siesta

EURO

CoE European Excellence in HPC Applications

MAX DRIVING THE EXASCALE TRANSITION

CASTIEL2 *"Code of the Month"* series with SIESTA

11/12/2024

Online

Co-funded by the European Union

EuroHPC Joint Undertaking

MaX - MAterials design at the eXascale has received funding from the European High Performance Computing Joint Undertaking under grant agreement no. 101093374 and Participating Countries in Project (Czechia, France, Germany, Italy, Slovenia and Spain).

MAX DRIVING THE EXASCALE TRANSITION

A. García (ICMAB) F. Pedron N. Wittemeier (ICN2) E. deF. Martins




SIESTA: a DFT code for large scale computational material science in HPC environments

Pablo Ordejón
Catalan Institute of Nanoscience and Nanotechnology – ICN2
Barcelona, Spain

siesta™

MaX "MAterials design at the eXascale" has received funding from the EuroHPC JU and participating countries (Czechia, France, Germany, Italy, Slovenia and Spain) under grant agreement No. 101093374.

Co-funded by the European Union

EuroHPC Joint Undertaking

MINISTERIO DE CIENCIA, INNOVACIÓN Y UNIVERSIDADES

Financiado por la Unión Europea NextGenerationEU

AGENCIA ESTATAL DE INVESTIGACIÓN

materials for…

batteries &
energy storage

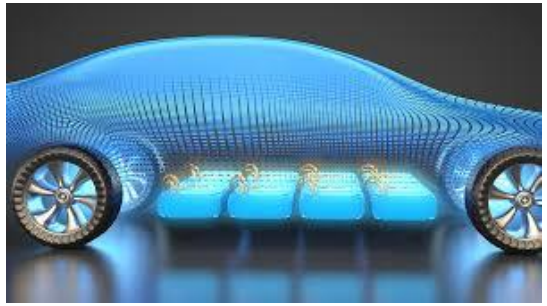

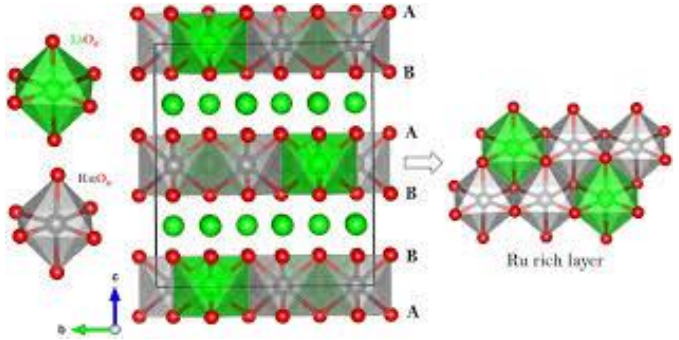

photovoltaics

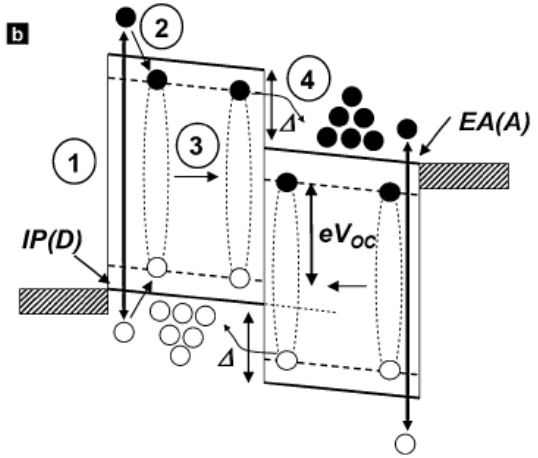

catalysis

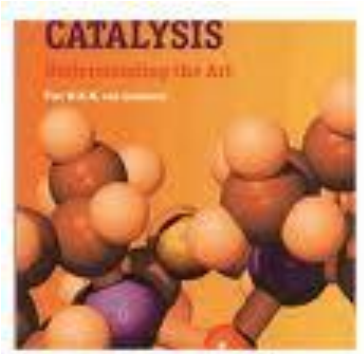

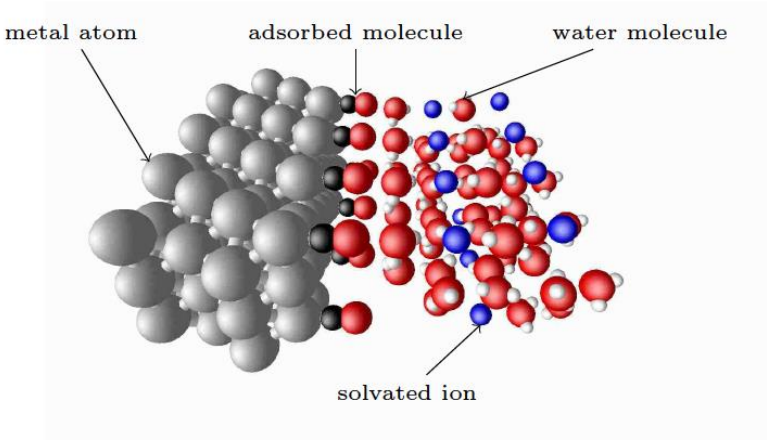

… and many more

computational materials design for…

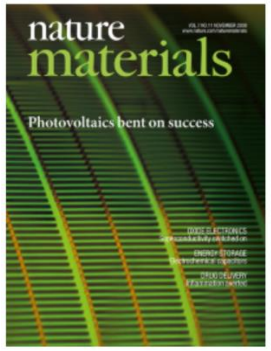

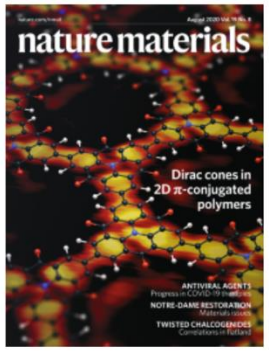

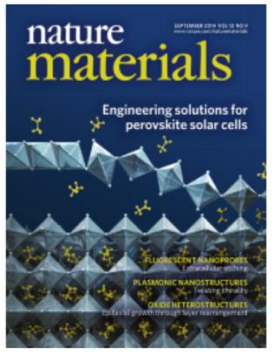

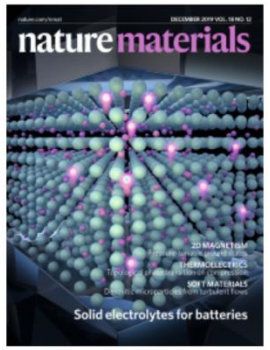

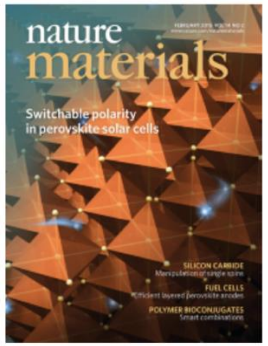

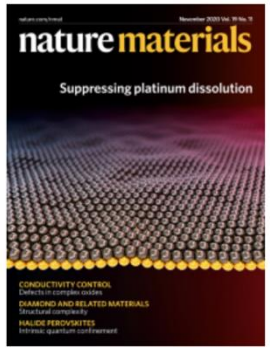

Nature Materials taken as example

Research areas covered in the journal

- Engineering and structural materials (metals, alloys, ceramics, composites)
- Organic and soft materials (glasses, colloids, liquid crystals, polymers)
- Bio-inspired, biomedical and biomolecular materials
- Optical, photonic and optoelectronic materials
- Magnetic materials
- Materials for electronics
- Superconducting materials
- Catalytic and separation materials
- Materials for energy
- Nanoscale materials and processes
- Computation, modelling and materials theory
- Surfaces and thin films
- Design, synthesis, processing and characterization techniques

- Materials are pervasive
- Connected to a number of **societal challenges**

accelerating discovery:

- **modelling**
- **codes**
- **HPC infrastructures**
- **data platforms**
- **AI**

THE RISE OF MATERIALS SCIENCE

3 Technologies That Could Create Trillion-Dollar Markets Over the Next Decade

By Greg Satell Updated April 21, 2019 9:00 a.m. ET

Yet today, we're in the midst of a materials revolution. Powerful simulation techniques, combined with increased computing power and machine learning, are enabling researchers to automate much of the discovery process, vastly accelerating the development of new materials

BARRON'S (April 2019)

Courtesy of N. Marzari

"There must be something wrong with your calculations."

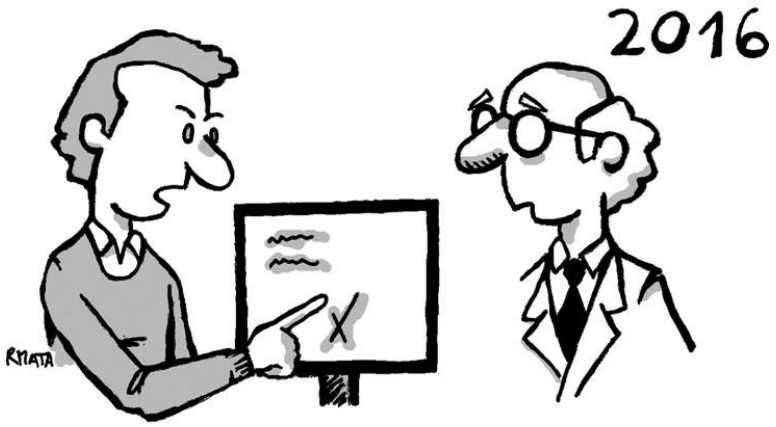

"There must be something wrong with your experiments."

Computational Materials Science

R. Mata and M. A. Suhm, Angew. Chem. Int. Ed. 2017, 56, 11011 – 11018

materials modelling

quantum mechanics based
atomistic modelling of materials
+
interfacing with **multiscale** approaches

Electronic Structure Methods

- highly accurate (predictive)
- computationally demanding
- **a case for HPC**

Schrödinger's Equations for the electrons

$$\nabla^2\psi(\mathbf{r}) + \frac{2m}{\hbar^2}\left[E - V(\mathbf{r})\right]\psi(\mathbf{r}) = 0$$

"Many body": r → $\{r_i\}$ for all electrons i

Newton's Equations for the nuclei

$$\vec{F}_i = -\frac{dE_e(\{\vec{R}\})}{dR_i}$$

materials modelling

quantum mechanics based
atomistic modelling of materials
+
interfacing with **multiscale** approaches

Electronic Structure Methods

- highly accurate (predictive)
- computationally demanding
- **a case for HPC**

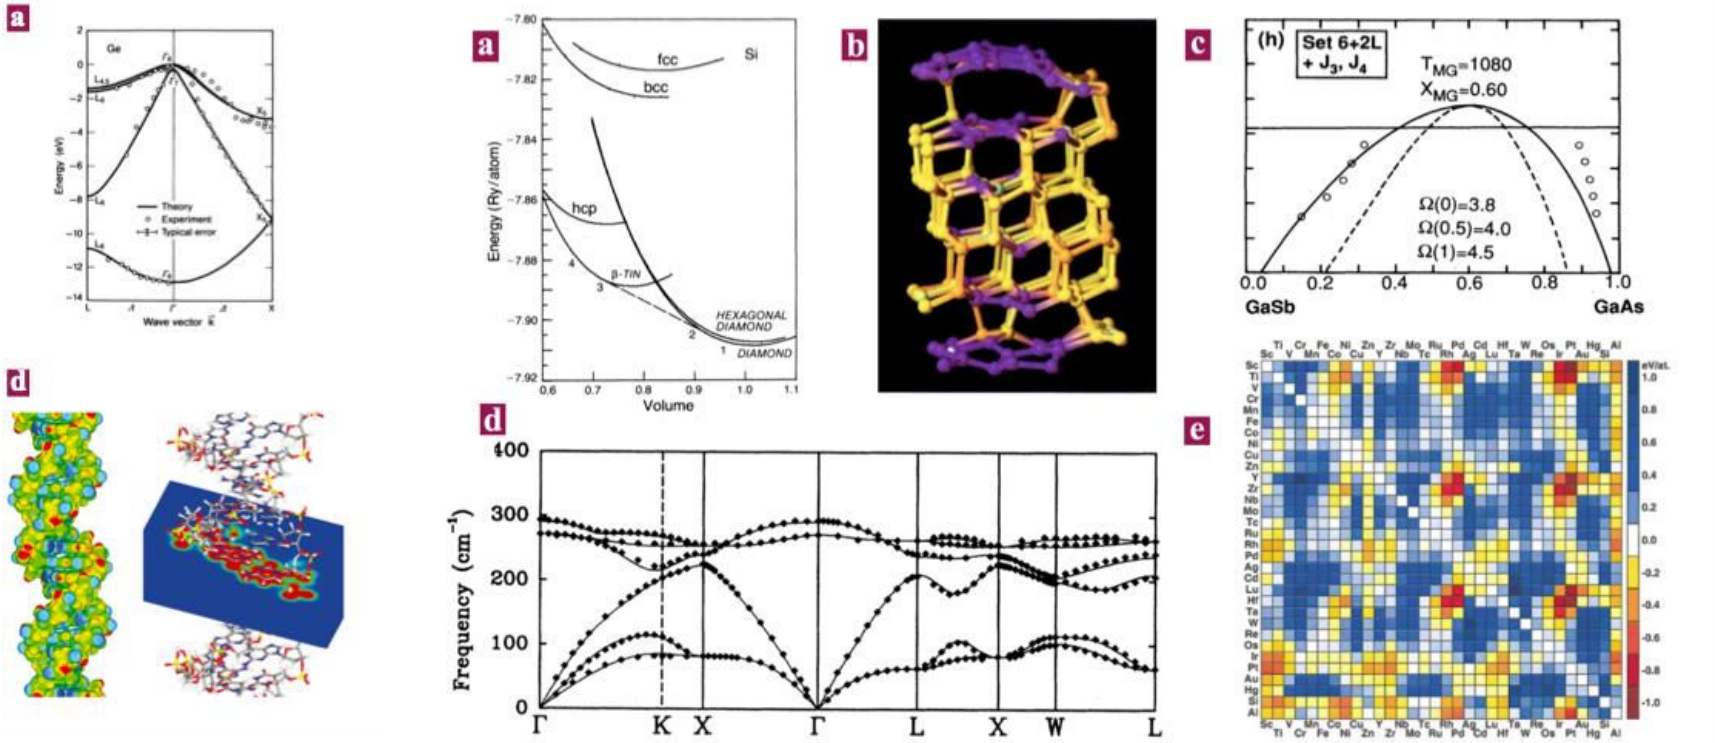

materials modelling

quantum mechanics based
atomistic modelling of materials
+
interfacing with **multiscale** approaches

Electronic Structure Methods
- highly accurate (predictive)
- computationally demanding
- **a case for HPC**

the **exascale** opportunity:

MAX

complexity

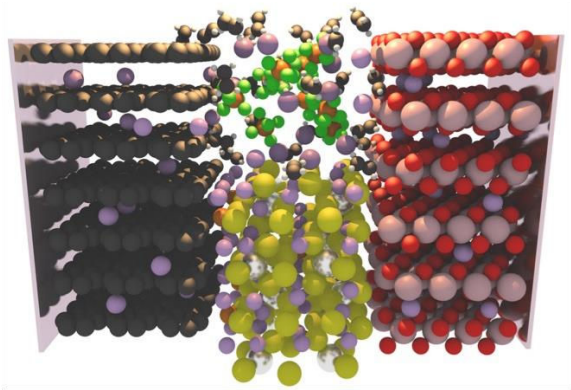

high-throughput screening

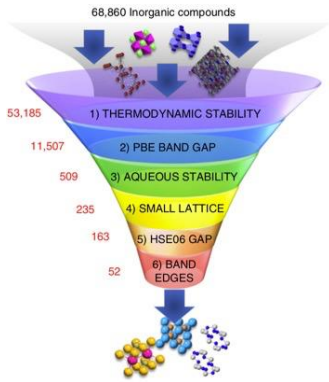

accuracy &
adv. properties

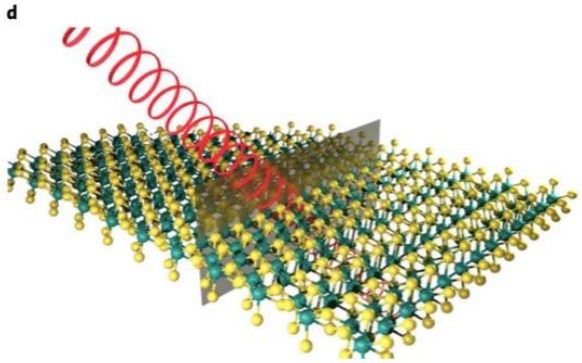

electronic structure

Density Functional Theory (DFT)

- **Density functional theory** (DFT) **widespread** well beyond its original electronic structure core-community
- applications ranging from **materials modelling, to quantum chemistry and drug design** (with about 30000 papers/y, arXiv 1701.00873)
- compatible with **high performance computing** and **high-throughput screening**

moving beyond DFT:

- **different hierarchies** can be climbed, including:
- many-body perturbation theory
- time-dependent and non-equilibrium methods
- quantum chemistry & wave-function based approaches

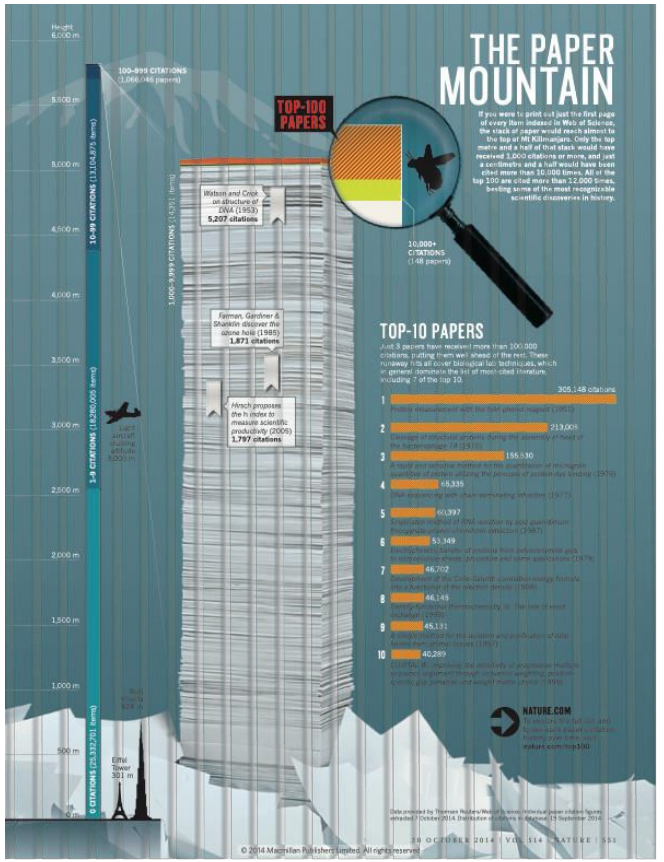

– van Noorden, Maher, Nuzzo, Nature **514**, 550 (2014)
– **2 papers about DFT in top10** (12 in top100) most cited papers over all disciplines

Density Functional Theory

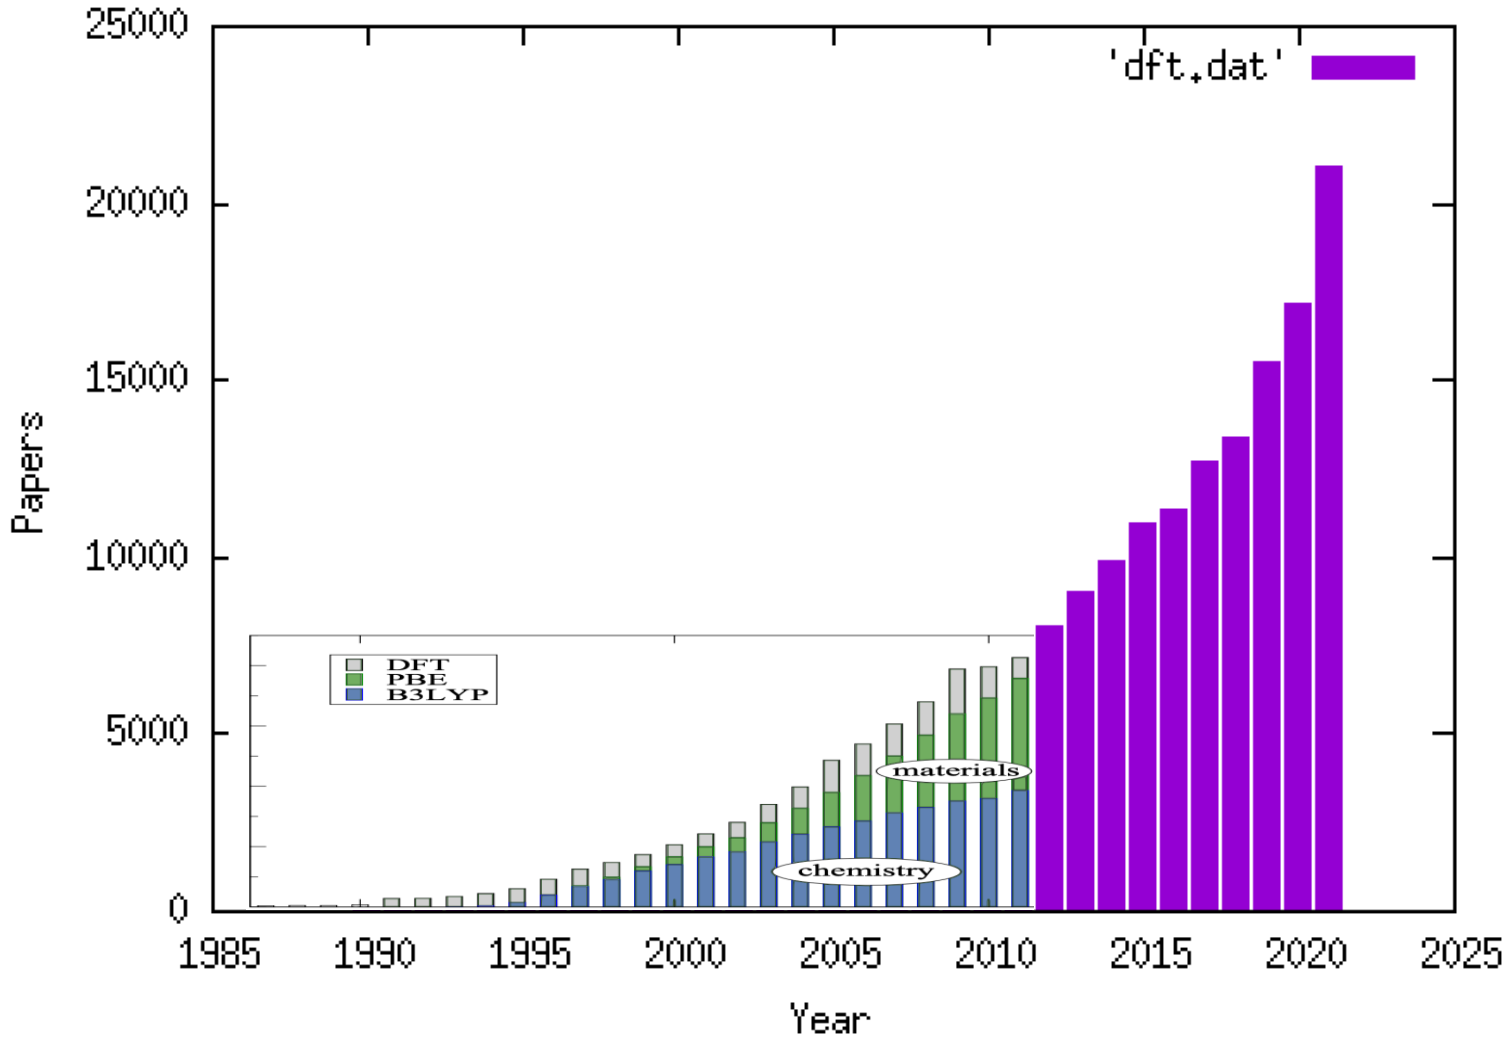

Perspective on Density Functional Theory
K. Burke, J. Chem. Phys. 136, 150901 (2012)

Density Functional Theory

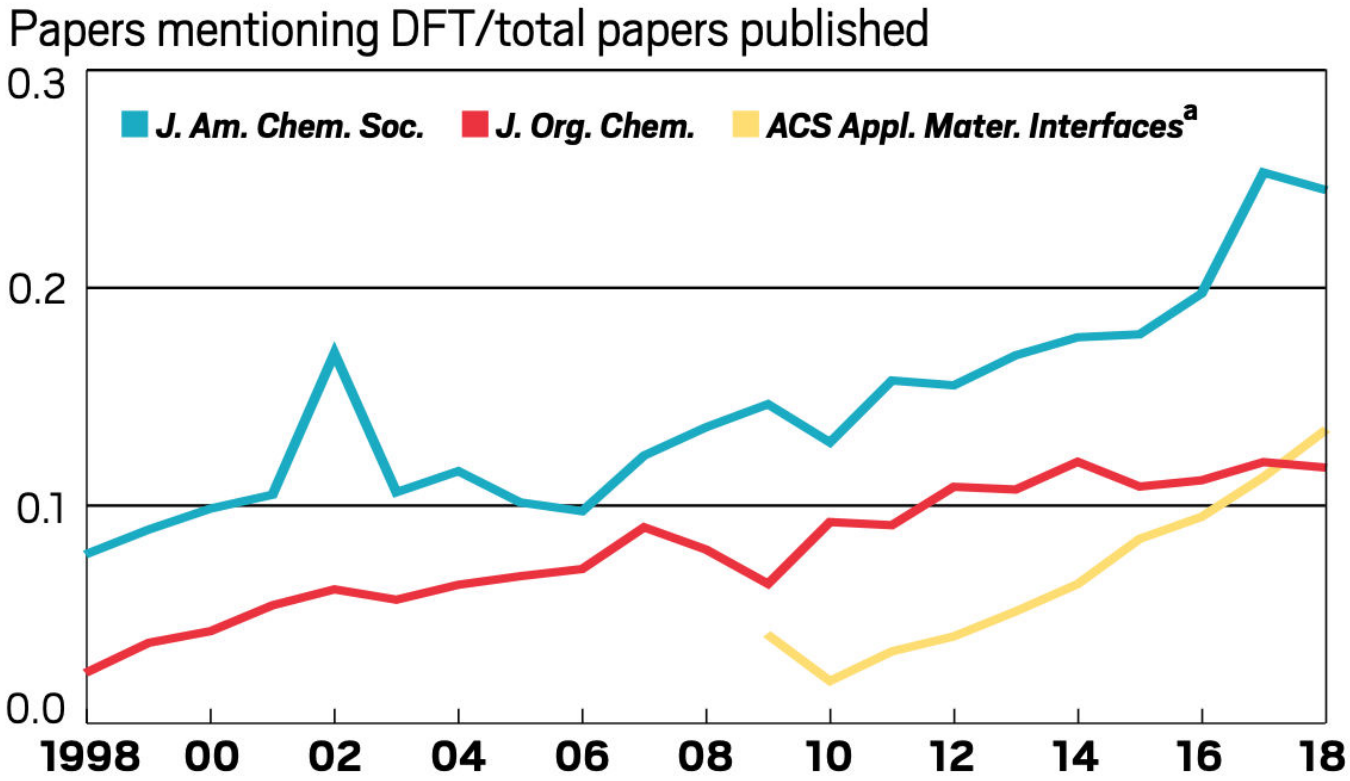

Sources: Data retrieved from the ACS Publications website on Aug. 29, 2019; Web of Science data retrieved on Sept. 3, 2019. **Note:** The data in this plot were normalized to account for any changes in the overall papers being published in these journals. **a** This journal launched in 2009.

S. Lemonick, C&EN, September 2019

The SIESTA code

SIESTA is a DFT code, density-functional theory

(like many others in many ways)

Aim from inception: EFFICIENCY - minimize time to solution

(independently of the computational platform)

Pioneer LINEAR-SCALING DFT code (or Order-N, O(N))

meaning: computational cost (CPU & memory) scaling linearly with number of atoms

The SIESTA code

DFT Shrödinger's equation for independent electrons

$$\nabla^2\psi(\mathbf{r}) + \frac{2m}{\hbar^2}\left[E - V(\mathbf{r})\right]\psi(\mathbf{r}) = 0$$

1. Choose a **basis set**

Plane Waves - APWs - LMTOs - Grids

Gaussians - Slaters

(Numerical) Atomic Orbitals with compact support

$$\psi_n(r) = \sum_\mu c_{n\mu}\phi_\mu(r)$$

$$\boxed{\phi_\mu(\vec{r}) = \varphi_\mu(|r|)\,Y_{lm}(\theta,\varphi)}$$

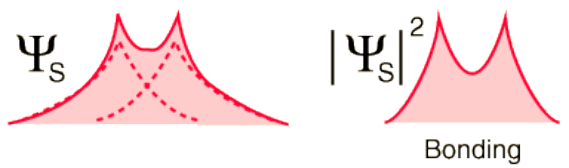

DFT Shrödinger's equation for independent electrons

$$\nabla^2\psi(\mathbf{r}) + \frac{2m}{\hbar^2}\left[E - V(\mathbf{r})\right]\psi(\mathbf{r}) = 0$$

2. Solve the Schödinger equation in the chosen basis set:

Building Hamiltonian matrix and **solving a generalized eigenvalue problem**

$$h_{\nu\mu} = \langle\phi_\nu|\hat{h}|\phi_\mu\rangle = \int d\vec{r}\ \phi_\nu^*(\vec{r})\hat{h}\phi_\mu(\vec{r})$$

O(N^3) commonly in DFT
O(N) in SIESTA

$$\hat{\boldsymbol{h}}_{\mu\nu}\boldsymbol{c}_{n\mu} = \varepsilon_n\,\hat{\boldsymbol{s}}_{\mu\nu}\,\boldsymbol{c}_{n\mu}$$

O(N^3) in general
O(N^3) to O(N) in SIESTA

The Numerical Atomic Orbitals Basis Sets

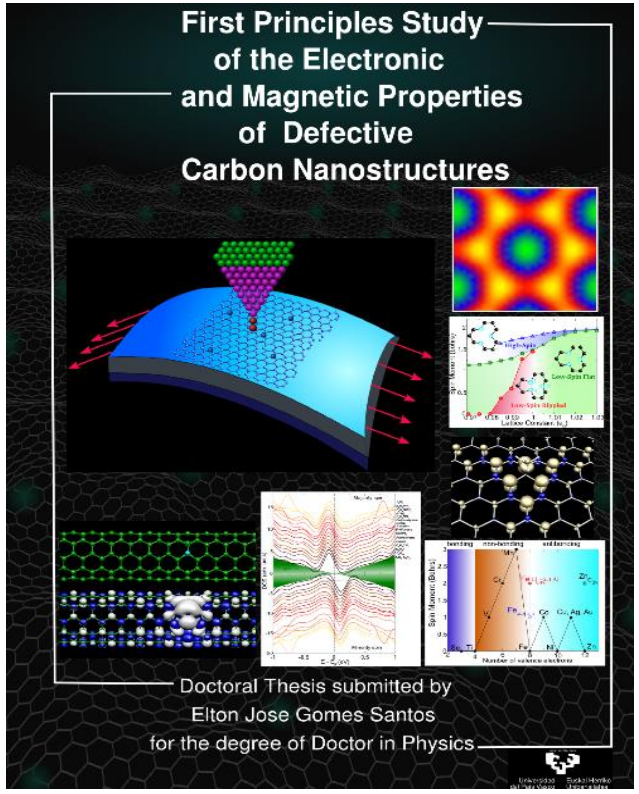

- DZP Basis
- E_{shift} = 50 meV
- r_c of TM increased to obtain converged E_B (variationally)

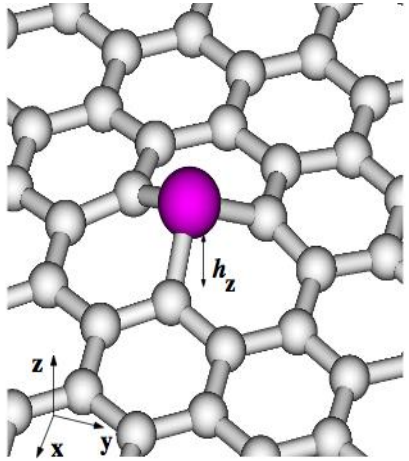

	r_{sp}^{NAO} (a_o)	r_d^{NAO} (a_o)
Co	8.00	4.73
Ni	10.94	6.81
Cu	8.87	5.52
Ag	10.48	6.52
Au	8.63	6.08
Zn	9.24	5.33

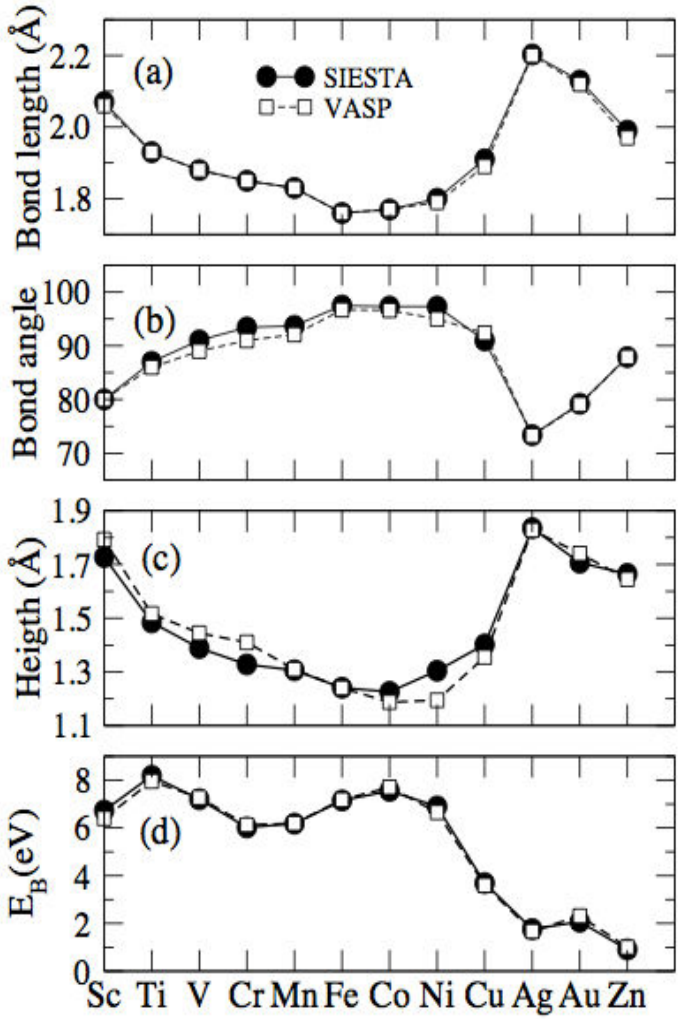

H-bonding in DNA base-pairs

Set of 30 DNA H-bonded basepairs

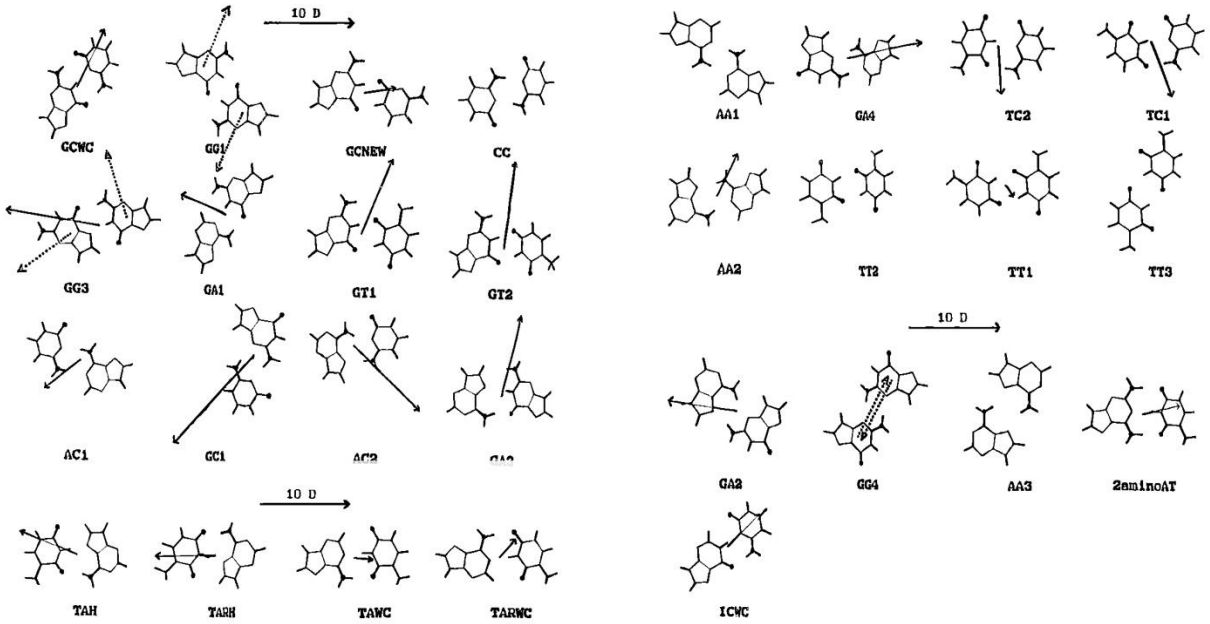

DFT with SIESTA vs MP2 (quantum chemistry)

Artacho et al. Mol. Phys. **101**, 1587 (2003)

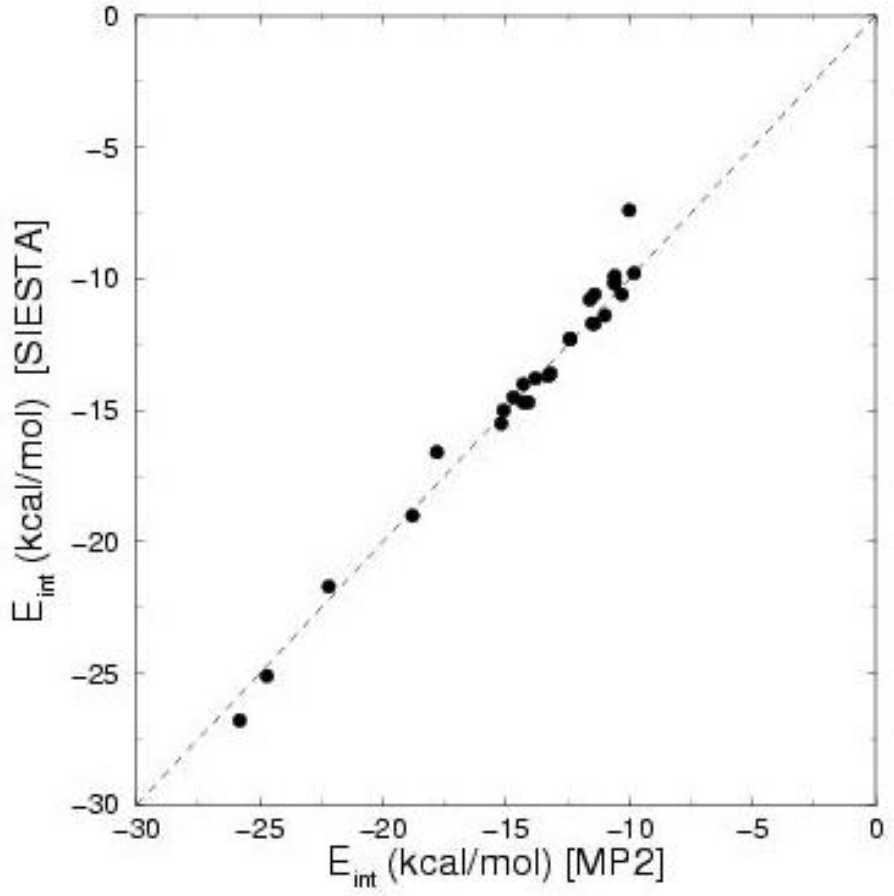

MP2 calculations: J. Šponer *et al*, J. Phys. Chem. **100**, 1965 (1996)

SIESTA: Enabling simulations of large systems

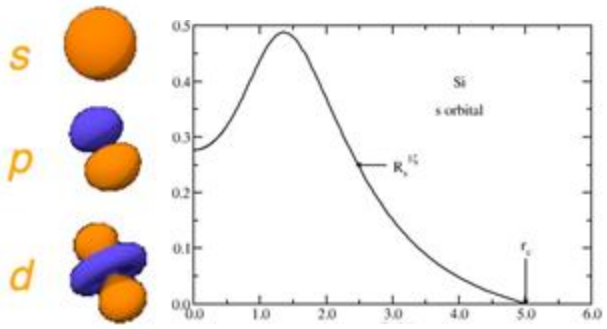

A DFT code using pseudopotentials and **finite-support atomic orbitals** as basis set.

- Much fewer basis elements than plane-wave codes
- Hamiltonian and overlap matrices are **sparse**
- **High efficiency for large systems, and small memory footprint.**
- Parallelized with MPI+OpenMP. GPU enabled

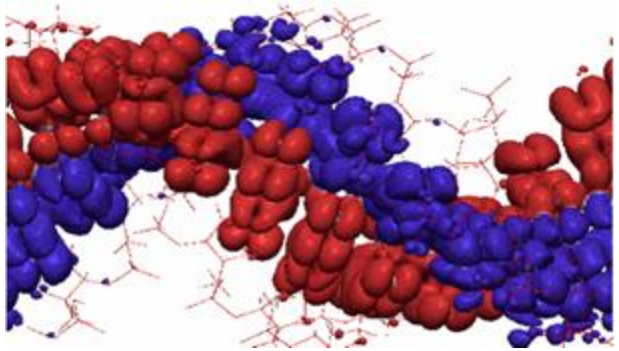

~800 atoms on a workstation (**1995**)

Used by **hundreds of groups worldwide** in various disciplines (**> 15000 citations** of the relevant papers)

- Electronic structure of molecules, solids, nanostructures, ...
- Total energy, forces, and stresses
- Molecular dynamics.
- **Very intuitive analysis tools**

- Electronic transport and electrochemistry with non-equilibrium Green's functions (**TranSiesta**)
- TD-DFT
- QM/MM

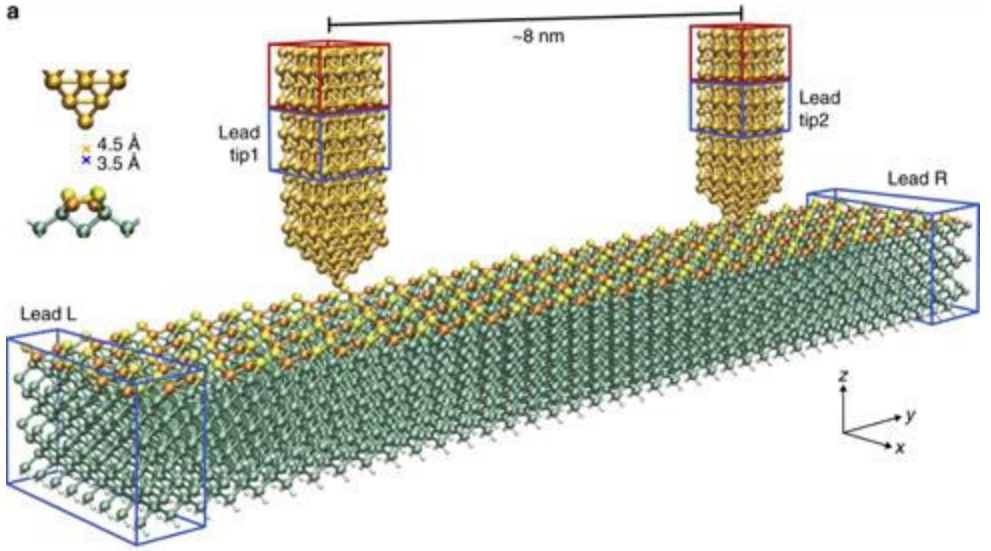

SIESTA development for HPC and scalability

Changing paradigm

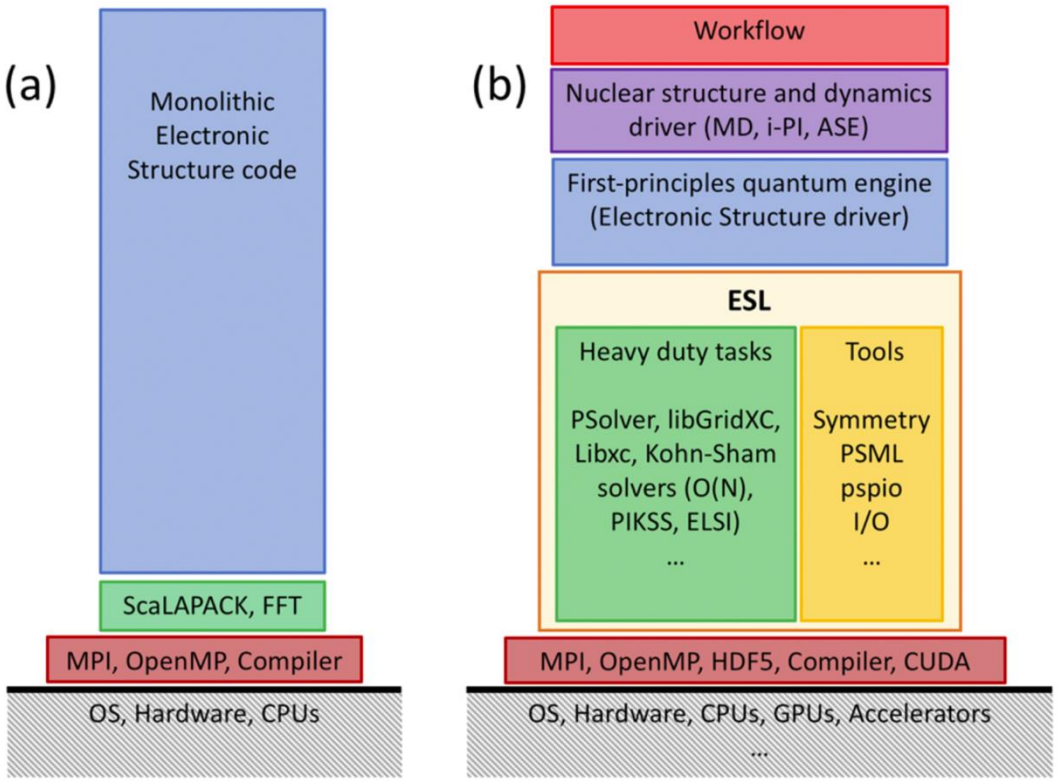

SIESTA development for HPC and scalability

ELSI (NSF-USA)

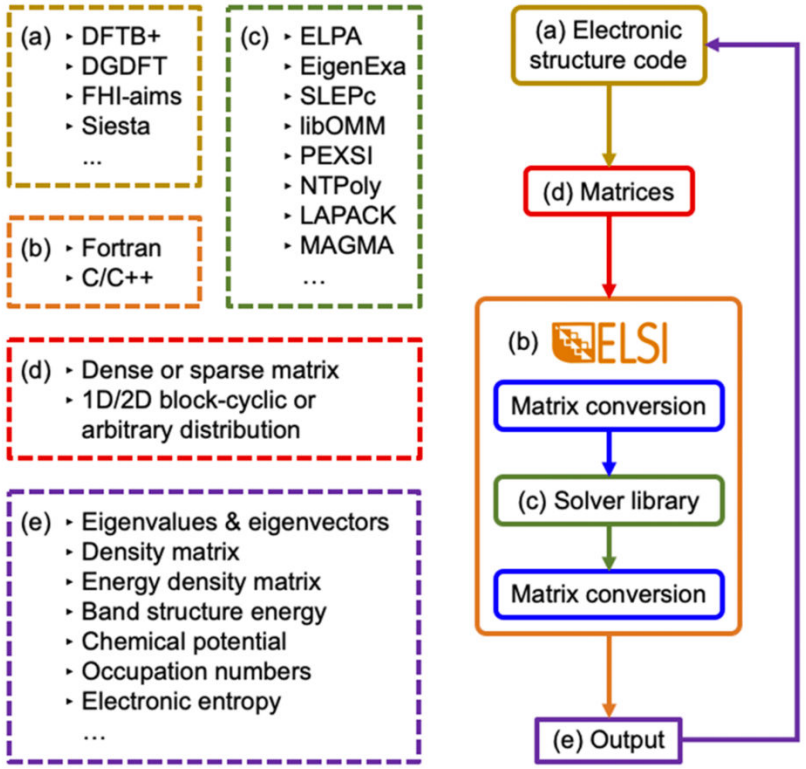

Electronic Structure Library (ESL) - CECAM

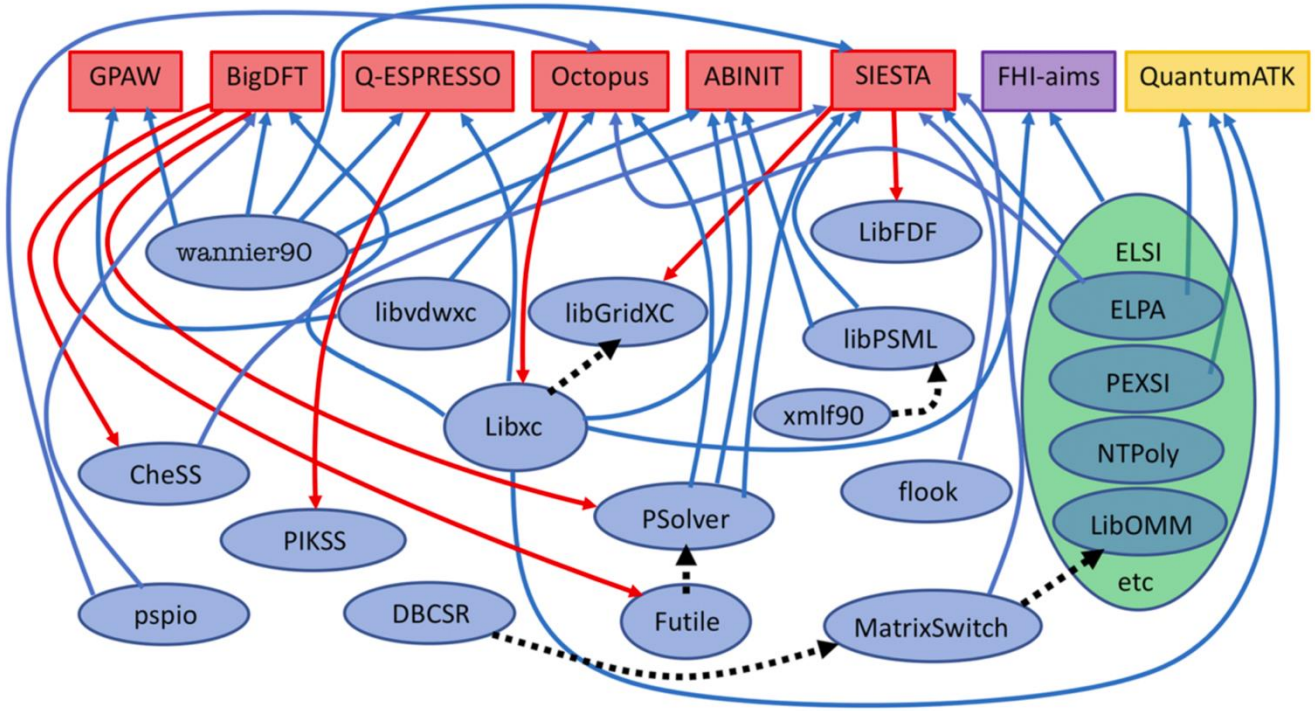

MJT Oliveira et al, JCP 2020

MAX

cam

MAX DRIVING THE EXASCALE TRANSITION

ESL BUNDLE FOR SIESTA

Consistent set of SIESTA dependencies to facilitate the deployment of all the involved packages, based on the ESL Bundle.

A significant effort of abstraction and interface building

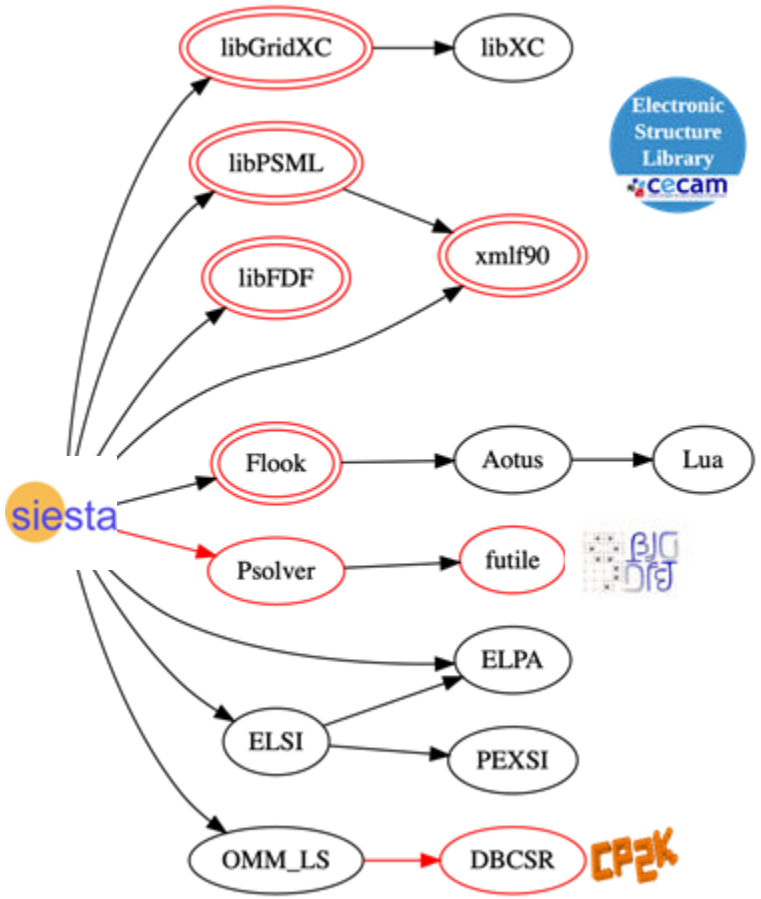

Domain-specific libraries originating in SIESTA itself or created for new functionalities in the code:

- libGridXC: laboratory for interface design
- xmlf90: used already by other community codes
- libPSML: enables pseudopotential interoperability

Scriptability via embedded interpreter with access to data structures of the code

New Poisson solver with flexible boundary conditions and optimized for hybrid architectures

Solvers: consolidated interfaces for continuously improved and performance-portable libraries.

New class of linear-scaling algorithms with efficient sparse-matrix library DBCSR as backend

Diagonalization-based solvers

In SIESTA, The **solver** stage typically takes 90-95% of the cpu time.

Need to use DIRECT solvers, as the number of desired eigenvectors is a substantial fraction of the matrix size

$$\sum_{\nu\beta}(H^{\alpha\beta}_{\mu\nu} - E_i S_{\mu\nu}\delta^{\alpha\beta})c^{\beta}_{\nu i} = 0$$

SIESTA uses pre-packaged libraries for this pure math problem:

- ScaLaPACK
  - pdsyev, pzheev and related drivers
  - MRRR
- ELPA: Alternative transformation sequence + optimizations
  https://elpa.mpcdf.mpg.de/

- Conversion of H and S to dense form
- Cholesky decomposition to reduce to standard eigenproblem
- Transformation to tri-diagonal form
- Solution of tri-diagonal problem
- Back-transformation

Cubic scaling with matrix size — Quadratic scaling for memory

Still competitive for low-cardinality basis sets

Solver strategies for performance and features: Use external libraries

In SIESTA, The **solver** stage typically takes 90-95% of the cpu time.

ELSI initiative to integrate solver libraries

https://elsi-interchange.org

Volker Blum, Duke

Lin Lin, Berkeley

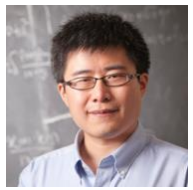

Jiangfen Lu, Duke

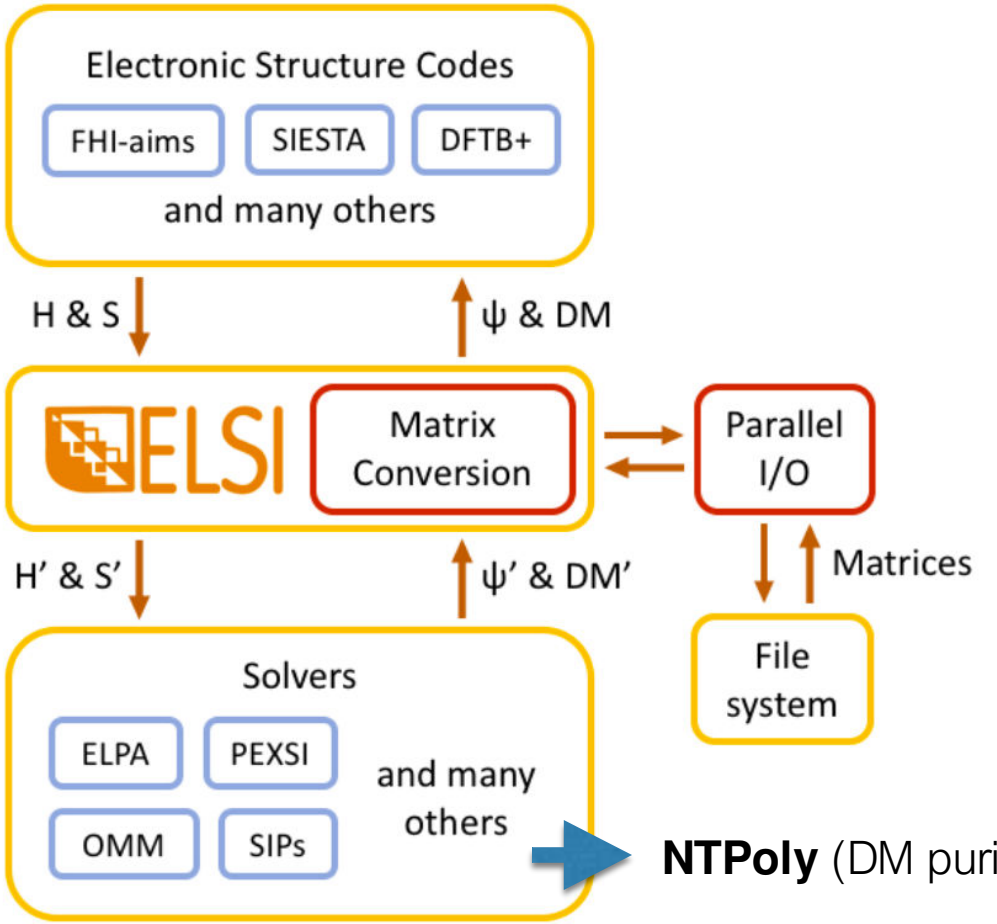

NTPoly (DM purification, O(N))

Interface in Siesta:
Alberto García (ICMAB)

Collaboration with
Victor Yu (Duke)

SIESTA: Strategic solver libraries

In SIESTA, The **solver** stage typically takes 90-95% of the cpu time.

Separation of concerns: Use of high-performing libraries, which can be ported to (pre)-exascale architectures:

- **ELPA** distributed diagonalization
- **ELSI** (includes ELPA, PEXSI, and other; in active development)

Support for new architectures: e.g. AMD GPUs

GPU acceleration with ELSI-ELPA in Marconi-100

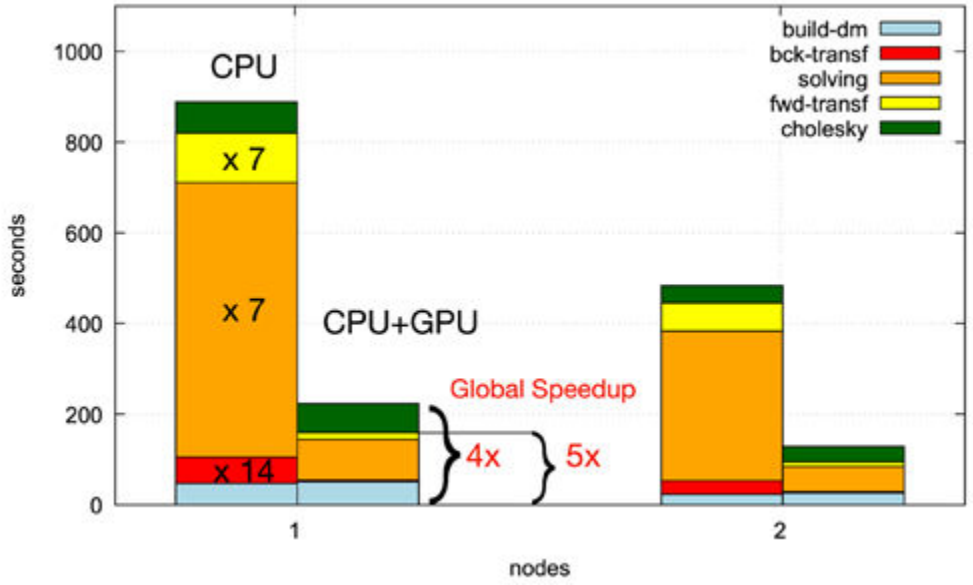

Performance and scalability for sars-cov-2 protein (8800 atoms)

Marconi-100 (CINECA): 32 CPUs+ 4 GPUs /node

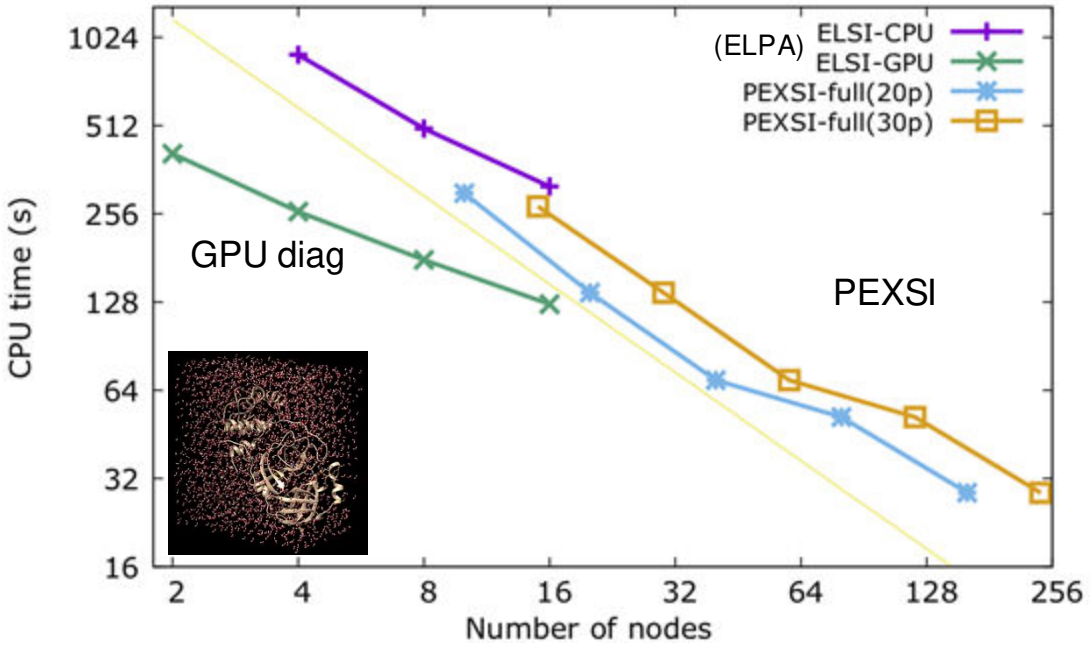

PEXSI: Multiple levels of parallelism: excellent scalability

For sufficiently big problems

(quasi-)1D:	$\mathcal{O}(N)$
(quasi-)2D:	$\mathcal{O}(N^{3/2})$
3D:	$\mathcal{O}(N^2)$

$$\hat{\rho} = Im\left(\sum_{l=1}^{P} \frac{\omega_l}{H-(z_l+\mu)S}\right)$$

MAX DRIVING THE EXASCALE TRANSITION

Very Large materials problems

DNA strand – 25 base pairs - 17,875 atoms
Label: DNA-25

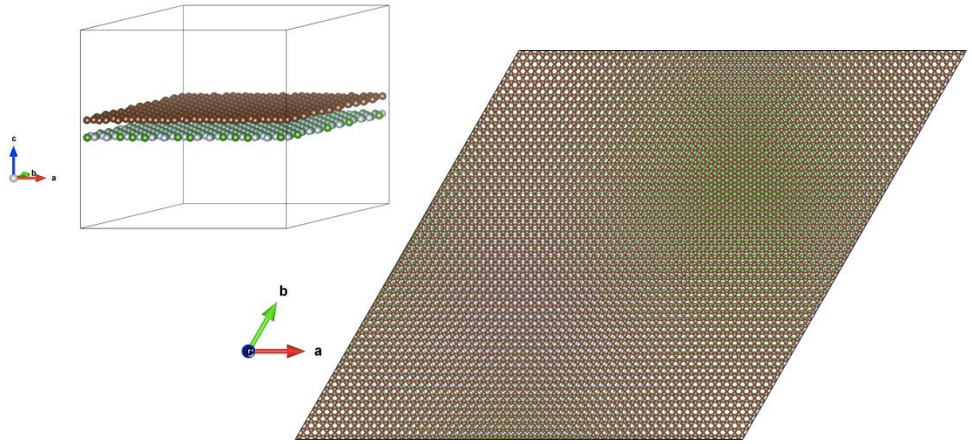

Graphene / BN (Moire pattern) - 12,770 atoms
Label: C-BN$_{0.00}$

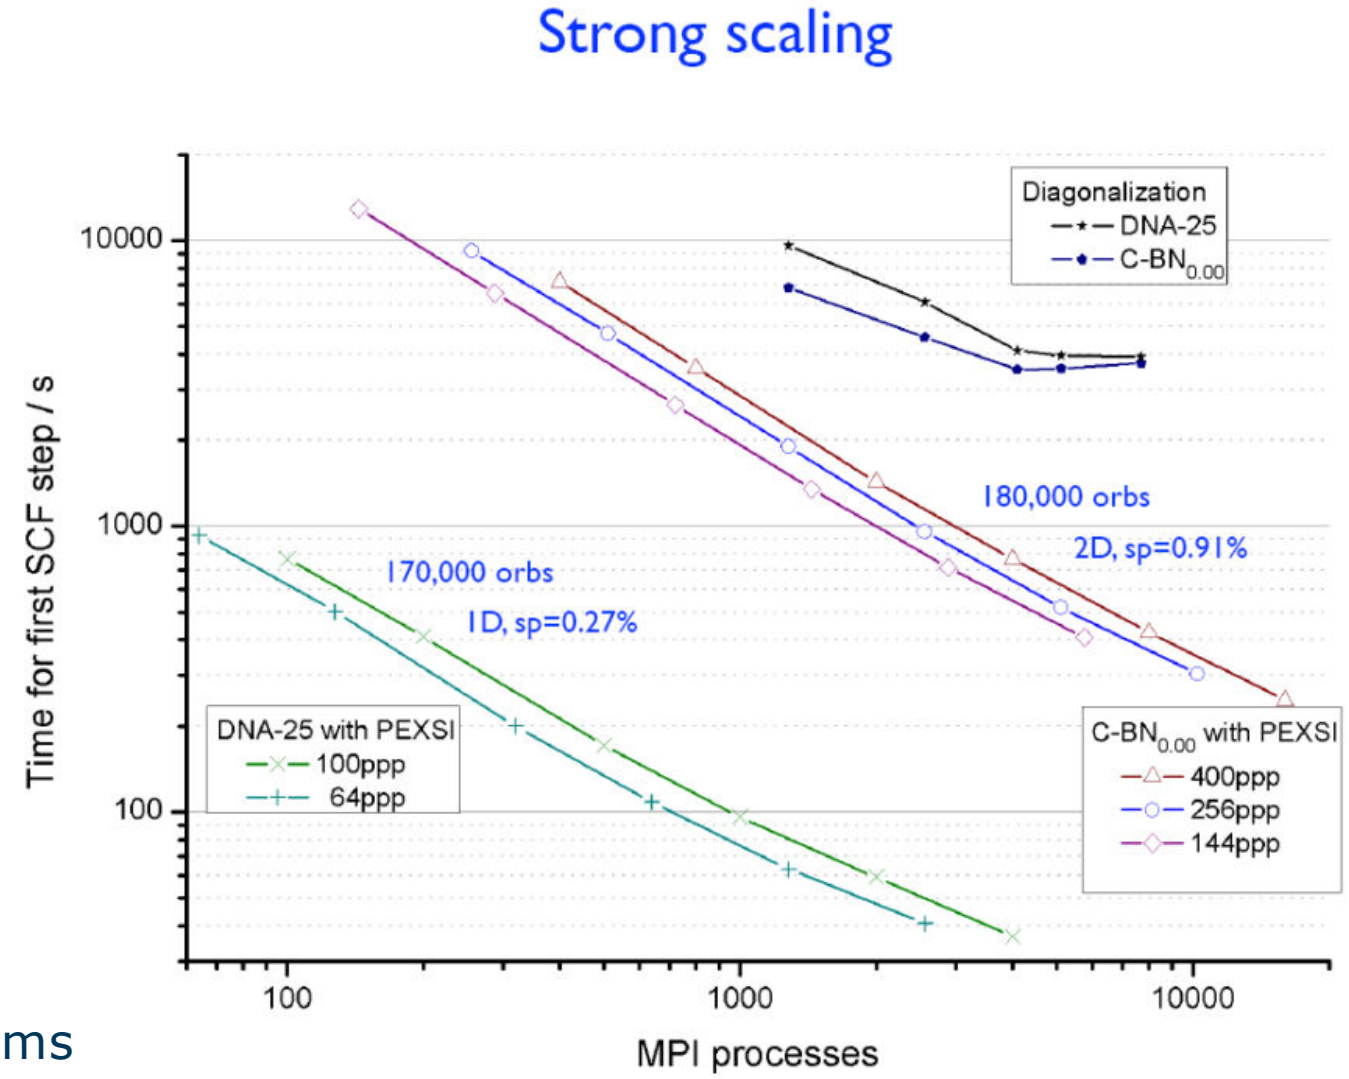

A. Garcia et al, JCP 2020

SIESTA development for HPC and scalability

SIESTA (diagonalisation & PEXSI)

DNA (1D)

Graphene – BN stack (2D)

sp: sparsity of Hamiltonian

ppp: Processes per pole

PEXSI: Pole EXpansion and Selected Inversion method
(Green's functions' based)

Diagonalisation: ScaLAPACK

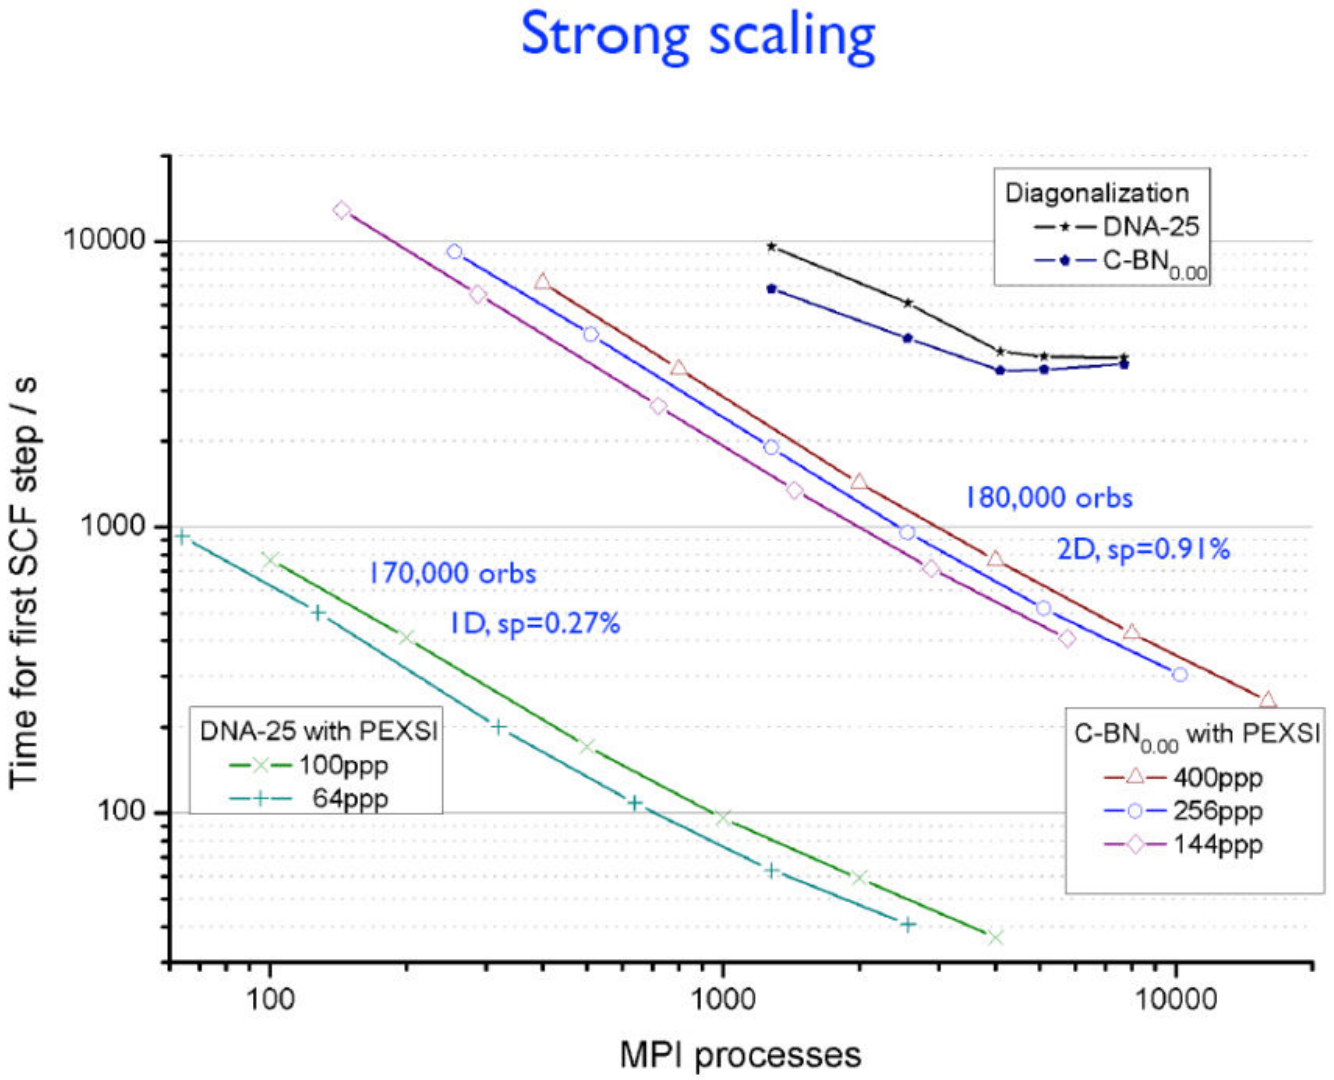

A. Garcia et al, JCP 2020

TranSIESTA – Open problems under an external bias (electronic transport, electrochemistry...)

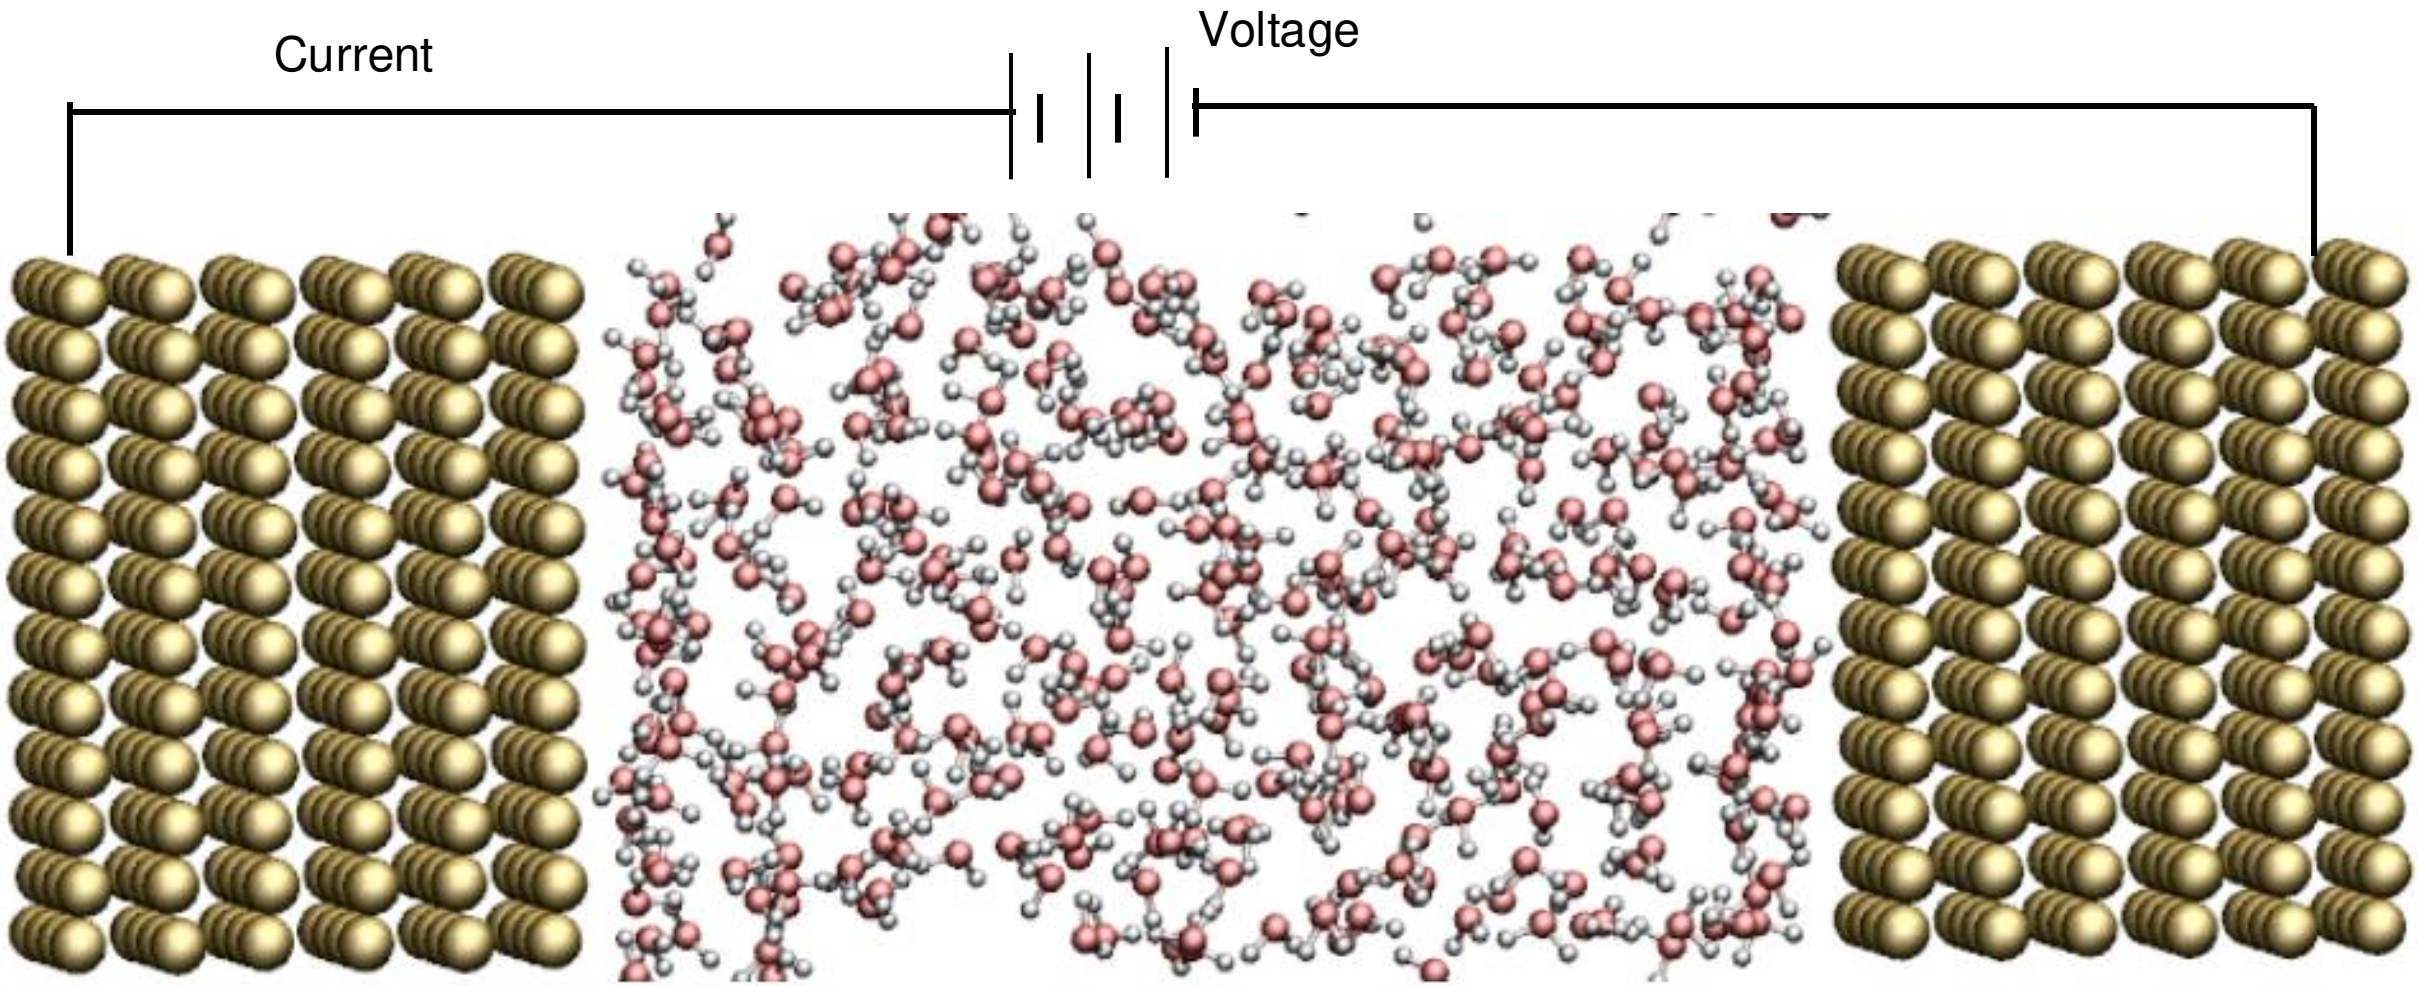

TranSIESTA – Open problems under an external bias (electronic transport, electrochemistry...)

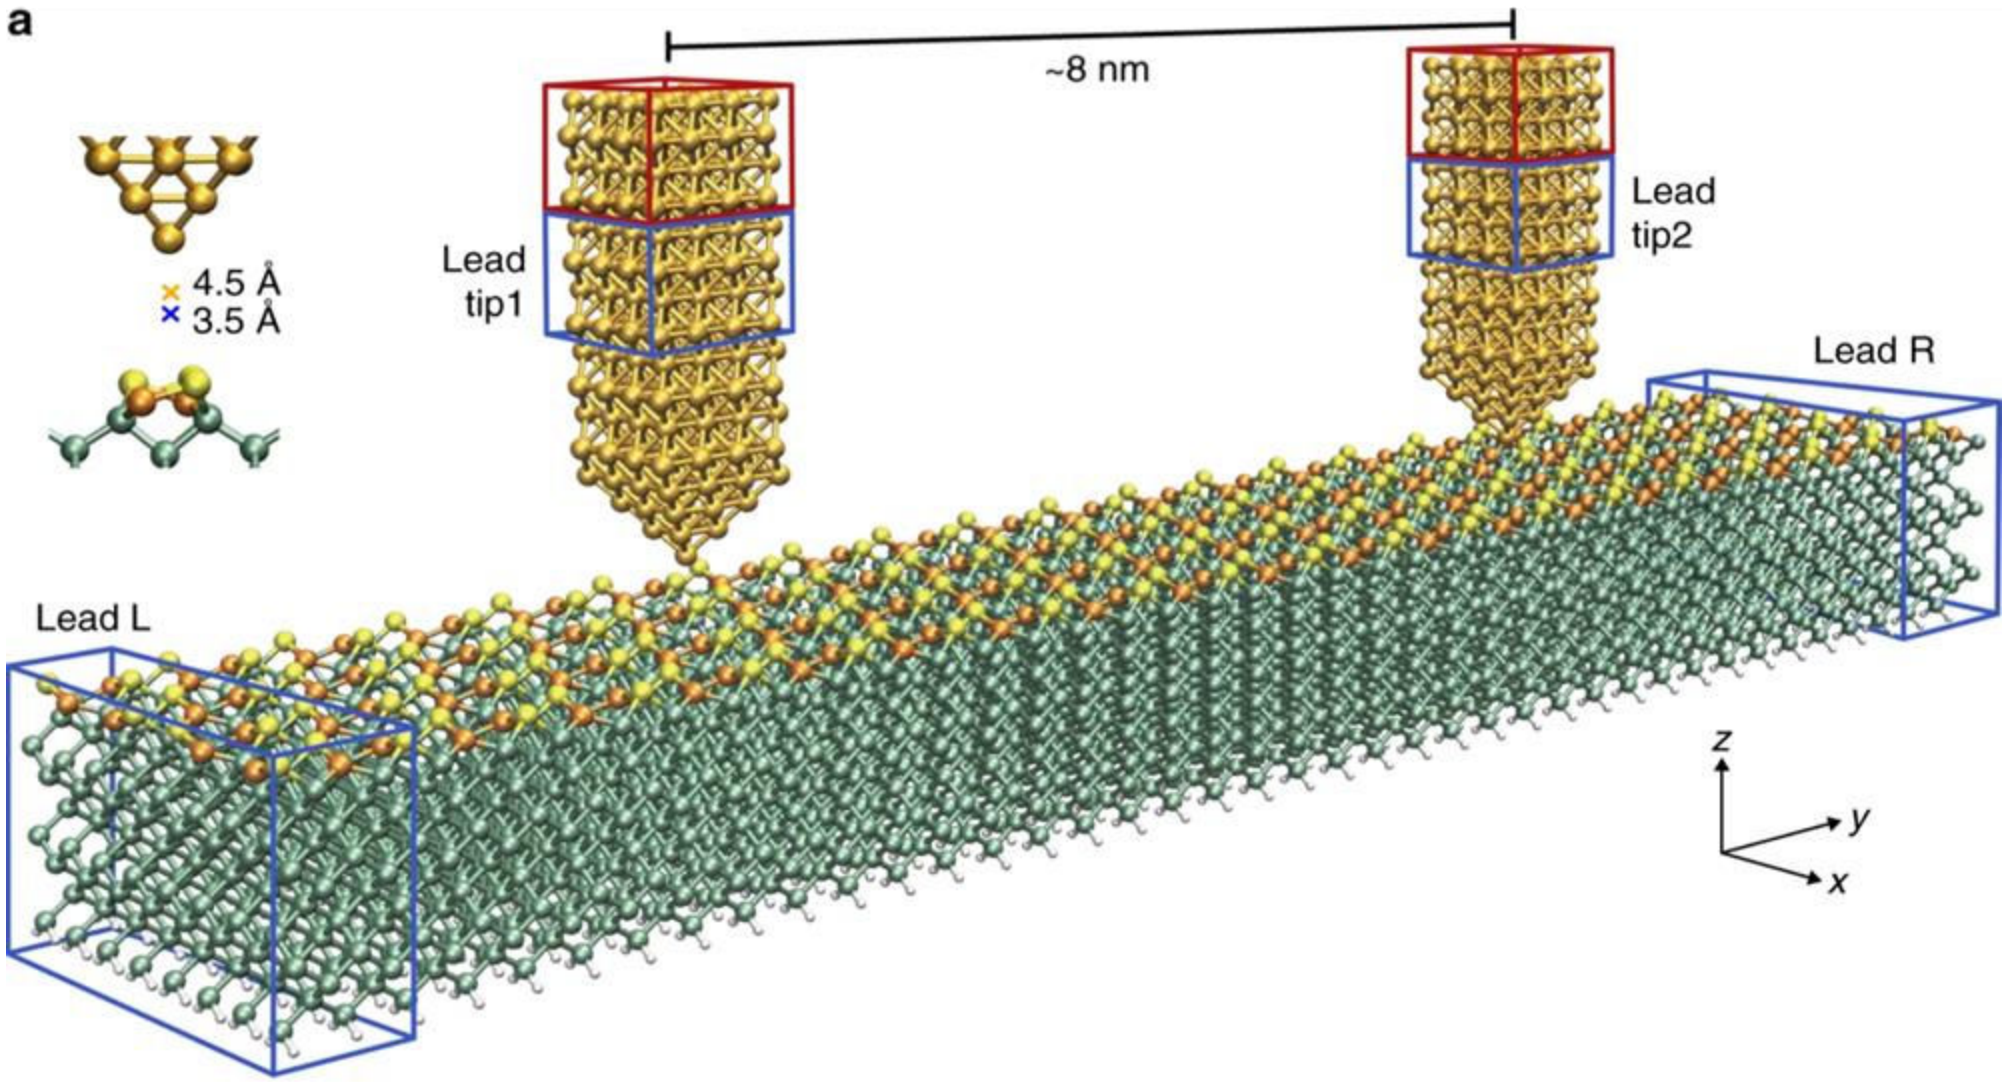

TranSIESTA – Open problems under an external bias (electronic transport, electrochemistry...)

Hybrid QM/MM simulations: electrolyte is MM, electrode is QM

$Au – (H_2O)_{236 \times n} – Au$ from n=1 to 10 Scaling: From 236 to 2360 H_2O molecules

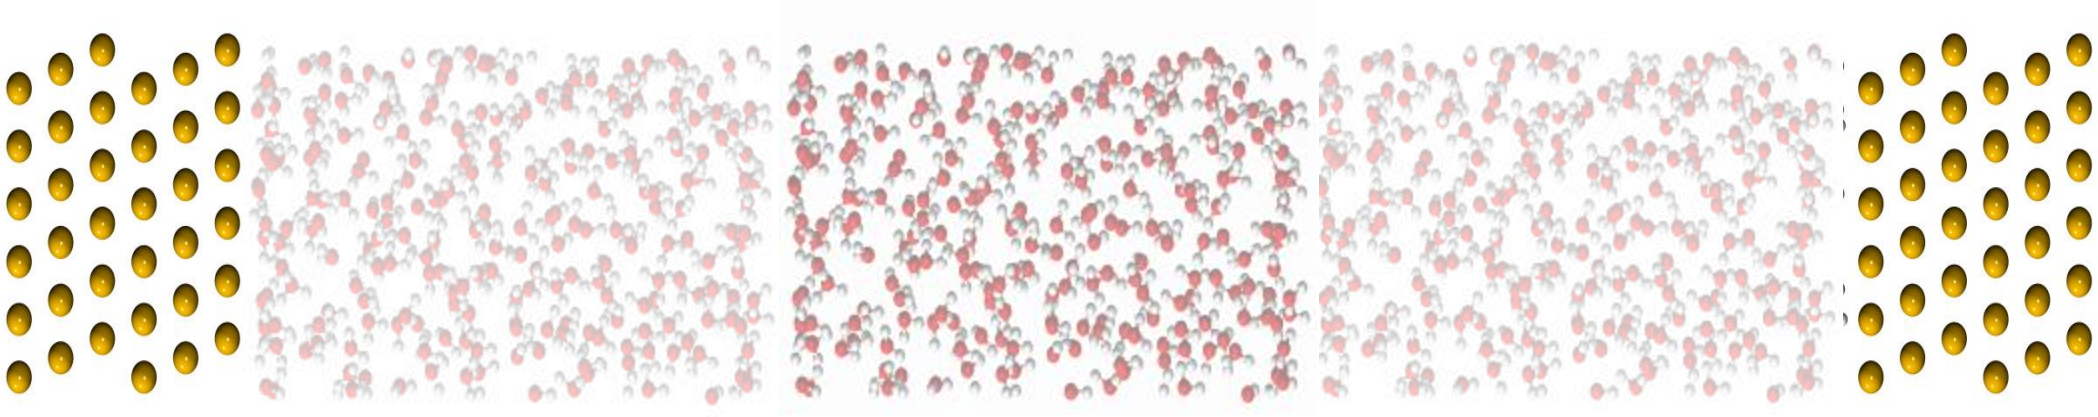

MareNostrum IV @ BSC
Intel Platinum 8160 @ 2.1 GHz

384 cores
96 MPI processes
4 openMP threads / MPI process

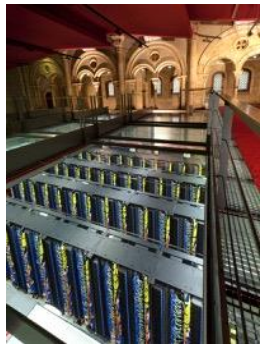

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

MAX DRIVING THE EXASCALE TRANSITION

Scaling vs system size

Scaling: From 236 to 2360 H_2O molecules

$Au - (H_2O)_{236 \times n} - Au$

Wall time for 1 MD step

In 384 cores, for n=1:
QM/MM: 1ps / day
QM (NEGF): 0.1ps / day

QM vs QM/MM

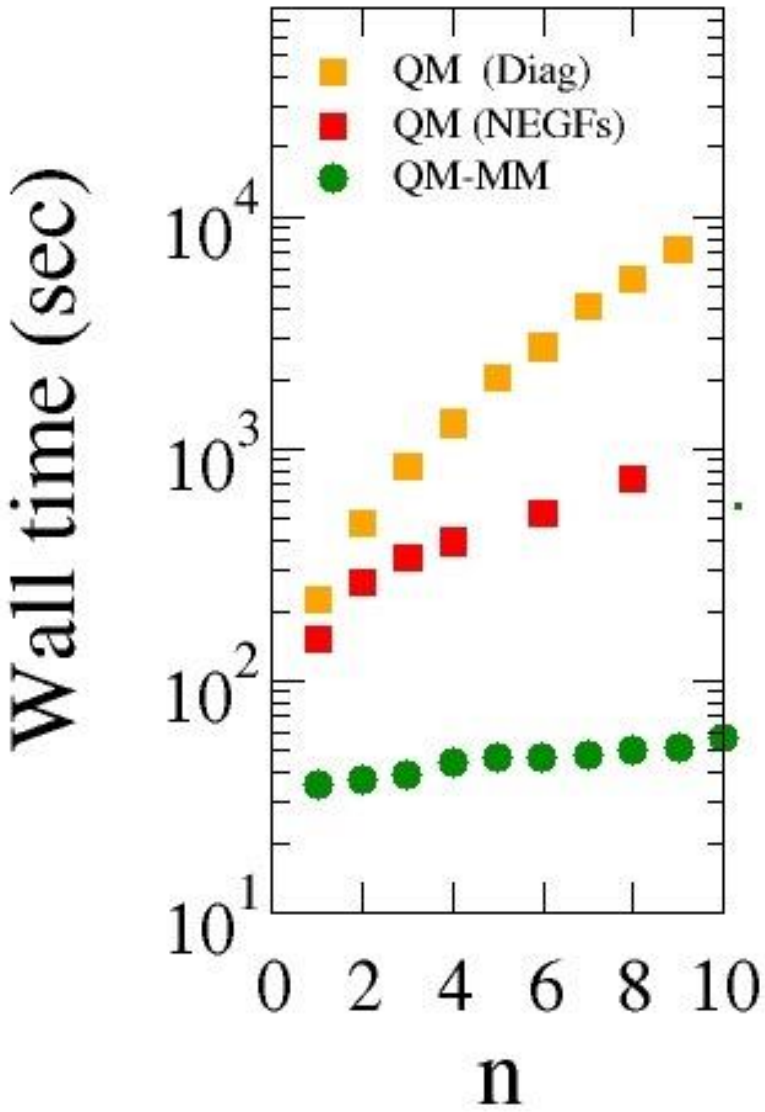

Multi-level parallelization in TranSIESTA (in development)

Old parallelization:

- over energy points (MPI)
- LAPACK (openMP)

New parallelization:

- over energy points (MPI)
- over k points (MPI)
- over matrix blocks (ScaLAPACK - MPI)

Different levels can be combined arbitrarily.

Hugely increased scalability. *Theoretical limit:*

*#energy points * #k points * #ScaLAPACKProcs*

In practical problems, this means about 5000 processors

System: Water between two Au slabs:
576 Au + 160 H_2O = 1,056 atoms (**12,320 orbitals**)
Parallel: over energy & matrix blocks

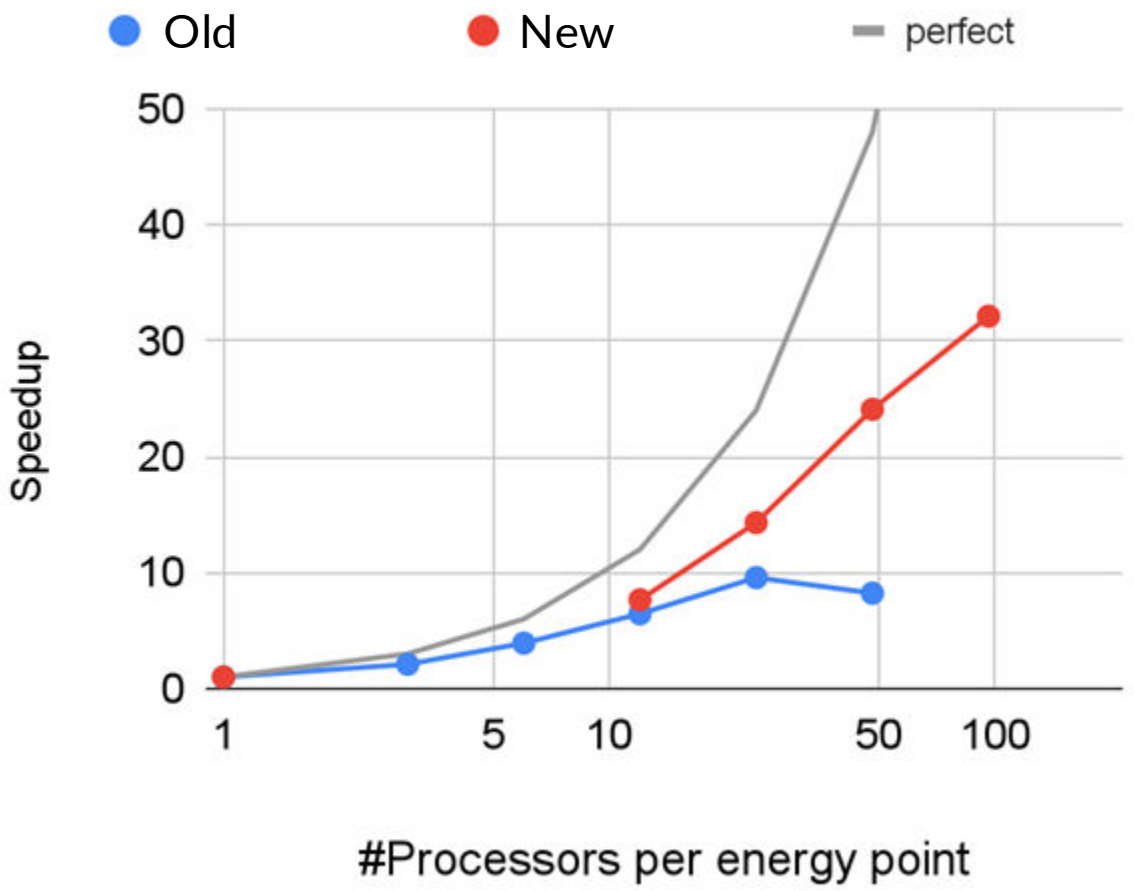

Deployment of MaX codes in EuroHPC supercomputers

Refined and more robust **CMake-based building** framework

- Deployment experience on many different machines.
- CMake enables good deployment on newer platforms.

This framework enables the automatic download of several core and optional dependencies (**ELSI**, **DFT-D3**, **Wannier90**, and internal libraries).

GPU support can be enabled if SIESTA is compiled against and existing **ELPA** installation (automatically detected if present).

This new CMake framework paved the way to provide the following options for installation in all EuroHPC centers:

- **Spack** recipes.
- **EasyBuild** recipes (upcoming in the next few months).

We also provide binaries via **Conda** for general usage.

Deployment of MaX codes in EuroHPC supercomputers

SIESTA 5.2 has been successfully compiled on **all non-ARM EuroHPC machines**, including **GPU support**.

GNU, Intel and Cray compilers are supported out of the box.

ARM support is underway and we expect to have it for **2025**.

Deployment of MaX codes in EuroHPC supercomputers

EuroHPC Machines		Architecture Type		Quantum ESPRESSO	Yambo	SIESTA	BigDFT	FLEUR
		CPU	GPU					
Deucalion	ARM	A64FX (ARM)		S	D	W	S	S
	CPU	AMD EPYC (x86)		S	S	D+B	S	S
	GPU	AMD EPYC (x86)	NVIDIA A100	S	S	D+B	S	S
Discoverer	CPU	AMD EPYC (x86)		M	S	D+B	S	S
Karolina	CPU	AMD EPYC (x86)		M	M	D+B	D	D
	GPU	AMD EPYC (x86)	NVIDIA A100	M	M	D+B	S	S
Jupiter	Cluster	SiPearl Rhea1 (ARM)						
	Booster	NVIDIA GH200 (ARM+H100)						
Leonardo	DCGP	INTEL SapphireRapids (x86)		M+B	M+B	M+B	M	M
	Booster	INTEL IceLake (x86)	NVIDIA A100+	M+B	M+B	M+B	M+B	M+B
LUMI	LUMI-C	AMD EPYC (x86)		M+B	M	D+B	S	S
	LUMI-G	AMD EPYC (x86)	AMD MI250X	M+B	D+B	D	D	W
MareNostrum5	GPP	INTEL SapphireRapids (x86)		M +B	D+B	M+B		D
	ACC	INTEL SapphireRapids (x86)	NVIDIA H100	M +B	D+B	D+B		D
	NGT GPP	NVIDIA Grace (ARM)						
	~~NGT ACC~~	~~INTEL EmeraldRapids (x86)~~	~~INTEL Rialto Bridge~~					
Meluxina	CPU	AMD EPYC (x86)		D	D	D+B	S	S
	GPU	AMD EPYC (x86)	NVIDIA A100	M	M	D+B	S	S
Vega	CPU	AMD EPYC (x86)		M	M	D+B	S	S
	GPU	AMD EPYC (x86)	NVIDIA A100	M	M	D(5)+B, M(4)	M	S

B	Benchmark
M	Deployed (module available)
D	Demonstrated (developers installation)
S	Supported Architecture (tested on similar architectures)
W	Work in Progress

MAX DRIVING THE EXASCALE TRANSITION

Where to find us:

- SIESTA project web: https://siesta-project.org/siesta/
- To get support, head to our Discord server: https://discord.gg/AqjX6aTNXR
- Tutorials and online documentation:

  https://docs.siesta-project.org/projects/siesta/en/stable/
- To help develop SIESTA, visit our Gitlab: https://gitlab.com/siesta-project/siesta
- For more general questions, we are also active on Stack Exchange:

  https://mattermodeling.stackexchange.com

Co-funded by the European Union

MaX "Materials Design at the Exascale" has received funding from the European Union under grant agreement no. 101093374.

EuroHPC

The project is supported by the EuroHPC JU and its members.

MINISTERIO DE CIENCIA, INNOVACIÓN Y UNIVERSIDADES

Financiado por la Unión Europea NextGenerationEU

Plan de Recuperación, Transformación y Resiliencia

AGENCIA ESTATAL DE INVESTIGACIÓN

Funded by the grant no. **101093374** from the **EuroHPC JU-EU**, and the grant no**. PCI2022-134972-2**, funded by the Spanish MCIN/AEI/10.13039/501100011033 and by the European Union NextGenerationEU/PRTR.

MAX DRIVING THE EXASCALE TRANSITION

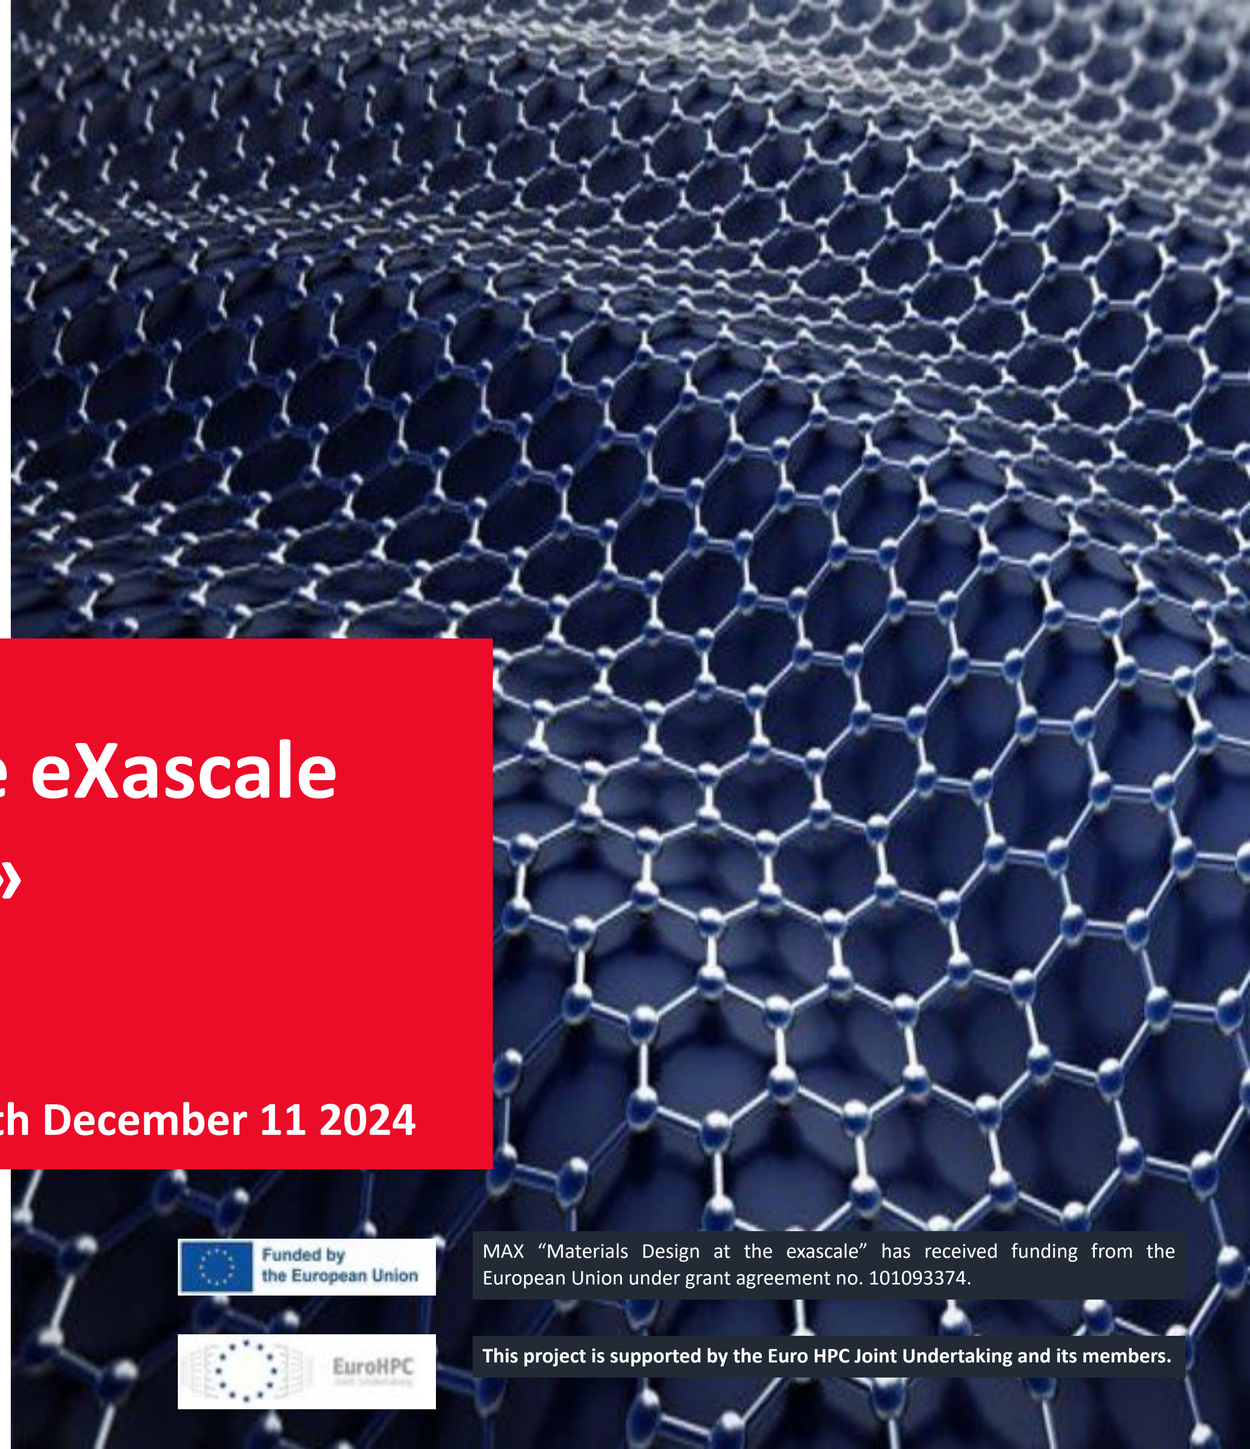

MaX – «Materials design at the eXascale European Centre of Excellence»

Daniele Varsano
S3 Centre, Istituto Nanoscienze CNR, Modena - Italy

Siesta: Code of the Month December 11 2024

Funded by the European Union

MAX "Materials Design at the exascale" has received funding from the European Union under grant agreement no. 101093374.

EuroHPC

This project is supported by the Euro HPC Joint Undertaking and its members.

Motivation: HPC at the exascale

the exascale challenge
in high performance computing

- 10^{18} Flops/s
- 10^{18} Bytes
- abrupt technology changes
- **action is needed** for full exploitation
- **heterogeneous** machines (multiple HW and SW stacks)

Jupiter: > 1 ExaFlops

MareNostrum V: Atos + NVIDIA H100 => 208 PFlops (estimated)

US DOE

El Capitan (@LLNL):
AMD EPYC+ AMD Mi300
=> 1742 PFlops

Switzerland

ALPS (@CSCS):
NVIDIA GH200
=> 435 PFlops

Leonardo: Atos + NVIDIA A100
(CUDA backend) => 239 PFlops

LUMI: CRAY + AMD cards
(ROCm, HIP) => 309 PFlops

ab initio materials modelling

quantum mechanics based
atomistic modelling of materials
+
interfacing with **multiscale** approaches

Electronic Structure Methods

- highly accurate (predictive)
- computationally demanding
- **a case for HPC**

the **exascale** opportunity:

Higher accuracy

High throughput screening

Improved modelling (complexity)

exascale opportunity: **complexity**

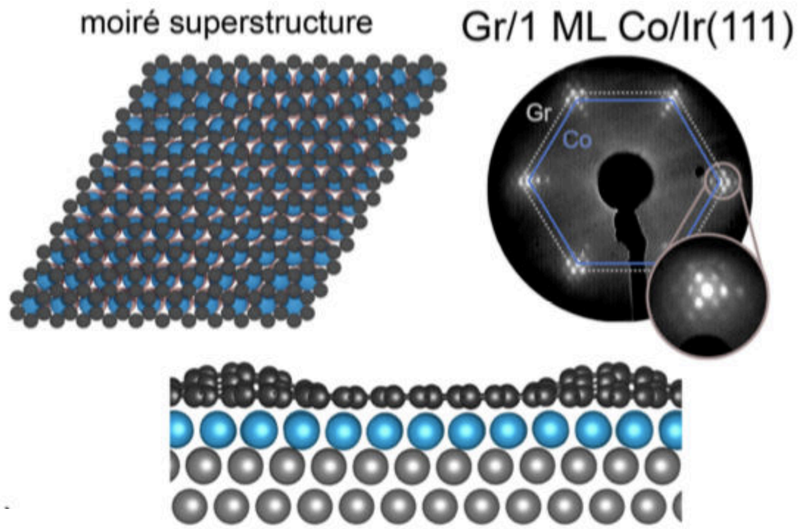

- Graphene / Transition Metal / Ir (111)
- **clear experimental evidence** for moire' pattern (**lattice mismatch**) and **Gr corrugation**
- 10x10 Graphene, 9x9 Iridium => 605 atoms / unit cell
- **Precise treatment of the structure** is important for modelling

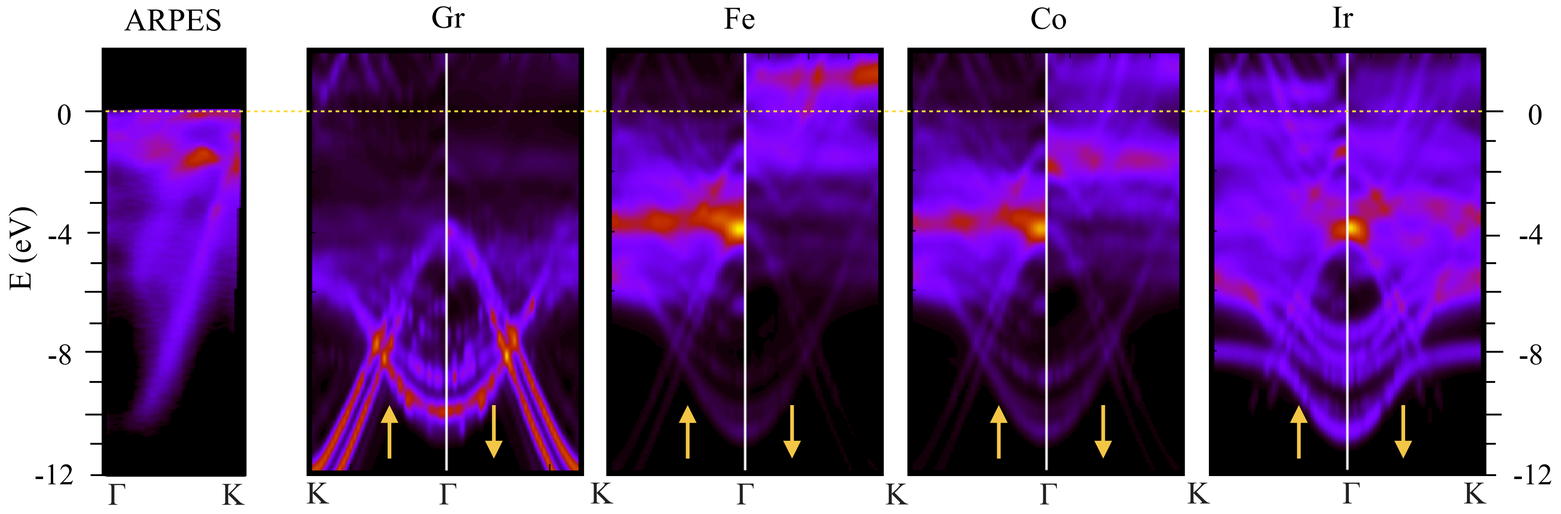

- Avvisati et al, J Phys. Chem. C **121**, 1639 (2017)

- Avvisati et al, Nano Lett. **18**, 2268 (2018)
- Calloni et al, J. Chem. Phys. **153**, 214703 (2020)
- Cardoso et al, Phys. Rev. Mat. **5**, 014405 (2021)
- Pacile' et al, Appl. Phys. Lett. **118**, 121602 (2021)

exascale opportunity: **high throughput screening**

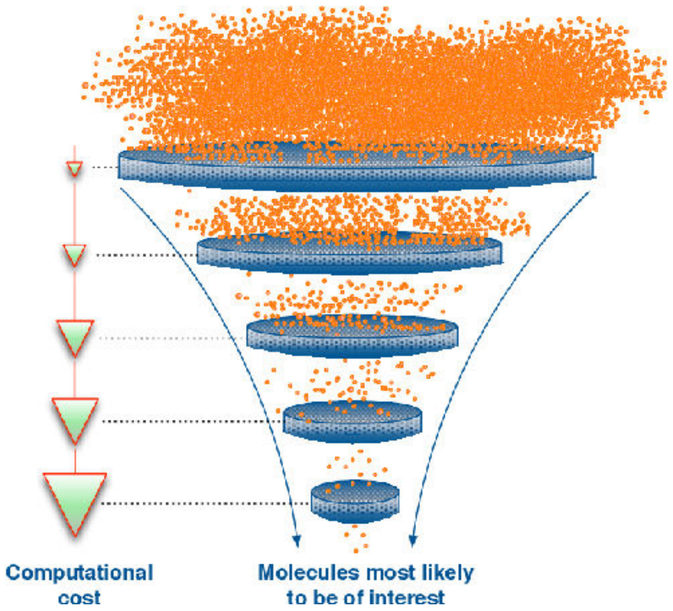

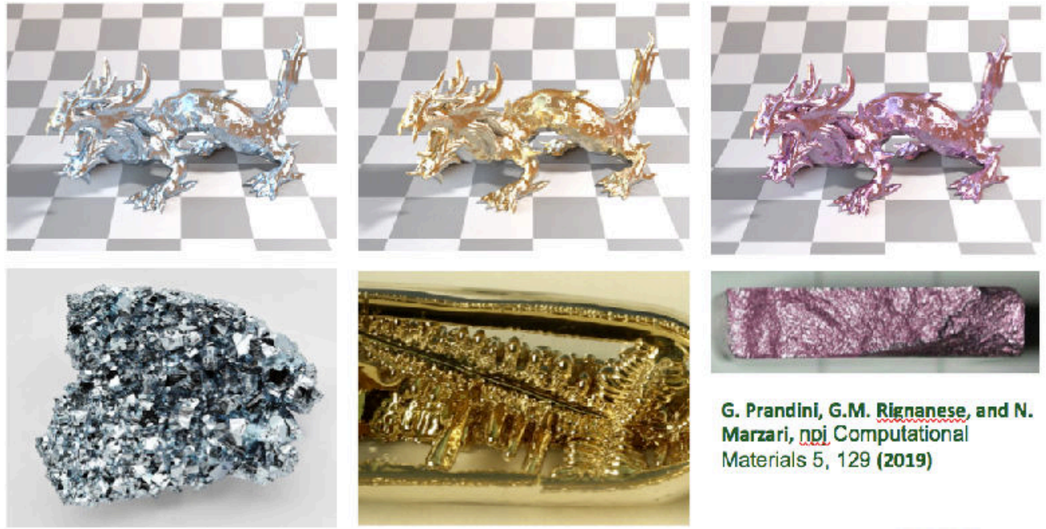

G. Prandini, G.M. Rignanese, and N. Marzari, npj Computational Materials 5, 129 **(2019)**

exascale opportunity: **high throughput screening**

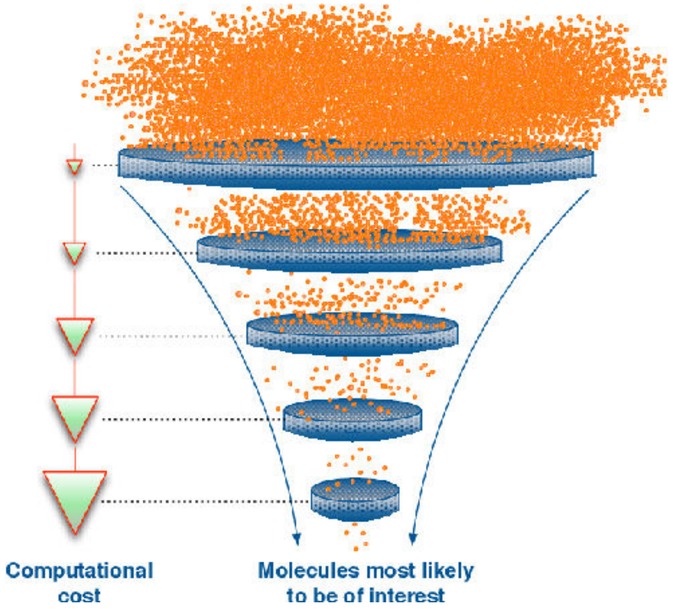

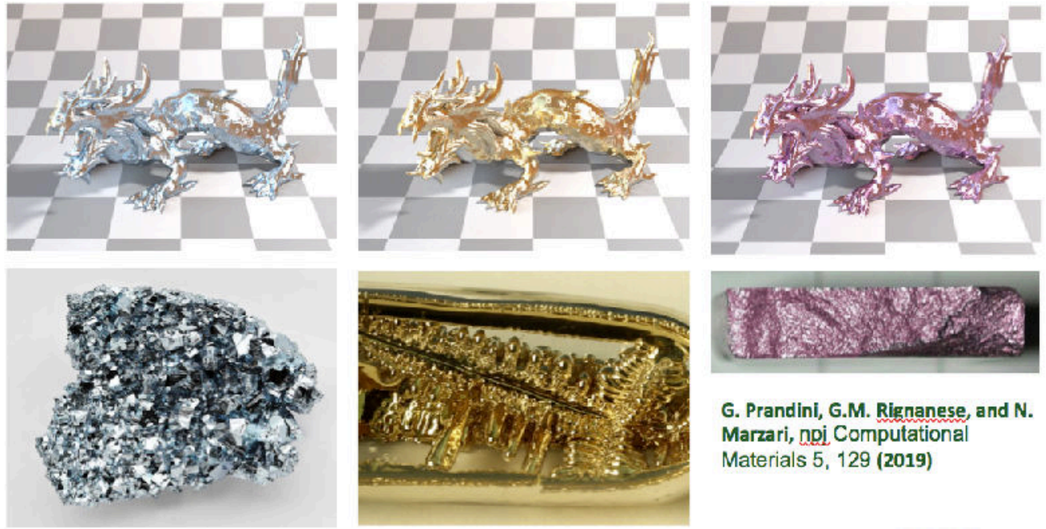

G. Prandini, G.M. Rignanese, and N. Marzari, npj Computational Materials 5, 129 **(2019)**

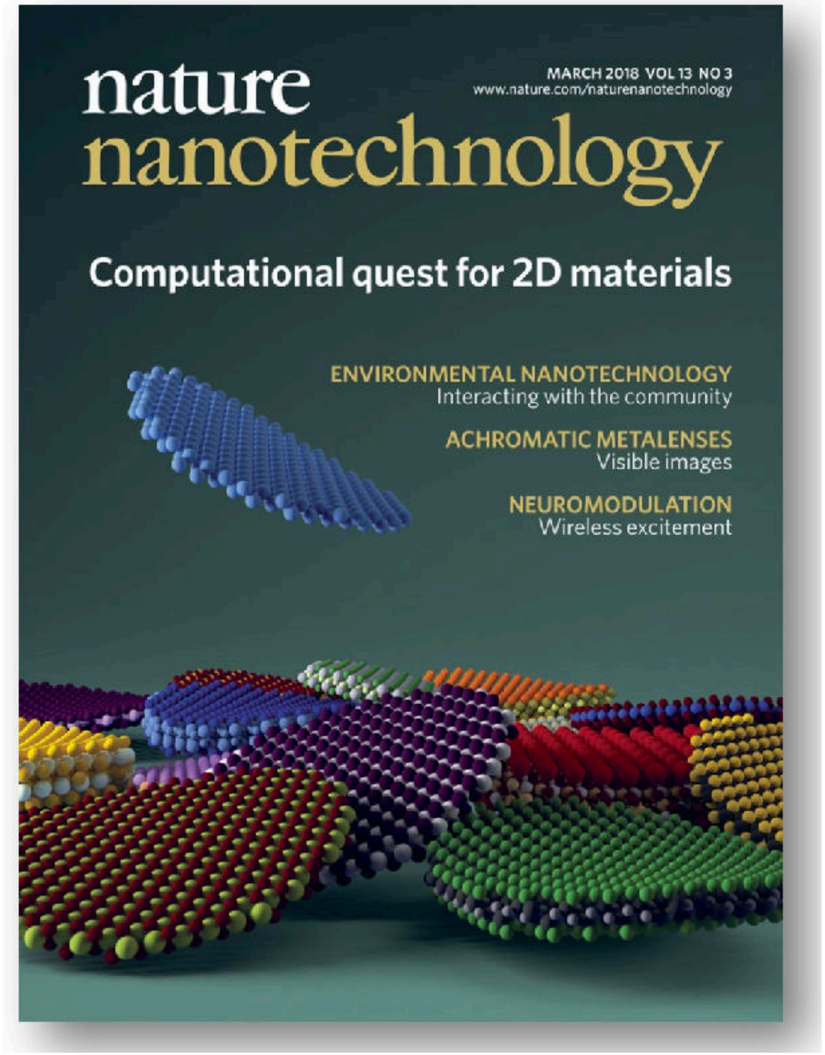

N. Mounet, ..., N. Marzari, Nature Nanotechnology 13, 246 (2018)
D. Campi, N. Mounet, M. Gibertini, G. Pizzi, and N. Marzari, *ACS Nano* 17, 11268–11278 (2023)

MAX DRIVING THE EXASCALE TRANSITION

Materials Design at the Exascale

- European centre of Excellence in **HPC applications**
- funded for **3 phases** (2015-2026)
- 16 EU partners, head-quartered at CNR (Modena, IT)
- focused on **electronic structure codes**

QUANTUMESPRESSO

siesta

Yambo

BigDFT

www.flapw.de Fleur

AiiDA

QUANTUMESPRESSO

www.flapw.de
fleur

BigDFT

siesta

Yambo

AiiDA

MaX

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

- widely used **open source, community codes** in electronic structure

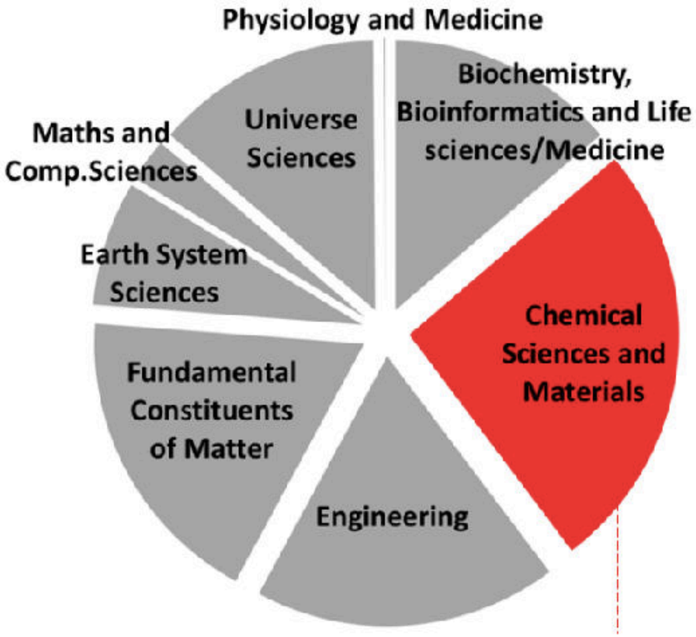

SELECTED ACTIVITIES

- **parallel optimization and performance portability** are key to keep exploiting HPC resources
- All MaX flagship codes released for **production with GPU support**

- Design of **exascale workflow** for selected scientific challenges

- hardware-software **codesign vehicles**
- **energy-efficiency** of codes

- large effort on **education and training**: hands-on schools and hackathons

A partnership with the required skills

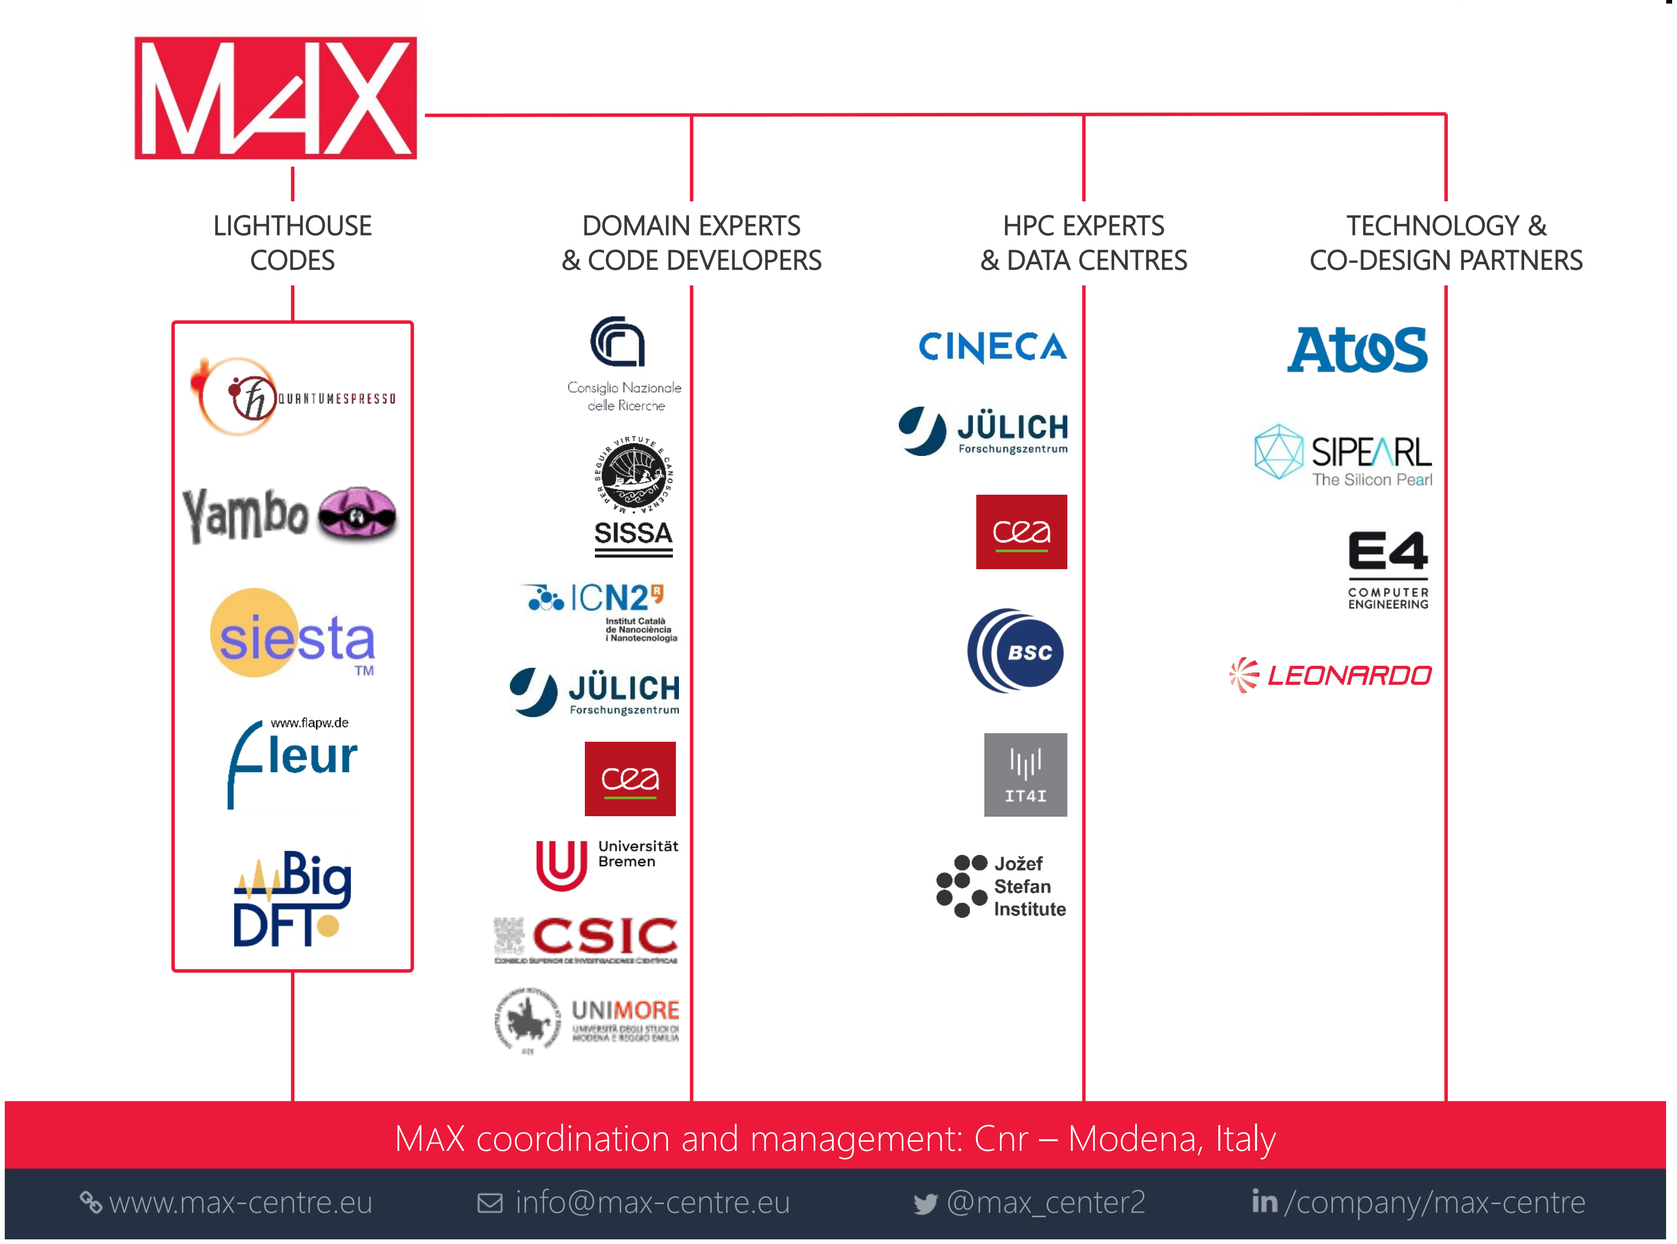

MaX flagship codes:

- electronic structure methods
- implement diverse computational approaches
- open source community codes
- large in terms of user base (~ 4000 citations/year)

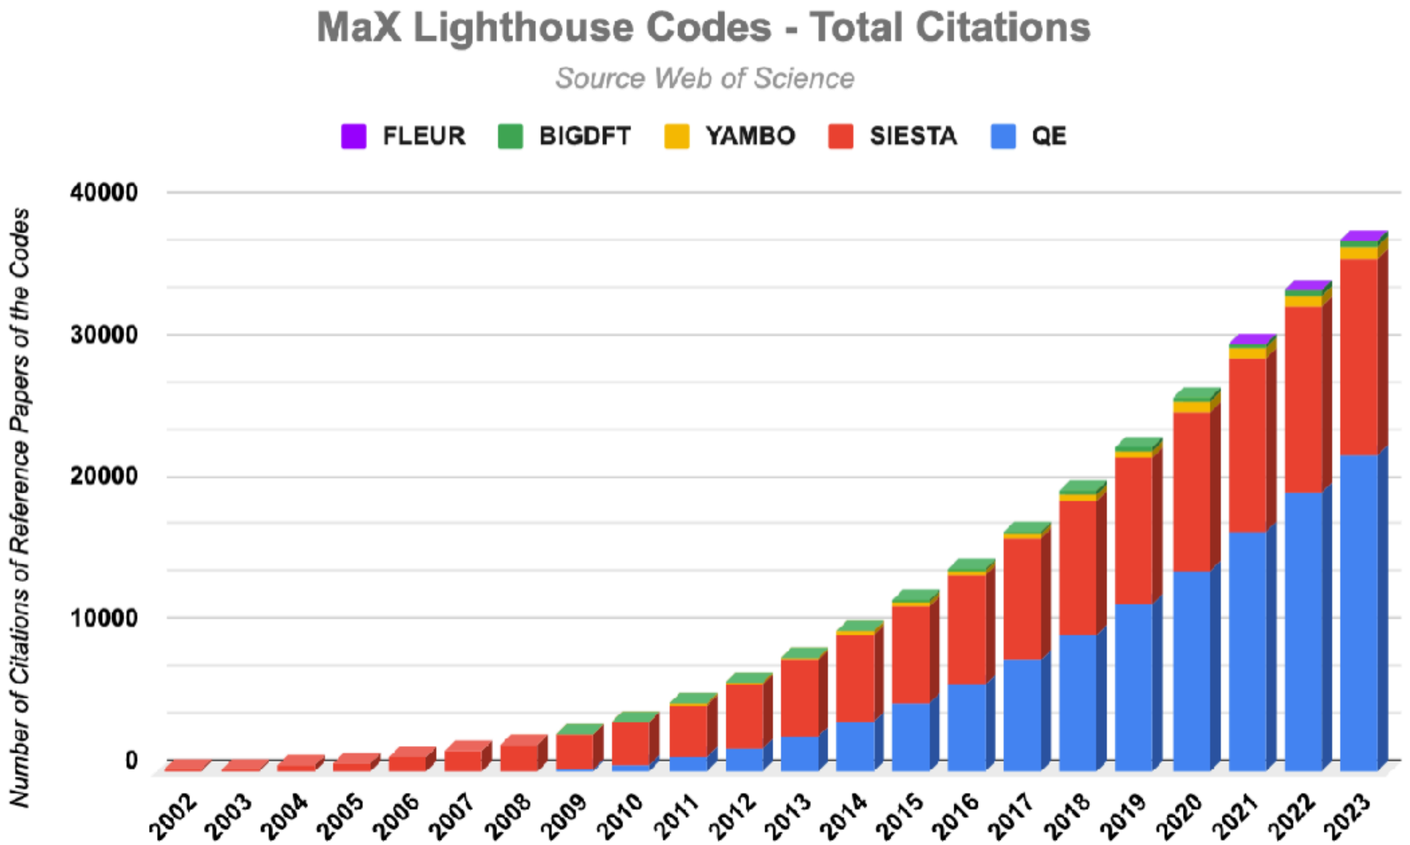

MaX flagship codes:

- electronic structure methods
- implement diverse computational approaches
- open source community codes
- large in terms of user base (~ 4000 citations/year)

Geographic distribution of authors' affiliation in peer-reviewed publications citing MaX lighthouse codes in 2023

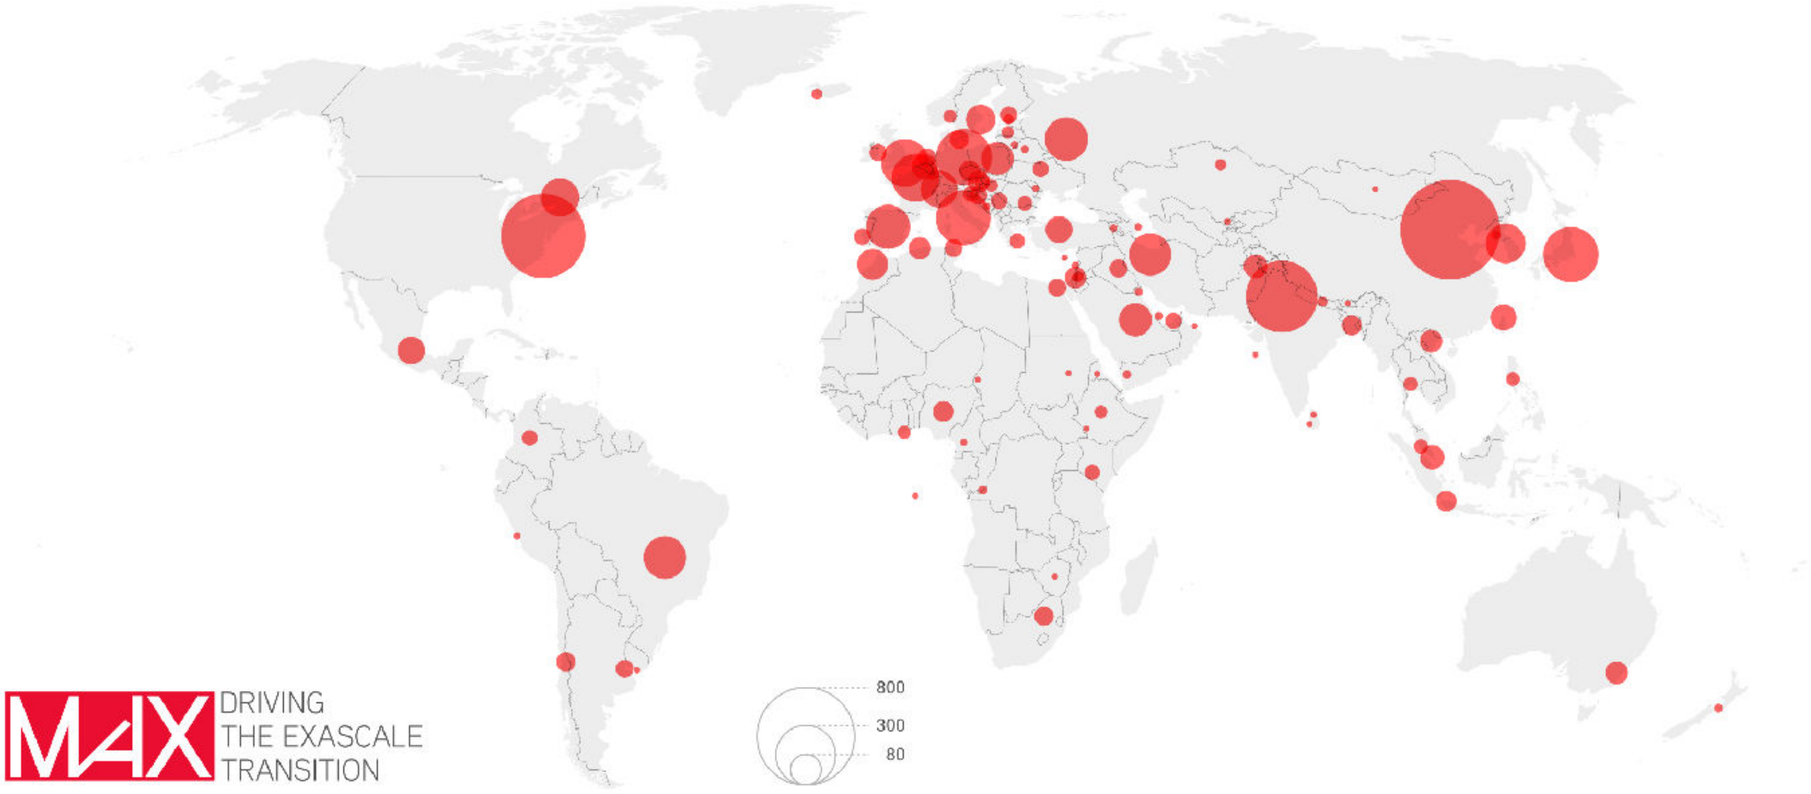

Source: Web of Science • Created with Datawrapper

Deployment on EuroHPC Systems

last updated: September 2024

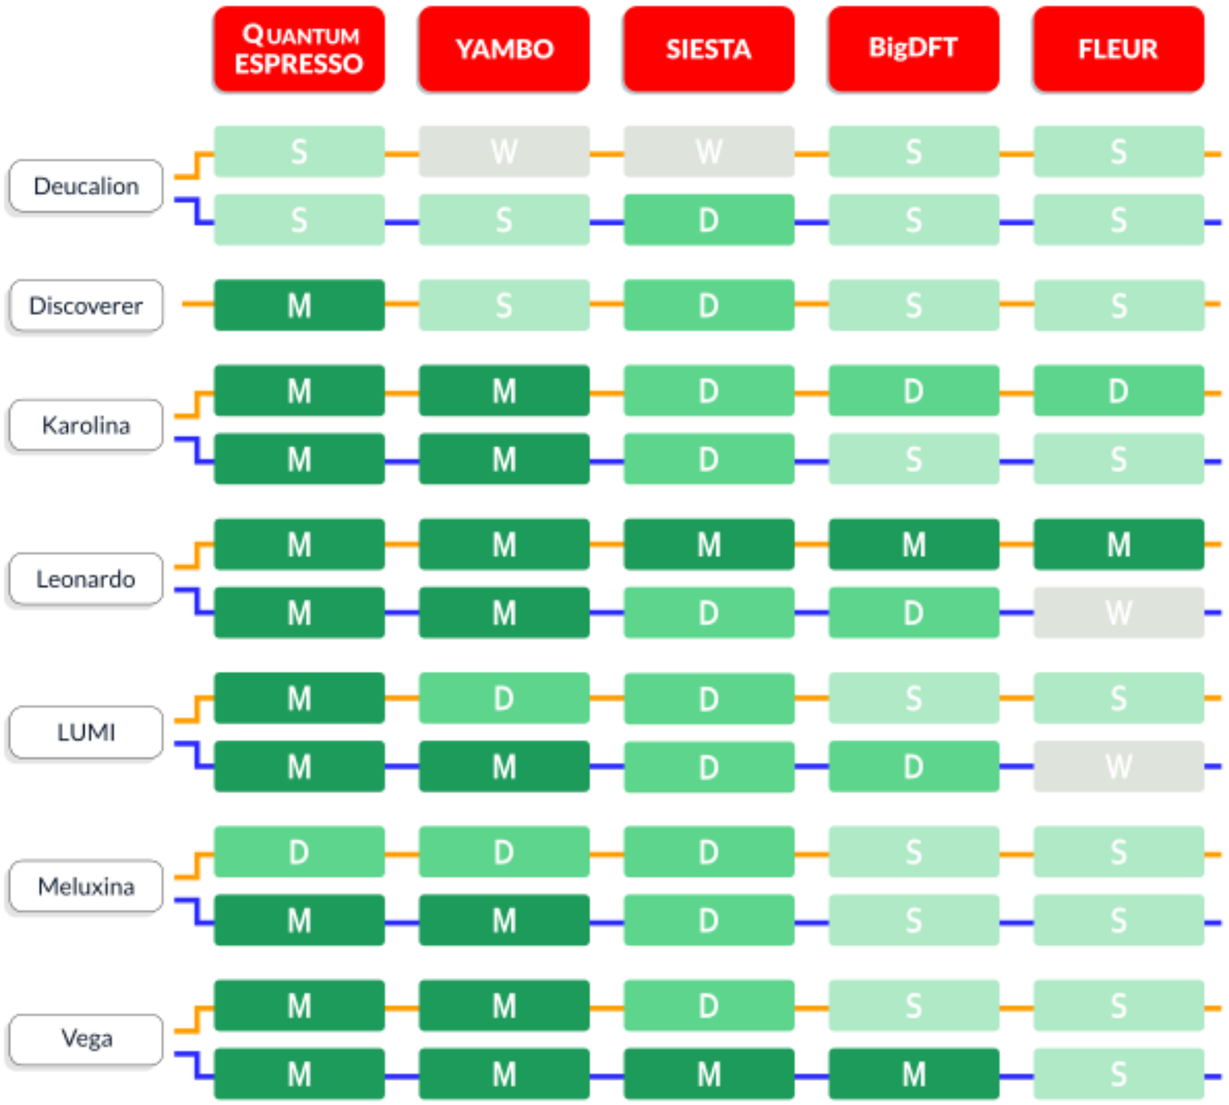

M
Deployed
module available

D
Demonstrated
developers installation

S
Supported architectures
tested on similar architectures

W
Work in progress

CPU
GPU

…and more:

- ARM Apple M1 (MacBook)
- ARM Graviton3 (AWS)
- ARM Fujitsu (Fugaku)

Exascale Workflows and Data

batteries & energy storage

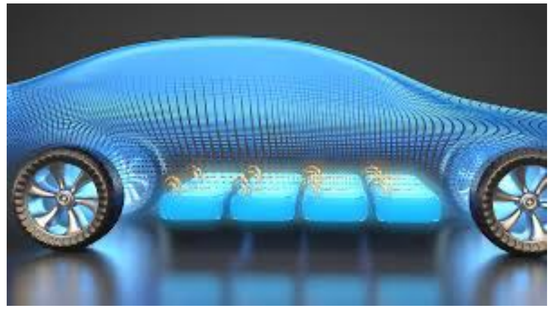

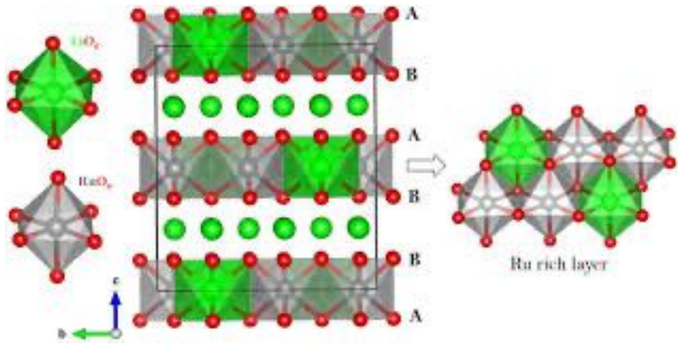

photovoltaics

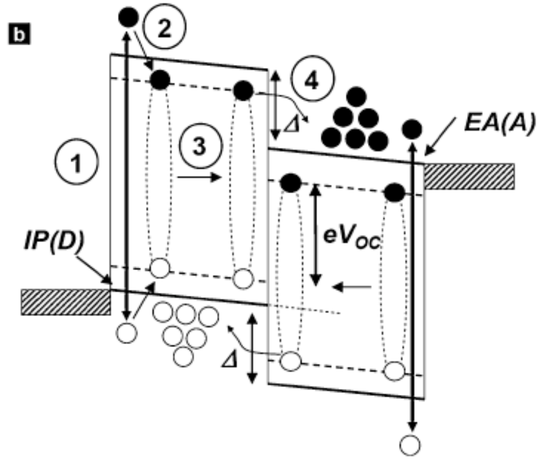

AiiDA HyperQueue

connectivity of very large data sets managed by AiiDA

Within MaX, **selected scientific grand challenges** are addressed, including:

- Electronic and heat conductivity from first principles
- Design and control of nanoscale magnetism
- Manipulation and control of coherent quantum states
- complex photocatalytic reactions and photovoltaic reactions
- biological and bio-mimetic materials

Training events: schools and Hands-on targeted to code users

Training in High Performance Computing (HPC) for Computational Materials Science

- **Content**: theoretical foundations, implemented algorithms using MaX flagship codes.
- **Hands-on Experience**: Code usage, best practices in HPC environments: parallel computations, efficient memory management, techniques for maximizing performance and scalability.
- **Format**: Typically 3/5 working days. In person, online, and hybrid formats

- **AiiDA-Yambo** tutorial: automating Green's function methods. (Feb. 2024)
- ENCCS/MaX Workshop (Efficient materials modelling on HPC with **QUANTUM ESPRESSO, SIESTA and Yambo**) (March 2024)
- PWTK: a scripting interface for **Quantum Espresso** Lubjana, May (May 2024)
- Workshop on Machine Learning for Materials (May 2024)
- **Quantum Espresso** school /NCC Czechia (June 2024)
- **Fleur** Hands-on tutorial 2024 edition (September 2024)
- MaX/HPC-NCC-Croatia QE workshop (November 2024)
- **SIESTA** School 2024 (November 2024)

Training events: schools and Hands-on targeted to code users

Training in High Performance Computing (HPC) for Computational Materials Science

- **Content**: theoretical foundations, implemented algorithms using MaX flagship codes.
- **Hands-on Experience**: Code usage, best practices in HPC environments: parallel computations, efficient memory management, techniques for maximizing performance and scalability.
- **Format**: Typically 3/5 working days. In person, online, and hybrid formats

- **AiiDA-Yambo** tutorial: automating Green's function methods. (Feb. 2024)
- ENCCS/MaX Workshop (Efficient materials modelling on HPC with **QUANTUM ESPRESSO, SIESTA and Yambo**) (March 2024)
- PWTK: a scripting interface for **Quantum Espresso** Lubjana, May (May 2024)
- Workshop on Machine Learning for Materials (May 2024)
- **Quantum Espresso** school /NCC Czechia (June 2024)
- **Fleur** Hands-on tutorial 2024 edition (September 2024)
- MaX/HPC-NCC-Croatia QE workshop (November 2024)
- **SIESTA** School 2024 (November 2024)

589 Participants in 2024

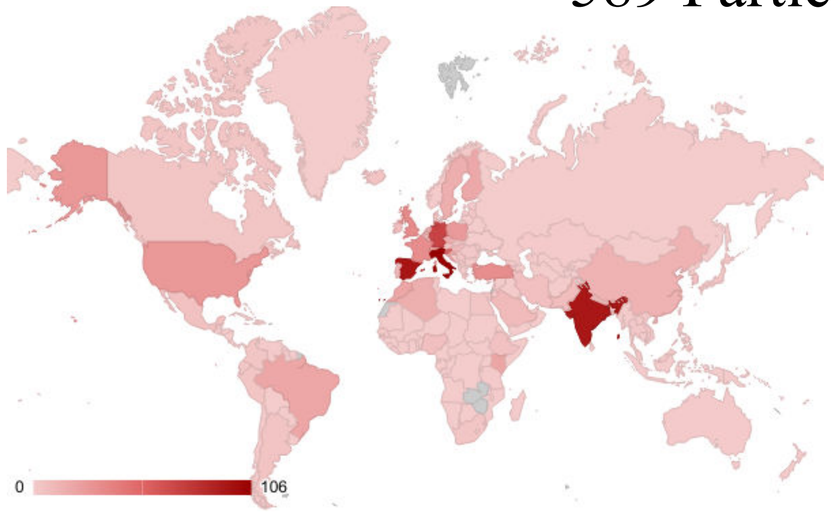

Lhumos Platform Lhumos www.alpha.lhumos.org

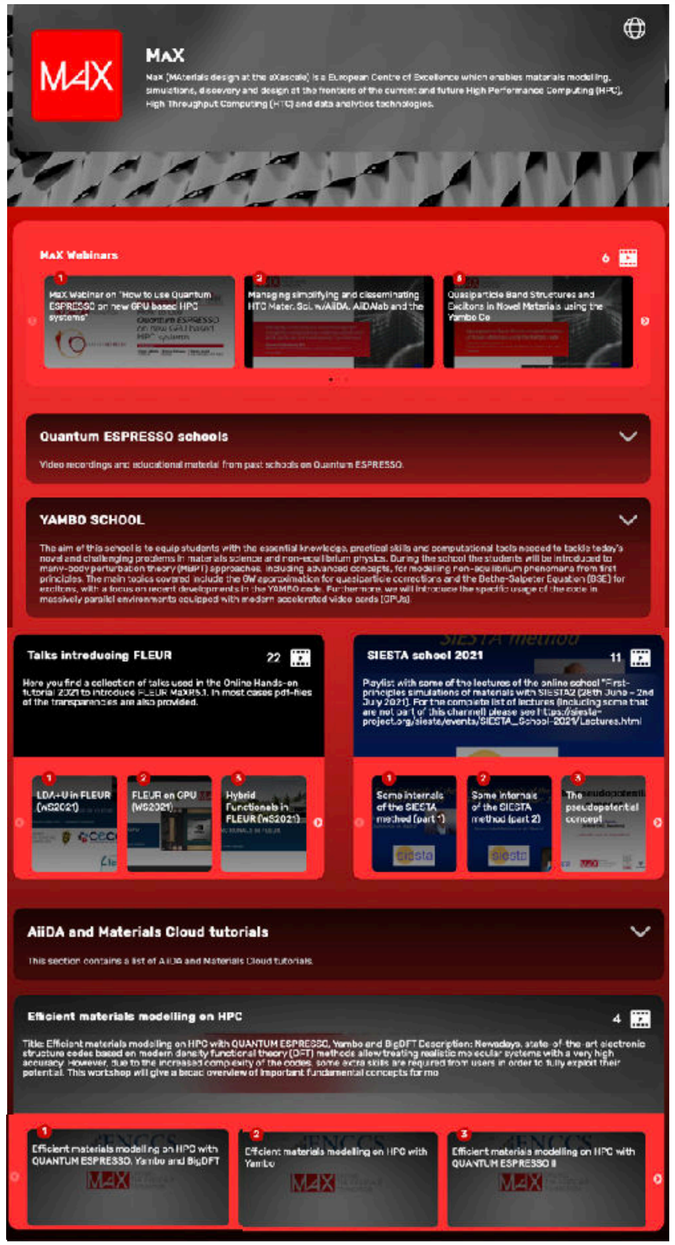

MAX DRIVING THE EXASCALE TRANSITION

Lhumos Platform

www.alpha.lhumos.org

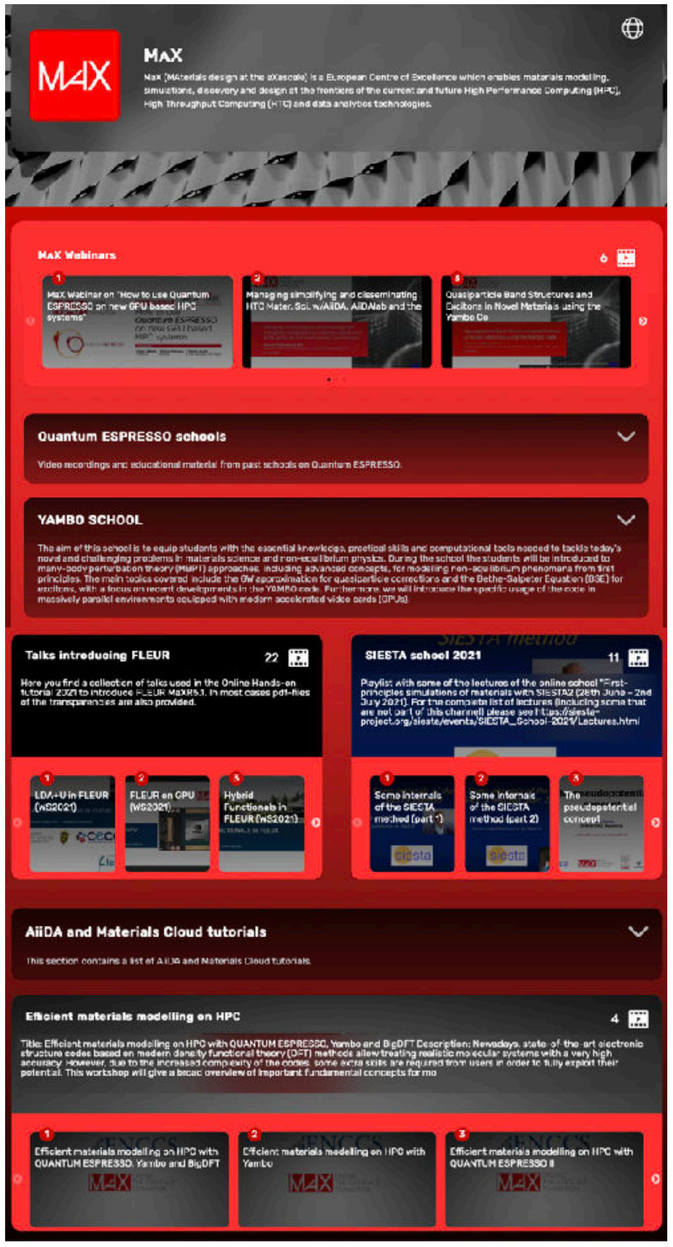

MAX DRIVING THE EXASCALE TRANSITION

Lhumos Platform

Lhumos

www.alpha.lhumos.org

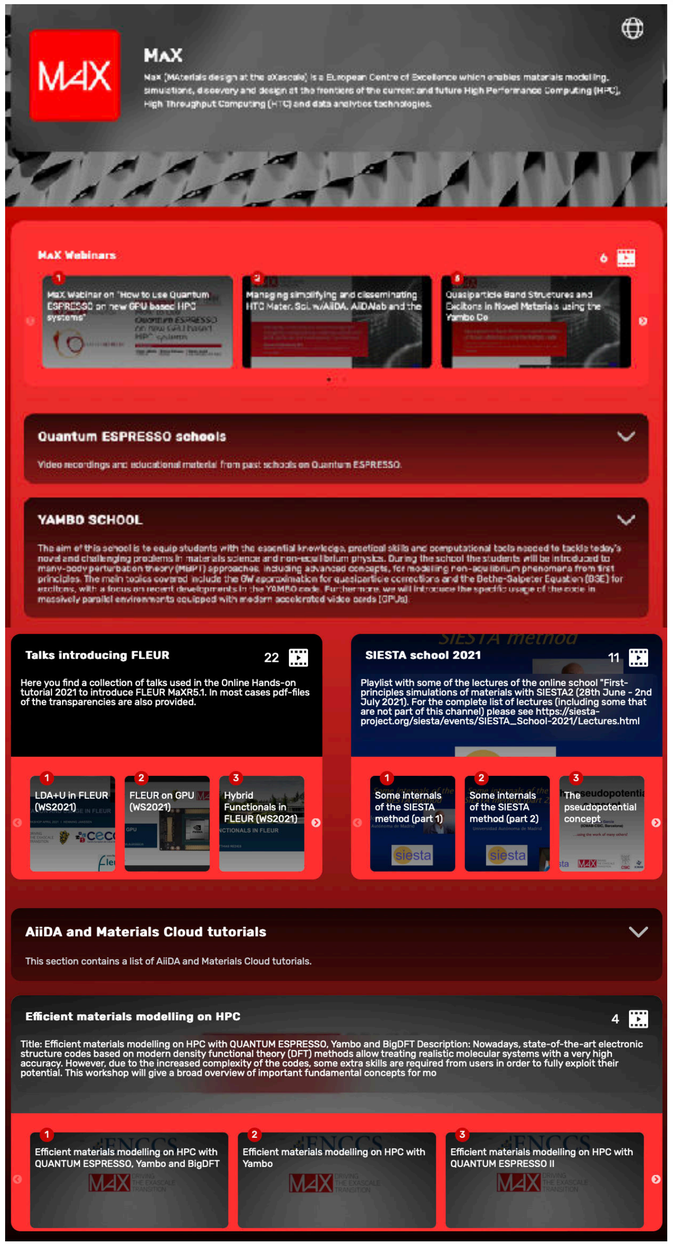

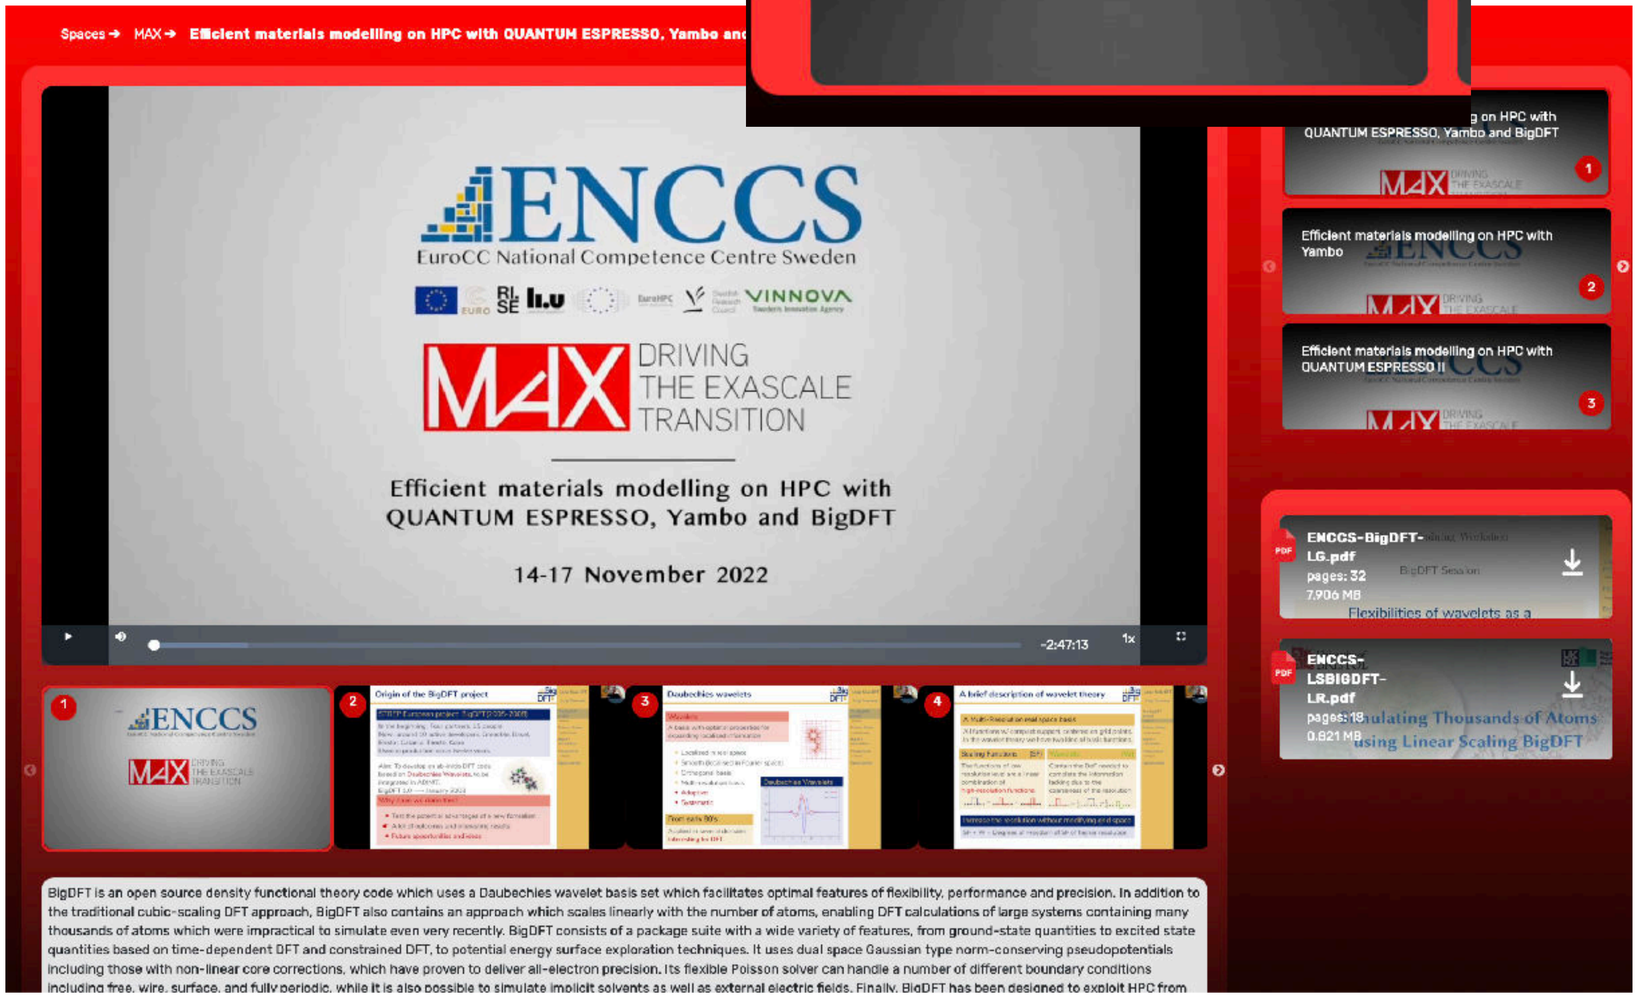

MAX DRIVING THE EXASCALE TRANSITION

Enjoy !

company/max-centre/

youtube/channel/MaX Centre eXascale

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