during design process, used for planning the space and evaluating the effectiveness of the graphic layout. Below the space during the exhibition with Mirko Cestari, Gabriella Scipione and Cinzia Zannoni (Cineca)





# SCIENTIFIC OVERVIEW

Cineca is involved in a wide scientific field thanks to EU projects and collaborations. From artificial intelligence to astrophysics, bioinformatics and digital humanities, HPC is used to support them all.

# AI agents that can see, speak and interact

Lorenzo Baraldi, Massimiliano Corsini and Rita Cucchiara  
Aimagelab, Università di Modena e Reggio Emilia, Italy

AImageLab is a research laboratory at the University of Modena and Reggio Emilia, focusing on AI, Computer Vision, Pattern Recognition and Machine Learning. At present, the staff comprises about 30 people, with 4 Full and Associate Professors, 3 Assistant Professors and 23 PhD students and Research Fellows (see [aimagelab.unimore.it](http://aimagelab.unimore.it)). Research at AImageLab focuses on the development of Deep Learning architectures and data-intensive Computer Vision approaches.

To satisfy the increasing computational needs, AImageLab has built a two-layer GPU-based computing environment which interacts with the facilities at Cineca. At the lowest level, researchers at AImageLab can build and train small-scale and medium-scale prototypes exploiting the on-site available GPU resources, while long-running experiments are transferred to Cineca. Together with CINI (Consorzio Interuniversitario Nazionale per l'Informatica), AImageLab is developing a partnership with NVIDIA for joint research activities on Deep Learning, Computer Vision and AI. This joint initiative aims to train students, nurture startups and spread adoption of the latest AI technology throughout Italy.

In the following, we provide a report of the most computationally intensive research activities carried out at AImageLab.


## Understanding videos and recognizing human actions

An important research line deals with the automatic comprehension of videos, with applications such as video understanding, classification and retrieval<sup>[1, 2]</sup>. In the past

years, AImageLab has developed innovative solutions for decomposing and summarizing the storytelling structure of videos, during a national PON-PNR research project “Città Educante” in which we cooperated with RAI on the analysis of broadcast videos, and then in collaboration with Facebook on investigating novel descriptors for video retrieval. More recently, AImageLab has focused on the design of novel architectures for spatio-temporal feature extraction and for the detection and classification of human actions in videos, in cooperation with Metaliquid<sup>[3]</sup>. Given the high dimensionality of video data, the role of high-performance computing is fundamental to investigate different solutions.

## Describing images in natural language

AImageLab interested in the development of image and video captioning algorithms which can automatically describe the visual content in natural language, thus connecting vision and language as our brain does. Most of the developed architectures rely on the combination of Convolutional and Recurrent Networks<sup>[4, 5]</sup>, although recently the group has investigated the use of self-attention and cross-attention architectures. Our M2 architecture<sup>[6]</sup> has recently surpassed the state-of-the-art image captioning algorithm on the most important benchmark, Microsoft COCO. For this line of work, high computational capabilities and high parallelism are required to investigate the possible design choices in reasonable time. For this reason, we have tested and exploited the D.A.V.I.D.E. platform, and we plan to continue on the upcoming Marconi 100 partition.



Example of human action detection across different frames of a video.





**Captions automatically generated from sample images, together with a visualization of the image regions on which the network focuses during the textual generation.**

Embodied AI: Agents that can see, perceive and interact

Recently, we have started a promising research direction which focuses on the development of mobile agents that can see, perceive and interact with the real world. On this line, we are developing robotic agents capable of navigating unseen environments exploiting natural language instructions<sup>[7, 8]</sup>, using both recurrent

and attention-based architectures. As this line of research needs simulated environments for faster training, we have also captured a virtual navigation environment in the Estense Gallery in Modena, as a new test bed for navigation agents. An expectable outcome of this research line is also the deployment of the algorithms on real robots that can navigate in the real world.



**Visualization of the virtual environment for navigation captured in the Estense Gallery in Modena.**

In the ambit of the Advanced School on Computer Graphics and Cultural Heritage organized by the Cineca (5-9th of October, 2019) we did a lesson entitled “Artificial Intelligent and Digital Humanities”. In this lesson we give some concepts and ideas behind Deep Learning and we presented something of the research previously described. The feedback received by the students was very good, even if they had an interdisciplinary background.

**References**

- [1] Baraldi, Lorenzo; Douze, Matthijs; Cucchiara, Rita; Jégou, Hervé “LAMV: Learning to align and match videos with kernelized temporal layers”, CVPR 2018
- [2] Baraldi, Lorenzo; Grana, Costantino; Cucchiara, Rita “Recognizing and Presenting the Storytelling Video Structure with Deep Multimodal Networks” IEEE Transactions on Multimedia, 2016
- [3] Tomei, Matto, Baraldi, Lorenzo, Calderara, Simone, Bronzin, Simone, Cucchiara, Rita “STAGE: Spatio-Temporal Attention on Graph Entities for Video Action Detection”. arXiv preprint, 2019
- [4] Baraldi, Lorenzo; Grana, Costantino; Cucchiara, Rita “Hierarchical Boundary-Aware Neural Encoder for Video Captioning”, CVPR 2017
- [5] Cornia, Marcella, Lorenzo Baraldi, and Rita Cucchiara. “Show, Control and Tell: A Framework for Generating Controllable and Grounded Captions.” CVPR 2019
- [6] Cornia, Marcella, Stefanini, Matto, Baraldi, Lorenzo, Cucchiara, Rita. “M2: Meshed-Memory Transformer for Image Captioning”. arXiv preprint, 2019.
- [7] Landi, Federico, Baraldi, Lorenzo, Corsini, Massimiliano, Cucchiara, Rita. “Embodied Vision-and-Language Navigation with Dynamic Convolutional Filters”. BMVC 2019
- [8] Landi, Federico, Baraldi, Lorenzo, Cornia, Marcella, Corsini, Massimiliano, Cucchiara, Rita. “Perceive, Transform, and Act: Multi-Modal Attention Networks for Vision-and-Language Navigation”. arXiv preprint, 2019

# Searching for water on Mars

Roberto Orosei

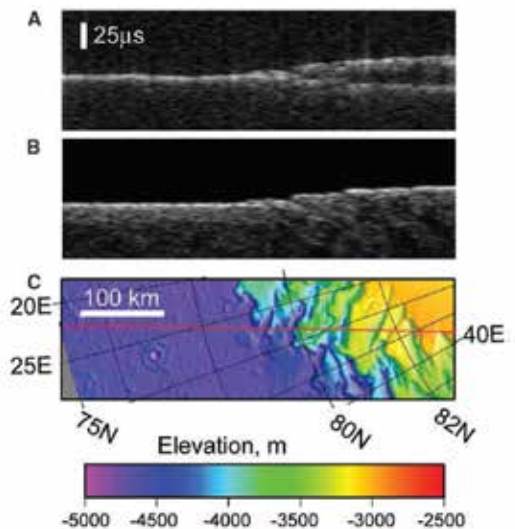
Istituto Nazionale di Astrofisica, Istituto di Radioastronomia, Bologna

The presence of water is considered a necessary condition for the existence of life as we know it. Even if Mars is today a cold and arid place, there is ample evidence that things were different in the past, and that rivers, lakes and perhaps seas existed for the first several hundred million years of its history. Such time span was sufficient for the emergence of life on Earth, and thus many have wondered if the same could have happened on our neighbor planet. The next question is then if life could have survived the dramatic changes in Martian climate and still be existing today, and where it could be found. The answer is summarized by the name NASA gave to its strategy for the exploration of the Red Planet: “follow the water”.

A fraction of Mars’ water was probably lost in space together with most of the planet’s atmosphere, which the weak gravity could not retain over the ages. More water is present as ice in the polar caps and in the ground. At some depth below the surface, ground ice would melt as temperature increases due to the heat coming from the molten interior of the planet, similarly to what happens on Earth. Currently, the only viable way to search for water over an entire planet is from orbit, and the only way to search for it at depth is by using ground penetrating radars (GPR’s), transmitting radio pulses at low frequency that penetrate into the ground and are reflected by changes in its structure and composition. Mapping such reflections over an area provides the local subsurface structure, while analyzing their characteristics can give information on its nature and composition.


On Earth, GPR has been used to probe polar ice sheets since the fifties of last century. Liquid water is the most reflecting natural material for radio waves, and because of this subglacial lakes have been identified in Antarctica, Greenland and, more recently, in the Canadian Arctic. Similar to terrestrial GPR’s, MARSIS was put on board the European Space Agency Mars Express spacecraft to look for water on Mars.

Radar pulses transmitted by MARSIS illuminate the entire surface of the planet, not just the portion directly beneath the spacecraft. Therefore, reflections coming from distant craters, mountains and similar features will be received after the echo coming from the surface directly below the spacecraft, and appear to be originating in the subsurface. The only way to identify true subsurface echoes is through numerical simulations of how the Martian topography reflects electromagnetic waves at the frequencies of MARSIS: any echo that appears both in real data and in simulations must be produced by surface topography, while one that is not present in simulated data has to be coming from the subsurface.



**Comparison between a real radar section of the Martian North polar cap (upper panel) and simulated data (middle panel): a secondary echo below the surface echo is visible only in real data, and originates from the bottom of the polar cap at a depth of about 1.8 km. In the bottom panel, topographic map along the ground track of the spacecraft (red line).**

**Picardi, G., and 33 colleagues 2005. Radar Soundings of the Subsurface of Mars. Science 310, 1925.**



**Thanks to Davide Coero Borga  
(INAF)**

A numerical code producing synthetic surface echoes was developed with the help of Cineca and run on its computers to simulate thousands of observations, over the course of a few years. The results were key in identifying subsurface structures, and led to the discovery of a  $20 \times 20 \text{ km}^2$  area of strong radar reflections coming from a depth of 1.5 km beneath the south polar cap, which could be explained only by the presence of liquid water. Because of the

very low temperatures at the Martian poles, it is believed that such water remains liquid because it contains salts that strongly depress its freezing point. Very cold and salty water isn't the ideal place for life, but there are microorganisms capable of thriving in similar environments on Earth. Thus, this discovery hints at the possibility that life could exist on Mars today, but to be certain it will be necessary to retrieve and analyze samples of Martian water.

# Binary Neutron Star Mergers

## Fully General Relativistic Simulations

**Riccardo Ciolfi** INAF - Padova Observatory, Italy  
**Wolfgang Kastaun** Albert Einstein Institute, Germany  
**Jay Vijay Kalinani** University of Padova, Italy  
**Bruno Giacomazzo** University of Milano Bicocca, Italy

In the last few years, observatories such as Virgo in Italy and LIGO in the USA were able to detect tens of gravitational waves emitted by binaries of compact objects, such as neutron stars and black holes. Neutron stars are particularly interesting objects since they are stars with masses slightly larger than the mass of the Sun, but with a radius of approximately 10 km (the radius of the Sun is around 700000 km). Matter contained in a neutron star reaches densities much higher than any density we can obtain in our laboratories on Earth. How matter behaves in such extreme conditions is still quite a mystery and therefore the study of neutron stars can allow us to understand those unknown regimes. Neutron stars are often found in binaries and some of these, after millions or billions of years of orbiting one around the other, can eventually merge into a single object. This process is accompanied by a strong emission of gravitational waves, as confirmed by detections made with our instruments (two confirmed cases up to now). Binary neutron star mergers may also produce bright electromagnetic signals and power highly energetic astrophysical phenomena such as gamma-ray bursts.

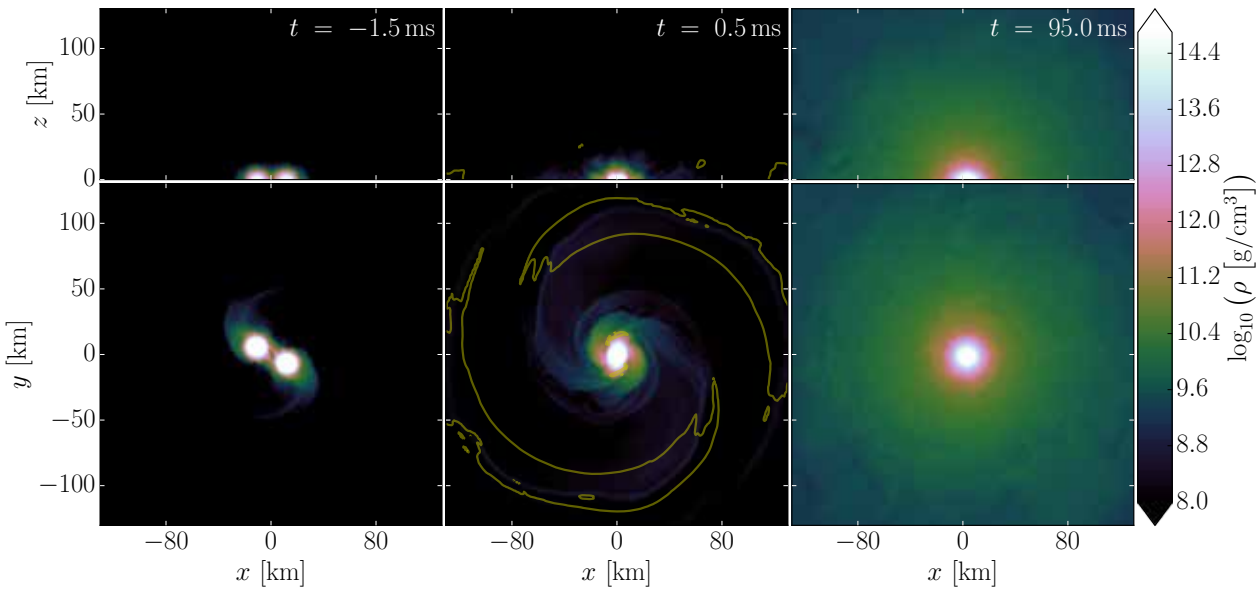
In order to extract significant information from observations of these objects via gravitational waves and electromagnetic signals it is necessary to model the dynamics of these systems solving the equations governing the movement of magnetized dense matter within the theory of general relativity via the use of parallel codes and HPC clusters. Our group studies what happens during a binary neutron star merger and the associated gravitational wave and electromagnetic emission to be compared with observations.

In 2019 we performed one of the longest simulations of the merger of two neutron stars with intense magnetic fields [1]. The simulations

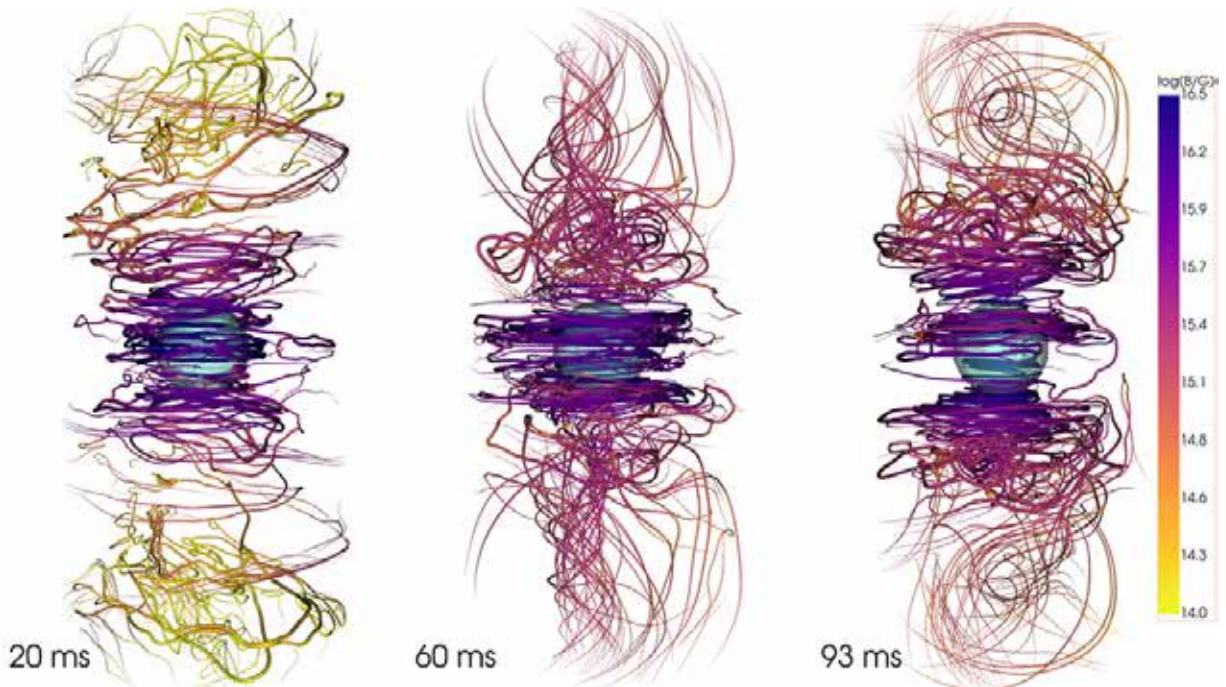
were performed using our WhiskyMHD code coupled with the publicly available Einstein Toolkit ([www.einsteintoolkit.org](http://www.einsteintoolkit.org)) and were run on the Cineca cluster Marconi A3 (Intel Skylake processors). Our simulations studied the possible formation of relativistic jets and short gamma ray bursts in case a long-lived neutron star or a black hole are the result of the merger of the two neutron stars. Gravitational wave detectors are indeed very sensitive to the last phases of inspiral, but they are not yet able to detect the signal emitted after the merger and therefore we do not know if a black hole or a massive neutron star is the result of such mergers. Simulations like ours point in favor of the formation of a black hole, since otherwise the system seems to be unable to produce the bright gamma ray emission that, for example, was observed together with the gravitational waves detected in 2017 (known as GW170817, the first neutron star merger ever detected). Our simulations also provide the full gravitational wave signals, including inspiral, merger, and post-merger phases. Such signals can be used to train data analysis algorithms in order to increase the chances of being able to detect also the post-merger emission in the future and hence being able to know which kind of objects these mergers form. Moreover, a comparison between data collected by Virgo and LIGO and our simulations may help shed light on the internal composition of neutron stars and therefore how matter behaves in such extreme regimes. This would allow for a significant advance in the fields of nuclear physics and high-energy astrophysics.

### References

[1] Ciolfi R., Kastaun W., Kalinani J. V., Giacomazzo B., 2019, *Physical Review D*, 100, 023005



**Evolution of the matter density during the last phases of inspiral and merger for a binary neutron star system. The top panels show the meridional plane while the bottom ones the equatorial plane. The simulation covered up to 100 ms after merger. Figure published in [1].**



**Magnetic field structure 20, 60, and 93 ms after merger. To give a scale reference, we added a (cyan) sphere of 10 km radius placed at the remnant centre of mass. Field line colors indicate the magnetic field strength. Figure published in [1].**

# HPC powers large scale de novo drug design

Federico Ficarelli and Carlo Cavazzoni

Cineca

Drug development is a long process: from the time a health threat is discovered, it takes years until a complete product is available on the market.

The process of looking for an effective treatment for a specific disease involves three subsequent phases, each one closer to the affected subject: in silico design, where a computer simulation pipeline is leveraged to refine experts' evaluation of potential candidate molecules; in vitro testing, where proper testing is performed on real biological samples and, lastly, in vivo testing, where actual live test subjects are involved.

Computer-aided drug design (CADD) techniques, the foundations of in silico design, are used to speed-up the early-stage drug development allowing the selection of new active compounds. Through virtual screening, large virtual chemical libraries can be screened to find active molecules for query targets.

The recent surge in open availability of protein structures, resolved by structural biologists, progressively raised the possibility to deploy structure-based drug design. Key for the success of the structure-based drug design is the evaluation of a chemical space big enough to allow the identification of chemical structures having the best complementary pattern of interactions with the biological target under investigation: in other words, the largest the search space that can be tackled by the computing platform considered, the higher the chances of success.

Dompé SpA, through its dedicated Drug Discovery Platform, invested since 2010 in a proprietary software for CADD. The most relevant tool is the de novo structure based virtual screening software LiGen™ (Ligand Generator), co-designed and developed in collaboration with Cineca. LiGen™, as a distinguish features, has been specifically designed from the ground up to run and scale

out efficiently on HPC platforms<sup>[1]</sup>.

The evolution of the LiGen™ platform has been selected as one of the industrial use cases of the ANTAREX project<sup>[2]</sup> aimed at the proposal of a holistic approach capable of controlling all the decision layers in order to implement self-adaptive applications optimized for energy efficiency.

To demonstrate the benefits of code optimization and scalability, Dompé selected the Zika pandemic crisis to support and promote the identification of novel drugs able to address the unmet medical need in terms of effective therapies. Dompé, in collaboration with Università degli Studi di Milano, selected 26 binding sites identified from the already resolved crystal structures of 5 Zika proteins and has created a virtual chemical space of 1.2 billion small molecular-weight molecules, the actual search space to be explored.

In order to be able to tackle the huge molecule set, teamwork was key: Dompé started by analysing the most computationally intensive kernels of LiGen™, the research team at Politecnico di Milano developed a runtime tuneable version of the molecular docking module while Cineca tackled the massive storage needs of the simulation by re-architecting the parallel I/O approach of the application.

The resulting tool was finally deployed and scaled out to the full size of the Marconi supercomputer at Cineca, a 10 PetaFlops system based on Intel Xeon Phi manycore processors, to screen the 1.2 billion ligand database (including all investigational and marketed drugs) targeting the Zika virus.

This represents the largest virtual screening experiment ever launched in terms of computational units (up to one million hardware threads) and size of the compound database (more than one billion molecules).


However, the collaboration between Dompé, Cineca and former ANTAREX partners continues in the ongoing development of the Exscalate platform<sup>[3]</sup>, a CADD integrated solution focused on future exascale architectures: soon, Exscalate will be able to count on the new pre-exascale system Leonardo hosted at Cineca. This platform, developed as a testbed for future exascale-ready applications in the frame of the EuroHPC project<sup>[4]</sup>, will be able to reach a power peak of 270 PetaFlops, which will make the screening and identification of new drugs even faster.

Furthermore, the Exscalate platform is going to be the key element on top of which the Exscalate4CoV consortium<sup>[5]</sup>, constituted by

Cineca, Politecnico di Milano, several European subjects and led by Dompé, is going to base its efforts toward providing an answer to the COVID-19 outbreak.

#### References

- [1] Beato, Claudia and Beccari, Andrea R. and Cavazzoni, Carlo and Lorenzi, Simone and Costantino, Gabriele, Use of Experimental Design To Optimize Docking Performance: The Case of LiGenDock, the Docking Module of Ligen, a New De Novo Design Program, *Journal of Chemical Information and Modeling*, 53, 6, 2013.
- [2] <http://www.antarex-project.eu>, retrieved 2020/02/27
- [3] <https://www.exscalate.eu>, retrieved 2020/02/27
- [4] <https://eurohpc-ju.europa.eu/>, retrieved 2020/02/27
- [5] <https://www.dompe.com/en/media/press-releases/exscalate4cov-supercomputing-platform-wins-3-million-euro-eu-funding-to-hunt-for-new-drugs-against-coronavirus-pandemic>, retrieved 2020/03/09



**Structure of a Zika Virus (ZIKV) polyprotein**

# 3D Saint Peter's Cathedral in Bologna Art, History and Digital Humanities

Fabio Massaccesi

Università degli Studi di Bologna - Dipartimento delle Arti (DAR)

The recent exhibition *Imago Splendida, capolavori della scultura Bolognese dal Romanico al Duecento* (curated by Massimo Medica and Luca Mor) made it possible to present some personal research which began in 2018, on the Cathedral of Saint Peter in Bologna.

The building has a foundation dating back to the 11th century and further renovations were undertaken over the following centuries, particularly between the 12th and 15th century. Its ancient medieval basilica, with a crypt and raised presbytery, would have been preserved until 1570, when it was completely rebuilt at the wish of the archbishop of Bologna, Cardinal Gabriele Paleotti, in its present form.

Those who enter the building will be able to admire an imposing Baroque architecture, symbol of the new Church that the Council of Trent, thanks to Paleotti himself, had imposed, but what we see today is totally different from that of the primitive episcopal seat. This is the great challenge of research: trying to reconstruct the ancient building through the sifting of archival documents.

A tortuous and unknown research path that focused on three key points.

The first centred on the attempt to find useful material (sources, archaeological data, drawings, archival documents ...), of which there was no certainty of existence, in order to reconstruct the building at the time of its demolition (1570), the second focused on the consequent triumphal furniture of the presbytery in which a central role was to be played by the big Cross (XII century) which still stands above the main altar and enriched by the presence of the two coeval Dolenti (the Virgin and John the Evangelist).


A type of historical-artistic research entirely

based on the restitution of no longer extant architectural structures that inevitably must overcome the even more difficult challenge of how to present it to an audience that is not exclusively expert (the scientific community and the general public).

As for the first two points, the research I conducted brought to light an enormous amount of data, sometimes apparently contradictory, which allowed us to reconstruct the structure of the ancient presbytery on which was imposed the "Corridore" (Rood Screen), with a good degree of accuracy. Records of its existence have been handed down to us in the Bolognese Chronicle (*Cronaca* 1563-1573) written by Giovan Battista Marescalchi in the 16th century. In Bologna, this denomination indicated a wall that ran along the entire width of the building and separated the part (the nave) in which the faithful had to be (two thirds of the length of the building), from the part designated to the officiants (canons and bishop), or the choir and the high altar (one third of the church).

This was a necessary architectural structure that responded to medieval liturgical practices and was completely different from the modern ones initiated by the Council of Trent, which systematically ordered their dismantling throughout Italy to make the liturgy of the solemn mass visible. On the Bolognese "Corridore" (Rood Screen) of Saint Peter's, which sources also attest with the word "Pontile" (from bridge), to demonstrate that it was walkable, there were altars (documented) and it is from the top of this structure that the priest gave the blessing to the faithful.

**"Corridore" (Rood Screen) between  
nave and chancel  
(Same Architecture)**





The geometric centre of medieval churches was the large “Crux de medio ecclesiae” (Cross in the middle of the church) which stood out as a visual and at the same time liturgical focus and was placed right on the Rood Screen (Pontile/Corridore).

In Saint Peter’s this important role was taken by the wooden Crucifix (today on the high altar) which, as research has clarified, must have been the object of a redevelopment by the Bolognese painter Jacopo di Paolo who replaced the cross in support of the wooden body of Christ somewhere between 1415 and 1417.

Research has thus clarified how until 1570 Saint Peter’s had an imposing, 5-meter high, wall septum which enclosed the church chancel and which served as a dividing point between the laity and the clergy.

Today these types of historical research can be truly effective thanks to the achievements of digital humanities.

On this occasion, thanks to the collaboration with Cineca and Same Architecture, the documentary research took the 3D reconstructions and made a video for the exhibition capable of telling



**Structure of the “Corridore” with the altars above it (Cineca)**

and testing the hypotheses that emerged on a documentary basis.



An interdisciplinary approach that opens up fascinating, new methodological challenges in which an ever greater role will be played in the restitution of the past by new digital tools in an association between the humanities and technological sciences which is increasingly indissoluble and necessary.

A meeting between disciplines that will give greater results in the difficult task of building our awareness of the past as a necessary basis for the enhancement of our historical-artistic heritage.

**References**

Fabio Massaccesi, I contesti architettonici delle croci trionfali bolognesi tra spazio e liturgia, in: *Imago Splendida*, capolavori della scultura lignea bolognese dal Romanico al Gotico, edited by Massimo Medica, Luca Mor, Milan 2019, pp. 53-73.  
 Idem, Il “Corridore” della chiesa agostiniana di San Giacomo Maggiore a Bologna: prime ipotesi ricostruttive, in: *Zeitschrift für Kunstgeschichte*, 2014, pp. 1-26.  
 Monika Schmelzer, *Der mittelalterliche Lehnstuhl in deutschsprachigen Raum: Typologie und Funktion*, Petersberg 2004.

**Crux de medio ecclesiae (Same Architecture)**



**3d reconstruction of the Christ and the Cross placed above the “Corridore” (Cineca)**

# Last Supper Interactive

Franz Fischnaller

Accademia Albertina di Belle Arti di Torino

With Last Supper Interactive, Cineca returns to collaborate with Franz Fischnaller after the experience of the Sarcophagus of the Spouses (<https://www.glietruschielaldila.it/il-sarcofago-degli-sposi>).

The project allows to experience the Last Supper by Leonardo Da Vinci from different points of view, even within the artwork itself. It was previewed at the Ars Electronica Centre in the Deep Space 8K during Ars Electronica 2019.

During the Researchers' Night held in Bologna in September 2019 (<http://nottedeiricercatori-society.eu/>) Franz Fischnaller presented a video in high resolution at the event dedicated to the future Leonardo supercomputer. Last Supper Interactive is a non-profit educational multimedia work on the experimental digital heritage derived from interdisciplinary research and entitled: "Technology and science in the arts, cultural heritage and digital humanities"

which proposes a method of creating cross disciplinary knowledge combining different forms of knowledge and shared practices.

The Last Supper Interactive Project (LSI) is a high resolution immersive VR storytelling that enables a 360 degree virtual access to Leonardo's Last Supper, drawing the visitors "between, inside and behind" the painting, and transporting visitors into the architectural environment where Leonardo created the artwork. Conceived for different types of viewing format/experiences, LSI provides a Stereoscopic 8K 3D immersive VR interactive storytelling application, a 3D digital Stereoscopic 8K animation, and a 4K UHD3D-VR 360-degree video. LSI is designed to run on different platforms and display systems, among which: Immersive virtual reality rooms, large scale multi viewer immersive spaces, immersive visualization environments and VR Head-Mounted Displays.



Different internal views of Immersive point cloud visualization in 8K of the Basilica Santa Maria Delle Grazie. LSI Project.© F.Fischnaller. Laser scanning: Politecnico di Milano.


Rendering: CINECA

For the LSI project the Visit lab team of Cineca has designed a new pipeline able to optimize the rendering of 8K stereoscopic 3D animation sequences together with the integration of special effects required by the artistic direction and the processing of massive point cloud data in order to create a flexible Blender open workflow according to the methodology followed in Cineca since the year 2011. The pipeline allows to easily reconfigure the workflow according to the requirements of the different output video content to be produced (8K stereo, 4K , VR scenes, etc.). To render a single scene of the LSI project using

a standard personal computer it would have taken 45 days. On the Galileo supercomputer instead it takes about ~ 10-12 hours (using only a small part of the computer: about 20 nodes). It was therefore possible to render all the video in 3 days.

**Credits:**

Video animation of “The Last Supper Interactive” Project (LSI project) - Rendering by CINECA  
 Author: Franz Fischnaller  
 Project Partners:  
 Cineca - Cineca Inter University Consortium (Italy)  
 Haltadefnizione (Italy)  
 Université De Rennes 1, Inria, Irisa, Insa (France)  
 Polytechnic University of Milan, Department of Mechanics (Italy)



**Tracking the Linear Perspective: experiencing a number of alternative virtualized views and interactive navigation allowing visitors to experience the art work from multiple viewpoints. LSI Project © F. Fischnaller**



**8192 \* 4320 pixels 8K (resolution of each image)**



**19060 8K frames have been rendered in order to produce also a stereo visualization**



**115 MB per frame for a total of 2192 GB (2 TeraByte)**



**95300 minutes 55600 core-h (computation time)**



**To visualize the interior of Santa Maria delle Grazie and the inner garden, the process manages a point cloud made up of 95 million vertices. The entire production pipeline is open**



# EUROPEAN PROJECTS AND COLLABORATIONS

Cineca is committed in many different EU-funded projects, both on digital infrastructures and on vertical thematic initiatives. In this context, we are working together with the most important HPC computing centres in Europe to help our scientific communities to leverage the most advanced available technologies.

# PRACE 6IP is looking forward to the exascale

Fabio Affinito and Lara Querciagrossa  
Cineca

In April 2019 the Sixth implementation project of PRACE has been triggered during a meeting in Bratislava. This new project, that will conclude at the end of 2021, will see the birth of the new three pre-exascale systems funded by the EuroHPC Joint Undertaking. For the first time, exascale supercomputers are at reach! In order to be ready for such a powerful hardware systems, also the software in the hands of the users needs to be ready. This is a very important challenge and the EU is putting a lot of effort in this direction, for example sustaining the development of new Centres of Excellence that aim to help the users' communities to face the exascale challenge. Also PRACE is playing a role in this match and, for the first time, in the PRACE 6IP project, a new Work Package has been created with the aim to design new "Forward Looking software solutions" to be deployed in the next generations of supercomputers, exploiting at the best their power. This Work Package, led by the Swiss National Supercomputing Centre (CSCS) and Cineca, consists of eight development projects, with 14 partners and a total number of more than 800 person-months. The projects in this Work Package are very ambitious and aim to address high-impact topics in the field of computational sciences, ranging from plasma physics to molecular modeling and to numerical libraries. All these projects have been selected with a competition whose outcome was decided by a peer-review process superseded by the PRACE Scientific Steering Committee. Cineca, in addition to the co-lead of this Work Package, is taking part to one of these project, "Performance portable linear algebra", together with Juelich Supercomputing Centre, CSCS and University of Ljubljana. In many scientific applications, for example in material sciences or computational chemistry, a large time of the simulation is spent in the resolution of complex linear algebra problems, such as the solution of linear systems of equations, the diagonalization of matrices, etc. The de-facto standard for the solution of these problems in distributed systems is the ScaLAPACK library. ScaLAPACK was developed in 1995: at that time most of the supercomputers were equipped with single-

core cpus. Now, after 25 years, supercomputers evolved and became completely different: each node has typically more than ten cores (for example, Marconi is equipped with Intel KNL nodes, each of them with 48 cores) and, likely, with one or more GPUs. The parallelism approach of the ScaLAPACK library, hence, does not perform well on the current architectures and the corresponding algorithms have not evolved to keep up with developments in hardware or modern solving strategies. Also, very well consolidate programming models for hybrid systems, like MPI+OpenMP, don't produce satisfactory results when projected to large scale machines as we expect to be the next pre-exascale architectures. The developers of ScaLAPACK have started, in order to keep the pace of the hardware evolution, a project to implement a new distributed linear algebra library, based on the concept of OpenMP tasks, called SLATE. The project to which Cineca is contributing is committed to collaborate with the development of the SLATE library, developing dense matrix eigensolvers.

Moreover, we are also contributing to implement a version of SLATE based on HPX. HPX is a library that permits to implement, in a modern C++ style, asynchronous tasks. The library produced in this project will be integrated in scientific applications, such as those that they can exploit the performance of modern architecture supercomputers. The first targets of the integration will be scientific applications such as Quantum ESPRESSO, SIRIUS and CP2K, which are also among the most used software in the Cineca users' community.



# PRACE SHAPE: projects for the SME

Eric Pascolo  
Cineca

SHAPE, the SME (Small Medium Enterprises) HPC Adoption Programme in Europe is a pan-European, PRACE-based programme supporting HPC adoption by SMEs. The Program aims to raise awareness and equip European SMEs with the necessary expertise to take advantage of the innovation possibilities opened up by HPC, thus increasing their competitiveness. SMEs can receive machine time on a PRACE system, and, most important, effort from a PRACE expert that will work alongside the SME in evaluating and/or developing the HPC-based solution. In return the SME provides in-kind effort, publicity of their participation in SHAPE, and a public White Paper on the work achieved in the project at its conclusion.

In 2019 the Axyon Fintech company worked with Cineca on SHAPE Project to “reshape” its AI workload management platform to use large scalable HPC systems.

The use of HPC in finance has become very important because of the advent of Big Data. The increasing use of AI algorithms and the consequently fast-growing need for computational power require the use of large scalable HPC systems able to handle and explore a wide array of Machine Learning models (e.g. deep neural networks) and to remain up-to-date with state-of-the-art methods.

Deep Neural Network classifies human tasks and forecasts mental workload parameters, based on eye-tracking and electroencephalogram data, combined with video analysis during driving test.

Over the years Cineca has hosted and supported 9 SHAPE projects and we are receiving more and more interest from SMEs to use and bring applications on HPC systems, both to optimize execution times and to use massively parallel systems to implement new business models.



Research To Business 2019: Eric Pascolo (Cineca) and Riccardo Folloni (Axyon)



# HPC for EUROfusion

Francesco Iannone

ENEA

Elda Rossi

Cineca

EUROfusion is the European Consortium for Development of Fusion Energy, born in 2014 in order to cement collaboration among fusion research bodies in Europe and Switzerland. EUROfusion supports and funds research activities on behalf of European Commission's Euratom programme.

The collaboration between Cineca and EUROfusion started in 2016, when a three years agreement was signed, jointly with ENEA, for hosting a Petaflop computer dedicated to large scale numerical simulations in plasma physics and materials science relevant in nuclear fusion reactors. The theory and the numerical modelling effort is a key for the experimental physics and design of the first nuclear fusion plant DEMO before 2050.

A new agreement was signed at the end of 2018 and, starting from January 2019, a dedicated partition of Marconi with a computing peak power of 8 PFlops was in operations. In addition to this, also two accelerated partitions are available to EUROfusion: 1 PFlops of Intel Xeon PHI Knights Landing nodes and 1 PFlops of GPU based nodes. Moreover, an interactive system for developers (Gateway) based on the same technology of the large conventional system, is used by the community.

The agreement will be operational for 5 years until 2023 and it is being deployed in two phases in order to take into account the latest new technologies of processors to replace the accelerated partition.

From the point of view of the quality of service,

availability of the computing systems, quality of system management, help desk, user support, high level support, Cineca offers a service to the EUROfusion users community with SLA which entails penalties in case agreed KPIs are not be matched. The Marconi operations are monitored on monthly base by a committee of computing scientists. It analyzes the status of the systems, as well as evaluates the issues occurred and propose the operational parameters to ensure the best exploitation of the HPC services. Next to the Operation committee a pool of experts composed of EUROfusion members and Cineca support team, handle the status of tickets submitted by EUROfusion users in order to fix them.

The nuclear fusion projects running on Marconi are selected by a dedicated committee on the base of their potential contribution to the activities of the EUROfusion work program and mainly concern:

- the plasma core physics: transport, turbulence and magnetic reconnection domains;
- the plasma edge: Scrape-Off-Layer and divertor physics;
- MHD stability and plasma disruption controls;
- Radiofrequency Heating: plasma-wave interaction in different electron frequency domains.



# SUPER

Debora Testi  
Cineca

SUPER (Supercomputing Unified Platform - Emilia-Romagna) is a regional project (POR-FESR-2014-2020) funded with the aim to extend the regional computing and storage infrastructure to support both scientific and industrial use cases.

The project is based on few main concepts:

- the geographical proximity of the major HPC and HTC centres in Italy, Cineca and INFN, respectively;
- the federation of computing resources which allows the data replication in an easy way, increasing resilience, and the availability and accessibility of data;
- security, flexibility, extensibility and scalability as key factors for the use cases evolving over time.

The first project objective is the update of the hardware platform with the acquisition of specific resources (both for computing and storage) to support to data intensive and data analytics applications. Then the partners will implement a set of federated basic and advanced services to easily and transparently access the resources of the different centres in the project. On top of these, use cases (UC) in different scientific domains (like genomics, regenerative medicine, and biobanks, advanced materials and innovative production systems) will be enabled to guide the implementation and serve as test beds also for other domains. While the UCs rely on different processes and data types, their needs for computation and data storage are similar and work will be carried out to define the minimum set of requirements and the resources needed from the infrastructure centres. In summary, the regenerative medicine UC will make available a platform that will use big data for patients' stratification, to improve diagnosis and consequently therapy; the genomics UC has high needs in computation and storage resources for the processing of the data: the SUPER infrastructure will make possible to process 200 genomes per week; the biobanks represent a strategic infrastructure

allowing to maintain tissue samples over time to further run analyses, like the omics ones, which are one of the innovative elements of health 4.0; the advance material domain is characterised by the increasing demand for HPC/HTC which makes possible the predictive simulations of complex systems and processes and the screening of materials with certain properties; the industry 4.0 activity is aimed to demonstrate how to tackle the challenges of different use cases from regional industries (i.e. e-maintenance for industrial machines, laser-based transport systems for material stocking, predictive maintenance of energy production from sustainable sources).

The SUPER federated integration of the infrastructure will allow:

- to tackle the big scientific challenges by integrating excellence HPC and HTC centres and services;
- to support future regional, national, EU and international projects;
- to move the frontier of research in fields of high socio-economic impact (like climate change, environment protection, artificial intelligence, and cyber security) besides those already explicitly supported by the project.

The consortium is composed by 12 partners and includes the major players in the region in above mentioned domains: CINECA, INFN, CMCC, ENEA, CNR, INAF, INGV, Università di Bologna, Università di Modena e Reggio Emilia, Istituti Ortopedici Rizzoli, Università di Ferrara e Università di Parma. The project will also see the involvement of industries from the region which will support the project with their use cases.

The project started in December 2019 and will last for 18 months.



# AIDA

**Giovanni Lapenta, Jorge Amaya** KU Leuven  
**Francesco Califano** University of Pisa  
**Francesca Delli Ponti** Cineca

Space is the final frontier and we are living a pioneer age of exploding interest for space exploration, from national and private enterprises. The project AIDA uses some of the most advanced data analysis techniques to study the data available from space missions and from space simulations. AIDA is based on three major trends.

First, there has been a tremendous accumulation of data from space studies. With the passage of time, technological advances allow to build spacecraft that carry faster and more powerful computers and instruments. These are capable of transmitting higher quality data at an accelerated rate.

Similarly, the computer simulations of space science are becoming ever more complex and their results are more and more detailed and rich, also leading to a rapidly increasing mass of data available.

AIDA is developing tools to make this data more readily accessible to the public, creating common standards for storing data when such standards do not yet exist while using community standards whenever possible. This advancement will be achieved by collecting existing tools and creating new ones into AIDApY, the main outcome of the AIDA project: a software tool to gather and analyse space-relevant data from observations and simulations.

The second trend is the vast amount of space data accumulated by past missions and simulation studies and by the new data generated every day by new and ongoing missions and by new numerical simulations. As a consequence, the amount of data to be scrutinised is becoming so large that a single human brain cannot grasp its entire complexity. Luckily, a trend is sweeping society to address this very same problem arising in many aspects of human activities: machine learning and artificial intelligence (AI).

AIDA is bringing the latest developments in AI to the space arena (AI in AIDA is indeed Artificial Intelligence for Data Analysis, AIDA). The competences and passions needed to be a space scientist are quite different from those of an expert in artificial intelligence. AIDA brings together expertise in computer science, AI and space science to create a new tool, namely AIDApY that will make the use of AI simpler for the space scientist and even the general public. Third, there is an ongoing trend in all big data analysis and in AI in particular to move towards a free open source language for computers called python. To truly democratize science, including the so-called trend in citizen science, to promote its spread in developing countries and to simplify access to the data and its analysis to all, python provides a free alternative. Python potentially has even more features than commercial alternatives and many science communities are currently making the transition to it. AIDA aims at fostering this same transition to python also for the space community, collecting in AIDApY (where py stands for python) tools previously based on non-free languages and developing new tools based on the most advanced tools in the state of the art.

The aim of AIDA is as much scientific as it is directed at transforming the way society can access the results of space missions and space simulations, making the data accessible to all and giving to all, including non expert citizens, access to the most modern tools to analyse the data. Just as amateur astronomers have in the past discovered moons, comets and astronomical processes, AIDApY will empower citizens to use the most advanced tools in data analysis and AI to explore space science.

<http://aida-space.eu/>




This project has received funding from the European Union's Horizon 2020 Research and Innovation programme under grant agreement No 776262.



AIDA partners during the second annual meeting.

## 1st Artificial Intelligence Data Analysis School for Heliophysicists

20 - 22 January 2020  
CINECA - Casalecchio di Reno, Bologna



# Highlander

Balasubramanian Chandramouli, Mara Chiara Liguori and Gabriella Scipione  
Cineca

Nowadays, still too often, the preservation of environmental resources, as well as the management of natural to semi-natural ecosystems and human sectors, are treated in a scattered way, with disciplines and decision makers proceeding separately and considering neither the chance of benefits nor the risk of damages potentially intergenerated.

Rather, an integrated and multi-actor approach at the landscape level, i.e. considering the interacting systems, sectors and resources in a mosaic of land uses, could optimize the efforts in addressing the challenges and seizing the opportunities for an overall smarter land management. This applies in particular to the Italian territory, that offers from north to south a huge variety of fragmented land covers, geomorphological settings, resources availability and human assets in abrupt and rapid transformation.

The Highlander project aims at improving existing data, indicators and tools, as well as at providing new ones, to support more efficient management of heterogeneous lands that suffer from pressures due to both natural and human factors.

The three main interrelated goals are: design and implement a framework of multi-thematic, last generation, harmonized data, indicators and tools, based on various complexity approaches, from remote and in-situ monitoring, analytical tools, numerical models up to machine learning algorithms, to investigate metrics originating from multiple domains (eg., microclimatic and genomic data) to establish inter-relations; fully exploit the Cineca HPC capabilities to generate, post-process, host, distribute and make accessible to and exploitable by multiple users both existing and newly generated data, indicators and tools into HPC-based services; ensure the long-term functionality of the created services thanks to the involvement of real users, so to preliminarily test and foster services' repeatability, transferability and scalability, up

to stimulating the design and development of a new generation of HPC.

The foreseen services are:

- a National Open Data Portal on climate and climate-related hazards;
- a set of illustrative “downstream applications”, aimed mainly at increasing the users’ awareness about the great potential of the information provided as well as their capacity in further elaborating and using data;
- an implementation of a chain of operations that directly generate the user-required information through a user-customized package.

The consortium of partners will define open data policies for the gathered or newly-generated data in order to define a national regulation and their re-use.

Hence, through the use of High Performance Computing, Highlander will make it possible to process data for generating climate forecasts and reducing the risks associated with climate change, for a more intelligent and sustainable management of natural resources and the territory.

<https://highlanderproject.eu/>



**Highlander**  
High performance computing  
to support smart land services



Co-financed by the Connecting European Facility Programme of the European Union  
Grant agreement n° INEA/CEF/ICT/A2018/1815462



# Mistral

Giuseppe Trotta and Gabriella Scipione  
Cineca

The project is funded under the Connecting Europe Facility (CEF) - Telecommunication Sector Programme of the European Union and started on 1 October 2018.

The core value proposition of MISTRAL platform is to create a National Meteorological Open Data Portal that provides to citizens, public administrations, and national and international private organizations meteorological data from observational networks, historical and real time weather analyses and forecasts, with high accuracy in terms of their availability gridded fields, probabilistic products, as rainfall forecasts for flash flood prediction and other data and products coming from the Italian operational forecast modelling chain.

The national portal involves the implementation of several services coming from HPC post processing activities that are summarized through the following use cases:

- Weather Data Downloader, that allows the users to download Observed Data (Current, Historical, Satellite, Radar) as well as Forecasted data on a regular grid interactive service (at high resolution and for different hours and with different updates, weather alerts);
- Italy Flash Flood: Gridded, high accuracy, probabilistic rainfall forecasts, for flash flood prediction. It is an application delivering probabilistic forecasts of 6-hour rainfall for Italy, created by ECMWF blending together post-processed ensemble forecasts, from ECMWF (“ecPoint”) and COSMO. Output shows pre-defined rainfall percentiles and probability thresholds for each COSMO grid point;
- Multimodal Forecasts: an application based on the post-processing of different meteorological Models, for issuing temperature and relative humidity forecasts over Italian weather stations. The Multimodel SuperEnsemble technique is performed by a weight calculation after a performance evaluation during a training period;
- Weather maps of observational and forecast data and post-processing layers.

The National Portal is a stable instrument for the dissemination of meteorological data, and acts as Generic Service of the Core Service known as “European Data Portal”. The high importance of this class of information, which many critical decisions are based upon, such as triggering an emergency intervention and/or evacuation of a specific area, represents the primary value for the long-term sustainability of the entire platform.

Cineca is the lead partner and the consortium is composed by Italian Civil Protection Department, the European Centre for Medium-Range Weather Forecasts, the Regional Agencies for the Environmental Protection of Emilia-Romagna and Piedmont and Dedagroup Public Services.

The goal of the MISTRAL portal is to facilitate and foster the re-use of the datasets by the weather community, as well as by its cross-area communities, to provide added value services through the use of HPC resources, turning it into the level of new business opportunities. The Hackathon organised at the Politecnico di Milano in May 2019 gathered dozens of participants from different sectors interested in the collection and management of meteorological data and proved a high level of interest in the opportunities offered by Mistral.

<http://www.mistralportal.it/>



Co-financed by the Connecting European Facility Programme of the European Union Grant agreement n° INEA/CEF/ICT/A2017/1567101.

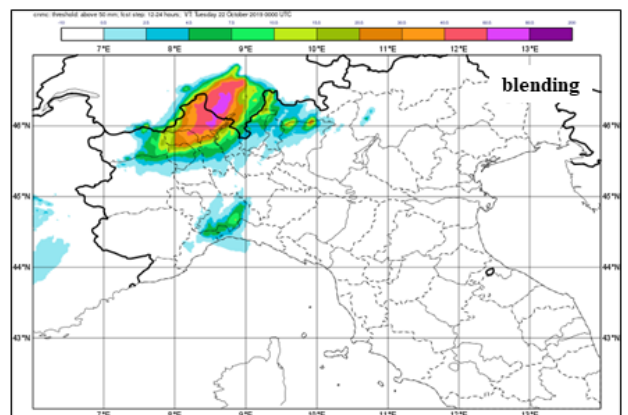
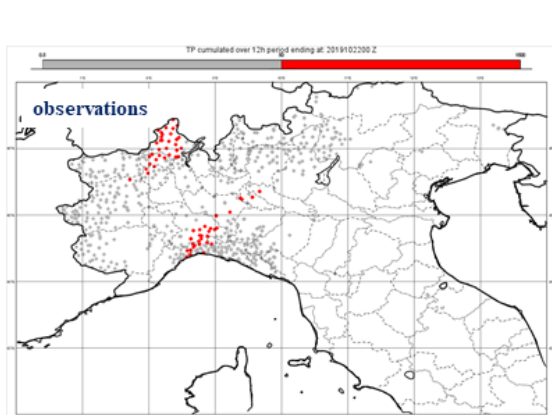
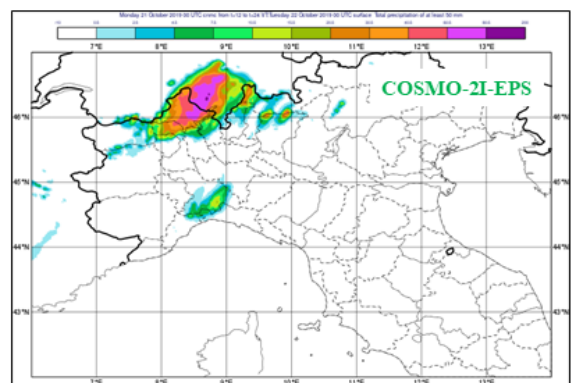
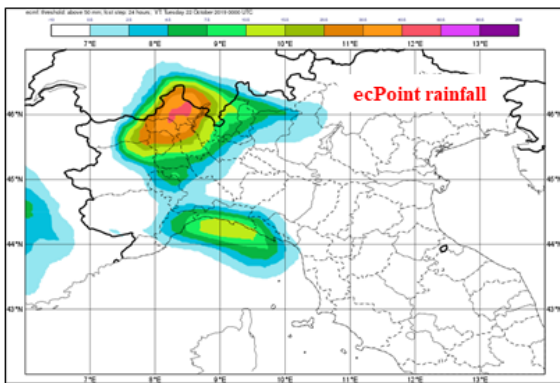
One of the main components of the Mistral platform is Meteo-Hub which is currently being implemented. Its main goal is to collect both observational and forecast data and provide a user interface for customized data selection and extraction. Starting from a list of available datasets specific filters can be applied as well as post-processing operations for data manipulation and transformation can be activated. Among post-processing currently available may include calculation of derived variables; cumulations, averages and other statistical elaborations over time; geographical elaborations, interpolations and cropping. The user who has been given access to the platform can directly submit a request or choose to schedule a plan of continuous extractions for updated data. Through a private area it is possible to manage

each request and in particular inspect the applied filters, monitor status and download the extraction result. In addition to the numerical representation of the data Mistral platform provides tools for graphical visualization of the observed and forecast data such as layered and choropleth maps and various plots to monitor the progress of a certain parameter over time.

An important objective for the Mistral project is to obtain authorisations from the regional functional centres to use the observational data of the ground stations in order to feed the Mistral portal. After the first year of work five regions have already joined: Emilia Romagna, Piedmont, Lazio, Calabria and Campania. Contacts with other regions are going on for presenting data in a unified manner, on a single portal.

## Case study of 21 October 2019: heavy precipitation over Piedmont

IC: 21/10/2019 00Z fc+12-24h; consider probability of tp > 50 mm/12h for ecPoint rainfall, COSMO-2I-EPS and the blending



**A result of the Italy Flash Flood Use case: Both ecPoint rainfall and COSMO-2I-EPS struggle to predict properly the locations possibly affected by heavy precipitation (see two images above). With respect to the observations, the best results are obtained by blending ecPoint rainfall and COSMO-2I-EPS (see blending image).**

Images obtained by



partner of Mistral project.

# DARE

Antonella Guidazzoli and Cinzia Zannoni  
Cineca

Digital transition is an unprecedented opportunity for the development of communities. It enables the connection of people, ideas and projects, enables the growth of human economic and social capital through new models of interaction between citizens and institutions.

DARE - Digital Environment for collaborative Alliances to Regenerate urban ecosystems in middle-sized cities - is a three year project under the 4th European grant for Urban Innovative Actions (UIA). DARE wants to demonstrate the effectiveness of an innovative digital-based and citizen-centred governance approach, aimed at facilitating, supporting and speeding up the implementation and evaluation of the regeneration process of the Darsena, a neighbourhood in the Italian city of Ravenna. DARE should become an example for other EU middle-sized cities involved in urban regeneration processes.

The project proposes the creation of a three layer digital environment, composed by a Data Management Platform (data layer), a Content Management System (editorial layer) and a so-called ViR-Virtual Realm (presentation layer), intended as the enabling technology needed to activate urban actors. The digital infrastructure

will allow the collection, management and creation of data related to the docks (Darsena), such as, among others, vehicular traffic, economic data about the activities of the area, environmental and population data. The platform, therefore, will make available a real-time snapshot of the environmental, social and economic situation of the Darsena. The role of the public administration should be innovated and it is called to act as process enabler, connecting political strategy, economic opportunities and citizens' needs. The main aim is to spread digital culture among the citizens in order to accompany them in becoming digitally aware city changers, that is, to put citizens at the centre by fostering their digital culture in order to benefit critically and actively from the chances that the digital society offers.

The approach will be collaborative: citizens will be involved through various solutions (portal, apps, information totems, events, workshops and courses) for a collective narration of the territory and the urban regeneration path.

At the end of the project it is expected that the Darsena area will be perceived by citizens, visitors, investors as a digital-based attractive, innovative, healthy, safe and inclusive urban ecosystem.

UN PROGETTO



**DARE**  
Digitale reale



**dare / avviso n.1**

**Cerchiamo**  
**persone**  
**curiose**



# CYBELE

Claudio Arlandini  
Cineca

CYBELE: Fostering Precision Agriculture and Livestock Farming through Secure Access to Large-Scale HPC-Enabled Virtual Industrial Experimentation Environment Empowering Scalable Big Data Analytics.

CYBELE generates innovation and creates value in the agri-food domain, and focus on the sub-domains of Precision Agriculture and Precision Livestock Farming in specific, as the real-life industrial cases supported, so the project empower capacity building within the industrial and research community.

Since agriculture is a high-volume business with low operational efficiency, CYBELE aspires to demonstrate how the convergence of HPC, Big Data, Cloud Computing and the IoT can revolutionize farming, reduce scarcity and increase food supply, bringing social, economic, and environmental benefits. CYBELE intends to safeguard that stakeholders have integrated, unmediated access to a vast number of largescale datasets of different types from a variety of sources, and they are capable of generating value and extracting insights, by providing secure and unmediated access to largescale HPC infrastructures supporting data discovery, processing, combination and visualization services, solving challenges modelled as mathematical algorithms requiring high computing power.

CYBELE develops large scale HPC-enabled test beds and delivers a distributed big data management architecture and a data management strategy providing: integrated, unmediated access to large scale datasets of

diverse types from a multitude of distributed data sources; a data and service driven virtual HPC enabled environment supporting the execution of multi-parametric agrifood related impact model experiments, optimizing the features of processing large scale datasets; also a bouquet of domain specific and generic services on top of the virtual research environment facilitating the elicitation of knowledge from big agri-food related data, addressing the issue of increasing responsiveness and empowering automation-assisted decision making, empowering the stakeholders to use resources in a more environmentally responsible manner, improve sourcing decisions, and implement circular economy solutions in the food chain.

Cineca's role in the project is mainly related to the hosting of the platform and helping to develop the required resource abstraction and management mechanisms to support the execution of data analytics workflows on HPC systems. It is also particularly involved in the development of a specific test case, aimed to develop an innovative tool for hail and frost forecast for fruit producers.

<https://www.cybele-project.eu/>



The CYBELE project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 825355.



# CYBELE

# Io Twins

Eric Pascolo  
Cineca

IoTwins: an unique and flexible platform for the creation of digital twins for industrial plants, a 20 million Euro European-funded project coordinated by Bonfiglioli Riduttori SpA in which Cineca participates to provide its know-how in the field of big data analysis and in the definition of complex workflows on cloud and HPC systems.

Within the project there will be 12 testbeds, each realizing a digital twin, the digitization of an industrial process, in order to optimize and improve its management.

Funded under Horizon 2020, IoTwins plans a total investment of 20 million € and is coordinated by Bonfiglioli Riduttori SpA, a world leader in the production of power reducers. The project benefits also from the scientific coordination of the University of Bologna, the computing infrastructures of Cineca and INFN, the National Institute for Nuclear Physics, and is under the aegis of the Emilia-Romagna Region through Art-ER, the regional company supporting innovation, attractiveness and internationalization of the territorial system. In addition to local partners, including Marpossa a world leader in the supply of precision instruments, the project will also include companies and scientific institutions of international importance such as Siemens, Barcelona Supercomputing Centre, the Barcelona Futbol Club, and the Fraunhofer Institute in Munich.

Digital twins are virtual copies of real processes and plants that can interact with each other

and, through the collection of large amounts of data, can simulate different scenarios to define corrective actions, optimize efficiency and diagnose anomalies before they occur.

In particular, 4 testbeds will be carried out in the manufacturing sector for predictive diagnosis in the numerous sectors of activity in which the consortium's industrial partners operate (wind turbines, machine tools for the production of automotive components, machines for the production of crankshafts, machines for the production of bottle caps). Another 3 testbeds will focus on the management and, in particular, will be applied to the management optimization of the Barcelona Camp Nou Stadium, on the so called Smart Grid, i.e. the set of an information network and a distribution network to manage the electricity grid in an "intelligent" manner, to monitor the quality of energy and the holistic management of supercomputing infrastructures. The remaining 5 testbeds - strength and value of the project for the entire regional industrial system - will cover the replicability and standardization of previous models to define new areas of application and to develop innovative business models in manufacturing and infrastructure management.



The IoTwins project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 857191.



# AI4EU – European AI On Demand Platform and Ecosystem

Roberta Turra  
Cineca

AI4EU has been funded by the European Commission under the topic ICT-26-2018-2020 - Artificial Intelligence, with 20M€ and started on the 1st January 2019. Coordinated by Thales, its objective is to efficiently build a comprehensive European AI-on-demand platform to lower barriers for innovation, to boost technology transfer and catalyse the growth of start-ups and SMEs in all sectors through Open calls and other actions. The platform will act as a broker, developer and one-stop shop providing and showcasing services, expertise, algorithms, software frameworks, development tools, components, modules, data, computing resources, prototyping functions and access to funding. Training will enable different user communities (engineers, civic leaders, etc.) to obtain skills and certifications. The AI4EU Platform will establish a world reference, built upon and interoperable with existing AI and data components (e.g. the Acumos open-source framework, QWT search engine ...) and platforms. It will mobilize the whole European AI ecosystem and already engages 80 partners in 21 countries including researchers, innovators and related talents. Eight industry-driven AI pilots will demonstrate the value of the platform as an innovation tool. In order to enhance the platform, research on five key interconnected AI scientific areas will be carried out using platform technologies and results will be implemented. The pilots and research will showcase how AI4EU can stimulate scientific discovery and technological innovation. An AI4EU Ethical Observatory has been established to ensure the respect of human

centred AI values and European regulations. Sustainability will be ensured via the creation of the AI4EU Foundation. The results will feed a new and comprehensive Strategic Research Innovation Agenda for Europe.

Cineca role in the Project is to provide technical environments for the pilot experiments, in collaboration with BSC. AI components orchestration and deployment will be tested with real life private data, in respect of privacy and IPR requirements through deployment of specialized Data Governance frameworks. Parallelization of the work and orchestration will be adapted and optimized for proper scalability. Cineca will benefit of its participation to the project in several ways: enabling the deployment of models and services that will be provided through the AI4EU platform, increasing market penetration for industrial edge/cloud services, making technological synergies with other AI related projects (e.g. Cybele, IoTwins, Highlander), making domain specific synergies with other pilots (e.g. in healthcare, precision agriculture, smart cities, industry ...), increasing the funding opportunities for Cineca industrial users through the AI4EU Open calls and intensifying the relationships with Big Data, edge/cloud and IoT ecosystem at the European Level.



The AI4EU project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 825619.





# FUTURE TRENDS

Cineca's mission is to support scientific research by providing access to advanced computational systems. Technologies such as Deep Learning, Quantum Computing, accelerated computing, will gradually revolutionize the way scientists conduct numerical experiments. It is part of our mission to design Cineca systems by incorporating technology changes and ensuring users to exploit them in a profitable manner.

# Completion of the MARCONI project for the evolution of the supercomputing system

**Sanzio Bassini**  
Cineca

At the peak of the use and exploitation of the supercomputing system FERMI project, 2011 - 2016, in 2014 the MARCONI project was developed for the evolution of the supercomputing system of Cineca allocating a budget of 50 million € to the project.

The MARCONI project was developed by defining two implementation phases. A first immediate phase, MARCONI-1 time frame 2016 - 2018, to increase the computing power by a 10x factor in comparison with what has been made available through the previous FERMI project. A second phase of upgrade, MARCONI-100 time frame 2019 - 2021, to bring the overall performance of Cineca supercomputing architecture to the order of 50 PFlops, the first step of a national roadmap towards exascale.

It is important to underline that, in light of evident elements of technological innovation of the supercomputing systems and constraints of economic sustainability and energy efficiency, the MARCONI project was already developed in 2014 considering the need to integrate in the Cineca supercomputing system a partition based on conventional components, let's say scalar, and a partition based on unconventional components, let's say hybrid or accelerated. Furthermore, it became increasingly important to start an architectural conceptualization that went beyond the single monolithic supercomputing system, but that it was necessary to have an overall architectural vision of the supercomputing system, which saw in single system image complementary and modular integration of diversified partitions within the architecture of the Tier-0 class system, and, on the other hand, a greater functional integration of the two systems, historically present in

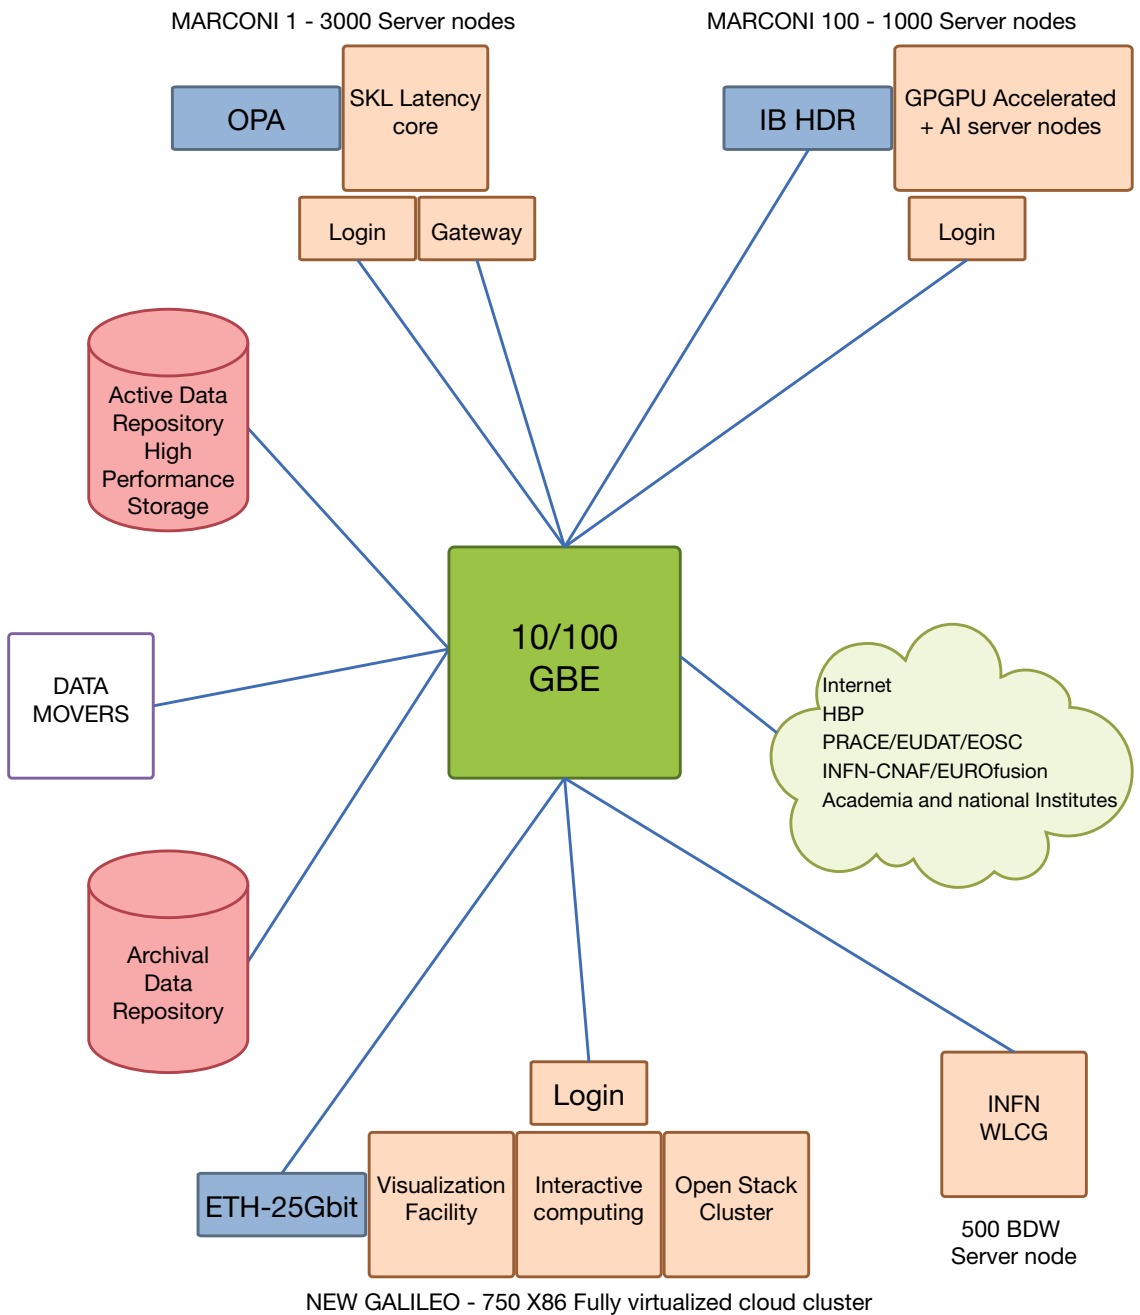
the Cineca configuration, the Tier-0 system, the scale out system, and the Tier-1 system, the system with the general purpose usage characterization.

Based on these design requirements, the MARCONI-1 system, also through adequate increases in the budget initially assigned, following the completion of the agreement with the EUROfusion community and the joint development projects with INFN, SISSA / ICTP and University of Milano-Bicocca, has reached the configuration of 7000 server nodes, half based on Intel Skylake microprocessors and half based on Intel Knight Landing microprocessors with a global performance in the order of 20 PFlops peak.

Last year the MARCONI project was completed with the upgrade planned for the second phase. The upgrade was also achieved thanks to the co-funding obtained through participation in the Procurement per Innovation action PPI4HPC. The upgrade consists in replacing the unconventional partition with a partition, logical name MARCONI-100, configured with hybrid server nodes with NVIDIA Volta GPGPU accelerators. The computational power of this partition is over 30 PFlops peak and brings the computing power of the MARCONI architecture, MARCONI-1 plus MARCONI-100, to about 50 PFlops peak. MARCONI-100 will remain in production until mid-2022 and will subsequently be replaced by the EuroHPC LEONARDO project pre-exascale system. With a longer-term view, already included in the budget allocation of the MARCONI project, the entire evolution path of the Cineca supercomputing system architecture will also include the updating of the Galileo Tier-1 system.

The updating of the Galileo system is part of the ICEI FENIX action, co-funded within the Human Brain flagship Project, for the creation of a European federated infrastructure of 5 European supercomputing centres, BSC, CEA, Cineca, CSCS, JUELICH. The new Galileo system will be configured with conventional scalar technology microprocessors, capable of offering cloud computing services.

In overall terms, the MARCONI project in the 2016 - 2021 period involved an investment of about 70 million €, of which 50 million € directly funded by Cineca and 20 million € co-funded by participation in projects for technological innovation of the European supercomputing infrastructure and through joint development actions with qualified Italian institutions.



# EuroHPC: a European and Italian success story

Gabriella Scipione  
Cineca

2019 has been a busy year for EuroHPC, the initiative leading the way in the European Exascale Supercomputing. Many of the planned goals have been reached thanks to the operative and effective collaboration put in place between the EU and the 32 European countries members of the Joint Undertaking.

## EuroHPC JU approved 840 million € funding for eight supercomputers

The 7th EuroHPC JU governing board of 4-5 June approved funding for 3 pre-exascale (capable of executing more than 150 Petaflop/s, that will be in the global top 5) and 5 petascale (capable of executing at least 4 Petaflop/s) supercomputers. The total value is 840 million € of which almost half will be paid by the EuroHPC budget. The other part will be paid by the bidding countries.

The pre-exascale systems are expected to provide 4-5 times more computing power than the current top supercomputing systems of the Partnership for Advanced Computing in Europe (PRACE). Together with the petascale systems, they will definitely increase the supercomputing resources available for European-level use, consequently many more users will have access to them.

The hosting sites of the 3 pre-exascale systems will be located in Bologna (Italy), Kajaani (Finland) and Barcelona (Spain). While the 5 petascale systems in Sofia (Bulgaria), Ostrava (Czechia), Bissen (Luxembourg), Minho (Portugal) and Maribor (Slovenia).

They will support the development of major applications in domains such as personalised medicine, drug and material design, bio-engineering, weather forecasting, and climate change.

## Leonardo Supercomputer and Consortium

Italy will host one of the pre-exascale class computers. Funded by the EU and the Ministry of the University and Research, conceived and managed by Cineca in collaboration with INFN and SISSA. The Italian pre-exascale system, Leonardo, will be one of the most powerful supercomputers in the world, able to project Italy towards the exascale class of HPC for research and innovation.

Leonardo was presented within the European Researchers Night, during an event with the directors and representatives of the most important Italian Research institutions, universities, researchers, and Roberto Viola Director General for Communications Networks, Content and Technology (conference call).

Leonardo will deliver between 220 and 250 peak PetaFlops, spread over 5000 compute nodes and laced together with a 200 Gb/sec interconnect. The system will also include 150 PB of storage and the entire machine is expected to consume 9 megawatts. Leonardo will foster the convergence of HPC, artificial intelligence (AI), and high-performance data analytics (HPDA).

Italy decided to fully co-fund 50% of the TCO required to host and maintain the system for the period indicated within the call (2019-2026). Italy was also very active in promoting and building a consortium for advanced computing with the Central-Eastern European Countries. During this year Austria, Slovakia, Slovenia and Hungary joined the Leonardo Consortium with the objective to improve the collaboration among them in order to foster the European HPC ecosystem and to collaborate and develop common activities of joined infrastructures (Slovenia will host 1 one of the five petascale systems), competence centres, research, development and innovation projects and initiatives in the national, bi-lateral and European context within EuroHPC.

Also Finland and Spain formed their respectively consortia, more focused on sharing the co-fund 50% of the TCO (Total Cost of Ownership) required. In total in EuroHPC, 21 of the 32 countries participating in the Joint Undertaking will be part of the consortia operating the centres, demonstrating the great success that this initiative had in Europe.

## Cineca signed the Hosting Agreement and EuroHPC launched the call for tender

The 26th of November 2019 in Strassbourg representatives of the eight centres representing the EuroHPC pre-exascale and petascale hosting entities signed the hosting agreements.



The 28th of November 2019 EuroHPC JU launched a call for tender for the procurement of three pre-exascale supercomputers in Italy, Finland and Spain with a close deadline set for the 3rd of January 2020. The total value of the tender is 416 million €. The machines will be owned by the EuroHPC JU, while the 5 petascale supercomputers will not be procured by EuroHPC JU but by the individual hosting entities.

The procurement is divided into three lots. One economic operator per lot will be selected for the component(s) acquisition, delivery, assembly, hardware and software installation and maintenance of these supercomputers.

The “Leonardo” EuroHPC supercomputer will be hosted in Bologna, Italy, in the premises of the new data centre of Cineca, at “Bologna Tecnopolo”, a site owned by the Emilia-Romagna and conceived for hosting also the ECMWF (European Centre for Medium Range Weather Forecast) data centre and the INFN Tier-1 system of the LHC experiment at CERN.

#### EuroHPC calls for proposal

EuroHPC has also the aim of supporting an ambitious research and innovation agenda, covering all scientific and industrial HPC value

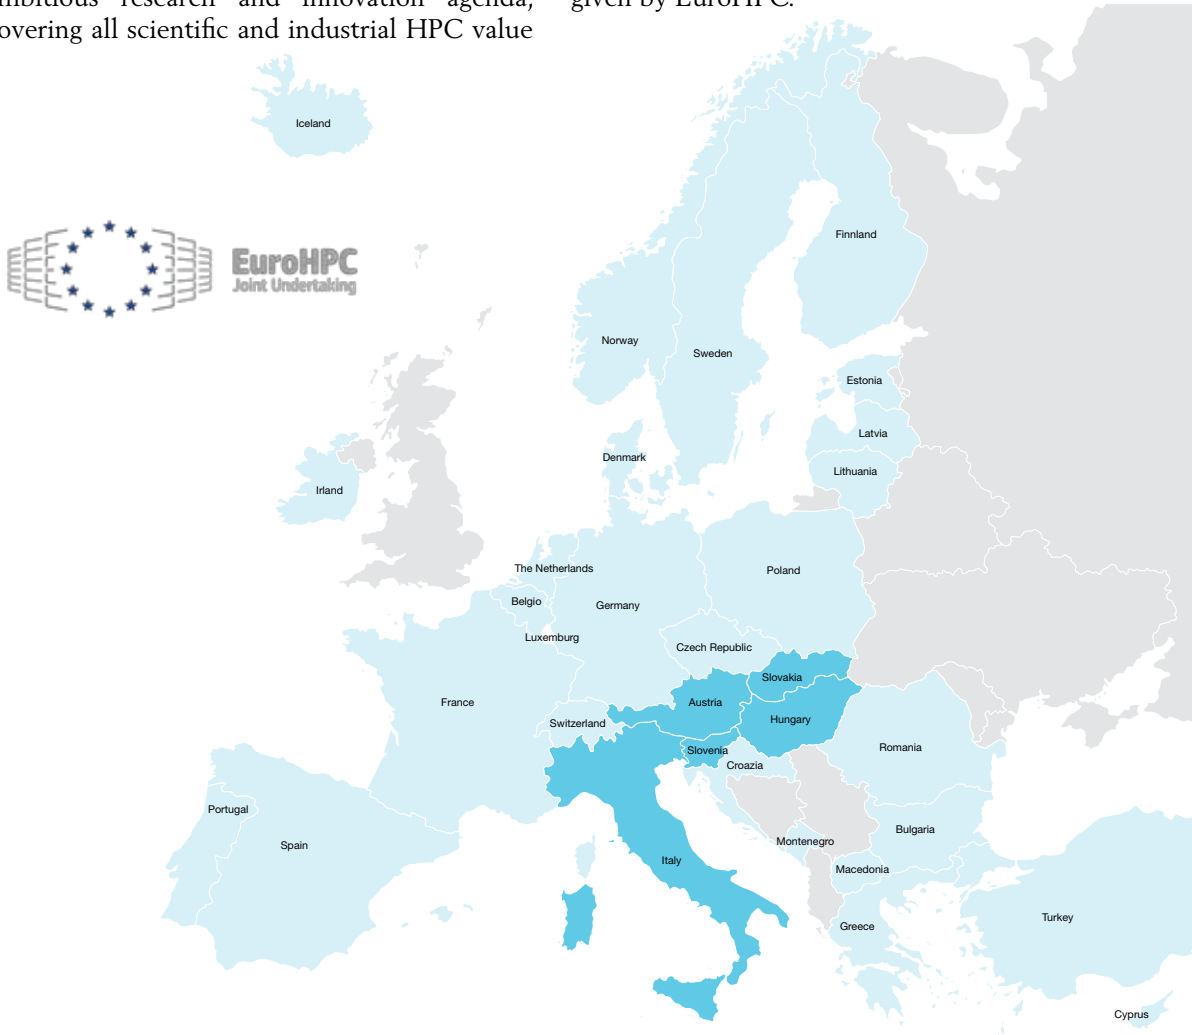
chain segments, including low-power processor and middleware technologies, algorithms and code design, applications and systems, know-how and skills for the next generation supercomputing era.

In 2019 the calls for proposals launched were two: Towards Extreme Scale Technologies and Applications and Innovating and Widening the HPC use and skills base.

The funds also in this case come from the EU and the respectively national ministries.

The Italian national funding authority in this case is the Ministry of Economic Development that contributed with 24 million €. The results of the calls will be published in the EuroHPC website during the course of 2020.

All these achievements are examples of a European and Italian success story. EU countries demonstrated a fruitful cooperation necessary to drive innovation and compete globally in order to boost Europe’s competitiveness in the digital area. Italy has seen the collaboration of its national stakeholders, MIUR , MISE, Cineca, INFN, SISSA and the Emilia Romagna Region, that worked all together to seize the opportunity given by EuroHPC.



# Accelerated High Performance Computing

Isabella Baccarelli, Luca Ferraro, Sergio Orlandini, Marco Rorro and Stefano Tagliaventi  
Cineca


Graphics Processing Unit (GPU) is a specialized hardware component originally developed for graphics applications. Surely unaware of the future ahead, and of GPU's impressive use in the nowadays high performance computing world, its design answered the need to process massive buffers of memory containing the bitmaps describing the texture (color pattern) of each object to be rendered. The involved computations were fairly simple, with no much branching with respect to typical CPU workloads, and entirely independent. This is a crucial aspect, in terms of application-oriented hardware: the GPUs were designed to have a high degree of parallelism with thousand of available cores, a high memory bandwidth, and a lighter context switch, at the cost of having a lower clock speed and less branching capability with respect to traditional CPUs. The first graphics accelerator marketed as "the world's first GPU" was the NVIDIA GeForce 256 in 1999 but its hardware was so specialized that it could only be used for graphics tasks. With the introduction of NVIDIA GeForce 8 Series, and the new generic stream processing unit, GPUs finally became a more generalized computing device, capable to execute arbitrary code (not just rendering routines) and programmers started making computational inroads against the CPUs. A fundamental role was played by NVIDIA CUDA programming language, which provided a way to write a GPU code in a C-like language (and Fortran-like, see PGI and IBM XL compilers extensions). More in general, different programming paradigms were inspired by this new architecture, and many

workable options alternative to CUDA became available for programming heterogeneous platforms such as OpenCL, an open standard for parallel computing. A different, more rapid and straightforward approach in accelerating application is the use of compiler directives given by OpenACC or OpenMP, or of the many GPU enabled and ready to use computational libraries. By now a subfield of research, GPU Computing or GPGPU that stands for "General Purpose Computing on GPU" has found its way into fields as diverse as: molecular dynamics, computational physics, quantum chemistry, computational fluid dynamics, weather modeling, statistics, bioinformatics, life science, medicine, not to mention industrial applications like oil exploration and stock options pricing determination or, last but not least, the trendy and extremely versatile machine and deep learning applications.

Many computing centres, universities and industries are moving towards accelerated solutions with GPU hardware. The importance that GPUs have achieved is also clear from the top positions of the TOP500 ranking. The first two positions are indeed occupied by clusters accelerated with GPUs and both with the same architecture, IBM Power-9 plus NVIDIA Volta V100. This promising architecture is the one that was chosen by Cineca for the next accelerated partition of the Tier-0 cluster Marconi, that will be equipped with about 1000 nodes and 4 NVIDIA Volta V100 per node. This outstanding computing power will be available to all the National and European academic communities in the early 2020.

In order to fully exploit the new hardware, the commitment of Cineca Consortium is to transfer our know-how in hybrid architecture and programming to the scientific community, by means of delivering courses concerning GPU architecture and hosting GPU Hackathons - sprint-like events where GPU programming specialists and academic researchers collaborate alongside and intensively on software projects for few days. Moreover, we are actively engaged

in giving support to academic community, as well as to Italian industrial companies of primary excellence, in developing technical and scientific applications. We consider this engagement a crucial factor in the present transition toward a new cluster technology, stating the pivotal mission of Cineca for 2020 in supporting this arising architecture in order to provide a real boost of the academic and industrial production.



# Supporting Eni HPC infrastructure

Cinzia Zannoni  
Cineca

For over 15 years, Cineca has been collaborating with Eni (the most important Italian Energy Company) in the set up of the company HPC ecosystem, implementing the most advanced algorithms developed by Eni, developing applications and managing the company HPC systems. From 2007 till 2012, Cineca hosted Eni supercomputers in its own datacentre. In 2013, Eni set up its own HPC infrastructure in its Green data Centre (GDC). Cineca continues to manage Eni's HPC infrastructure in GDC.

On Cineca side, the collaboration involves many people with different skills, able to work in research projects, and in the development and management of HPC applications and systems. In March 2020, Eni set up its fifth generation High Performance Computing System, HPC5, located in the Eni's Green Data Centre (GDC). The system has a peak performance of 52 PFlops, is expected to be ranked in the fifth position of the June 2020 TOP500 List, emerging as the most powerful HPC system World Wide installed by an industrial company.



HPC5, associated to the other HPC4 system, takes the whole Eni's HPC infrastructure to the availability of 70 PFlops.

This infrastructure provides a strategic support to Eni's digital transformation process along its whole value chain, from exploration to reservoir management to the analysis of all the data generated during the operations.

In 2019 Cineca collaborated with Eni in the evaluation of novel HPC technologies and on the evolution and maintenance of HPC applications, as well as on new Eni research projects developing new applications for reservoir simulation, data analysis and interpretation.

In the course of 2019, 40 projects were developed with Eni:

- 5 projects concerned the parallelization, optimization and portability of hybrid codes on different platforms,
- 2 project concerned the evolution of the HPC production environment,
- 13 projects concerned the development and maintenance of production applications,
- 20 projects concerned the development of 5 new applications and 6 new components and new solvers, to be integrated into the production applications.

The collaboration with Eni allowed Cineca staff to face the problems and the typical needs of industrial production, to grow in the understanding of the specific domain and in the ability to work in highly specialized and effective multidisciplinary teams.

Courtesy of [www.eni.com](http://www.eni.com)



# National ecosystem for Industry 4.0

Claudio Arlandini  
Cineca

Advanced production systems are based on a multitiered pyramid of enabling IT technologies, going from a primary level including Cloud Computing and cybersecurity, up to Big Data, Simulation, Digital Twin, Cognitive technologies built on top of those. Cineca infrastructure and competences, reinforced through a network of qualified partners to constitute a real innovation ecosystem, is the ideal partner for industrial innovation projects that lead to tangible Return-on-Investments and ultimately new jobs creation.

The action of Cineca as a Digital Innovation Hub goes through four support pillars:

- enabling test-before-invest proof of concepts projects;
- supporting the individuation of funding opportunities for innovation enablement;
- introducing the partner in a valuable innovation and networking ecosystem;
- enhancing skills and providing tailor-made training.

At the national level, the ecosystem is developed with the close collaboration with two important innovation enablers, the Associazione Big Data<sup>[1]</sup> and the BI-REX Competence Centre<sup>[2]</sup>.

The Big Data Association collects all the stakeholders involved in supercomputing and big data production and management in Italy to map and analyse their potential impact on scientific, social and business domains. The

aim is to interconnect and jointly exploit the knowledge, capacities, research and innovation potentials of this community to leverage the effects of actions and investments made so far and to maximize their impacts, locally but also at national, EU and international levels.

BI-REX is one of the six Italian Competence Centres for Industry 4.0 established in 2018, bringing together 57 actors in partnership between universities, research centres and 45 companies of excellence to assist companies and especially SMEs in the adoption of enabling technologies Industry 4.0, with a special focus on Big Data, additive Manufacturing and new related business models.

At the international level, Cineca offers support to industries through initiatives like the PRACE SHAPE<sup>[3]</sup> (SME HPC Adoption Programme in Europe) or the BDVA (Big Data Value Association) i-Spaces<sup>[4]</sup> programmes, or in projects financed by the European Commission through the H2020 Framework Programme and the EuroHPC Joint Undertaking. In particular, EuroHPC is building a network of HPC Competence Centres, one for each EU country, providing HPC services to industry (including to SMEs), academia and public administrations, delivering tailored /modular solutions for a wide variety of users, with an aim to ease and foster the transition towards wider uptake of HPC in Europe.

## References

[1] <https://associazionebigdata.it/>

[2] <http://bi-r-ex.it/>

[3] <http://www.prace-ri.eu/hpc-access/shape-programme/>

[4] <http://www.bdva.eu/1-Spaces>

## National ecosystem



# Quantum Computing

Daniele Ottaviani  
Cineca

Beginning in the middle of the last century, when the first rudimentary computers made their appearance in the world, humanity seemed to be heading towards a technologically thriving future. Inexorably, the power of digital computers grew exponentially day after day, letting humanity to imagine a bright future: everything that is impossible today, tomorrow will be tackled, thanks to the new computing power available. This actually went on and on until the beginning of the new millennium, when suddenly the miniaturization technology of the CPUs, which was the first driver of the unbridled technological evolution of computers, had to face a fundamental problem: the rules that manage the quantum world.

Laws different from those that regulate the macroscopic world: the laws of quantum mechanics. Overcoming the stumbling block of these laws is still a gigantic problem for digital computers: when the miniaturization of the CPUs reached the quantum world, it was thought to continue the evolution of classic computers with multi-core technology, increasing the number of processes parallel in a single computer. In this way the total blockage of the evolution of digital computers was temporarily averted, even slowed down a lot. Towards the end of the last century a further idea was proposed to circumvent the gigantic problem: trying to find a way to exploit the unpredictable quantum behaviours to our advantage.

To do this, however, scientists immediately realized that they had to think about something completely different from the classic computers to which we were all used. Thus was born the idea of the quantum computer: unlike its classical counterpart, it revolutionizes the concept of computer from the basis: its basic logical unit, the qubit, is a bit capable of exploiting the peculiarities of quantum mechanics to its advantage. In particular, the qubit is able to exploit the quantum superposition, a phenomenon whereby a microscopic particle manages to assume several classical states simultaneously. At the computing level, for example, this property is exploited to allow intrinsically parallel calculations; a quantum computer, in fact, is able to process multiple

states simultaneously, making this machine exponentially more powerful than its analogous counterpart. Today we are experiencing the 1960s of quantum computers: they exist, they work, but they are still prototypes, not yet able to support classic computers in production calculations.

We at Cineca have been following the evolution of this very interesting new field since 2018. In these two years we have been able to see the speed of evolution of the new quantum machines. Although still far from being used at their full power, the new calculators require skills completely different from the classic programming; this implies, among other things, the necessity to be ready when one day not too far away they will become mature enough to appear in the world of scientific production.


For this reason Cineca is continuing to monitor their evolution both by employing people in the field and by organizing scientific dissemination events, such as the workshop “Quantum Computing and High Performance Computing”, with its second edition in 2019. It is Cineca’s intention to become a point of reference for Italian quantum computing, providing spaces of visibility, hours of calculation and specialist support for the research projects of Italian universities.

In 2020, Cineca is going to make a concrete commitment to realizing this vision. In recent months we are drafting documents for a tender that will collect the best quantum service vendors to buy the resources we need. Among the big names we have in mind, stand out IBM (for general purpose quantum computing) and D-Wave (for quantum annealing), even if the public nature of the tender makes it accessible to all those who are able to satisfy our prerequisites. In this first phase we decided to buy services to be redistributed to Italian research bodies: subsequently, this service will turn into the probable purchase of a real quantum computer. It is hoped that by the end of the year Cineca will be able to effectively distribute these resources to the public; the scientists who will take advantage of our services will then tell, if they like, their experiences in the next edition of our annual workshop, which will probably be held a few days before Christmas 2020.

# Quantum Computing and High Performance Computing

CINECA - Casalecchio di Reno, Bologna

See you again soon!



CINECA











# CINECA

Headquarters:  
Via Magnanelli, 6/3  
40033 Casalecchio di Reno,  
Bologna, Italy  
Tel. +39 051 6171411

Via R. Sanzio, 4  
20090 Segrate, Milano, Italy  
Tel. +39 02 269951

Via dei Tizi, 6/B  
00185 Roma, Italy  
Tel. +39 06 444861

Via Medina, 40  
80133 Napoli, Italy

Via dei Vestini, 31  
66100 Chieti, Italy

ISBN 9788886037419