

**Prague, 9-10 May 2016**



<http://www.max-centre.eu>

# **Materials design at the exascale**

**EU Centre of excellence - a H2020 infrastructure**

## **Andrea Ferretti**



- ❑ **MaX in a nutshell**: why, what, who, and how
  - ❑ MaX flagship **codes**
  - ❑ Materials informatics ecosystem: **workflows and data**
  - ❑ the **exascale challenge**
  - ❑ **service** catalogue
  - ❑ **pilot user cases**
  - ❑ **web services**, intranet, user portal
  - ❑ training & education
  - ❑ organisation & **governance**
  - ❑ **networking**, communication, dissemination
- **main ideas**
  - **MaX @ Month 8**
  - **MaX within the EU HPC ecosystem**



Led by users...

<b>CNR Modena</b>	E. Molinari
<b>SISSA Trieste</b>	S. Baroni
<b>ICN2 Barcellona</b>	P. Ordejon
<b>FZ Juelich</b>	S. Bluegel
<b>EPFL Lausanne</b>	N. Marzari

... strongly supported by HPC centres

<b>CINECA Bologna</b>	C. Cavazzoni
<b>ETH Zurich</b>	T. Schulthess
<b>FZ Juelich</b>	D. Pleiter
<b>KTH Stockholm</b>	E. Laure
<b>BSC Barcelona</b>	J. Cela

... and by 3 technology / business / education partners

**E4 SpA    CloudWeavers ltd    ICTP**



- ❑ implementing a **sustainable programming platform** in the domain of quantum materials simulations (**WP1**)
- ❑ building a dynamic **data framework** to manage the automation of high-throughput calculations, **data storage and exchange**, provenance, re-usability, etc (**WP3**)
- ❑ enabling the **exascale transition** (**WP4**)
- ❑ developing a catalogue of **user-oriented services** (**WP5**)
- ❑ dedicated integrated actions with **industrial lead users** (**WP2**)
- ❑ addressing the skills gap offering **training and education programs** (**WP6**)



# flagship codes

WP1-WP2-WP3-WP4-WP5-WP6-WP7

## quantum engines:

- legacy community codes in electronic structure
- large share of electronic structure codes **using most of the PRACE resources** in Materials Science (22%).
- overall, > **6500 subscribed users**, ~2000 journal citations since sep 2015
- part of the **PRACE benchmark suite**
- already able to exploit ~ 1 PFlop/s and **exascale oriented**



PRACE

# flagship codes

WP1-WP2-WP3-WP4-WP5-WP6-WP7



## SIESTA IS NOW RELEASED UNDER GPL OPEN SOURCE LICENCE

Following the agreement of moving the code under a new fully Open Source compliant license (GA, IPR section), the SIESTA developers took the action to change the licensing conditions for the code. Starting in the Spring of 2016, SIESTA will be released under the GPL open-source license ([https://en.wikipedia.org/wiki/GNU\\_General\\_Public\\_License](https://en.wikipedia.org/wiki/GNU_General_Public_License)). In the past SIESTA had always been free for academic use, but re-distribution was not permitted. It is expected that the move to an open-source license will result in a more flexible and dynamic development for the program. On the one hand, SIESTA will be able to incorporate functionalities already existing in other GPL codes. On the other, the barrier for contributors will be lowered, as new developments will be more easily re-distributed. With this change, all the flagship codes in the MaX consortium are open-source.


7

max-centre.eu

# workflows and data

WP1-WP2-WP3-WP4-WP5-WP6-WP7

- ❑ built around the 4 pillars of materials informatics: **Automation, Data, Environment, Sharing**
- ❑ low level: **automation** of calculations and **data storage** (workflows, provenance, re-use, search)
- ❑ high level: intuitive research **environment** to accelerate scientific discovery (high-throughput), **sharing capabilities** for collaborative research
- ❑ defines an **ecosystem** of codes, workflows, data, HPC resources, etc



# exascale challenge

WP1-WP2-WP3-**WP4**-WP5-WP6-WP7

- advanced programming models** to deal with complex architectural changes
- novel algorithms** aimed to cope at best with massive parallelism
- domain specific libraries:** performance, software sustainability, HW hiding
- HW-SW co-design, tight interaction with HW manufacturers (2+ dedicated partners)



 **ETP4HPC**

# exascale challenge

WP1-WP2-WP3-**WP4**-WP5-WP6-WP7



## CO-DESIGN: GETTING STARTED AND BUILDING THE NETWORK

Within the scope of the co-design activities of WP4, MaX started discussing with several players (INTEL, IBM, ARM, etc.) involved in developing technologies along the exascale roadmap. Discussions started at SC15, and continued in the first months of 2016, pointing out that HW manufacturers prefer to work on co-design using less complex software with respect to full featured applications (the latter are difficult to manage and their performance and profiling is often too dependent on the input dataset). In reply to that, MaX partner CINECA has developed two mini-app's concerning domain specific FFT and linear algebra (not originally included in the DOW, but adding effort on top of what was already committed), easy to be ported to new hardware, eliminating all the complexity of the full application but being accurate in reproducing numerical functionalities. The mini-app work was presented at the Computing on Low-power Architectures (COLA) workshop (February 2016, Ferrara). Based on the mini-app initiative, a successful collaboration started with the Mont-Blanc project. MaX representatives (CINECA) proposed the possibility to use the newly developed mini-app's to contribute to the application co-design and validation of the Mont-Blanc exascale architecture within the Mont-Blanc project. The discussion with Mont-Blanc PIs started at the COLA workshop and continued a few weeks later in a dedicated meeting with the Mont-Blanc team working on application porting and benchmarking. Mont-Blanc is now using the MaX mini-apps the results were presented at the 2016 Exascale Applications & Software Conference (EASC 2016) in April in Stockholm.

At the same time, co-design initiatives are promoted by MaX within several partnerships. For example:

- a FET-PROACTIVE proposal was submitted by CINECA together with IBM on variable precision arithmetics (co-design on materials science applications planned);
- a formal contract is foreseen to be signed by CINECA and INTEL on co-design themes;
- Quantum-ESPRESSO was selected for the benchmark suite of the PRACE (Pre Commercial Procurement): it is thus used to design, test, and validate the involved prototype systems (about 1 1PFlop each) targeting a European architecture with top energy efficiency.



mont-blanc

Intel

IBM

# service catalogue

WP1-WP2-WP3-WP4-**WP5**-WP6-WP7



## MAX CODE DOWNLOADING SERVICE

New high-quality releases of its open-source flagship codes are a key output of MaX. The Center offers to all users:

- Downloading of last releases and documentation
- Access to a repository of benchmarks [planned].

← PRACE



## MAX BASIC SUPPORT

MaX ensures basic support for free to all users of its codes:

- Forum/ mailing- list support for codes related issues;
- Forum support for materials science and HPC issues [planned];
- Knowledge base navigation [planned].



## MAX HELP-DESK

Direct contact help-desk service for:

- Code utilization (focus on performances and HPC optimization);
- Input preparation and output analysis;
- Code personalization (small modifications).



## MAX ADVANCED SUPPORT


Direct contact help-desk service dedicated to HPC users (specifically for proposal submission):

- Specific technical issues (scaling, CPU and memory requirements);
- Guided choice of appropriate codes for specific research context.

← PRACE

# service catalogue

WP1-WP2-WP3-WP4-**WP5**-WP6-WP7



## MAX ADVANCED CONSULTING

Dedicated on site advanced consulting focused on industrial needs:

- Scientific case, workflow definition;
- Dedicated code development;
- Code porting;
- Full support to complex research issues.



## MAX DEDICATED TRAINING

Education and training initiatives

- Schools and Courses organization;
- Specific modules;
- Training through research in the MaX Labs;
- Dedicated courses and tailored programs.



## MAX FORGE ENVIRONMENT (PLANNED)

Access to a 'MaX integrated development environment (IDE)' open to developers of scientific software in the field.

# service catalogue

- ❑ some services are free of charge
- ❑ others are for pay (different levels of pricing)
- ❑ MaX partner **CloudWeaver Ltd** dedicated to analyse the market and identify suitable **business models** to achieve self-sustainability

**CLOUDWEAVERS**



## QUANTUM ESPRESSO goes to market

Doing business with open-source software is a considerable challenge. The Quantum ESPRESSO developers' group has taken up the challenge by signing an important commercial agreement with a leading scientific software company, Schrödinger LLC (<http://www.schrodinger.com>), under the umbrella of the Quantum ESPRESSO Foundation (<http://foundation.quantum-espresso.org>), which coordinates the efforts of several MaX members (SISSA, CINECA, EPFL, ICTP, among others) also leveraging the consultancy expertise of another MaX member (CloudWeavers). Following this agreement, Quantum ESPRESSO will be included in the molecular and materials modeling suite commercialized by Schrödinger LLC, bundled with proprietary software. The coexistence in a same product of open-source and proprietary software components is requiring a considerable technical effort, which we believe will be of great benefit for the entire community of (non-commercial) users of the Quantum ESPRESSO suite of codes, and will constitute an example of original business model of great interest for the open-source scientific-software community at large. Within the MaX project, CloudWeavers helped the group since the initial phases of negotiation, by drafting a proposal, consulting on legal matters and licensing, by providing effort and cost estimates, and giving advice on the content and finalization of the commercial agreement. The contract between Schrödinger and the Quantum Espresso Foundation was signed in spring 2016.

← ISV's

# pilot user cases

WP1-WP2-WP3-WP4-WP5-WP6-WP7

## COMPUTATIONAL PROTOCOL FOR FRICTION AND TRIBOCHEMISTRY

(Key team: M.C. Righi, CNR)

INDUSTRIAL PARTNER: **Total SAS and Toyota Europe**

INDUSTRIAL INTEREST: Understanding and control of friction are key topics in manufacturing and in any application with parts in relative motion, with huge implications, e.g., in the automotive sector. Stress-assisted chemical reactions, occurring between two surfaces in relative motion, play a central role in tribological phenomena such as wear, which is ruled by bond breaking/forming across the interface, and boundary lubrication, where interfacial molecules react with the contacting surfaces and modify them.

This is the basic principle for the friction and it also explains the influence of lubricants, such as diamond/diamond-like carbon.

TARGET: Developing a computational protocol for the simulation of friction occurring in solid and boundary lubrication. Initial simulations and a parameter study to extend time/length scales.

PILOT CASE 1

## Computational protocol for understanding thermal energy storage in molten salts and nanofluids

(Key team: P. Ordejon, ICN2)

INDUSTRIAL PARTNER: **Abengoa Research**

INDUSTRIAL INTEREST: Innovation is a key factor for realising the potential of novel, renewable and clean energy sources to compete with fossil fuels. Improving the efficiencies and thus lowering the cost of energy generation is essential for the economic viability of novel, renewable and clean energy sources. Solar thermal energy is one of the most competitive alternatives: the heat generated can be stored in a heat storage medium by means of an intermediate heat transfer fluid, for further transportation and use.

Molten salts are one of the preferred options used by industry. Optimising the thermal storage and transport properties of molten salts is therefore key for the global efficiency of these technologies. One of the options for improving these properties, in particular thermal conductivity and specific heat, is to introduce nanoparticles dispersed in the molten salt. Simulation plays a very important role in this research and innovation process, as producing reliable experimental data for these systems is difficult and expensive. Within MaX, we will collaborate with Abengoa Research to elaborate computational protocols for the determination of the key thermal parameters of molten salts containing nanoparticles. The protocols will then be made public, and will be tested and used by Abengoa Research, within their proprietary activities.

TARGET: Elaborating computational protocols for the determination of the key thermal parameters of thermal storage fluids like molten salts, and to predict and explain the modification of these properties by the addition of nanoparticles (the so-called nanofluids).

PILOT CASE 2

# pilot user cases

WP1-WP2-WP3-WP4-WP5-WP6-WP7

## Computational protocol for simulating the colour optical properties of natural dyes for food industry

(Key team: S. Baroni, SISSA )

INDUSTRIAL PARTNER: **Mars Chocolate**

INDUSTRIAL INTEREST: In order to comply with EU regulations on food safety, many colorants commonly employed in the food industry will be substituted with natural ones. However, the colour palette from natural sources is still incomplete, and the understanding of the relation between the chemical structure and the optical properties of natural pigments inadequate, thus making the chase of molecules with custom-designated chromatic properties a blind search.

TARGET: Designing and implementing a simulation protocol to assist industrial research in the identification of the most suitable natural dyes to express custom-designated hues of colours (particularly in the purple-blue gamut), and to apply it to the screening of anthocyanins.

PILOT CASE 3

## Design and discovery of optimal solid-state electrolytes

(Key team: N. Marzari, EPFL )

INDUSTRIAL PARTNER: **Bosch Research**


INDUSTRIAL INTEREST: Improving performance and safety of electrochemical storage systems to reduce amount of greenhouse gas emitted by vehicles.

TARGET: Search for new materials for solid-state electrolytes for lithium ions and protons in a high-throughput screening approach, by understanding diffusion processes and finding accurate descriptors for ionic conductivity.

PILOT CASE 4


# web services, user portal

WP1-WP2-WP3-WP4-WP5-WP6-WP7




# training

WP1-WP2-WP3-WP4-WP5-WP6-WP7



□ bridge the **skills gap** for computational materials science @ HPC

# organisation & governance



## International Advisory Board Members

Alessandro Curioni, IBM Research Zurich & IBM Europe

Giulia Galli, University of Chicago and MICCoM

Kersti Hermansson, Uppsala University and European Materials Modeling Council (EMMC)

Boris Kozinsky, Bosch R&T Center Boston

Jens Norskov\*, Stanford University

Marie-Christine Sawley, Intel Paris Exascale Laboratory

Erich Wimmer, Materials Design and EMMC


TBC, Graphene Flagship

← international (US)

← EMMC

← Graphene Flagship

# organisation & governance



## International Advisory Board Members

Alessandro Curioni, IBM Research Zurich & IBM Europe

Giulia Galli, University of Chicago and MICCoM

Kersti Hermansson, Uppsala University and European Materials Modeling Council (EMMC)

Boris Kozinsky, Bosch R&T Center Boston

Jens Norskov\*, Stanford University

Marie-Christine Sawley, Intel Paris Exascale Laboratory

Erich Wimmer, Materials Design and EMMC

TBC, Graphene Flagship

← IBM

← Bosch

← Intel

# networking (EU)

- E-CAM, NOMAD CoE's**: a strategy meeting (sept 2015) and a joint “psi-k highlight” paper
- EoCoE CoE**: kick-off meeting attendance
- PoP CoE**: collaborative work on SW profiling
- European Materials Modelling Council (**EMMC**)
- NFFA-Europe**: we aim at a synergy on service provisioning concerning simulations
- Graphene Flagship**: ongoing conversation to establish a formal collaboration to access MaX services
- PRACE**: collaboration on training initiatives; MaX service is supporting PRACE applications
- MARVEL**: a long term Swiss NSF project on materials, and MaX partner on workflows and data activities
- IPCEI-HPC-BDA**: coordination with leading stakeholders from Italy, France, Lux, Spain
- SESAMenet and EDISON**: coordination in view of joint training activities and proposals
- Region EmiliaRomagna** BigDATA and HPC initiative


# networking (int'l)

- ❑ Participation to a **materials science coordination forum** (chair G. Galli, Chicago) including newly established centres from **US, EU, CH, and JP**. Among other: MICCoM, Berkeley, Rutgers, ECAM, NOMAD, MARVEL, NIMS. A meeting is planned in Paris during spring 2017.
- ❑ Scientific collaboration (CSCS) with **ORNL and DoE Exascale Computing Project** on domain specific libraries. Further collaborations (CSCS) with **BNL, ORNL** concerning QE.
- ❑ Collaborations with **UC Berkeley and LBNL** (prof Lin Lin and Chao Yang) on the development of efficient algorithms for massively parallel architectures.
- ❑ MaX is especially active (training & education) in **developing countries** thanks to its collaboration with the **ICTP partner**.

# MaX in numbers



5 research centers



5 HPC centers



3 education & business partners



> 70 researchers



~6.500 users (subscribers)




4 flagship codes



2 data, workflow & interoperability suites

## ONGOING ACTIVITY (SEPTEMBER 2015—APRIL 2016)




**thanks!**



# MaX in numbers



5 research centers



5 HPC centers



3 education & business partners



> 70 researchers



~6.500 users (subscribers)




4 flagship codes



2 data, workflow & interoperability suites

## ONGOING ACTIVITY (SEPTEMBER 2015—APRIL 2016)





**MAX**