


# CINECA

2017







Cineca HPC Report 2017  
website: [www.hpc.cineca.it](http://www.hpc.cineca.it)  
mail: [info-hpc@cinca.it](mailto:info-hpc@cinca.it)

Contents are © CINECA 2017 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  
Monograf s.r.l.  
Via Collamarini 5  
40138 - Bologna - Italy

Cover image:  
David Radice - article on page 50  
*Mass ejecta, disk and remnant neutron star from a binary neutron star merger.*  
Simulations performed on Marconi with the THC code.

Thanks to:

Fabio Affinito; Paola Alberigo; Andrea Bartolini; Andrea Benassi; Sebastiano Bernuzzi; Emanuela Bianchi Janetti; Luca Biferale; Gianluigi Bodo; Francesco Bonelli; Pietro Bonfà; Lidia Bressan; Gianbattista Bussi; Claudio Cacciari; Claudio Cafiero; Francesco Califano; Sandro Capetti; Carlo Cavazzoni; Ines Cerenzia; Silvio Sergio Cerri; Mirko Cestari; Giovanni Chillemi; Letizia Chiodo; Gabriele Coccia; Patrizia Coluccia; Ciro Cottini; Grazia Cottone; Maurizio Cremonesi; Alberto Cuoci; Martina Datteo; Remo De Donno; Alberto Debernardi; Cristiana Di Valentin; Tim Dietrich; Andrew Emerson; Gianluca Fazio; Matteo Ferri; Lara Ferrighi; Giuseppe Fiameni; Alessio Frassoldati; Guido Fratesi; Daniela Galetti; Davide Emilio Galli; Roberto Gaspari; Luca Govoni; Francesco Grasso; Robert Evaristo Grisenti; Massimiliano Guarrasi; Gael Raymond Guedon; Antonella Guidazzoli; Stefano Lovato; Vinicio Magi; Therese Malliavin; Francesco Mambretti; Luca Maragliano; Mario Martina; Silvano Massaglia; Andrea Michelotti; Andrea Mignone; Davide Modesti; Silvia Monfardini; Nicola Munerini; Alessandro Nagar; Giovanni Onida; Paolo Orlandi; Gianfranco Pacchioni; Eric Pascolo; Albino Perego; Daniele Perilli; Sergio Pirozzoli; David Radice; Costanza Ronchi; Paola Rossi; Riccardo Rossi; Philomena Schlexer; Daniele Selli; Ivan Spisso; Marco Tavoso; Debora Testi; Sergio Paolo Tosoni; Claudia Truini; Roberta Turra; Piero Ugliengo; Annarita Viggiano; Riccardo Zanella; Filippo Zanetti; Cinzia Zannoni.

# CINECA HPC REPORT

# 2017



Dear Colleagues,

In preparing for writing this introduction of the annual report 2017 on the outcomes and achievements of the year, I realized the real difficulty is represented in finding a good compromise between the presentation of a long list of ordinary duties and the overemphasis of some events that have happened in the previous months. These may represent small steps in a long journey, but in reality constitute a potential stronghold for further crucial strategic development.

The main achievements are represented by the consolidation of the supercomputing infrastructures and by the collaboration agreements signed with qualified national partners – results that are strictly connected. The National Institute of Nuclear Physics (INFN) joined the Cineca Consortium during 2017 and together with the International School of Advanced Studies (SISSA) in a joint partnership with the International Centre of Theoretical Physics (ICTP) and Universities of the Consortium, signed a collaboration agreement with Cineca to co-fund a substantial upgrade of the Marconi supercomputing system.

Marconi was ranked number 14 in the last November 2017 Top500 and represents one of the main Tier-0 systems in PRACE and a world-class system in Europe. Such a valuable infrastructure constitutes the hardware infrastructure for user communities not only in the traditional domains but also in the fields of precision medicine, artificial intelligence, neural network training and big data analytics as well.

Two events from the past year in particular must be mentioned. The first is the EuroHPC declaration, signed by seven Member States – France, Germany, Italy, Luxembourg, the Netherlands, Portugal and Spain in occasion of the Digital Day, organized in Rome on March 23rd 2017 as part of the 60th anniversary celebrations of the Treaty of Rome, subsequently joined by others that brings to 14 the number of Member States that signed the agreement. The EuroHPC declaration has the aim of acquiring and deploying by 2022-2023 a pan-European integrated exascale supercomputing infrastructure, with an intermediate step of a pre-exascale performance that should be reached by 2020-2021.

The second event concerned the ECMWF (European Centre for Medium Range Weather Forecast) decision to locate the new data centre in Bologna by 2019. ECMWF Member States approved the proposal by the Italian Government and the Emilia Romagna Region in June 2017 as part of an international competition.

These two successes are fostering a very ambitious program to set up in Bologna one of the main concentrations of HPC infrastructures in the world. In the same city there will be the data centre of ECMWF, the world-level HPC centre of Cineca, and the INFN Tier-1 system of the LHC experiment at CERN. Moreover, the last two centres will be connected in order to allow the different communities to access both data repositories in a transparent way from the computing facilities hosted in their data centres. The results of all these developments are leading to a physical infrastructure able to create in Bologna a major European hub for big data processing and computing. Cineca is part of this ecosystem and will continue to maintain its commitment and spend its energies and efforts to keep the Italian digital infrastructure at the top level of the international panorama.

With kindest regards,

  
Sanzio Bassini

Director of High Performance Computing Department of Cineca

## **GENERAL OVERVIEW**

<b>2</b>	<b>2017 in a nutshell</b>
<b>4</b>	<b>Infrastructures and services</b>
<b>6</b>	<b>Usage report</b>
<b>8</b>	<b>Users' statistics</b>
<b>10</b>	<b>Training</b>
<b>11</b>	<b>Projects in numbers</b>
<b>15</b>	<b>Staff</b>
<b>16</b>	<b>Events</b>

## **NEW PROJECTS**

<b>20</b>	<b>HPC-Europa3</b>
<b>22</b>	<b>OPRECOMP</b>
<b>23</b>	<b>European Open Science Cloud</b>
<b>24</b>	<b>Max Center of Excellence</b>
<b>26</b>	<b>PPI4HPC and ICEI</b>



## EMERGING TECHNOLOGIES

HPC infrastructure evolution	30
Trends on energy-efficient supercomputer management	32
Deep Learning	34
Big Data	36

## USERS' PROJECTS

PRACE	40
LISA	41
ISCRA	42
Exemplary projects	44



# GENERAL OVERVIEW

Cineca HPC is the main center 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.

# 2017 in a nutshell

Carlo Cavazzoni and Giuseppe Fiameni  
Cineca

This has been a very important year for the HPC in Cineca, both from the point of view of the infrastructure upgrades and the launch of new initiatives at European level with Cineca as one of the main protagonists. From the point of view of the infrastructure, during 2017 the last partition of Marconi has been installed, bringing the entire computational capability above 20 PFlop/s, that places Marconi within the top systems in the world in the rank of Top500 ([www.top500.org](http://www.top500.org)). Marconi is ranked number 14 in the world and keeps the leadership in the EU area.

The increase of the Marconi configuration was possible also thanks to a very important collaboration with some of the major Italian research bodies (INFN, SISSA, ICTP and Bicocca University, Milan) that contributed in the acquisition of the final part of the Marconi system, with more than 1.000 nodes of the new SkyLake technology by Intel.


At the level of the EU ecosystem, 2017 was a year dense of new challenges and opportunities for Cineca. EU reinforced the commitment to expand the investment in the HPC domain and create a new value chain launching a very ambitious program to fill the gap with other economic areas in the world, such as US, China, Asia/Pacific. Many new initiatives started, and Cineca was among the main institutions involved and representing Italy. In particular for the first time EU decided to co-invest in the system procured by PRACE members, in order to expand the resources available for all EU researchers. Cineca with other important HPC centers in Europe participates to a EU project that aims to manage a point procurement of innovative solutions in HPC (<https://www.ppi4hpc.eu/>).

A similar initiative, ICEI, relates to the Human Brain Project to deploy a supercomputing infrastructure for the neuroscience community. The two projects are presented in the following pages in details.

We are proud of our “women in HPC” (<https://www.womeninhpc.org/>).



# SUPERCOMPUTER MARC:NI



"Every day sees humanity more  
victorious in the struggle with  
space and time."

*Guglielmo Marconi*



# Infrastructures and services

Elda Rossi and Antonella Guidazzoli  
Cineca

In year 2017 the main platform available to researchers for scientific computing was Marconi: the so-called Cineca Tier-0 system.

Marconi was co-designed by Cineca on the Lenovo NeXtScale platform. It is based on the next-generation of the Intel Xeon Phi™ product family alongside with Intel Xeon processor E5-2600 v4 product family, and offers to the scientific community a technologically advanced and energy-efficient high performance computing system.

Marconi was gradually completed in about one and a half year, between mid-2016 and January 2018, according to a plan based on a series of updates:

- July 2016: Marconi-A1 is put in production, equipped with Broadwell chips, with a computational power of 2PFlop/s peak
- January 2017: Marconi-A2 is added, equipped with the Knights Landing nodes, with an additional computational power of approximately 11PFlop/s peak.
- July 2017: Marconi-A3 is added, in a preliminary partition of 1512 nodes, based on SkyLake technology, with an additional power of 5 PFlop/s peak
- January 2018: the Marconi-A3 partition is increased with additional 792 nodes, reaching the final configuration of the system and a total computational power of 20 PFlop/s peak.

Marconi takes advantage of a new Intel architecture for the internal network (OmniPath), which provides the high performance interconnectivity required to efficiently scale the system's thousands of servers.

A high-performance Lenovo GSS storage subsystem provides data storage capacity.

Marconi allows the access to state-of-the-art processor technology, enabling an extremely high-performance system with a 'green' soul. Indeed one of the parameters of the project developed by the Cineca team is to gradually increase the computational power up to 50PFlop/s without exceeding, at any stage, the limit of 3MWatt power consumption.

A second system (Galileo), was available to researchers in 2017 at the Tier-1 level. It is an Intel base cluster dedicated to projects with less parallel requests.

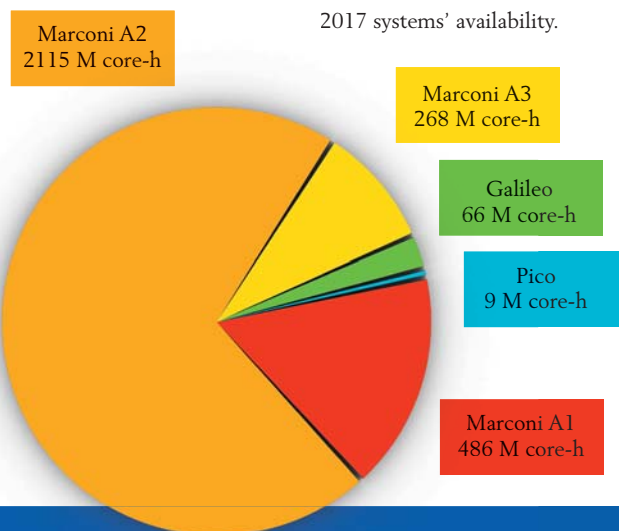
Finally an additional platform, Pico, dedicated to BigData and bioinformatics applications, offers also a cloud-based environment for innovative applications.


The computational capacity of these systems is usually expressed in terms of core-h, which means the number of computing hours that can be exploited for each single core of the system. In the picture below, the total annual capacity of each system is reported, for a grand- total value of about 3 billion core-h.

Cineca hosts also a Virtual Theater that can accommodate fifteen spectators and is composed of: one BARCO video projector Galaxy NH12; a surround sound system; a cylindrical projection screen of 9.4 x 2.7 meters covering an angle of 120°, made of glass fiber.

It is an integrated immersive three-dimensional display environment with sense of presence perception. By wearing appropriate glasses, viewers can experience virtual reality environments thanks to stereoscopic vision.

Moreover in the same space two Head mounted displays (Oculus DK2) are available for individual immersive navigations.



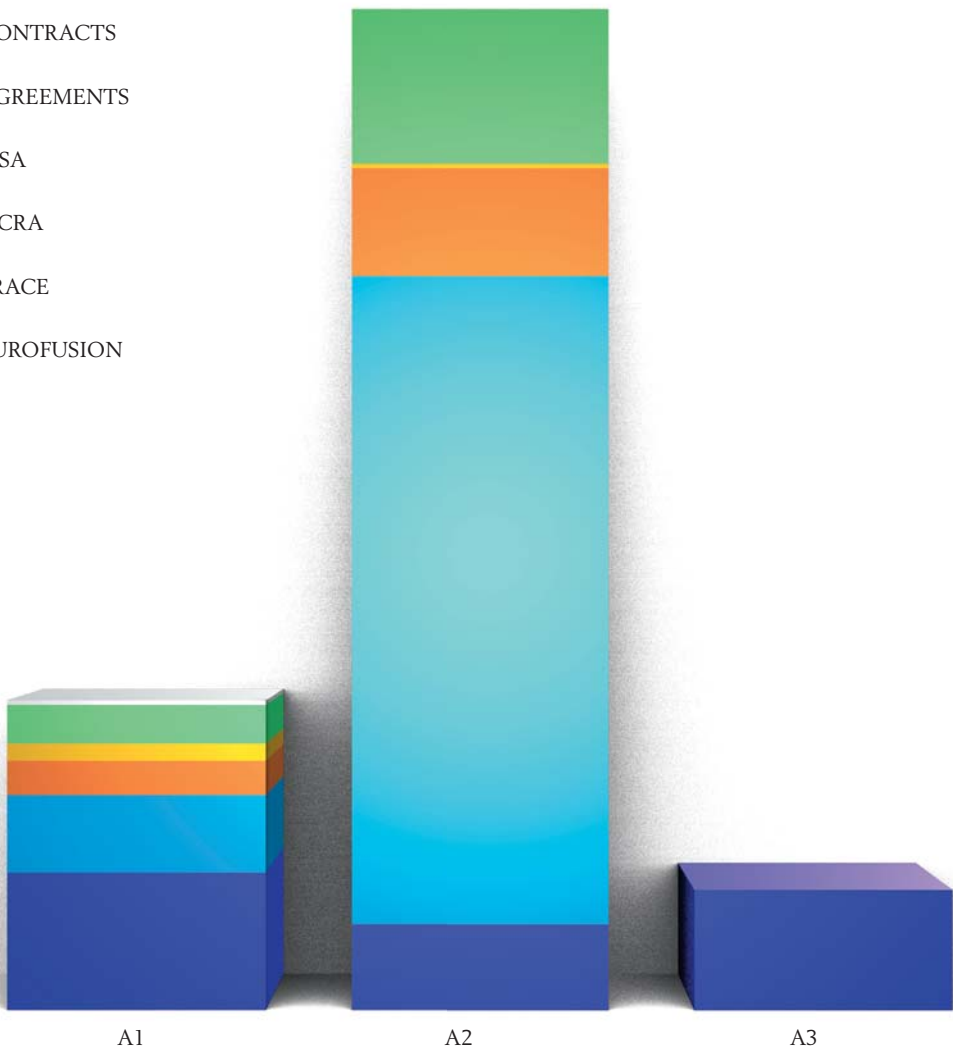



EUROfusion delegates visiting Cineca infrastructure.

# Usage report


In 2017, as shown below, computing resources have been distributed to Italian and European researchers via peer-review based granting programmes (PRACE, ISCRA, LISA). PRACE projects were accounted for 50% of the resource usage, followed by ISCRA (10%). Another important type of projects refers to “Agreements”, special collaborations with the main national research institutions in Italy. “EUROfusion” refers to resources used by the European Consortium for the development of fusion energy; “contracts” to collaborations with Italian companies.

In figure on the right we show how resources were used on Marconi by different application fields: 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). Nevertheless, the most traditional disciplines of Condensed Matter Physics and Astrophysics are still well represented with 27% of usage.




Core-h of Marconi accounted for project categories.

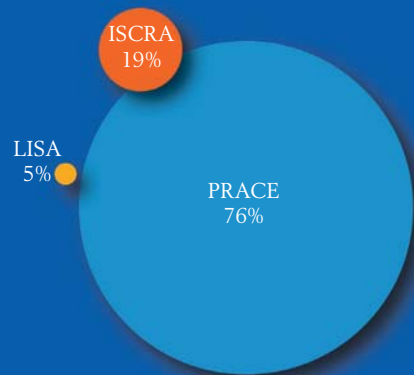




Publications



Number of citations per year



Distribution of core-h on HPC infrastructures via call PRACE, IS CRA and LISA.

Information extracted from the reference database "Web of Science Core Collection" (Font: Web of Science - Thomson Reuters), analyzing the publications using HPC technologies in all the science and engineering fields, citing Cineca.

# Users' statistics


During 2017 we acquired about 1200 new users on our HPC systems, for a total number of about 3.500 active users.

The great majority of the users are males (82%) and working for Italian institutions (66%), mainly Universities or other public research bodies in Italy or abroad (80%).

The large Italian cities, clustering multiple institutions, are well represented: the Milan

area (587 users, also thanks to a dedicated programme funded by the regional authority), followed by Rome (312), Bologna (242), Trieste (198) and Turin (117).


Among the more represented foreign nationalities: Germany and France (thanks to the EUROfusion community), Spain and United Kingdom.




Courtesy of Alessio David et al. - Politecnico Milano  
Polymer viscoelasticity from atomistic molecular dynamics simulations.




Number of Cineca users by gender and nationality.



Nationalities of Cineca users.



Classification of users' institutes.



# Training

Claudia Truini  
Cineca

Training has always been a hallmark of the support activities carried out by Cineca for the Italian research community and users.


Starting from 2012 it was extended to European researchers, Cineca being recognized as a PRACE Advanced Training Center in HPC.

During 2017 we reached quite impressive numbers: 28 courses, distributed across the three sites of the consortium, Bologna, Milan and Rome; 6 schools and 2 workshops. 50 people in the HPC department contributed as teachers, for a total of 120 days of lectures and highly specialized training. In total, over 670 Italian and European researchers took advantage of our training program. Students appreciated these courses and the surveys show high satisfaction levels, reporting an average rating of more than 8/10. Moreover, teaching collaborations have been activated in 3 academic courses (Masters and Doctorate schools), held in three Italian

Universities.

Besides traditional lectures, many other initiatives attracted young promising scientists to Cineca to spend short or extended visits under the mentoring of our experts, who help them to maximize the output of their research or enhance their expertise in HPC techniques.

Also this year we took part in Summer of HPC (SoHPC), a PRACE initiative offering young University students from all over Europe the chance to do internships at major computer centers in the months of July and August. In summer 2017 we hosted two students, and, again this year as in 2016, one of the two students, Arnau Miró Jané, won the PRACE Summer of HPC Best Visualisation Award for his project “Web visualization of the Mediterranean Sea” [<https://www.youtube.com/watch?v=5KGYMsK1JmY>].



## Average ratings:

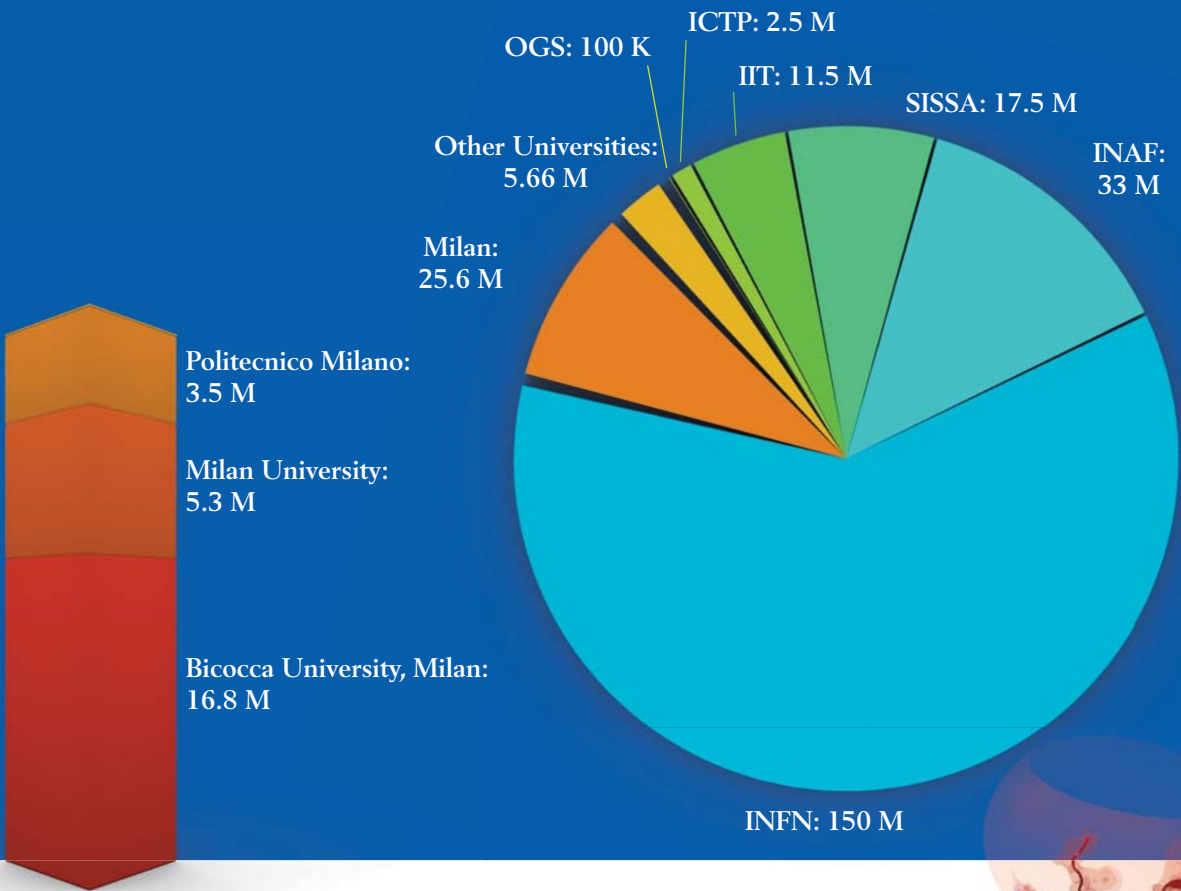
- for courses 8/10
- for teachers 9/10

# Projects in numbers

Roberta Turra, Daniela Galetti and Cinzia Zannoni  
Cineca

Our users get access to Cineca HPC systems by means of projects. Some of them are based on peer-review mechanisms (Prace, ISCRA, Lisa) and are described later in this report. In addition to them we offer other possible agreements, for collaboration and computing power provisioning, with the main Italian research organizations as well as with industrial companies in Italy.

## Agreements with Italian Universities and Research Centers (~ 240 M core-h)



## Other agreements (~ 5 M core-h)

ELIXIR (Distributed infrastructure for life-science information) : 2,6 M core-h

NIG – Network for Italian Genomes : 250 K core-h

Telethon foundation : 2 M core-h



## Collaborations with Industrial Companies

Altran: Engineering and R&D services (Torino)

Amet: Engineering Company (Torino)

Chiesi farmaceutici: Innovation in healthcare (Parma)

Dallara Automobili (Parma)

Elica hoods (Ancona)

Ferretti Yachts (Forli)

Nolan Group: motorbike helmets (Bergamo)

OlsaGroup: Optical Lighting Systems (Torino)

R&D CFD: Computational Fluid Dynamics (Modena)

RED Fluid Dynamics: Research & development in Fluid Dynamics (Rimini)

Unipol Group: insurance (Bologna)



Courtesy of  
RED Fluid Dynamics (Image and simulations)  
Board: Scifi model by Slater Design

## Collaboration with ENI

For over 15 years Cineca has been collaborating with Eni (the most important Italian Energy Company) on the development of applications and on the management of HPC systems for Exploration & Production. 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 2017 Eni replaced HPC1, the first system located in its own Green Data Center (GDC), with a more powerful hybrid cluster (HPC3) managed by Cineca like its predecessors. In June 2017 Cineca was in charge to manage the new Eni PetaByte storage facility in GDC. Meanwhile Eni started a new project with Cineca to make the production software portable on different hybrid architectures and to address, in 2018, the replacement of HPC2 with a new hybrid cluster, foreseen as one of the most powerful industrial HPC clusters in the world.

In 2017 Cineca collaborated with Eni on the evolution and maintenance of production applications, as well as on new Eni research projects developing new applications for data analysis and interpretation.

In the course of 2017, 40 projects were carried out with Eni:

- 4 projects concerned the parallelization, optimization and portability of hybrid codes on different platforms
- 17 projects concerned the development and maintenance of production applications
- 19 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 people 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.

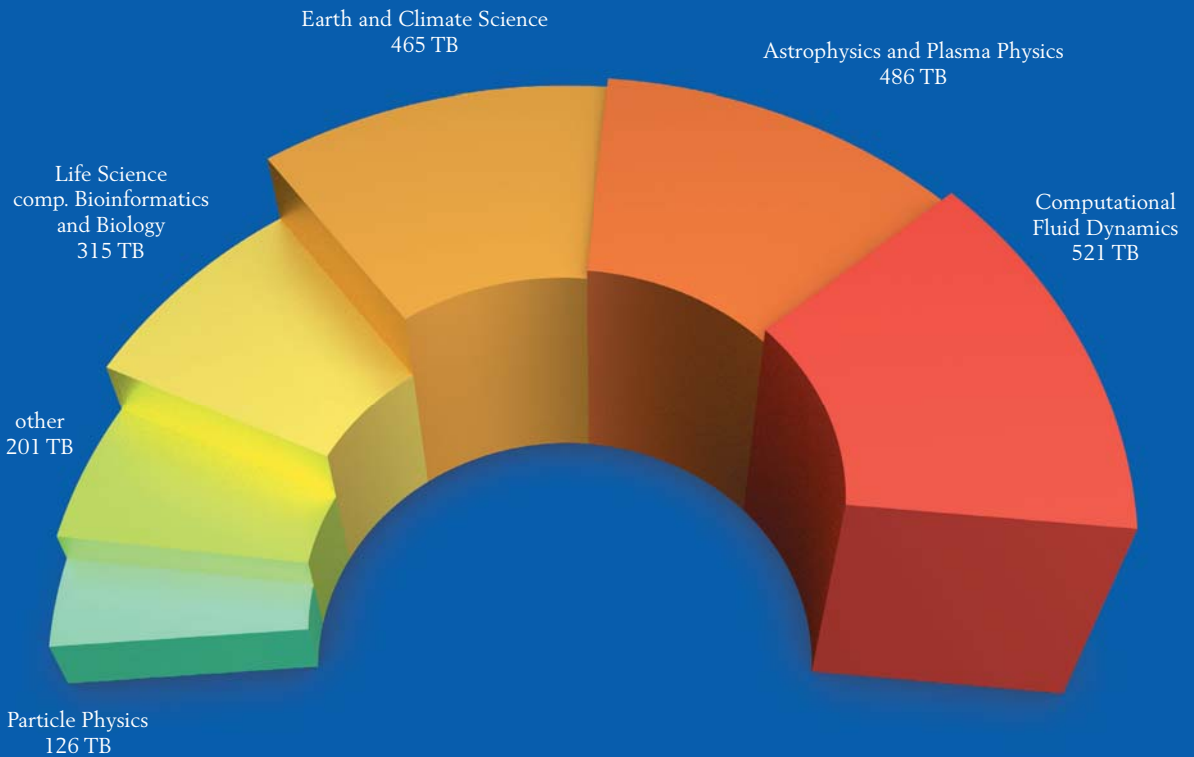


Courtesy of [www.eni.com](http://www.eni.com)

## Data as a resource

From a traditional point of view in HPC applications, data are generated from simulations and are stored in a “archive” for further analysis or post-processing.

Cineca manages a long-term archive where data can be maintained and preserved. In 2017 it was as large as 8.2 PB (raw) and connected to the tape archive in a transparent way via LTFS (Linear Tape File System), an IBM software for multi-level storage.



Storage consumption by scientific domain.

## Big data activity

The term “big data” refers to data sets that are voluminous, complex and rapidly growing and that require the use of predictive, or advanced, analytics in order to extract value from them. The term commonly refers to real data, collected by transaction systems, sensors, social media, IoT, GPS, web apps, scientific instruments (eg NGS, telescopes, seismographs).

While most of the HPC storage infrastructure still hosts high volumes of data generated by simulations, the observational data are increasing in volume and demand specific tools and services, beside the storage service, that have the potential to enable the data-driven science and innovation. EU funded projects, in the big data domain, and ad-hoc research agreements, provide the requirements for developing such tools and services and integrating them with the open source software already available.

Big data projects with HPC staff directly involved:

- Bioinformatics 15 projects
- Industry 4.0 17 projects
- Digital Humanities, Insurance, e-Gov, Media, Energy 6 projects



# Staff

In the HPC department of Cineca work about 86 people, distributed into the three premises: 48 in Bologna, 10 in Milan and 28 in Rome.

Two colleagues have retired in 2017. They were two real pillars of our department, protagonists of the history of supercomputing in Italy and Europe.



**Giovanni Erbacci** started working in Cineca in 1978 and since the beginning he was involved in supercomputing. He was responsible for the HPC Projects Division, and the director of the Cineca's Summer School on Parallel Computing. He actively participates in different EU projects, and his main involvement was on PRACE (Partnership for Advanced Computing in Europe), specifically focused on SHAPE, and on the "Operational services for the HPC ecosystem" activity.

He has been a member of the PRACE Technical Board and has been appointed Italian Delegate in the PRACE Research Infrastructure.

Giovanni holds a degree in Computer Science from the Pisa University.

**Sergio Bernardi** joined Cineca in 1978, a young researcher with a degree in Computer Science from the University of Torino. Sergio has started his activity in Cineca as system analyst and system manager of the top computing system at the time (CDC CY76).



He then joined the small internal R&D group focused on computer graphics and CAD applications.

In the early '90 Sergio got back to system management as head of Systems and Operation department.

In 2007 started his involvement in EU projects. Since then his activity focused on PRACE and he was member of the PRACE Board of Directors from 2010 until 2017.

New colleagues were enrolled in 2017.

**Neva Besker** got a degree in Chemistry at Sapienza University of Rome and PhD in Medicinal Chemistry. She worked in the Theoretical and Computational Physical Chemistry field on classical and ab-initio methods on macromolecules and molecules of biological and pharmaceutical interest. In Cineca she is working in community support.

**Pietro Bonfà** received a degree in Physical Engineering and the PhD in Physics from University of Parma. He has a background in solid state physics and he is actively involved in the development of computational chemistry codes.

**Alessia Marruzzo** has a degree and a PhD in Physics. Her interest was on glassy physics: Molecular Dynamics Simulation, Statistical Mechanics, Spin Glass Models, Inference and Optimization Problems. Alessia joined Cineca last May; she is working on simulations of sedimentary basins in collaborations with ENI geological teams.

**Eric Pascolo** is a member of HPC User Support team. His work is focused on monitoring HPC production and creating innovative solutions to optimize and increase the production. He has a degree in Physics from University of Bologna.

**Riccardo Zanella** has got a Master degree in Telecommunications Engineering and a PhD in numerical analysis. He worked at the University in the field of numerical minimisation of differentiable functions. Now he is in Cineca in charge of studying Deep Learning applications based on open source software on HPC systems.

# Events



## SC17

November 12-17, 2017

Cineca was at booth #771 and presented perspectives of development of HPC in Cineca and the development of the supercomputers Marconi and D.A.V.I.D.E.

SC17 brings together the most respected minds in high performance computing, networking, storage and analysis to debate on research and innovation that will open the door to new scientific and economic opportunities.

September 12 - 15, 2017

The International Conference on Parallel Computing was hosted this year by Cineca in cooperation with the University of Pisa, Dept. of computer Science and ISTI/CNR of Pisa.

The ParCo conferences stimulate the development and application of parallel computers on a world-wide scale.

The conference hosted more than 100 delegates from more than 20 countries from all the 5 continents. 6 invited speakers, and more than 30 parallel sessions presented recent advances and discussed new research topics in the area covered by the conference, which is focused on all aspects related to parallel computing. The invited talks covered themes related to robotics, quantum computing, Internet of Things, climate and weather forecast, HPC computer installation trends as well as new programming models for HPC. As usual in ParCo (which is run every two years in Europe) the delegates had good opportunities to discuss new research results and to interact each other to enforce new cooperations and possibilities for research interactions.

## ParCo2017





# Discover supercomputers at the European Researchers' Night!


September 29, 2017

At Cineca booth it was possible to meet our researchers, understand why parallelism is the basis of supercomputing, how complicated it is to design and use a supercomputer and see how useful it is in many fields: science, industry and even cultural heritage.

Children, but also adults, amused themselves thanks to games, smart glasses, mobile applications, so to figure out how supercomputing can help us in our everyday life and why it is our friend.

<http://nottedeiricercatori-society.eu/>

<https://www.frascatiscienza.it/pagine/notte-europea-dei-ricercatori-2017/>



## Summer of HPC

July and August 2017

Summer of HPC is a PRACE (Partnership for Advanced Computing in Europe) programme that offers summer placements at top HPC centers across Europe, in a PRACE partner country. This programme offer to late stage undergraduates and early stage postgraduate students the opportunity to spend two months of the summer at a HPC center. Twenty students go to ten PRACE member centers around Europe, where they work on PRACE related projects for 2 months during the summer.



Cineca in 2017 hosted two students:

Arnau Mirò worked on Viewing the Mediterranean Sea project;

Petr Stehlik worked on Web visualization of Energy load of an HPC system project.



# NEW PROJECTS

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 centers in Europe to help our scientific communities to leverage the most advanced available technologies.

# HPC-Europa3

Debora Testi  
Cineca

HPC-Europa3 is a multi-disciplinary programme for short collaborative research visits with transnational access to some of Europe's biggest supercomputing resources.

In the last years numerical simulation methods have come to full prominence and are at the base of new gains in knowledge and progress in science. These methods are essential for understanding the theoretical and experimental data related to the fundamental laws of physical phenomena. At the same time, the panorama of High Performance Computing (HPC) in Europe has reached a remarkable level of consistency, efficiency and harmony, and it could reach further levels of completeness and cohesion with the Transnational Access model that is offered by the HPC-Europa initiative. Thanks to HPC-Europa3, researchers are supported to become completely self-sufficient users of the HPC facilities included in the pan-European initiative for supercomputing, leading to maximum exploitation of the experience gained and of the scientific network created through the programme.

HPC-Europa3 is an EC-funded partnership of 9 leading HPC centers and 1 center of excellence which aims to implement a programme of Transnational Access visits, giving European-based researchers the opportunity to spend a period of time at one of the HPC centers project partners, collaborating with one of the associated scientific research teams, and to benefit the use of some of the most relevant HPC research infrastructures in Europe complemented by strong scientific and technical on-site support. Visitors can work closely with a "host" research group working in a similar field of research: many research groups are already associated with the programme as "host" research groups, but new ones can join at any time; SMEs are welcome to participate too.

Applicants can work in any discipline, but must require the use of HPC resources for their research project. HPC-Europa3 covers travel and living costs, and visits can last between 3 weeks and 3 months. The programme is open to researchers of all levels, from postgraduate students to senior professors.

Following EC guidelines, applications are encouraged from the new EU member countries and those who do not have access to similar computing facilities at their home institute. Under the new "Regional Access" sub-programme, applicants in the Baltic states and South East Europe who have little or no HPC experience are particularly encouraged to visit host departments in Sweden and Greece respectively.

Calls are open all the time but evaluation takes places 4 times per year with dates reported on the project website.

In parallel, HPC-Europa3 Networking Activities promote co-operation and co-ordination between the participating infrastructure and other European initiatives HPC-Europa3 Joint Research Activities address technological issues that will improve the overall performance of these research infrastructures.

In brief, HPC-Europa3 Transnational Access programme goals are:

- access to world-class HPC systems
- scientific collaboration with host researchers in any field
- technical support by the HPC centers
- travel and living expenses fully reimbursed
- more than 1200 visitors supported in the next four years
- about 100 million CPU hours offered

For further information and contact details please see our website: <https://www.hpc-europa.eu>  
HPC-Europa3 receives funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No.730897.

Infrastructure on High Performance Computing

Pan-European Research

# HPC *Europa*3

HPC-Europa3 receives funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No.730897



**HPC-Europa3** Transnational Access is a multi-disciplinary programme for short collaborative research visits with access to some of Europe's biggest supercomputing resources.

Courtesy of  
Ivan Marri (University of Modena and Reggio Emilia, Department of Sciences and Methods for Engineering and CNR/nano) and Marco Govoni (The University of Chicago, Institute for Molecular Engineering)  
Electronic Coupling Effects in Silicon Nanocrystals: Ab-Initio Results

**HPC-Europa3** offers access to first-class HPC platforms for researchers throughout the whole of Europe in all domains of computational science, providing advanced computational environments, technical support and training.

The programme runs until 2021  
Applications may be submitted at any time

Please submit your application at  
<http://www.hpc-europa.eu/>



# OPRECOMP

Andrew Emerson  
Cineca

Guaranteed numerical precision of each elementary step in a complex computation has been the mainstay of traditional computing systems for many years. This era, fuelled by Moore's law and the constant exponential improvement in computing efficiency, is at its twilight: from tiny nodes of the Internet-of-Things, to large HPC computing centers, sub-picojoule/operation energy efficiency is essential for practical realisations. To overcome the "power wall", a shift from traditional computing paradigms is now mandatory.

OPRECOMP investigates the theoretical and practical understanding of the energy efficiency boost obtainable when accuracy requirements on data being processed, stored and communicated, can be lifted for intermediate calculations. While approximate computing approaches have been used before, in OPRECOMP for the first time ever, a complete framework for transprecision computing, covering devices, circuits, software tools, and algorithms, along with the mathematical theory and physical foundations of the ideas will be developed that not only will provide error bounds with respect to full precision results, but also will enable

major energy efficiency improvements even when there is no freedom to relax end-to-end application quality-of-results.

The mission of OPRECOMP is to demonstrate using physical demonstrators that this idea holds in a huge range of application scenarios in the domains of IoT, Big Data Analytics, Deep Learning, and HPC simulations: from the sub-milliwatt to the megawatt range, spanning nine orders of magnitude. In view of industrial exploitation, we will prove the quality and reliability and demonstrate that transprecision computing is the way to think about future systems.

OPRECOMP is a 4-year research project funded under the EU Framework Horizon 2020 - Future and Emerging Technologies (FET), started in January 2017. It involves the following organisations: IBM Research Zurich (Project co-ordinator, Switzerland), ETH Zurich (Switzerland), CEA (France), University of Perugia (Italy), University of Bologna (Italy), Cineca (Italy), University Jaume (Spain), Queen's University (UK), GREENWAVES TECHNOLOGIES (France) and Technical University Kaiserslautern (Germany).

For further information and contact details please see our website: <http://oprecomp.eu/>





# European Open Science Cloud

Claudio Cacciari  
Cineca

The European Open Science Cloud (EOSC-hub) project creates the integration and management system of the future European Open Science Cloud that delivers a catalogue of services, software and data from the EGI Federation, EUDAT CDI, INDIGO-DataCloud and major research e-infrastructures. This integration and management system (the Hub) builds on mature processes, policies and tools from the leading European federated e-Infrastructures to cover the whole life-cycle of services, from planning to delivery. The Hub aggregates services from local, regional and national e-Infrastructures in Europe, Africa, Asia, Canada and South America.

The Hub acts as a single contact point for researchers and innovators to discover, access, use and reuse a broad spectrum of resources for advanced data-driven research. Through the virtual access mechanism, more scientific communities and users have access to services supporting their scientific discovery and collaboration across disciplinary and geographical boundaries.

The project also fosters skills and knowledge among researchers and service operators by delivering specialised trainings and by

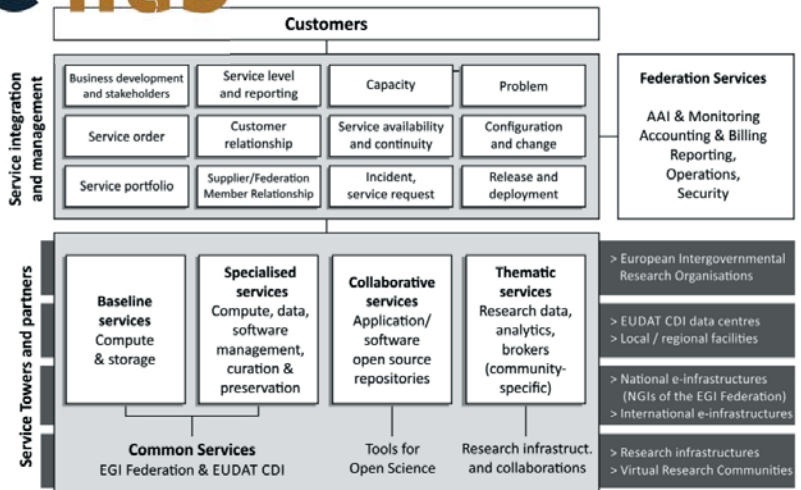
establishing competence centers to co-create solutions with the users. In the area of engagement with the private sector, the project creates a Joint Digital Innovation Hub that stimulates an ecosystem of industry/SMEs, service providers and researchers to support business pilots, market take-up and commercial boost strategies. EOSC-hub builds on existing technology already mature and addresses the need for interoperability by promoting the adoption of open standards and protocols. By mobilizing e-Infrastructures comprising more than 300 data centers worldwide and 18 pan-European infrastructures, this project is a groundbreaking milestone for the implementation of the European Open Science Cloud.

Cineca joined the EOSC-hub project as partner of the EUDAT2020's infrastructure and has an important role in the project, it is member of the Project Management Board and leader of Work Package about "Thematic Services".

This is a three years project started late 2017. It represents an important opportunity for Cineca to gain knowledge about new technologies, to consolidate relationships with international partners and to establish new links with research communities.



## EOSC-hub



Conceptual overview of the EOSC-hub project.

# MaX Center of Excellence

Carlo Cavazzoni and Pietro Bonfà  
Cineca

MaX - Materials design at the Exascale is a European center of excellence funded under the FET-HPC (Future Emerging Technologies in HPC) program in the H2020 workplan.

MaX works at the frontiers of the current and future High Performance Computing (HPC) technologies, to enable the best use and evolution of HPC for materials research and innovation.

MaX is creating an ecosystem of capabilities, ambitious applications, data workflows and analysis, and user-oriented services. At the same time, MaX enables the exascale transition in the materials domain, by developing advanced programming models, novel algorithms, domain-specific libraries, in-memory data management, software/hardware co-design and technology-transfer actions.

At the core of MaX there are four flagship codes, widely-used applications, based on rather diverse models, mainly oriented to structural, electronic, magnetic properties and to spectroscopies of materials from first principles:



<http://www.quantum-espresso.org/>



<http://departments.icmab.es/leem/siesta/>



<http://www.flapw.de/pm/index.php>



<http://www.yambo-code.org/>

The four codes are complemented by a workflow management system:



<http://www.aiida.net/>

AiiDA allow the implementation of specific workflow for material science application, that can be stored and reused many times for different materials. The workflows can have different degree of complexity, from a single application, to a complex graph with many steps and forks. AiiDA also allows storing and sharing data.

MaX is working on these flagship codes on the present HPC platforms, by implementing new capabilities and algorithms for the study of complex materials, properties and processes in realistic condition, far beyond the current realms.

At the same time, MaX is radically enhancing the performance of the flagship codes in terms of scaling, robustness, and usability, and will make them ready for the forthcoming exascale hardware architectures.

In this way, MaX is designing and implementing a sustainable approach ready for further use beyond its core codes and field. A detailed Flagship codes description is given at <http://www.max-centre.eu/flagship-codes/flagship-codes-2/>

For further information and contact details please see our website:  
<http://www.max-centre.eu>

## Natural Food Colors simulated by Quantum Espresso



Ab-initio methods are used by food-industry research to discover natural replacements for artificial colourants. One of MaX industrial pilots combines classical and quantum approaches to design molecules with desired color and improved stability.

See more:


<http://www.max-centre.eu/pilot-case-3/>

## Advanced capabilities for materials modelling with Quantum ESPRESSO

A recent paper reports advances in the MaX flagship code Quantum ESPRESSO. It has been published on "Journal of Physics: Condensed Matter". In the manuscript, "Advanced capabilities for materials modelling with Quantum ESPRESSO", authors discuss recent extensions and improvements of the code, covering new methodologies and property calculators, improved parallelization, code modularization, and extended interoperability both within the distribution and with external software.


See more:

<http://www.max-centre.eu/2017/10/25/qe-new-paper-on-quantum-espresso-advances/>



## A new scalability record in a materials science application

Another step towards the exascale was made by MaX researchers at CNR who have run a multi petaFlop simulation with the MaX flagship application Yambo. A single run reached 3 petaFlop/s on 1000 nodes (68000 cores) on the KNL partition of Marconi@Cineca within a Prace project. The real-life GW calculation, performed on a realistic polymer used as chemical precursors of chevron-shaped graphene nanoribbons, is part of an effort on computational spectroscopies in collaboration with experimental teams in Austria, Italy and Switzerland.



<http://www.max-centre.eu/2017/04/19/a-new-scalability-record-in-a-materials-science-application/>

# PPI4HPC and ICEI

Carlo Cavazzoni and Giuseppe Fiameni  
Cineca

The Cineca HPC infrastructure will undergo a large upgrade in 2019 - 2020 leveraging different EU initiatives. According to the press release of January 11th 2017, the European Commission is ready to invest EUR 1 billion jointly with the Member States in building a world-class European supercomputers infrastructure.

*"Supercomputers are considered fundamental to process large amounts of data and bring benefits to the society in many areas from health care and renewable energy to car safety and cybersecurity."*

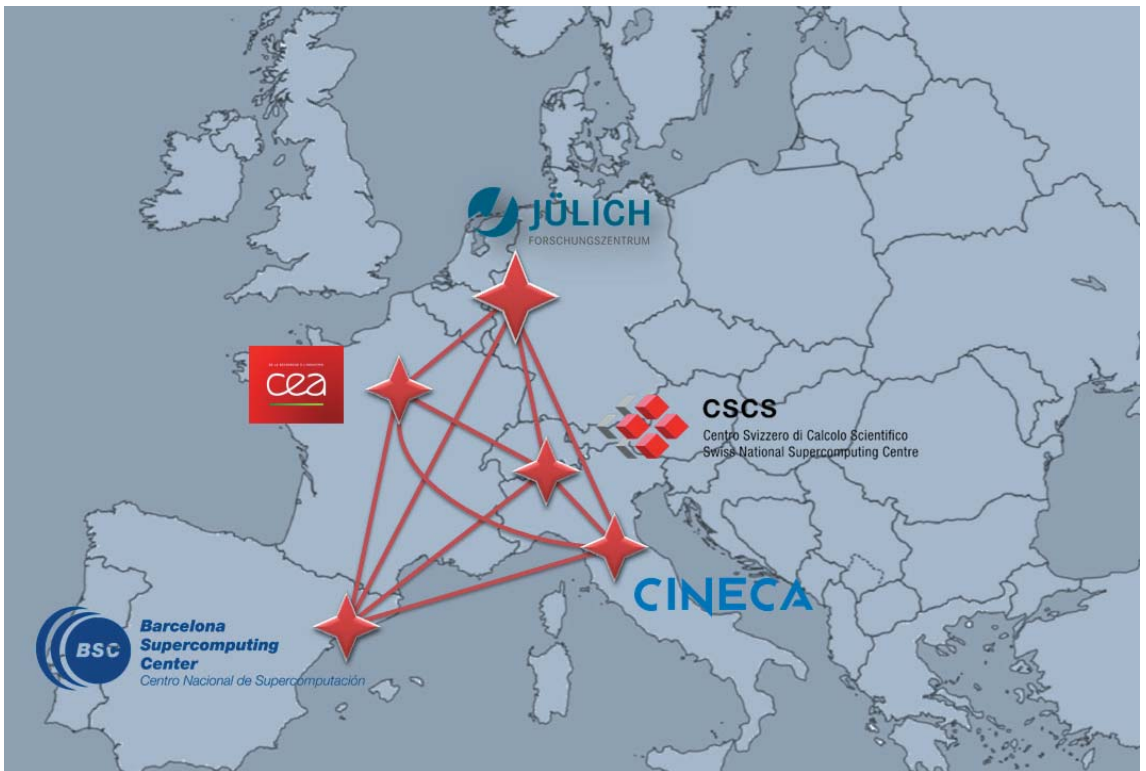
Cineca is involved in two major initiatives supported by EU in the field of HPC. PPI4HPC (Public Procurement of Innovative Solutions for HPC - [www.ppi4hpc.eu](http://www.ppi4hpc.eu)) is a

project of leading European supercomputing centers, including Cineca, that aims to execute a joint public procurement for the first time in the area of High Performance Computing. The co-funding by the European Commission will allow for a significant enhancement of the planned pre-exascale HPC infrastructure from 2019 and pave the path for future joint investments in Europe. The total investment for all the centers is planned to be about € 73 million and it will bring in Cineca a new system of approximately 30PF/s of peak performance. The participants will work together on coordinated roadmaps for providing HPC resources optimised to the needs of European scientists and engineers.



One to one meeting with vendors. Barcelona, 4-6 October 2017.





The second initiative is ICEI (Interactive Computing e-Infrastructure) with the aim of providing stable computational resources to the Human Brain Project, the neuroscience community and beyond.

The final goal is to realise the key elements of the planned infrastructure through the procurement of equipment and R&D services. To this purpose, a coordinated procurement process is foreseen among the partners. Furthermore, significant additional parts of the infrastructure and R&D services will be realised through in-kind contributions from the participating supercomputing centers.

As part of this collaboration, Cineca will deliver a system around 500 nodes equipped with new memory technology (i.e. NVRAM) to support interactive scientific computing. Interactive computing refers to the capability of a system to support distributed computing workloads while permitting on-the-fly interruption by the user.

The real-time interaction of a user with a program runtime is motivated by various factors, such as the need to estimate the state of a program or its future tendency, to access intermediate results, and to steer the computation by modifying input parameters or boundary conditions.

Within neuroscience applications (i.e. brain activity simulation, large image volume rendering and visualization, connectomics experiments), the runtime can be modified interactively so that the user can gain insight on parameters, algorithmic behaviour, and optimization potentials.



# EMERGING TECHNOLOGIES

Cineca's mission is to support scientific research by providing access to advanced computational systems. Technologies like Deep Learning, Non-volatile Memory, Quantum Computing, High Performance Converged Network 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.

# HPC infrastructure evolution

Carlo Cavazzoni  
Cineca

## Finalization of Marconi as a convergent infrastructure

During 2017 we completed the configuration of Marconi with the upgrade of the SkyLake partition. Now Marconi is a system with a total computational power of about 20 Pflop/s, made of a conventional partition (latency cores) and an accelerated partition (Intel KnightLanding cores).

For the near future, we plan a reworking of the HPC infrastructure, where the other two clusters, Galileo and PICO supporting Tier-1 and BigData workloads, will be integrated in a “convergent” architecture within Marconi itself. The same hardware components will address different needs: Tier-0 highly parallel (scale out) applications, Tier-1 and industrial workloads, visualisation workloads and cloud services.

## D.A.V.I.D.E. heterogeneous system

D.A.V.I.D.E. (Development for an Added Value Infrastructure Designed in Europe) is an innovative system co-funded by the PRACE-3IP (PCP initiative), designed to evaluate possible solutions for the next-cycle exascale system. The same architecture has been chosen by the Department of Energy in US for the new HPC flagship system that should grant

US a new leadership in the HPC. For Cineca users it will represent a great opportunity to become acquainted with the architecture, to adapt the codes and keep a leadership in HPC competitiveness. D.A.V.I.D.E. represents one of the largest testbed architecture for Artificial Intelligence available in Europe.

D.A.V.I.D.E. has been designed by an Italian company, E4 Computer Engineering, in collaboration with IBM and the University of Bologna. It features an innovative monitoring system (conceived by the team of Luca Benini) that can be used to optimize the workload and maximize energy efficiency. It will be available for production early in 2018.

Logical Name	DAVIDE
Model	E4 Open RACK
Architecture	OpenPower with NVlink and NVIDIA GPUs
Processor	8 core Power 8+
Accelerator	NVIDIA P100
Number of nodes	45
Processors and Accelerator per node	2 processors and 4 accelerators
Memory per node	256GByte
Peak Performance	1PFlop/s

### Tier-0 MARCONI

Conventional partition:

- A1 (720 BDW nodes)
- A3 (2300 SKL nodes)

Accelerated partition

- A2 (3600 KNL nodes)

### Tier-1 Cloud/HT

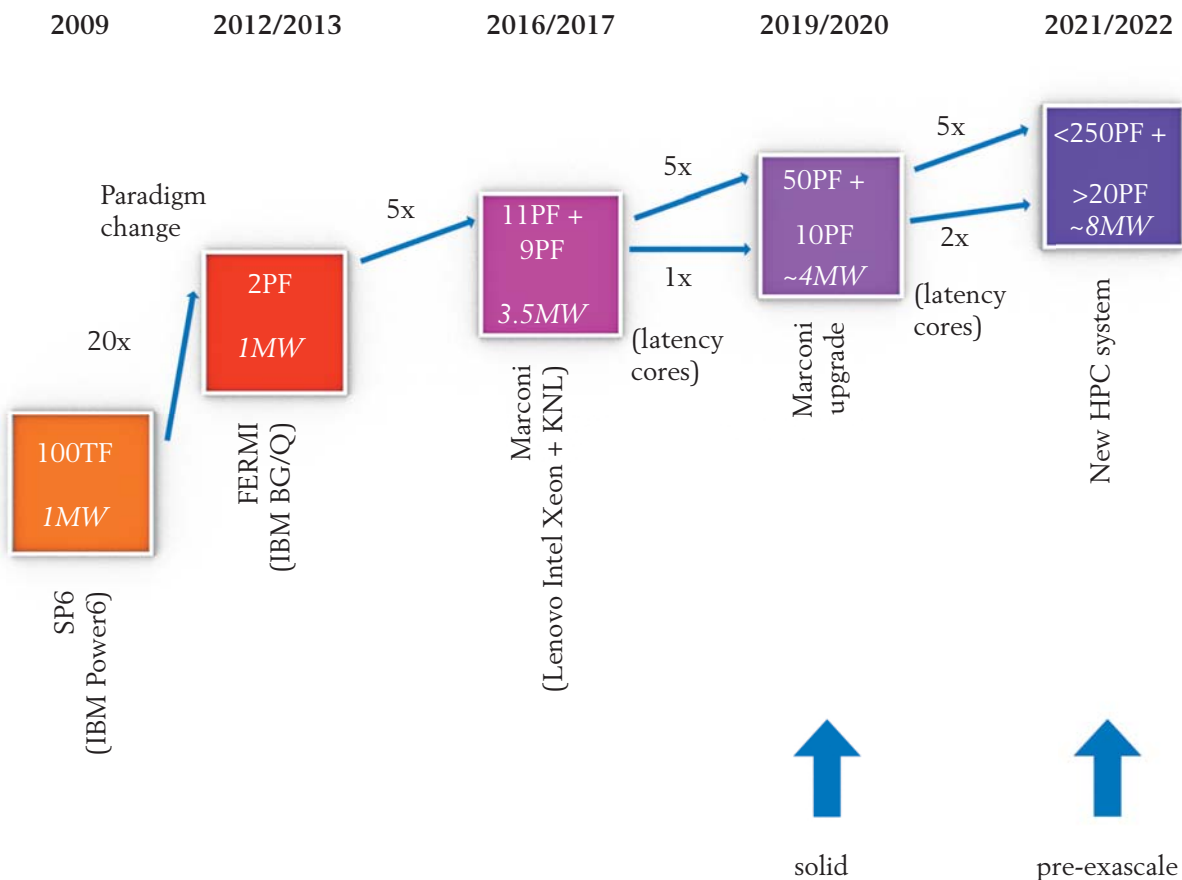
792 BDW nodes

- Cloud HT for INFN
- Cloud HPC
- Tier-1 HPC (newGALILEO)

Lenovo GSS + IBM Flash  
> 30 PByte

The “convergent” configuration of Marconi.





### Medium-term evolution of the infrastructure

The installation of Marconi is part of a more global plan in two steps, approved by the Cineca governance in 2016, with an expected duration of 5 years. Now the first step is completed and we are approaching the second one, planned in 2019/2020.

Looking at the evolution of Cineca infrastructure since 2009, when IBM SP6 was our top system, we see an increase of about 200 times in computational power with a corresponding increase of only 3.5 in power consumption.

It has to be noted that Marconi is made of two partitions, a conventional one for high productivity workload (9PFlop/s) and a second one (accelerated) for scale-out workload (11PFlop/s). This characterization will be

maintained in the next generation, which could include a new increased partition, based on special architectures/processors, for scale-out workload, and a partition based on more conventional servers (latency cores). This second partition, more suitable for high productivity and cloud-based workflows, will probably grow at lower pace with respect to the scale-out partition, since it is limited by the power consumption. In fact if we consider the last step of our current road-map, where a pre-exascale system is expected, we estimate that the scale-out partition will be five times more powerful than the previous generation. It will be most probably based on specific hardware for low-power/high-efficiency HPC architecture. On the contrary, the conventional partition will be only two times faster than the previous generation.

# Trends on energy-efficient supercomputer management

**Andrea Bartolini**

**Department of Electrical, Electronic, and Information Engineering “Guglielmo Marconi” (DEI) – University of Bologna**

Future exascale computing systems will feature thousands of nodes each embodying hundreds of cores, deep memory hierarchies and have a wide range of parallel accelerators, as well as a complex inter-node communication system. In addition, as the effect of technology scaling, the effect of hardware variability will become visible at the macroscale, making homogeneous nodes heterogeneous in practice.

This effect is already visible in today's systems and already cause large efforts to system administrators and supercomputer owners to guarantee efficiency and fairness across all the users and at the datacenter level<sup>[1][2][3]</sup>. Moreover, the increased number of components inevitably grows the hardware failure risk putting serious concern about exascale system reliability and MTTF (Mean Time To Failure)<sup>[4]</sup>.

With this increasing complexity in terms of computing resources, infrastructure and software components, it is getting harder to operate supercomputing systems at the most efficient point. Performance is not the only concern in today's supercomputers design and operation, but also energy and power consumption as well as resiliency.

Indeed, as an effect of the quest towards the exascale supercomputer the main actors (datacentre owner, system administrators and final users) are becoming obsessed by the scalability of applications, effective usage of available resources as well as control of the energy and power consumed by the applications and by the cooling infrastructure. All these parameters need to be checked in real-time.

Always on monitoring of key performance and cost parameters of the datacenter, supercomputer, application and users can help in creating awareness around these keys factors. However, being capable of connecting these parameters with Big Data and Artificial Intelligence technologies can cause the quantum leap towards automated approaches of the supercomputer and datacenter administration,

maintenance and optimization.

Cineca is pioneering in this field. Since several years, a strategic research agenda has been signed with researchers at the Energy-Efficient Embedded System Laboratory at the department of Electrical and Information Engineering “Guglielmo Marconi” (DEI) of the University of Bologna to research and prototype innovative solutions in this domain.

During 2017 an integrated monitoring framework has been designed and deployed in Cineca. The framework is composed by a scalable and extensible software infrastructure and leverages big data and internet of things technologies.



Courtesy of E4.



Web Service Front End snapshot. The web service designed for EXAMON (EXAMON WEB) enables users to query energy and performance metrics based on their job\_ID and system administrators to inspect cluster behavior.

The framework was deployed on the Galileo supercomputer collecting more than 440 performance metrics every 2 seconds from each compute node as well as tens of IoT (Internet of Things) nodes dislocated in the compute room for monitoring the thermal evolution and the cooling subsystem.

During the same period a web service has been designed for easy visualization of the collected data, by aggregating it per job and cluster. EXAMON WEB (EXAscale MONitoring on the Web) is the first attempt to provide to the different actors awareness of supercomputer's efficiency.

Finally, a new prototype system has been co-designed for energy-efficiency. D.A.V.I.D.E. (Development for an Added Value Infrastructure Designed in Europe), is an innovative High Performance Computing cluster designed by E4 Computer Engineering for PRACE (Partnership for Advanced Computing in Europe).

D.A.V.I.D.E. on top of best-in-class components HW components features custom hardware and an innovative system middleware software designed for fine grain power and performance monitoring and cluster power capping.

**References**

- [1] Ashby, S. et al. (2010), [http://science.energy.gov/media/ascr/ascac/pdf/reports/Exascale\\_subcommittee\\_report.pdf](http://science.energy.gov/media/ascr/ascac/pdf/reports/Exascale_subcommittee_report.pdf)
- [2] Dongarra, J., et al. Int. J. High Perform. Comput. ppl. (2011), <http://hpc.sagepub.com/content/early/2011/01/04/1094342010391989>
- [3] Shalf, J., et al. International Conference on High Performance Computing for Computational Science (2010)
- [4] Cappello, F., Int. J. High Perform. Comput. Appl. (2009), <http://hpc.sagepub.com/content/23/3/212>

# Deep Learning

Riccardo Zanella and Giuseppe Fiameni  
Cineca

The world of Data Science is in continuous transformation both on the technological and machine learning side and the demand for predictive analysis is growing ever more. With the data doubling every year, data intensive applications are increasing as well as the demand of high-end resource capacity to analyse collected data sets. The explosion of analysis applications has become a major driver for revising system architecture and tools leading to the proliferation of software components and frameworks which may require multi-node and multi-core systems to scale-up and provide high performance. In this context, Machine learning and Deep learning are steadily proving to be successful methods for a variety of use cases, and their popularity has resulted in numerous open-source software tools becoming accessible to the public and popular across different scientific disciplines.

Cineca is actively supporting users to achieve state of the art efficiency on Deep Learning tasks in HPC systems. During this year, we carried out a preliminary comparative study of state-of-the-art deep learning software tools and benchmark them on Cineca HPC facilities providing insights that would be useful for ISCRA (Italian SuperComputing Resource Allocation) users.

To ensure consistency and reproducibility of benchmark results, the comparison is based either on convnet-benchmarks<sup>[1]</sup> open source project or on test code directly available on frameworks source. We considered both training and inference phases, and we granted the access to HPC resources in Cineca (Nvidia GPUs or Intel Xeon accelerators), by exploiting latest optimizations implemented on current versions of TensorFlow, BVLC Caffe, Intel® Caffe, MILA Theano, Intel® Theano, Intel® Nervana™ Neon. Results are easily summarized as: TensorFlow exhibits comparable results both on Intel® Knights Landing and Nvidia

K40 GPU, allowing uniform user experience in Marconi A2 and Galileo infrastructures.

We were able to measure Nvidia P100 GPU performance efficiency on Deep Learning tasks, especially in conjunction with Intel® Nervana™ Neon: this certifies Intel® efforts on the development of a Deep Learning tool capable to reach top performances on all hardware. Results have been presented at the PARCO2017 Conference and during an all-day workshop co-organized with Intel® Nervana™ AI Academy.

The workshop comprised a number of relevant scientific use cases, a description of optimization strategies suggested by Intel® and implemented both in TensorFlow and Caffe, and an introduction to Intel® Nervana™ Neon and Ngraph.

Stefano Cozzini and Rossella Aversa (CNR-IOM) gave an overview of their recently proposed method, based on modern Deep Learning techniques, capable to classify Scanning Electron Microscope (SEM) images at the nanoscale.


Ludovico Silvestri (LENS) and Matteo Roffilli (Bioretics) underlined the importance of Deep Learning methods on image analysis in the context of the Human Brain Project by exposing their experience in analyzing extremely large (TeraByte) datasets containing human and mouse brain information produced within the European Flagship Human Brain Project.

Rita Cucchiara and Lorenzo Baraldi (AImageLab UNIMORE) described results achieved in automatic video analysis (semantic annotation of content, tagging, similarity search, saliency detection and textual captioning) with specifically designed convolutional neural networks and recurrent networks such as LSTM, trained over very large datasets of broadcast video in sport, movies and culture.

#### References


[1] <https://github.com/soumith/convnet-benchmarks>

LENS in the HBP: strategic micron-scale resolution mapping of the whole brain.



Whole mouse brain cell census.

Whole mouse brain vasculature.



Human brain cell census.




Image courtesy of Ludovico Silvestri, European Laboratory for Non-linear Spectroscopy (LENS).

A view of the audience attending the Cineca-INTEL workshop.



# Big Data

Roberta Turra  
Cineca

With a public announcement at the EBDVF in Versailles, on the 21.11.2017 Cineca Big Data Platform has been granted the i-Space label by the Big Data Value Association (BDVA), the private counterpart to the EU Commission for the implementation of the Big Data programme. Data innovation spaces (i-Spaces) are trusted platforms targeted to accelerate take up of data driven innovation in commercial sectors like Manufacturing 4.0, Logistics, e-Commerce, Media, Aerospace, Automobile, Energy, Agriculture, Pharmacy, as well as in non-profit sectors (e-Government, Environment, Public Health, Smart Cities). They host closed as well as open data from business and public sources for the purpose of the execution of pre-competitive Proof-of-Concept projects and offer trusted and secure environment allowing Research, Education and Innovation stakeholder to innovate with data and acting as hubs to connect different stakeholders.

With their infrastructure, tools and skill offerings, i-Spaces help the surrounding (industrial and public) ecosystem to overcome high-risk investment burdens for faster adoption of what new BigData technologies can offer to them.

Excellence of the infrastructure, quality of services, evidence of projects already being supported in different sectors and of an existing vibrant ecosystem together with a clear growth and business strategy are the key factors for achieving the i-Space label and each candidate undergoes an evaluation process in order to assess its level of fulfillment.

Cineca gained the silver label and this achievement has been possible thanks to the ever growing number of big data projects, the


flexibility of the big data platform and the skills and competences that enable high quality services, most notably the training and user support services.

Regarding the big data projects, the main domains supported by the platform are Bioinformatics and IND4.0. In bioinformatics, Cineca has been directly involved in 15 big data projects in 2017 (while providing the infrastructure to another 105 IS CRA projects). In the IND4.0 domain, Cineca collaborated in 2017 to 17 projects. Other domains are Digital Humanities, Insurance, e-Gov, Media, Energy, globally accounting for another 6 big data projects in 2017. These projects are mostly EU funded and provide the requirements for developing our big data platform, tools and services. Thanks to that, the Cineca big data platform has high flexibility and accommodates different kinds of users with different usage models (from GPU intensive workloads on the HPC infrastructure in batch mode, to the more interactive containerized usage of the resources, to the platform as a service available on cloud computing). Different options are also available for data storage: from the data lake (unstructured data), to data base with metadata management and to repositories (structured data).

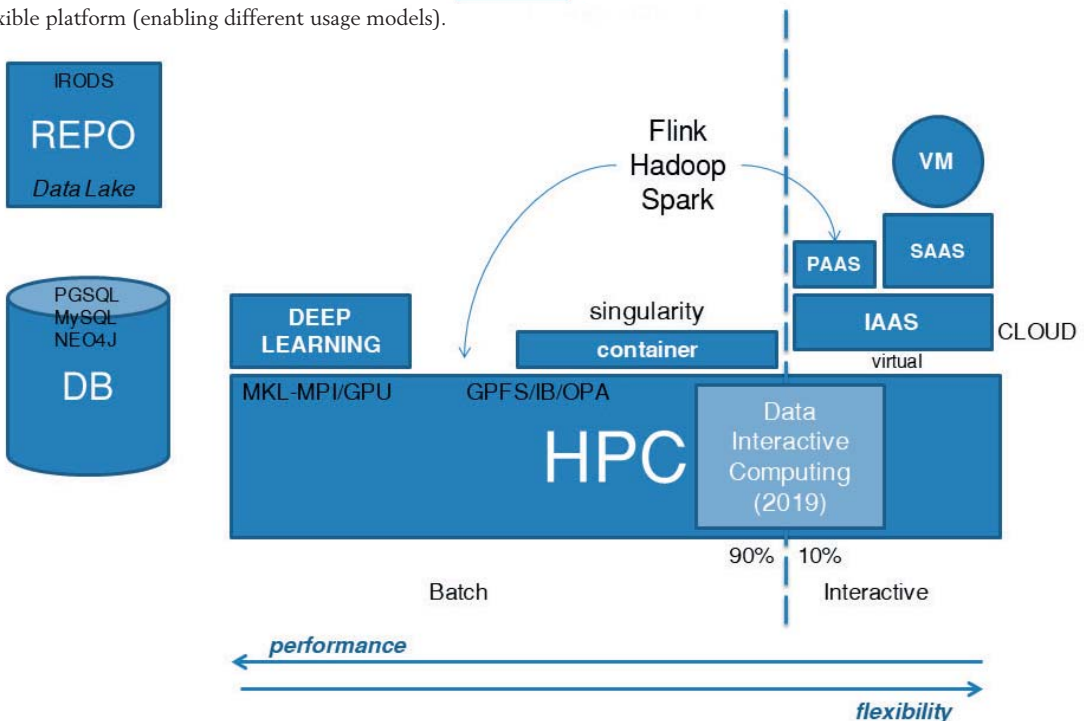
What makes this platform different from all the other big data platforms (and the other i-Spaces) is the high computational power which is available and can be used to reduce the time-to-market, or time-to-science, of deep learning solutions and the possibility to couple real data and virtual data, to improve simulation models.



BDVA labelled i-Spaces.



A flexible platform (enabling different usage models).





# USERS' PROJECTS



Three different programmes are available to offer computational resources to researches, based on a peer-review access: Prace, at the European level; Iscra, at the national level; Lisa, funded by Regione Lombardia and restricted to researchers working in the area.

Moreover we offer other possibilities to get access to HPC resources, specifically to Italian researchers.

In the following, some exemplary projects granted in 2017 are presented.

# Astrophysics and Plasma Physics

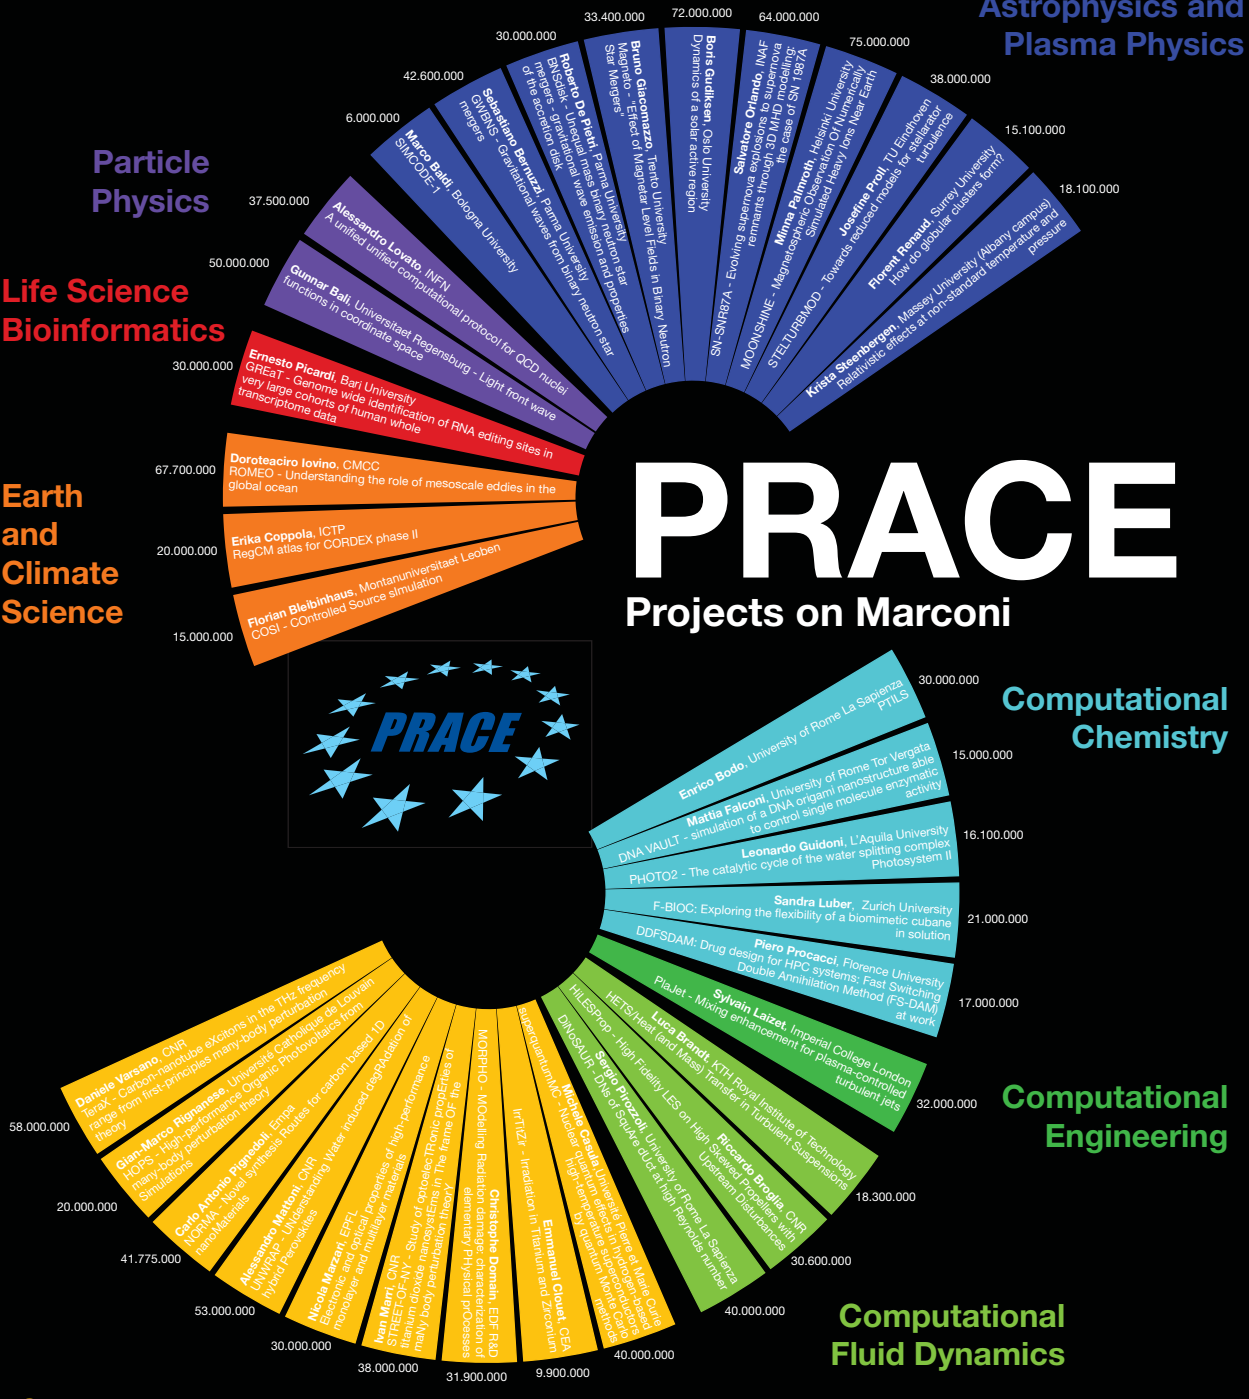
# Particle Physics

# Life Science Bioinformatics

# Earth and Climate Science

# PRACE

## Projects on Marconi



# Computational Chemistry

# Computational Engineering

# Computational Fluid Dynamics

# Condensed Matter Physics

# PRACE

## Italian PI

## Computational Chemistry

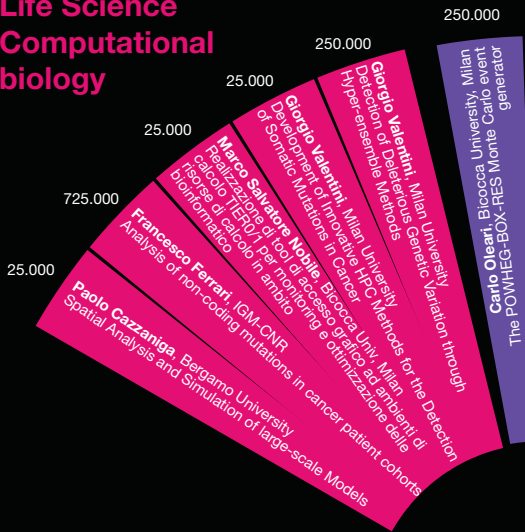
## Earth and Climate Science

## Life Science Bioinformatics

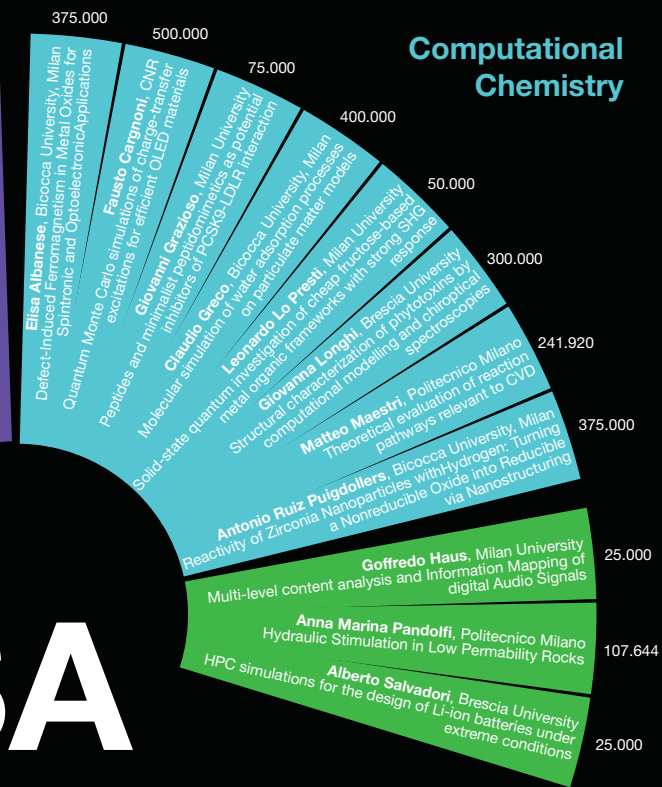
Pino D'Amico, CNR - SIDEX - Size Dependence of the Excitonic properties of MoS2 nanoribbons from many-body perturbation theory	Piz Daint 59,000,000
Maddalena Damore, Turin University - NANOCAT-ZN Modeling mechanism of reactions and active sites in Ziegler-Natta nano-systems, looking for reactivity indicators in heterogeneous catalysis	SuperMUC 30,000,000
Adriano Pietropaolo, Catanzaro University - NANOCASSEMBLY-targeting the polymer assembling in supramolecular nanoparticles for in cell imaging devices	SuperMUC 15,200,000
Alessandra Sabina Lanotte, CNR - MIXER: How stratification, rotation and confinement impact on the turbulent mixing of passive and active particulate matter	MareNostrum 25,000,000
Michele Migliore, CNR - LACEHIP: Large scale CELLular model of the HIPpocampus	Juqueen 21,000,000

# Particle Physics

# Life Science Computational biology



# Computational Chemistry



# Life Science Bioinformatics

# LISA

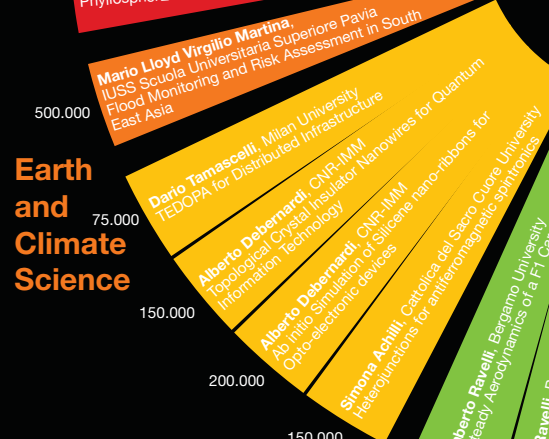


## Regione Lombardia


# Computational Engineering



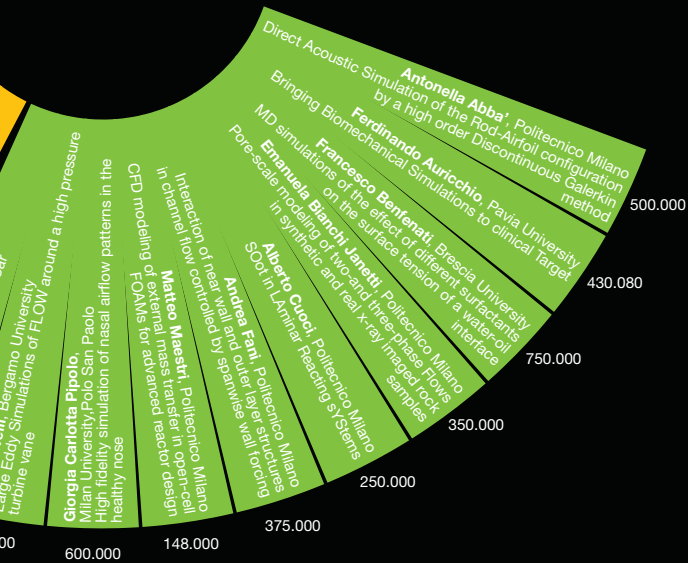
# Earth and Climate Science




# Condensed Matter Physics



# Computational Fluid Dynamics

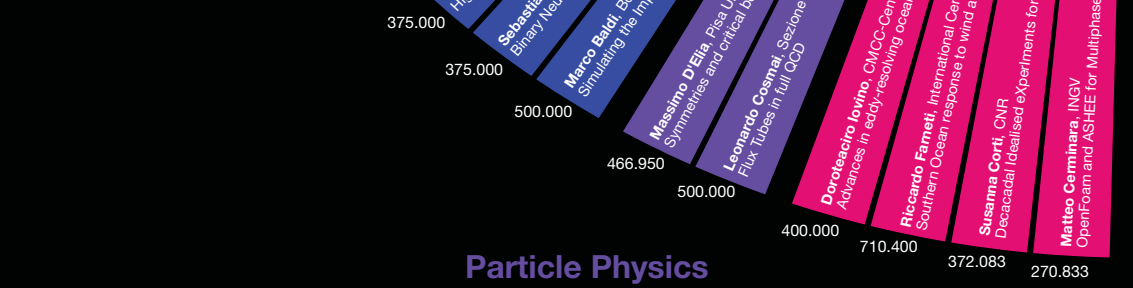


# Computational Chemistry



# ISC

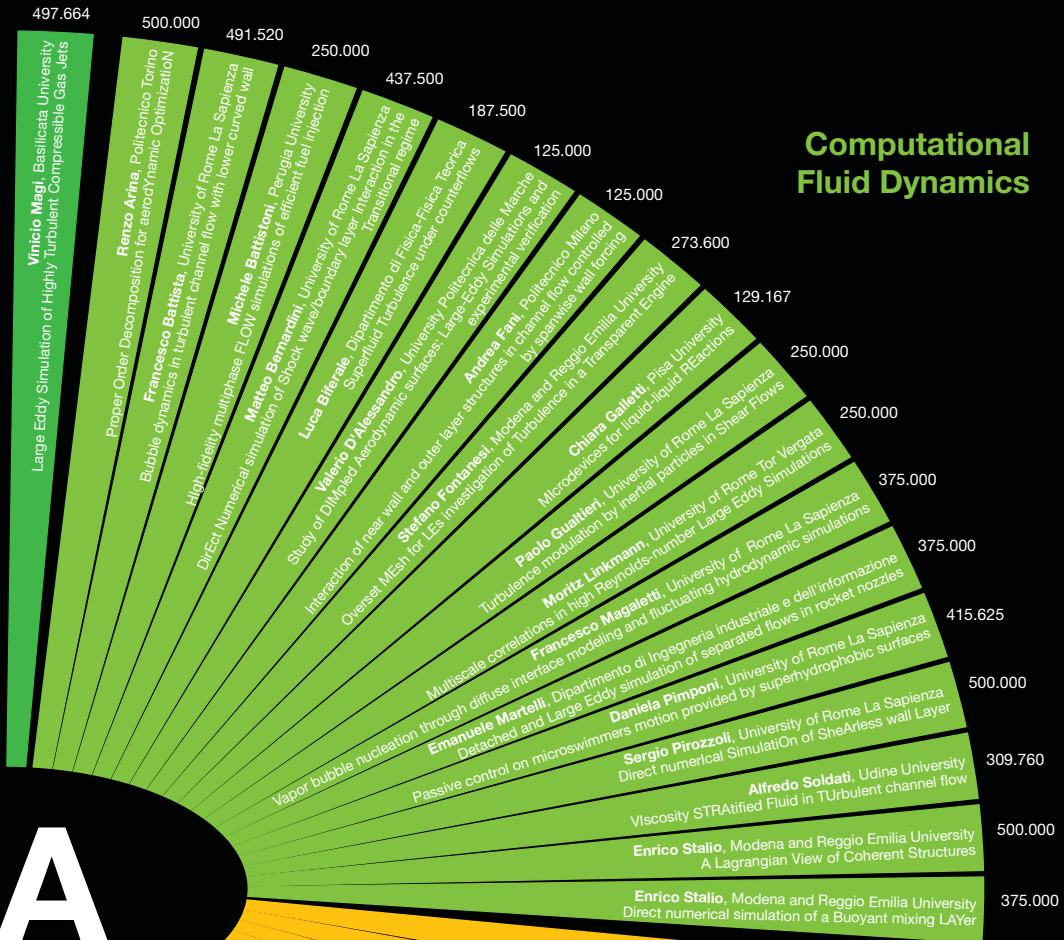
# Astrophysics and Plasma Physics



# Particle Physics

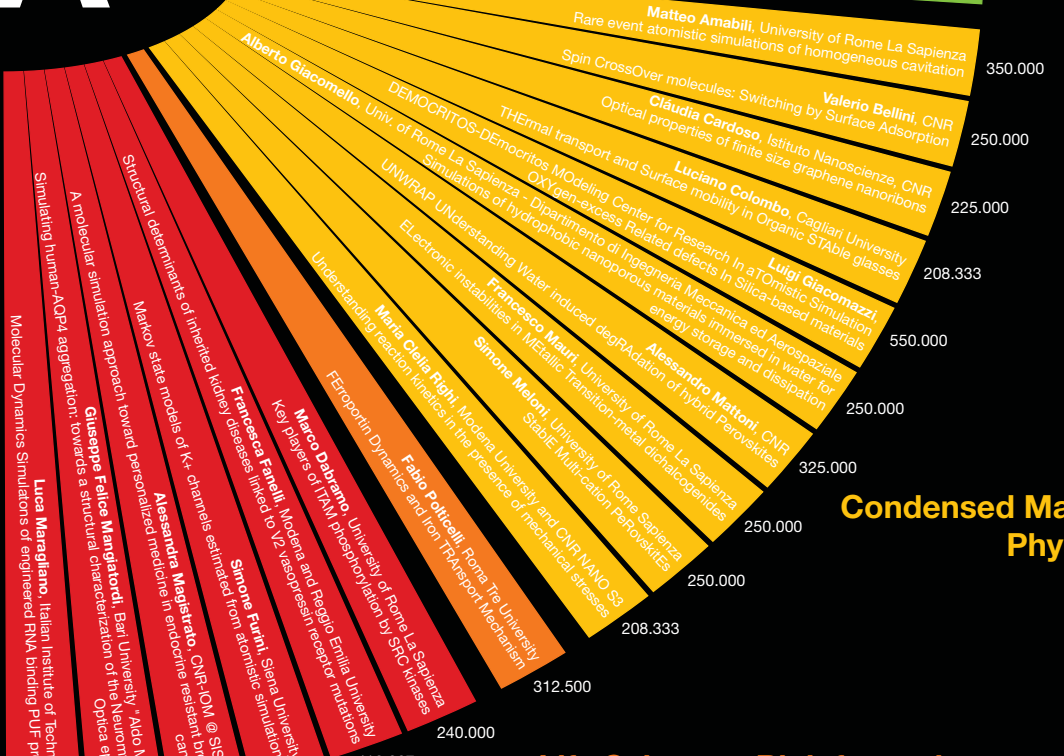
# Earth and Climate Science

# Computational Engineering



# Computational Fluid Dynamics

# RA



# Condensed Matter Physics

# Life Science - Bioinformatics

# Life Science - Computational biology

# The Race UP Team

Andrea Michelotti, Filippo Zanetti, Nicola Munerlin, Marco Tavoso, Stefano Lovato  
Race UP Team, Department of Industrial Engineering, University of Padua

The Race UP Team is a team of students from the University of Padua, working to design and produce two competition cars every year that will take part in the Formula SAE and Formula Student competitions. The team is divided into two divisions, one producing a combustion car and the other an electric car. The Electric Division last year participated at Formula SAE Italy, held in Varano de Melegari, during which the car was evaluated by experts based on economic and design choices. We won the design event and the business plan presentation, together with the overall competition for our category.

This was achieved through careful design, choice of adequate materials and precise manufacturing of every part of the car. In particular, in the years 2015-2016, the aerodynamic department of the Electric Division collaborated with Cineca, using the supercomputer Galileo HPC located in Bologna, in order to simulate the behaviour of the car in the best possible way. The main goals were to validate the aerodynamic package developed, to obtain useful data for the control of the vehicle and to understand the interaction between the chassis, the suspensions and the aerodynamic surfaces.

This was done by means of a CFD (Computational Fluid Dynamics) simulation, performed using the software Star CCM+. These simulations use the finite volume method to solve numerically the differential equations that control the behaviour of a generic fluid, in this case air.

Starting from a simplified CAD model, a mesh is generated around the car (i.e. a grid of points over which the equations will be solved approximately). The mesh is denser over the areas close to the car, where there are high


gradients of pressure and velocity, and is sparser away from the car. The total number of cells of our mesh was over 34 million, so as to correctly simulate all the parts of the car and to consider all the complicated phenomena happening. The construction of the mesh is very delicate: a poor mesh will lead to a poor simulation that may not converge or may give wrong results.

Once the mesh is ready, the next step is to set the simulation in terms of parameters (air density and viscosity), type of simulation (compressible or incompressible...), turbulence model (we used the k-w model), boundary conditions (e.g. velocity of the air approaching the car, speed of the wheels) and many other (specific model for the radiator...).


Now the simulation is ready to be solved: its enormous complexity required 32 CPUs, 160GB of RAM, 287 hours of computation and 3000 iterations. The highest residual (related to turbulent kinetic energy) was close to  $10^{-2}$ , which indicates that the result was close to convergence. The data we got from the simulation were the forces acting on every aerodynamic surface (both downforce and drag), the distribution of pressure and velocity around the car (that can be useful for example to detect the areas where the flow detaches), the position of the center of pressure (important for the vehicle control) and many other information. We performed four simulations, with different velocities varying from 7.5 m/s to 30 m/s.

This simulations with the whole car have been very useful during the design event, since it is very much appreciated by the judges, because it allows a better understanding of the behaviour of air fluxes around the car, with respect to simulations where only some parts are involved.


Turbulent kinetic energy of the flow after the impact with the car



Velocity magnitude of the flow around the car seen using streamlines



Velocity of the flow after the impact with the car



# DEM simulation of pharmaceutical powders behaviour during manufacturing operations

Andrea Benassi, Claudio Cafiero, Roberto Gaspari, Ciro Cottini

Chiesi Farmaceutici, R&D headquarter, Parma (Italy)

Drug Product Process Development & GMP Manufacturing Department




Mirko Cestari Cineca

In the pharmaceutical form known as Dry Powders for Inhalation (DPI) the active molecule, in the form of crystalline powder grains, is directly inhaled by patients breathing from a medical device with a powder reservoir or fed by a single dose capsule/blister. In order to reach the alveoli in the peripheral lung airways the powder grain diameter must be smaller than 5  $\mu\text{m}$ . Such small particles are usually mixed with larger excipient particles, typically sugar powders such as Mannitol or Lactose with particle size of several hundreds of microns working as a carrier. The drug product to be manufactured and partitioned into medical devices is thus constituted by a complex poly-disperse powder blend whose mechanical behaviour (e.g. flow ability, mechanical resistance, tendency to compaction etc.) can be extremely different depending on the ingredients physico-chemical properties, the grain morphology and the dosage strength. Low strength powder blends are usually non-cohesive exhibiting a high flow ability, no hysteresis effects, low compaction capability, however they are prone to segregation phenomena. In high strength powder blends cohesive forces are dominant over gravity and segregation is almost completely suppressed, however such blends do not flow easily, they show a pronounced tendency to compaction and rat-holing, they show hysteresis (i.e. their local structure and density retain memory of the past mechanical stresses suffered). This high variability from product to product makes the design of products, production and filling processes quite complex, time and resources consuming. To make the design work faster and more effective we are developing a Discrete

Element Modelling (DEM) platform able to predict the behaviour of our powder blends in a mixer, filling machine or inside a medical device. As a proof of principle we run on Marconi a set of DEM simulations of a device filling equipment partitioning the DPI powder with an auger (volumetric partitioning system) while a stirrer keeps the powder continuously in motion. With such a tool we can easily study the occurrence, during filling manufacturing, of phenomena such as rat-holing, segregation by shear or percolation, powder compaction and caking and preventing them by properly choosing process parameters (e.g. auger or stirrer speed) or by redesigning the components geometry. The free parameters of the DEM contact potentials have been set by reproducing in-silico few simple experiments probing the static and dynamics rheological properties of our DPI powders, namely the angle of repose and the mass flow rate measurements. The DEM code of our choice is LIGGGHTS, an open source code based on the LAMMPS molecular dynamics engine developed by DCS Computing. With the public version of the code a typical simulation reproducing few partitioning cycles of the whole manufacturing run over 710 cores for 24 hours. The premium version features a dynamic load balancing tool which should speed up the simulations and improve their scalability.

Through the CFDEM Coupling software LIGGGHTS can be coupled to OpenFoam to simulate fluid-powder combined systems, an attractive scenario for a company developing DPI products as it enables for the possibility to simulate the aerosolization of powders into a medical device upon patient breathing or the active powder dispersion into the patients lungs.





# Coherent Structures and Extreme Events in Rotating Multiphase Turbulent Flows

Luca Biferale


Department of Physics and INFN, University of Rome, Tor Vergata

Fluids are ubiquitous on the earth as in the sky - inside blood arteries, in the atmosphere, inside planets, in the extra-galactic space, etc... Often flows become turbulent, developing highly irregular spatial and temporal chaotic evolutions. Theory lags behind. Fortunately, supercomputers can be used to simulate turbulent flow in fine detail. In our group at the University of Rome Tor Vergata, we have been carrying out state-of-the-art simulations of turbulent flow under rotation, which are relevant to real life situations, but can also be used to extract useful insights about the Navier-Stokes equations themselves.


Understanding the dynamics of rotating fluids is key due to the fact that the Earth spins on its axis, meaning that the oceans, the atmosphere and the inner mantle are all subjected to Coriolis force. Moreover, many engineering applications, e.g. turbomachines, do need to cope with highly turbulent flows under strong rotation. The advantage of numerical simulations is that "in silico" we have more freedom than in a laboratory, we can remove boundaries and/or modify the equations of motions to ask new questions. In a few words, we can simplify or complexify the problem and then perform numerical experiments in idealised setups. In our simulations we have changed the rotation

rate in a turbulent flow to study what happens under extreme conditions. We observed that at some point there is a transition in the flow leading to the formation of strong and large structures where the energy condensates: a sort of tornado-like structures oriented in the direction of the rotation. What is interesting is that the transition is connected to a reversal of the energy flux, from large to small scales at low rotation to the inverse when spinning is large enough.

A unique aspect of our study is connected to the possibility to study also Lagrangian turbulence, thanks to the presence of small particles advected by the flow. We provided the first systematic analysis of the preferential concentration and advection induced by the presence of large tornado structures. With particles lighter/heavier than the fluid being mainly attracted inside/outside the vertical structures (see Figure). The supercomputing power has been provided by PRACE and allowed us to look at billions of particles as they flow through the volume, providing quantitative data at unprecedented levels. This study is achieved within the activity of the ERC NewTURB project and was recently published in the Physical Review X in L.Biferale et al. PRX 6 (4), 041036 (2016).



Time evolution of three bunch of particles injected along one vertical axis on a turbulent rotating flow. Colours identify particles that are heavier (blue), lighter (green) than the flow. Red trajectories correspond to evolution of tracers (same density of the advecting fluid). Below we show one typical Eulerian configuration of the rotating flow. Notice the tendency of light particles to accumulate inside vertical structure and of heavy particles to be dispersed horizontally.



# Numerical relativity in the era of multi-messenger astronomy

**Sebastiano Bernuzzi** Parma University & INFN, **Tim Dietrich** Nikhef, **Alessandro Nagar** Centro Fermi & INFN, **Albino Perego** INFN, Bicocca, Parma University **David Radice** Princeton University & Princeton Astrophys. Sci. Dept.

The August 2017 observation of gravitational and electromagnetic radiation from the collision of two neutron stars marked the beginning of multi-messenger astronomy. The coincident detection of different types of radiations provided the first direct evidence that linked together some of the the most extreme events in the Universe. The interpretation of observations strongly relies on the study of the dynamics of binary systems in General Relativity. Numerical relativity, the art of solving Einstein's equations with computational methods, is the main tool for theoretical physicists to connect the strong-field dynamics of the source with the emitted radiations using a first-principles and ab-initio approach.

The simulation of binary mergers is a 4D, multi-scales and multi-physics problem that requires the solution of nonlinear PDEs in complex geometries. Our group develops numerical methods and codes for such simulations and computes theoretical predictions of gravitational and electromagnetic radiation emitted by compact binaries. The most challenging aspects in these simulations are the dynamical interaction between hot matter at supra-nuclear density, the modelling of neutrino radiation, the resolution of magnetohydrodynamical instabilities, and the production of accurate


gravitational waveforms for many different binary parameters. A typical simulation needs HPC techniques with computations distributed on hundreds up to few thousands of CPUs and weeks of runtime for covering the dynamical timescale of the merger (tens of milliseconds). Simulations give us access to physics that cannot be studied in any other way.

Simulations performed on Marconi Cineca with PRACE and ISCRA allocations allowed us to build precise models of the gravitational waveforms that were crucial for the analysis of the August event. The simulations were also employed for the interpretation of the electromagnetic transient powered by the radioactive decay of the heavy elements synthesized in the merger event. Moreover, we showed that a joint analysis of gravitational and electromagnetic observation can set strong constraints on the unknown equation of state of neutron star's matter, i.e. at conditions that no laboratory on Earth can explore. Key for our analysis was the use of numerical relativity strong-field information to link the observations.

## References

Bernuzzi et al, Phys.Rev.Lett. 115 (2015)  
Dietrich et al Phys.Rev. D96 (2017)  
Perego et al ApJ Lett. 850 (2017)  
Radice et al ApJ Lett. 852 (2017)

This study is also supported by: ERC H2020, Centro Fermi, CIPE/INFN, NSF. S. Bernuzzi, A. Perego, T. Dietrich and A. Nagar are Virgo members.



Gravitational waves from a binary neutron stars merger.

[BAM code, LRZ/SuperMUC, Gauss Large-Scale project 16th call, allocation pr48pu]

# Large Eddy Simulation of Highly Turbulent Compressible Gas Jets


Francesco Bonelli, Annarita Viggiano, Vinicio Magi  
School of Engineering, Basilicata University

In this project, a comprehensive investigation of the mixing properties of fully turbulent hydrogen round jets with high density ratios has been carried out in the near- and intermediate-field by using an in-house 3D Large Eddy Simulation (LES) model, named FLEDS. The investigation of hydrogen jets is a relevant subject both for the design of new propulsion systems and for safety issues, such as the sudden release of hydrogen from a high pressure vessel in the case of failure. Specifically, the aim of this project is to untangle the question whether the density ratio between the surrounding and the jet affects the turbulence mixing properties. This question has been addressed by several authors from theoretical, experimental and numerical points of view, but the outcomes are not yet conclusive. A significant issue of these jets is the presence of sharp density gradients in the flow field that can numerically produce non-physical spurious oscillations even under low-Mach number conditions. In order to avoid these instabilities, a recent localized artificial diffusivity scheme has been implemented. First, the model has been validated by comparing both the centerline mean velocity properties and the turbulence statistics of a low-Mach air jet into air with experimental data. A


qualitative view of the jet coherent structures and a quantitative analysis of the turbulence scales have also been provided. Then, LES has been used to investigate two hydrogen jets with very high ambient to jet density ratios, up to 50, and high Reynolds number, up to about  $10^5$ , with the same momentum flux. As expected, the centerline velocity decay rate is much larger than that of the air jet, however in the near- and intermediate-field the slopes of the velocity profiles are not the same if based on the classical definition of the effective diameter. Indeed, the results show that a new effective diameter, based on an effective density, is required. Furthermore, the spreading rates of the two hydrogen jets are less affected by the density ratio compared to the velocity decay rate.

#### References

BONELLI F., VIGGIANO A., MAGI V. (2016). How does a high density ratio affect the near- and intermediate-field of high-Re hydrogen jets? INTERNATIONAL JOURNAL OF HYDROGEN ENERGY, vol. 41 (33), p. 15007-15025, <http://dx.doi.org/10.1016/j.ijhydene.2016.06.174>



$Y_{H_2}$ : 0.01 0.02 0.03 0.04 0.05 0.06 0.07 0.08 0.09 0.1 0.2 0.3 0.4 0.5 0.6 0.7 0.8 0.9



Q isosurface equal to  $2.5 \times 10^{-4}$ , coloured by the hydrogen mass fraction, in the case of a jet of hydrogen issuing into C4H10 with density ratio equal to 28.8, centerline velocity at the jet exit equal to 225.29 m/s, Reynolds number equal to about  $10^5$  and Mach number equal to about 0.17. A blow-up near the injector is also shown.

# Vlasov simulations in phase space of plasma turbulence

Francesco Califano and Silvio Sergio Cerri

Department of Physics “E. Fermi”, Pisa University


Nearly all astrophysical plasmas are in a turbulent state. Investigating the properties of turbulent fluctuations along the fluid and kinetic spectrum down to the dissipation scale is a crucial step to understand how turbulence feeds back on the large-scale evolution of such systems and the transport properties. In this context, space plasmas are probably the best laboratory for the study of collisionless plasma turbulence, the Earth’s environment being accessible to increasingly accurate direct measurements. In situ solar wind and magnetosheath observations have provided basic constraints to theoretical models by determining the typical values for the slopes of the electromagnetic energy spectra, revealing the presence of breaks in the turbulent cascade at kinetic scales, characterizing intermittency and revealing the presence of coherent structures. The transition from the magnetohydrodynamic “inertial range” spectrum to the kinetic spectrum observed around the proton gyroradius gives clear evidence of a change in the physics underlying the cascade process, and its understanding in terms of kinetic processes is today a matter of a strong debate.

In order to investigate the transition of collisionless plasma turbulence from MHD scales to the kinetic regime and to shed light

on the nature of turbulent fluctuations in the ion kinetic range we have developed a hybrid Vlasov-Maxwell (HVM) model where fully-kinetic ions are coupled to fluid electrons<sup>[1]</sup>. Thanks to recent additional MPI/openMP parallelization and further optimization, we have been able to perform pioneering Eulerian simulations of a six-dimensional phase space, i.e. three-dimensional space and three velocity dimensions (3D-3V), using a  $384^2 \times 64 \times 51^3$  grid on Marconi KNL. For the first time<sup>[2]</sup>, we extended previous 2D-3V HVM studies to simulations that simultaneously span more than one order of magnitude in wave numbers parallel to the background magnetic field and more than two decades in the perpendicular wave numbers. The resulting turbulent fields as they appear in real space are shown in the figure, exhibiting the typical anisotropic behaviour of magnetized plasma turbulence and the sheet-like structures.

## References

- [1] S.S. Cerri, F. Califano, Reconnection and small-scale fields in hybrid-kinetic turbulence, *New Journal of Physics* 19, 025007 (2017); L. Franci, et al., Magnetic Reconnection as a driver for a Sub-Ion Scale Cascade in Plasma Turbulence, *ApJ* 850:L16, 2017; S.S. Cerri, L. Franci, F. Califano, S. Landi, P. Hellinger, Plasma turbulence at ion scales: a comparison between PIC and Eulerian hybrid-kinetic approaches, *Journal of Plasma Physics* 83, 705830202 (2017); S.S. Cerri, F. Califano, F. Jenko, D. Told, F. Rincon, Subproton-scale cascades in driven hybrid-kinetic plasma turbulence, *Astrophysical J. Lett.* 822:L12 (2016)
- [2] S.S. Cerri, S. Servidio, F. Califano, Kinetic Cascade in Solar-Wind Turbulence: 3d3v Hybrid-Kinetic Simulations with electron inertia, *Astrophysical Journal Letter* 846:18, 1-6 (2017)



Three-dimensional rendering of current density,  $J$ , and magnetic field fluctuations,  $\delta B$ . Magnetic field lines with magnetic field strength,  $|B|$ , are also shown, along with embedded current density iso-surfaces (purple).

# A full atomistic computational study of the inactive states of the human $\alpha 7$ nicotinic receptor

Grazia Cottone Palermo University  
Letizia Chiodo University Campus Bio-Medico of Rome  
Therese Malliavin Institut Pasteur, Paris  
Luca Maragliano IIT Genova

Nicotinic acetylcholine receptors are ligand-gated ion channels (pLGICs) that regulate signal transmission at the neuromuscular junction<sup>[1]</sup>. The determination of the conformations corresponding to their diverse functional roles is subject of investigation and debate. Even with a certain number of X-rays structures becoming recently available, only few major structural features clearly distinguishing conductive channels from the non conductive resting, closed-locked or desensitized ones have been highlighted. The characterization of the desensitized conformation is especially complex, as experimental data provide conflicting information for different channels. Only very recently, a closed-locked structure has been crystallized bound to strychnine<sup>[2]</sup>, giving some hints to discriminate among the non conductive states. Structures of human channels are still lacking. In the past we provided a model for the open active human  $\alpha 7$  nicotinic receptor<sup>[3]</sup>. Recently, we refined via MD a conformation non-conductive though still bound to full agonists<sup>[4]</sup>. The analysis of several structural descriptors and of the protein-ligand modes of binding suggests that this conformation could correspond to the desensitized  $\alpha 7$ . The protein presents a closed pore, and a modified interface structure, with specific loops moving to provide a conformation compatible with a closed channel, while keeping the binding pockets as in the active state. Based on the observation of one single constriction point in the middle of the pore, the obtained structure should correspond to a fast desensitized. We also built and refined, via MD, a model of  $\alpha 7$  bound to antagonists. We identified structural properties that characterize a closed-locked conformation with respect

to the open and desensitized states of  $\alpha 7$ . Comparison with available experimental data for non-conductive states of various pLGICs confirms the quality of our model. Simulations of the receptor embedded in a lipid bilayer/water system have been performed with NAMD on the hundreds of nanosecond time scale, with the available resources provided by Cineca, projects IS CRA C CLOS NICO and IS CRA B ION LGIC.

## References

- [1] A.Taly et al., Nat Rev Drug Discov. 2009, 8,733
- [2] X.Huanget al, Nature 2015,526, 277
- [3] L.Chiodo et al., PloS One 2015,10, e0133011
- [4] L.Chiodo et al., Biophys. Chem., 2017, 229, 99

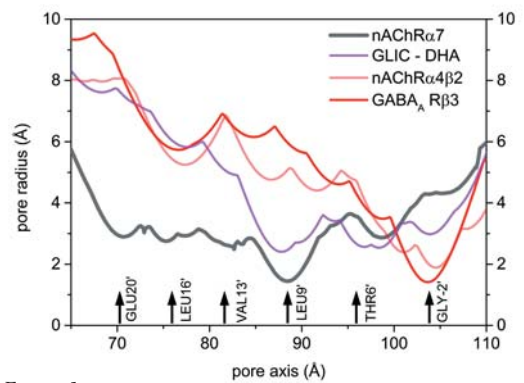



Figure 1 Pore profile in the modeled desensitized structure, compared with experimental structures claimed in desensitized state.

## Figure 2

Dot surface of the channel pore in the open, desensitized and closed-locked conformations modeled by the authors, compared to the experimental structures of a desensitized GLIC (PDB entry: 5J0Z) and of the Gly receptor bound to strychnine (PDB entry: 5CFB). Two subunits are sketched in cartoon representation and colored according to the secondary structure color code. Green and blue dots indicate where the pore radius allows room for a single water and where the radius is double the minimum for a single water, respectively. Regions in red are inaccessible to water.



# ABOUT THE FR I / FR II DICHOTOMY

Silvano Massaglia, Gianluigi Bodo, Paola Rossi, Sandro Capetti, Andrea Mignone  
Departments of Physics, Turin University

Extragalactic radio sources have been classified into two categories based upon their radio morphology: a first class of objects, named Fanaroff-Riley I (FR I), which is preferentially found in rich clusters and hosted by weak-lined galaxies, shows jet-dominated emission and two-sided jets at the kiloparsec scale that smoothly extend into the intracluster medium, where they form large-scale plumes or tails of diffuse radio emission. The second class, named Fanaroff-Riley II (FR II, or classical doubles), found in poorer environments and hosted by strong emission-line galaxies, presents lobe-dominated emission and one-sided jets at the kpc scale that abruptly terminate into hot-spots of emission. In addition to morphology, FR I and FR II radio sources have been distinguished based on power: objects below  $10^{25}$  W/Hz/str at 178 MHz were typically found to be FR I sources.

The distorted, diffuse, and plume-like morphologies of FR I sources led researchers to model them as turbulent flows, while the characteristics of FR II, such as their linear structure and the hot-spots at the jet termination, are associated with hypersonic flows.

By means of hydro-dynamical 3D simulations of supersonic jets, we investigated how the displayed morphologies depend on the jet parameters. Bow shocks and Mach disks at the jet's head, which are probably responsible for the hot spots in the FR II sources, disappear for a jet kinetic power lower than  $10^{43}$  erg/s. This threshold compares favorably with the luminosity at which the FR I/FR II transition is

observed.


The problem was addressed by numerical means carrying out 3D HD simulations of supersonic jets that propagate in a non-homogeneous medium with the ambient temperature that increases with distance from the jet origin, which maintains constant pressure.

The jet energy in the lower power sources, instead of being deposited at the terminal shock, is gradually dissipated by the turbulence. The jets spread out while propagating, and they smoothly decelerate while mixing with the ambient medium and produce the plumes characteristic of FR I objects.

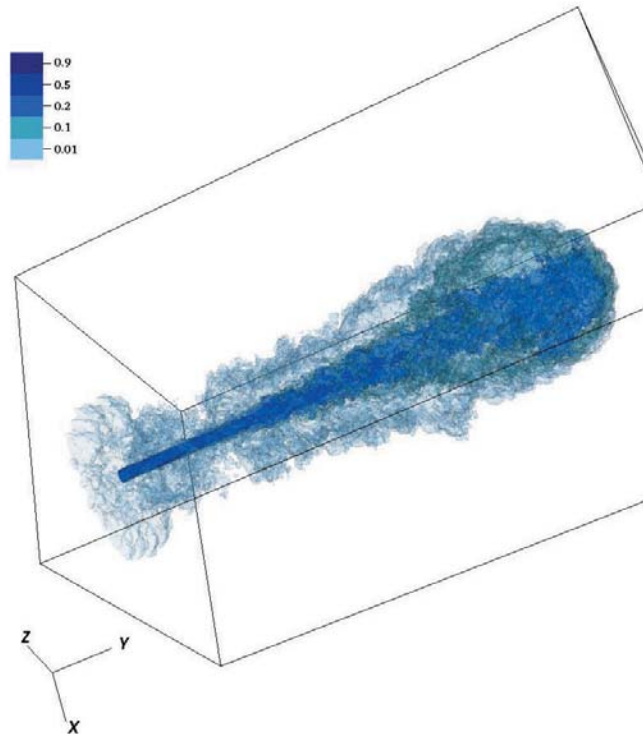
Three-dimensionality was shown to be an essential ingredient to explore the FR I evolution because the properties of turbulence in two and three dimensions are very different, since there is no energy cascade to small scales in two dimensions, and two-dimensional simulations with the same parameters lead to FR II-like behavior.

## References

S. Massaglia, G. Bodo, P. Rossi, S. Capetti, A. Mignone  
"Making Fanaroff-Riley I radio sources. I. Numerical hydrodynamic 3D simulations of low-power jets"  
A&A, Vol. 596, id. A12, 10 pp.



3D volume rendering of the tracer distribution and 2D longitudinal cut of the density distribution for case of a jet with Mach number = 4, jet-to-ambient density ratio = 1/100, kinetic power =  $1.1 \times 10^{42}$  erg/s at a life time of  $5.7 \times 10^7$  yrs.





# Turbulence and secondary motions in square duct flow

Sergio Pirozzoli, Paolo Orlandi

Department of Mechanical and Aerospace Engineering, University of Rome La Sapienza

Davide Modesti, Francesco Grasso

CNAM-Laboratoire DynFluid, Paris

Turbulent wall-bounded flows are typically studied in canonical idealizations, mainly the plane channel and circular pipe flow. However, ducts in practical industrial contexts frequently have more complex shape. As a first step in geometrical complexity, we have studied the flow in a duct with square cross section, in a wide enough range of Reynolds numbers to reach flow conditions which are representative of fully developed turbulence. Direct numerical simulations of the Navier-Stokes equations are carried out, with sufficiently fine mesh to resolve all of the relevant flow scales. The problem is made computationally challenging by the presence of secondary (cross-flow) motions which do not occur in canonical flows, and whose typical turnover time is much longer than the typical time scales of wall turbulence. Figure 1 shows instantaneous streamwise velocity contours in the cross-stream and near wall-parallel planes (colour scale from blue for low speed, to red for high speed), with superposed streamlines of the time-averaged cross-stream flow, for the highest Reynolds number case which we managed to simulate ( $Re_{\tau} \approx 1000$ ). The secondary motions are found to take the classical form of eight counter-rotating eddies which modulate the near-wall, small-scale turbulence.

The intensity of the secondary eddies is found to be of the order of 1-2% of the bulk flow velocity, and approximately unaffected by Reynolds number variation. Despite their effect of redistributing the wall shear stress along the duct perimeter, we find that secondary motions do not have large influence on the bulk flow properties, and the streamwise velocity field can be characterized with good accuracy as resulting from the superposition of four flat walls in isolation. As a consequence, we find that parametrizations of the duct resistance based on the hydraulic diameter concept are successful in predicting duct friction. Figure 2 shows the friction factor as a function of the bulk Reynolds number based on the hydraulic diameter ( $Re_{D_h} = u_b D_h / \nu$ ), where the solid line denotes the reference friction factor curve for smooth pipe flow and symbols indicate present DNS data (circles), previous DNS data of Pinelli et al. (2010) (triangles) and Vinuesa et al. (2014) (squares), experiments by Jones (1976) (stars), circular pipe flow DNS data by El Khoury et al. (2013) (right-triangles) and circular pipe experiments by McKeon et al. (2004) (daggers). Excellent collapse of the various distributions is obtained, especially at high Reynolds number.




Figure 2

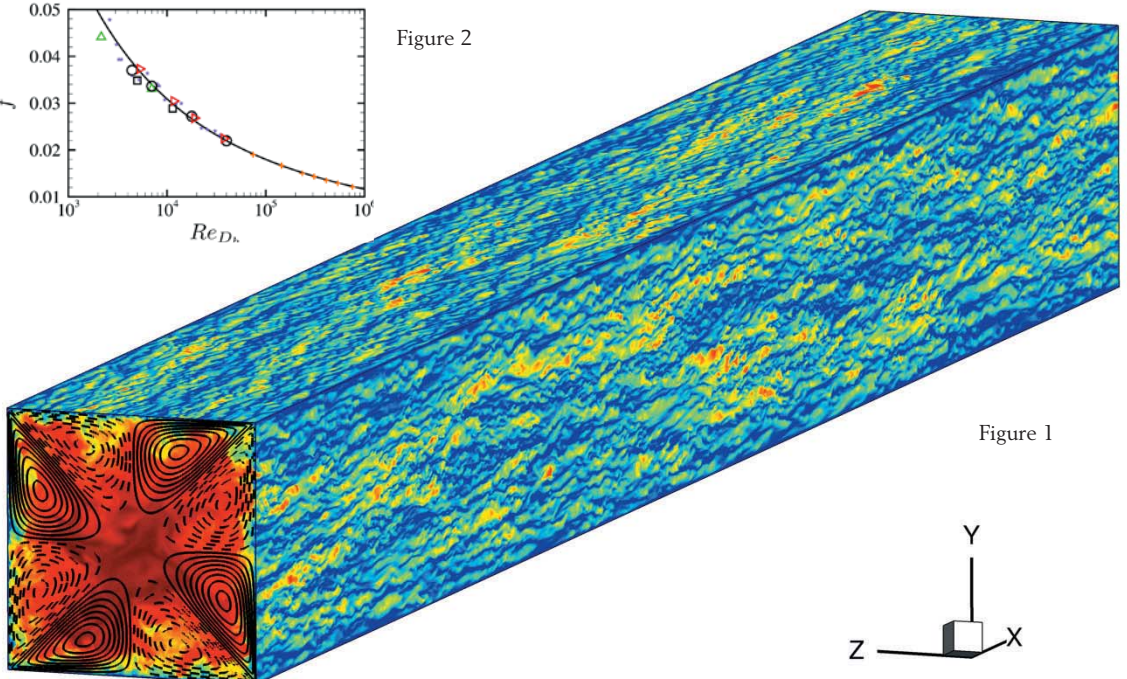


Figure 1

# Electronic Structure and Reactivity of Oxides Thin Films

Sergio Paolo Tosoni, Gianfranco Pacchioni and Philomena Schlexer  
Bicocca University, Milan

This project focused on chemical and physical properties of metal-supported oxide thin films and, more in general, metal-oxide interfaces. We studied both physical (charge transport mechanism) and chemical (charge transfer and adsorption) properties by means of density functional calculations.

We performed a detailed investigation in the frame of the density functional theory of structural, electronic and vibrational properties of freestanding and Mo(100)-supported CaO films, as functions of the film thickness and intensity of the external static electric field.<sup>[1]</sup> The phonon frequencies negligibly depend on the external electric field. A small gradual increase of the energy of CaO phonons upon increase of the film thickness was found. The effect of Mo support was observed in the systematic decrease of the energy of phonons. The applied electric field showed a minor effect on the structure of CaO films, whereas electronic properties of the oxide were significantly affected.

The adhesion of gold on oxide surfaces is a very important parameter in the design of heterogeneous catalysts. Calorimetric experiments show a trend in gold binding which is not fully clear: in particular, the Au adhesion on an ionic non-reducible oxide (MgO) is similar to what reported for a reducible oxide such as TiO<sub>2</sub>. In turns, CeO<sub>2</sub> show a surprisingly strong binding capability toward Au. In order to clarify this trend, we performed a systematic series of calculations on Au adhesion on magnesia,<sup>[2]</sup> titania, and ceria, spanning from an isolated Au atom to an Au<sub>20</sub> cluster and extended gold/oxide interfaces. It turned out that two aspects may contribute in clarifying this trend, namely: a) the match between oxygen lattice positions on the oxide surfaces and gold adatoms, which

is excellent for magnesia and ceria but poor for titania and b) the role of oxygen vacancies.


ZnO bilayers supported on Cu(111), Ag(111) and Au(111) were studied, showing a small charge transfer from the ZnO to the Au and Ag support.<sup>[3]</sup> In the case of Cu(111), the charge flow has an opposite sign. The adsorption of a probe molecule, CO, was simulated to size the effect of the different supports on the ZnO surface properties. The CO stretching frequency is red-shifted on ZnO/Cu, blue-shifted on ZnO/Au and almost unaltered on Ag, correlating well with the work function of the ZnO-metal system.

## References

[1] Kuklin et al. Phys Rev B 2017, 95, 165423.

[2] Tosoni et al. J. Phys. Chem. C 2017, 121, 28328

[3] Tosoni et al. J. Phys. Chem. C 2017, 121, 27453



DFT calculations applied to oxide thin films and surfaces: effect of an external electric field on the vibrational properties and electronic structure of supported CaO thin films (top), Au adsorption on MgO, TiO<sub>2</sub> and CeO<sub>2</sub> (middle) and ZnO thin films supported on Cu, Ag and Au (bottom).


# How does collagen interact with hydroxyapatite surface?

Piero Ugliengo  
Turin University

Bone has a hierarchical structure based on the mineralized fibril, an organic matrix envisaging collagen protein in tight interaction with hydroxyapatite mineral (HAP) and stabilized by water molecules. The tremendous complexity of this natural composite material hides the extraordinary features in terms of high compressive strength and elasticity imparted by the collagen protein. Clearly, understanding the nanoscale interface and mechanics of bone at atomistic level can dramatically improve the development of bio-composite and the understanding of bone related diseases. We tried to shed some light on the interaction between a model of a single-collagen-strand (COL) with the most common dried P-rich (010) HAP surface. The methods of choice are static and dynamic simulations based on density functional theory. Collagen is made in large extent by proline (PRO) and derivatives and PROs side chain is known to affect the

collagen triple helix stability dramatically. However, the role of PRO side chain in the COL/HAP interface has never been studied so far at quantum mechanical level. To decrease the enormous structural complexity of collagen itself, we employed a simple collagen model, e.i. a single strand based on the poly-L-proline type II polymer (PPII), which, for its composition, nicely suites our purposes. We discovered that during the HAP adsorption process the polymer deforms to create a relatively strong electrostatic interaction between the PRO carbonyl C=O group and the most exposed Ca ion of the HAP surface. Both dynamic and static simulations agree that the HAP surface guides the formation of PPII conformers which would be un-stable without the support of the HAP surface. This work is the first step towards the development of a full collagen model envisaging a three-interlocked helical polymer interacting with the HAP surfaces.

Adsorption of the collagen model on the hydroxyapatite surface. Ca ions, P and O atoms as cyan, yellow and red spheres; carbon, nitrogen and hydrogen atoms in green, blue and light gray.



# Pore-scale modeling of multiphase flow in synthetic and real x-ray imaged rock samples

Emanuela Bianchi Janetti, Gael Raymond Guedon  
Politecnico Milano

Multiphase flow in porous and fractured media is ubiquitous in environmental and industrial applications. The current state of understanding and modeling of these flow dynamics is still largely incomplete. This constitutes a critical barrier to scientific and technological advancements. In this context, we aim to develop an entirely original conceptual and theoretical framework to characterize two- and three-phase flow at the pore scale in synthetically generated media and in real rock samples. The final goals of this study are (a) to provide high quality information to be employed in groundwater and reservoir modeling, including the estimation of relative permeabilities which accounts for the mutual interaction of multiple flowing phases and (b) to quantify the uncertainty associated with such information.

To achieve these objectives, periodic pressure driven two- and three-phase flow mimicking experimental procedures performed during water and alternating water-gas injection are numerically solved using the open-source OpenFOAM code. Here we summarize the preliminary results performed with capillary number  $Ca = 10^{-5}$  within six water-wet limestone rock samples characterized by very similar macro properties (connected porosity and absolute permeability of about 10% and 250 mD, respectively).

i) The residual saturations (end-points of relative permeability curves) in both oil/water and gas/oil/water environments is strongly affected by the geometry of the pore space.

ii) A consistent reduction of residual oil saturation is observed for all the samples analyzed passing from two-phase (oil/water) to three-phase (gas/oil/water) flow. The mechanism for such enhanced oil recovery is ascribed to the replacement of trapped oil in large pores by gas.

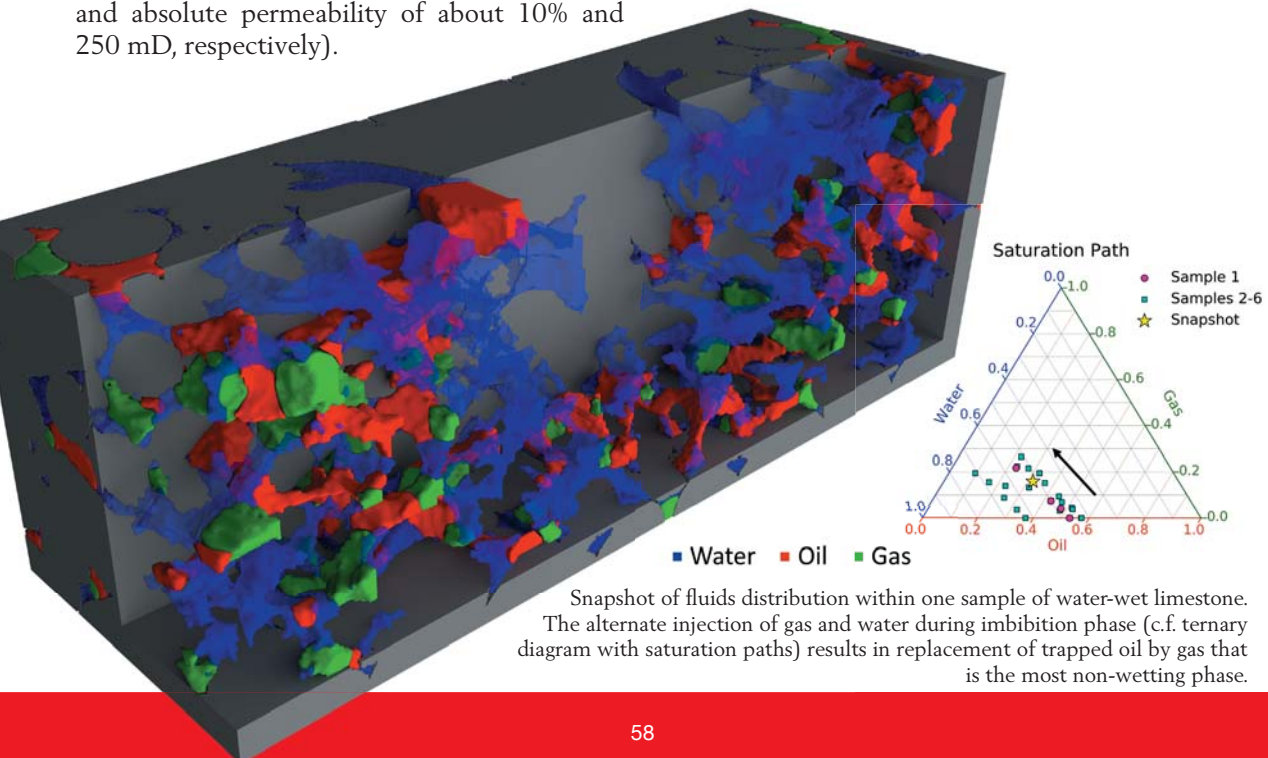
iii) Larger oil relative permeability values are predicted in three phase (gas/oil/water) flow with respect to two phase (oil/water) flow.

iv) Water relative permeability does not change significantly passing from two- to three-phase environment.

Although the development of the computational framework is in its early phase, our results demonstrate the feasibility and the capabilities of our modeling-based approach. In particular, it enables us to progress beyond current experimental constraints, with reference to exploring a variety of operational settings, pore-space structure, fluid rheology, as well as rock wettability.

## Acknowledgments

This project was supported by a LISA program (PoreFlow).



Snapshot of fluids distribution within one sample of water-wet limestone. The alternate injection of gas and water during imbibition phase (c.f. ternary diagram with saturation paths) results in replacement of trapped oil by gas that is the most non-wetting phase.

# Assessing the risk of coastal flooding caused by hurricanes in the Caribbean and Central America

Gianbattista Bussi, Lidia Bressan, Ines Cerenzia, Gabriele Coccia, Mario Martina  
IUSS Pavia & RED spa


Hurricanes can trigger localised rises in the sea level, called storm surge. Coastal flooding due to storm surge is a pressing issue in the Caribbean and Central America, as recent hurricanes such as Irma and Maria showed. A thorough assessment of the coastal flooding hazard is a key factor for national-scale planning and resources allocation. Furthermore, the insurance sector is increasingly interested in investing in new risk pricing tools for insurance products against hurricanes. Defining the areas that are most likely to be affected by coastal flooding (hazard mapping) is a necessary step towards the achievement of a scientifically sound hurricane risk assessment.

Given the complexity of the problem, hydrodynamic models are needed to estimate the response of the sea level to tropical cyclones. These mathematical tools which solve the shallow water equations based on finite volume or finite element numerical schemes, taking into account external forcings (pressure and wind). In this project, the Geoclaw storm surge model was used to reproduce the response of the sea level to tropical cyclones in the Caribbean and Central America. The model was set up using the most updated information (1-km bathymetry, 250-m emerged land elevation) and calibrated

against sea level gauge data.

Due to the vast extension of the study area, a very long historical record is needed to estimate low-probability hazards such as the 1-in-100-years sea surge. The historical records, which start in 1850, do not cover exhaustively the area. Thus, a synthetic catalogue of more than 60000 tropical cyclones was generated using stochastic regressions, maintaining the same statistical characteristics of the actual observed hurricanes, such as path variability, pressure distribution and radius distribution. Geoclaw was used to compute the sea level rise caused by each of these stochastic hurricanes.

The very large number of simulations and the complexity of the model required the use of Cineca's high-performing computing facilities to simulate all the stochastic events. The results were 1-km sea surge maps for a very high number of events. These maps were downscaled to a resolution of 100 m for all the locations where an asset or group of assets was identified (more than one million locations), and for each location a distribution function of the sea level surge was built. Based on this distribution functions, several hazard maps were produced, including the 1-in-100-years map showed here.



Storm surge hazard map for 100 year return period in the Caribbean and Central America. Every dot represents a specific asset or group of assets (e.g., a neighbourhood, a city or a road).

# Formation of soot and nitrogen oxides in laminar, non-premixed flames


Alberto Cuoci and Alessio Frassoldati

Department of Chemistry, Materials and Chemical Engineering “G. Natta”, Politecnico Milano

The problem of pollutant emissions from combustion devices is gaining rising importance due to their negative effects on human health and environment. In particular, soot (i.e. carbonaceous particles) is today recognized as a responsible of cancer and lung disease, and is the second-biggest human cause of global warming. Similarly, nitrogen oxides (NOX) contribute to the formation of smog and acid rains, as well as tropospheric ozone. Numerical simulations of combustion devices can be successfully used to help identifying the conditions reducing the formation and the emissions of such pollutant species and to design more environmentally sustainable devices, without sacrificing the efficiency and the safety.

Formation of pollutants (especially soot) in combustion is governed by complex, intricate chemical reaction paths, involving hundreds of species and thousands of reactions. Moreover, reaction rates are strong nonlinear functions of temperature and concentrations and the evolution of species occurs over a wide range of characteristic chemical times. In this Project, the laminar SMOKE CFD code, specifically conceived for solving laminar reacting flows described by detailed kinetic mechanisms, was adopted, in which special numerical

algorithms were implemented for facing the issues mentioned above. In particular, the code, based on the OpenFOAM framework, includes stiff ODE solvers for chemistry and detailed calculation of transport properties. It was used for simulating and analyzing a large set of 2D coflow flames, which can be considered a simplified system sufficiently representative of real devices in which non-premixed combustion is carried out. The analyses allowed highlighting the interactions between the mixing and the chemistry leading to the formation of pollutants (NOX and soot). The most relevant chemical paths were identified, together with the operating conditions able to reduce the pollutant emissions. Thanks to the adoption of a Discrete Sectional Method for describing the formation and evolution of soot, it was also possible to reconstruct the PSD (Particle Size Distribution) function of carbonaceous particles. Since most recent experimental evidences demonstrated that the effects of soot particles on human health depend on their size distribution, this represents a significant result, paving the way for more advanced analyses and fundamental studies on formation of soot in combustion.



Instantaneous maps of mass fractions of NOX in a naturally flickering laminar coflow flame fed with pure propane and regular air (fuel velocity: 10 cm/s, air velocity: 7 cm/s, burner diameter: 9 mm). Kinetic mechanism from CRECK Modeling Group (210 species, 10800 reactions) available at: <http://creckmodeling.chem.polimi.it/>.

# Optical Properties of Silicene Nano-ribbons

Alberto Debernardi

CNR-IMM, Agrate

Matteo Ferri

CNR-IMM, Agrate, Milan University

Guido Fratesi, Giovanni Onida

Milan University


Silicene, the silicon analogous of graphene, is a honeycomb two-dimensional (2D) lattices constituted by Si. As a free-standing material, silicene is predicted to present electronic band-structure having the Dirac's cone and to be stable in a buckled structure. The traditional scaling of the electronic devices had reached limitations of electrostatics and short-channel effects, threatening the continuance of Moore's law. The industry has already moved to thin film channel devices and silicene is a good candidate to reach the ultimate limit of channel thickness. Indeed, silicene can benefit of its compatibility with the current industrial processes based on the silicon technology. The recent experimental measurement of a 2D silicene sheet on metallic substrate has triggered further attention on these materials. However, the strong coupling between the epitaxial silicene lattice and its metal hosting substrate inhibits the required semiconducting properties and represents a drawback in the fabrication of functional electronic devices. The opening of an energy gap is also necessary for on/off switching. A possible way to achieve this in silicene is to restrict the geometry into one dimension by constructing nanoribbons, i.e. silicene stripes of finite width. Ab initio calculations predicted that free-

standing silicene nanoribbons with zig-zag edges have an antiferromagnetic semiconducting ground-state, which may find some possible applications in nanoelectronics and spintronics. We have computed ab initio the structural, electronic and optical properties of zig-zag oriented silicene nanoribbons. The minimal width for structurally stable planar structure having zig-zag edges corresponds to a 4-chain ribbon, shown in the Figure, whose ground state presents spin polarized edges antiferromagnetically coupled, and a lattice parameter along the nanoribbon axis contracted ( $\sim 5\%$ ) with respect to the bulk value. After having determined the dependence of structural and electronic properties as a function of ribbon widths, we computed optical spectra for nanoribbons with shorter width, in which excitonic effects are relevant due to the confinement in a quasi-1D structure. For light polarized parallel to the ribbon axis, we calculated significant differences in optical absorption spectra and peak positions, which constitute the fingerprint of ferro- or antiferromagnetic coupling of the electronic states at the ribbon edges.

#### Acknowledgments

This project was supported by LISA (ASSO) and ISCRA C (OPSiN) programs.

Absorption spectra of antiferromagnetic silicene nanoribbon computed within the independent particle random phase approximation (IP-RPA) and including local fields (LF). The relaxed structure and the isosurface corresponding to a density of 6% of the maximum value in the opposite spin populations (red, green) of the valence density are displayed in the inserts.



# Design of centrifugal pumps by CFD simulations and surrogate-based shape optimization

Remo De Donno

Department of Mechanical and Industrial Engineering, Brescia University

Centrifugal pumps are widely used in many applications, in particular in the industrial, automotive, agricultural and domestic field. They must therefore be designed for a wide range of pressure ratios and flow rates. The design and the performance prediction of this class of turbomachines is however far from trivial due to the high number of free geometric parameters to be determined, whose effect on pump performance can not be easily evaluated. Nowadays, the availability of ever increasing computational resources makes the coupling of CFD and shape optimization algorithms, either evolutionary or gradient based, a viable path toward an automatic, robust, and fast design strategy for turbomachinery. This approach allows to drastically reduce the time required by the trial-and-error procedure traditionally employed by designers, since the design methodology is completely automatized.

A fully automated surrogate-based optimization method has been developed to improve the centrifugal pump efficiency, completely based on open-source software: the Scilab software is used for the geometry definition, the OpenFOAM toolbox for meshing the fluid domain and solving the 3D-RANS equations and the DAKOTA software for governing the

surrogate-based optimization.

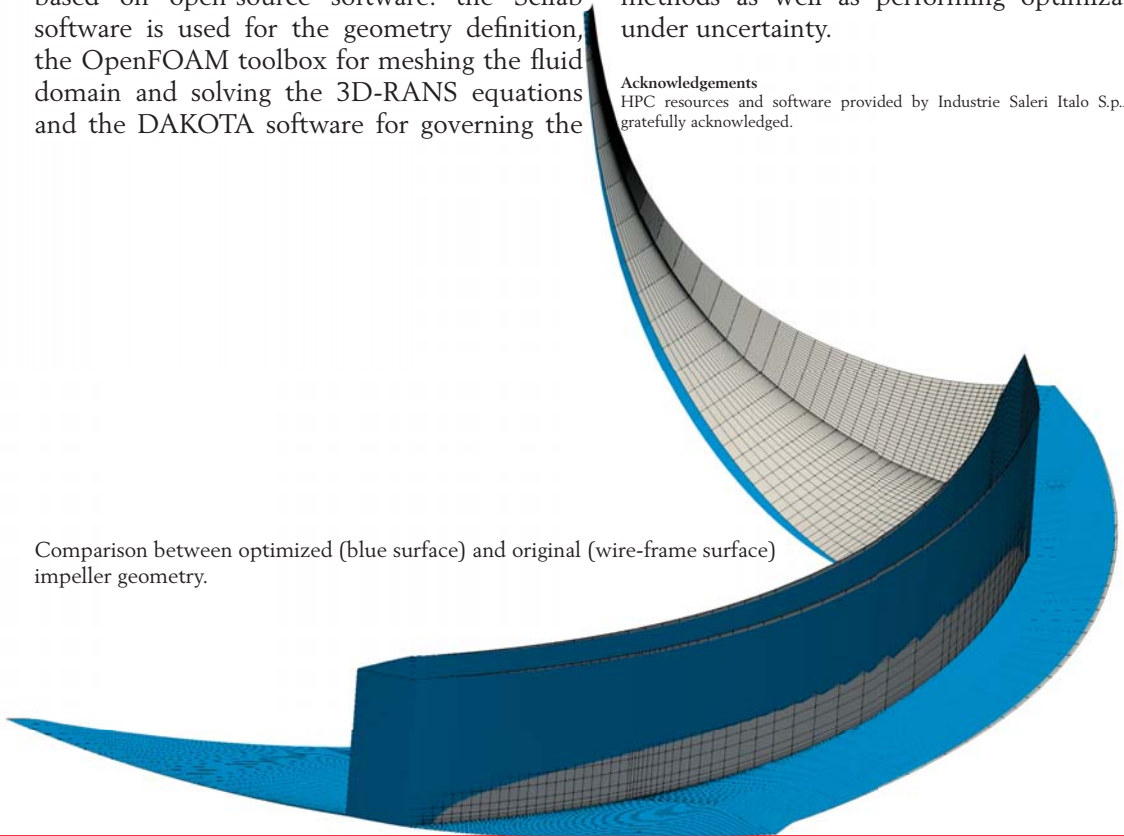
The method has been tested and validated on the ERCOFTAC centrifugal pump, where the impeller blade has been represented with Bezier polynomials and 16 control points have been used as design variables for the optimization.

The Kriging surrogate model has been adopted for this work and trained on computer experiments in order to connect accurately the impeller geometry with the pump performance, predicted by computational fluid dynamics. A single objective genetic algorithm has been set in order to maximize the pump efficiency coefficient, while keeping constrained the pressure rise coefficient, to ensure the initial operating condition is not changed during optimization. The results of this work show an improvement in the pump efficiency of about 3.4% with respect to the initial design and, therefore, demonstrating the effectiveness of the proposed surrogate-based optimization strategy. Future work will be devoted to increase the CFD accuracy by using hybrid RANS-LES methods as well as performing optimization under uncertainty.

#### Acknowledgements

HPC resources and software provided by Industrie Saleri Italo S.p.A. are gratefully acknowledged.

Comparison between optimized (blue surface) and original (wire-frame surface) impeller geometry.





# Microscopic investigation of the crystallization slowdown in supercooled liquid mixtures

Francesco Mambretti, Davide Emilio Galli

Department of Physics, Milan University

Robert Evaristo Griseni

Institut für Kernphysik, J. W. Goethe-Universität, Frankfurt am Main, Germany

The microscopic study of the stability of supercooled liquids with respect to crystallization is a fundamental open problem and is also of high technological interest. In these out-of-equilibrium condensed matter systems the crystallization competes with glass formation; therefore, elucidating the mechanisms that govern crystal growth could provide a significant step towards the understanding of the nature of the glass transition.

Mixing two components is a common strategy to hinder crystallization, favouring the persistence of disorder. In particular, it is apparent that geometric frustration due to the different atomic sizes of the two components could play a key role in preventing the crystal growth.

Crystallization slowdown has recently been observed by the group of R. Griseni (J. W. Goethe-Universität, Frankfurt am Main) via X-ray scattering measurements on supercooled Argon-Krypton liquid microjets, exploiting an innovative technique<sup>[1]</sup> which probes the inner structure of the microjets at various composition ratios.

Nowadays, computer simulations play a leading role in giving insight into microscopic scale processes during crystallization. We are supporting these experiments with out-of-equilibrium Molecular Dynamics simulations

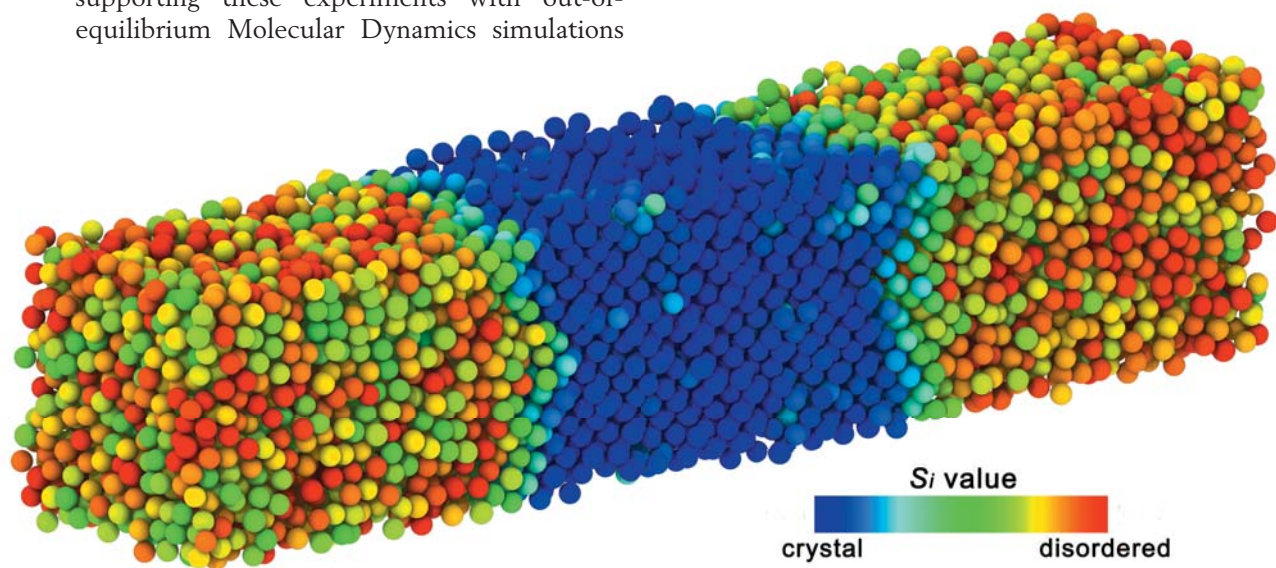
of about  $10^4$  particles in periodic boundary conditions using LAMMPS (<http://lammps.sandia.gov>). We have had access to Cineca HPC resources thanks to the LISA PUMAS (Powerfully Unveiling Microscopic Atomic Structures) project. Our simulations of solid growth from a crystal seed closely mimic the experimental conditions and qualitatively reproduce the observed crystallization slowdown. We are analysing the microscopic trajectories with Local Bond Order (LBO) parameters<sup>[2]</sup> which bring information about the orientational order surrounding each particle. We have evidence that by increasing the concentration of the minority component the local structural order parameter<sup>[3]</sup>  $S_i$  manifests much larger fluctuations during the crystallization process. Above 30%-70% composition ratio the time scale of the crystal growth rate decreases so much that it is no more possible to observe the crystallization process on computationally reasonable time scales. This is a possible signal of an arrested crystal growth dynamics.

## References

[1] R.E. Griseni et al., Phys. Rev. Lett. 90, 234501 (2003)

[2] P. Steinhardt et al., Phys. Rev. B 28, 784 (1983)

[3] J. Russo et al., Sci. Rep. 2, 505 (2012)



Snapshot of an Ar(85%)-Kr(15%) supercooled binary mixture during the crystallization process from a crystal seed placed in the middle of the simulation box. For both components, the color map refers to the structural LBO parameter  $S_i$  measured for each particle, clearly revealing the interfaces between the crystal and the disordered phase.

# Investigating the Reactivity of Defective Graphene Supported on Metallic Surfaces

Daniele Selli, Gianluca Fazio, Daniele Perilli, Martina Datteo, Costanza Ronchi, Lara Ferrighi, Cristiana Di Valentin

Department of Material Science, Bicocca University, Milan

The development of inexpensive electrocatalysts to be used in electrochemical devices is of central interest considering the worldwide impelling energy-related issues. Enabling reactivity in graphene (G) by chemical doping or by defects introduction is presently considered a viable approach to obtain active, durable, and selective electrocatalytic materials.


Graphene is usually grown or deposited on a substrate, which has been usually considered as a mere support. Recently, however, the possibility to work with atom-thick 2D layers has opened to a new emerging type of catalysis: catalysis “under cover”, i.e. within the space between the substrate and the graphenic overlayer.

In the first part of our work we investigated the geometrical, electronic and magnetic properties of a carbon monovacancy in the free-standing G sheet. Despite the large technological interest in defective G, the understanding of magnetism in such system is far to be clear. With our work, we show that self-interaction corrected hybrid functional methods (B3LYP-D\*) are capable of correctly provide a  $\pi$  magnetization of 1  $\mu$ B. The crucial role played by the exact exchange is highlighted by comparison with PBE-D2

results. Our results are supported by recent STS experiments [Phys. Rev. Lett. 2016, 117, 166801].

In a second part, we focus on the interaction between G and the Cu(111) and Pt(111) surfaces and how this affects the chemical reactivity towards water. While pristine G interacts only via dispersive forces with the substrate, defective G forms strong organometallic bonds with metal atoms because the unsaturated C atoms at the vacancy site have a high affinity towards the surface metal atoms. However, this organometallic bonds cause a decrease in the chemical reactivity observed for the free-standing G towards water dissociation. In the case of the Cu substrate, the large energy gain associated with water dissociation at the C vacancy is sufficient to over-balance the energy cost for breaking of the C-Cu bonds. On the contrary, in the case of the Pt substrate, the C-Pt bonds are too strong and water dissociation cannot be observed.

The results mentioned above have been reported in two publications: J. Phys. Chem. C, 2017, 121, 8653–8661 and ACS Appl. Mater. Interfaces, 2017, 9, 29932–29941.



Defective graphene supported by a metal surface. Water molecules can penetrate the space between the two materials where they exhibit a different reactivity.





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

ISBN 978-88-86037-38-9