Materials design towards the exascale: Codes, Workflows, and data

Andrea Ferretti

[CNR-NANO, Modena, Italy]

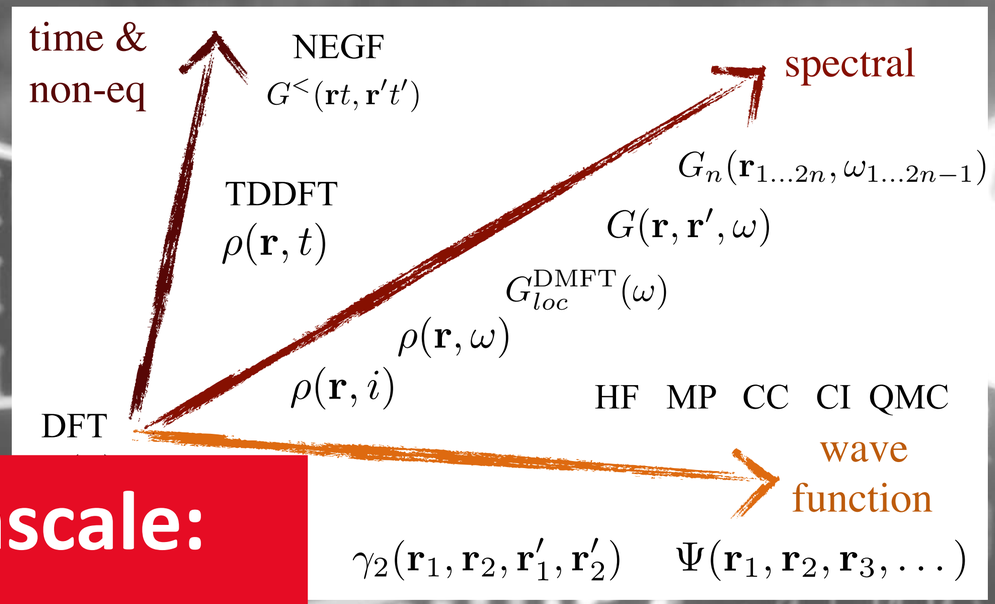

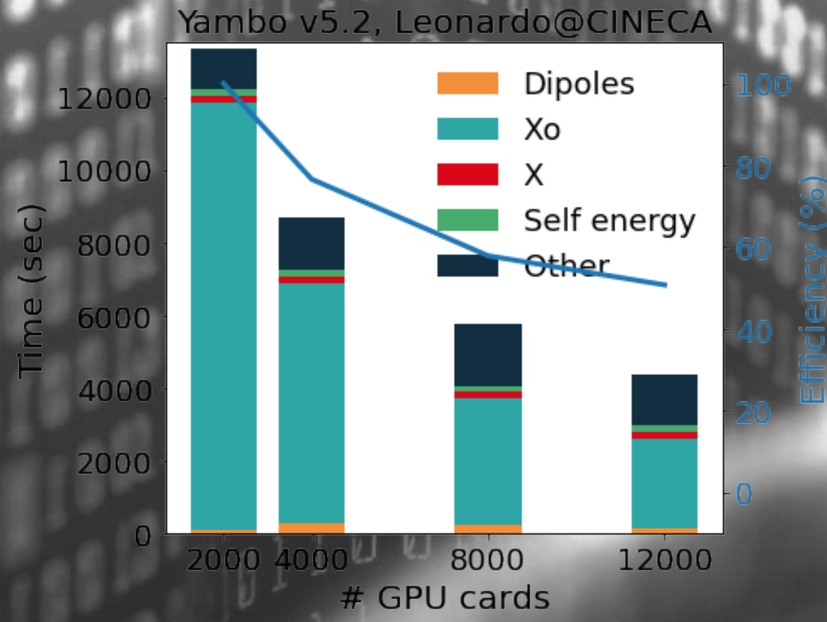

outline

Materials Science & MaX

- Exascale
- MaX: Materials design at the exascale

Activities on MaX Codes

- Parallel performance
- Porting and performance portability
- New features

```
74   !$omp parallel default(shared), private(ir)
75   !$omp do
76   do ir = 1, fft_size
77     isc%rho_tw_rs(ir) = cmplx(isc
78   enddo
79   !$omp end do
```

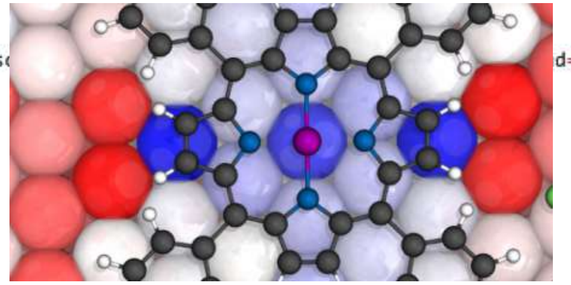

More on Workflows and data

- Exascale workflows
- Tools & Examples
- Data infrastructure

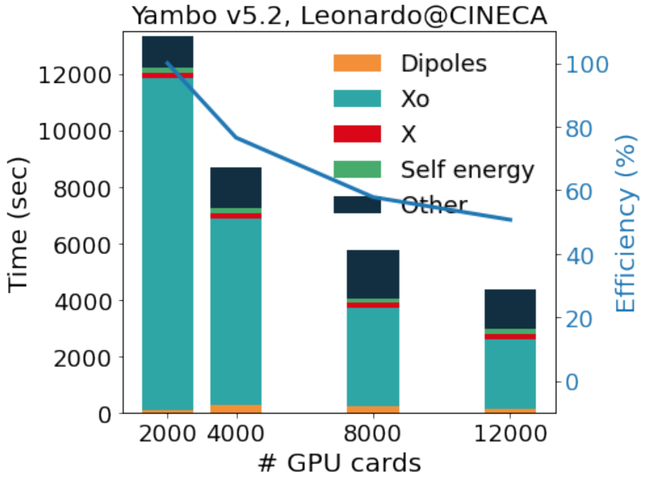

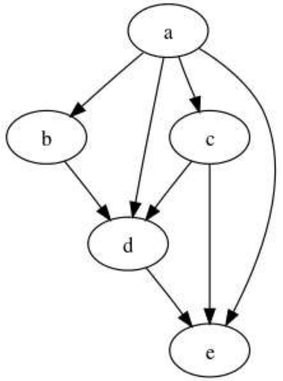

today: HPC at the exascale

MAX DRIVING THE EXASCALE TRANSITION

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)

EuroHPC Joint Undertaking

Jupiter: NVIDIA GH200
> 1 ExaFlops

MareNostrum V: Intel + NVIDIA H100 => 180 PFlops

US DOE

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

Switzerland

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

Leonardo: Intel + NVIDIA A100 => 241 PFlops

LUMI: CRAY + AMD Mi250
=> 380 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)

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

Machine Learning &
Artificial Intelligence

High throughput screening

Improved modelling (complexity)

exascale opportunity: **complexity**

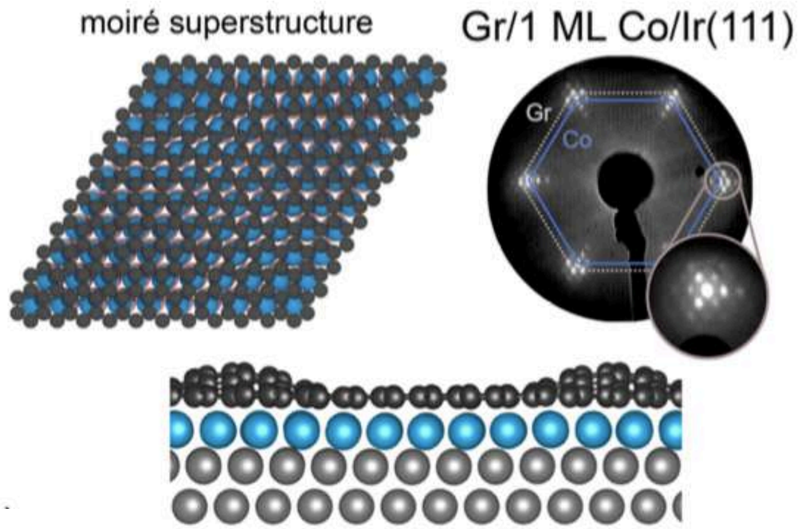

- Graphene / Transition Metal / Ir (111)
- **clear experimental evidence** for moire' pattern (**lattice mismatch**) and **Gr corrugation**
- 10x10 Graphene, 9x9 Iridium => 605 atoms / unit cell
- **Precise treatment of the structure** is important for modelling

- 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)

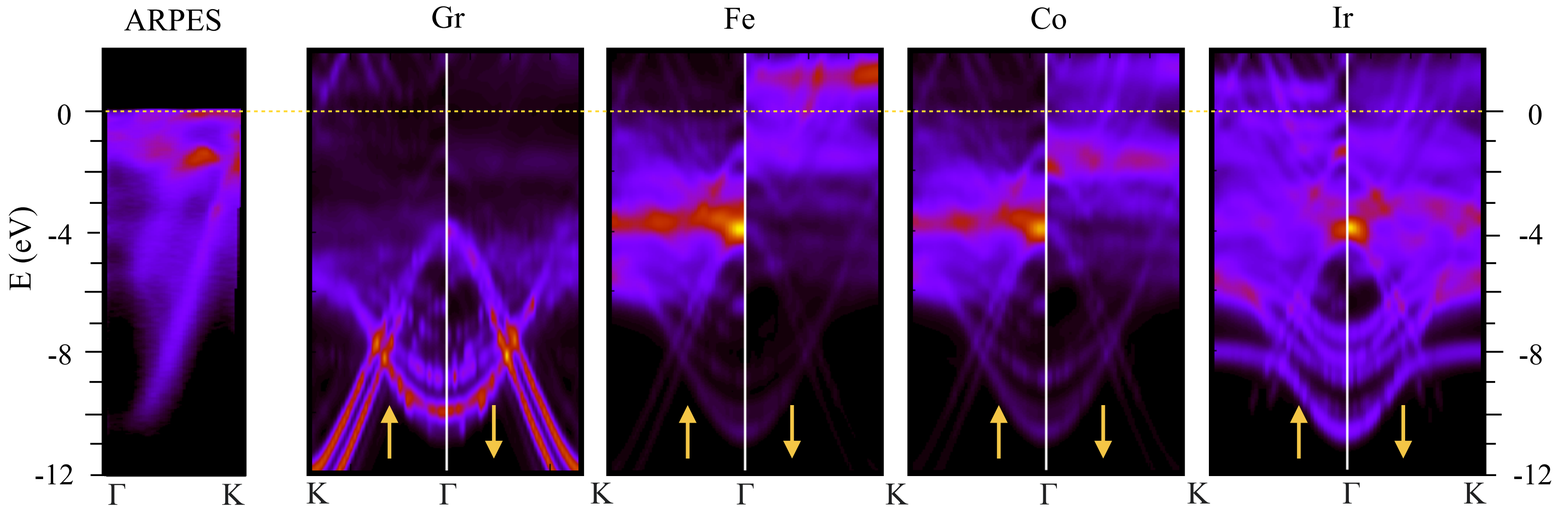

exascale opportunity: **high throughput screening**

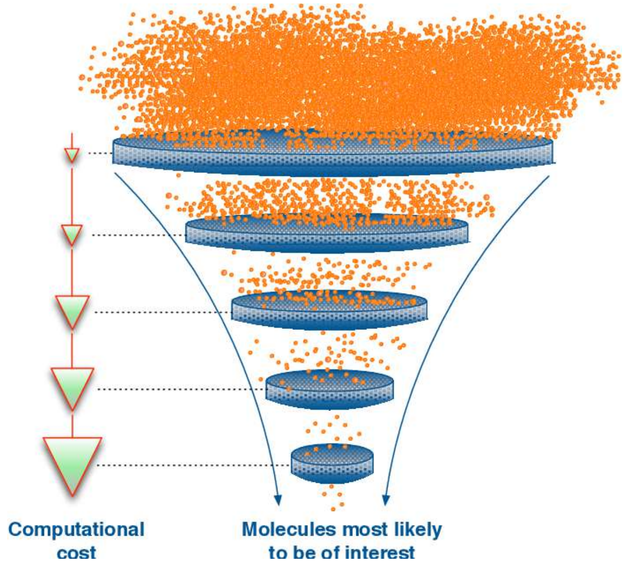

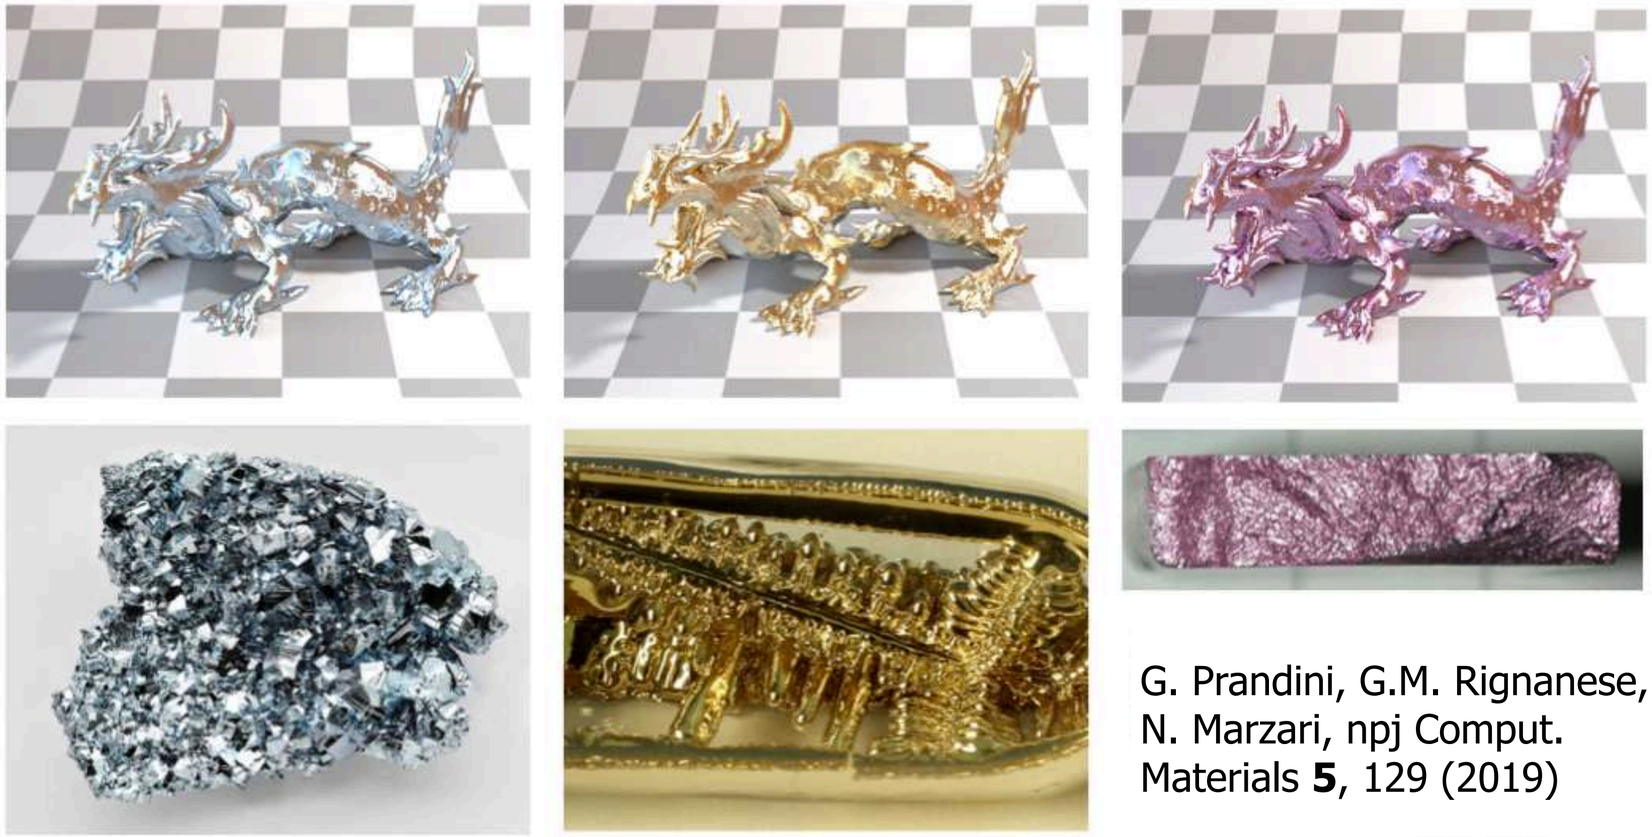

G. Prandini, G.M. Rignanese, N. Marzari, npj Comput. Materials **5**, 129 (2019)

Adv methods in electronic structure

nature materials

INSIGHT | REVIEW ARTICLE

https://doi.org/10.1038/s41563-021-01013-3

Check for updates

Electronic-structure methods for materials design

Nicola Marzari [1]✉, Andrea Ferretti [2] and Chris Wolverton [3]

Density functional theory (DFT):

- applications ranging from **materials modelling, to quantum chemistry and drug design**
- compatible with **high performance computing** and **high-throughput screening**

beyond DFT:

- **multiple hierarchies** can be climbed
  - wavefunction-based methods
  - **many-body perturbation theory and spectral methods**
  - time-dependent and non-equilibrium methods
  - **ensembles**

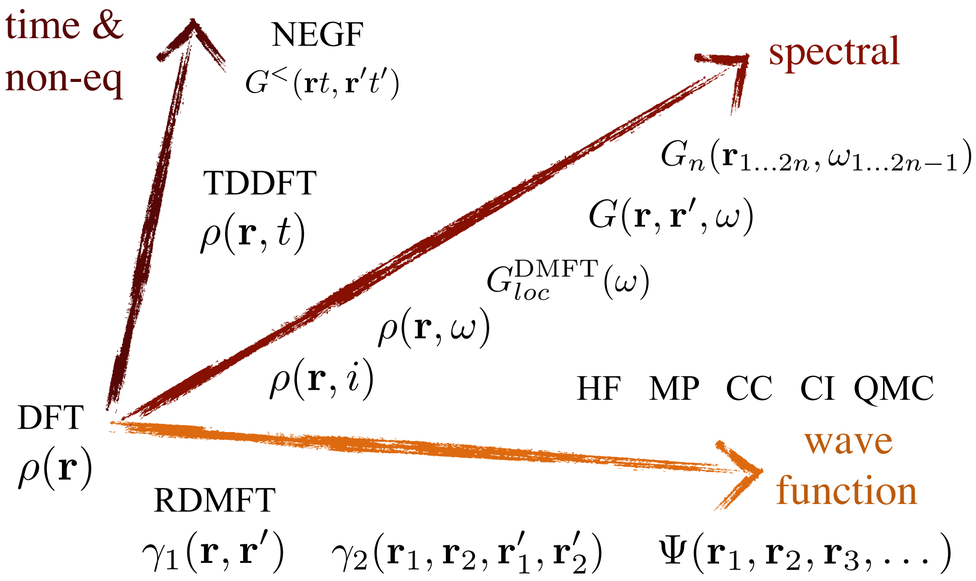

Materials Design at the Exascale

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

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 **open source, community codes** from the **electronic structure** domain

QUANTUMESPRESSO

siesta

Yambo

BigDFT

www.flapw.de
Fleur

AiiDA

Most examples in this talk

QUANTUMESPRESSO

www.flapw.de
fleur

BigDFT

siesta

Yambo

MaX

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

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

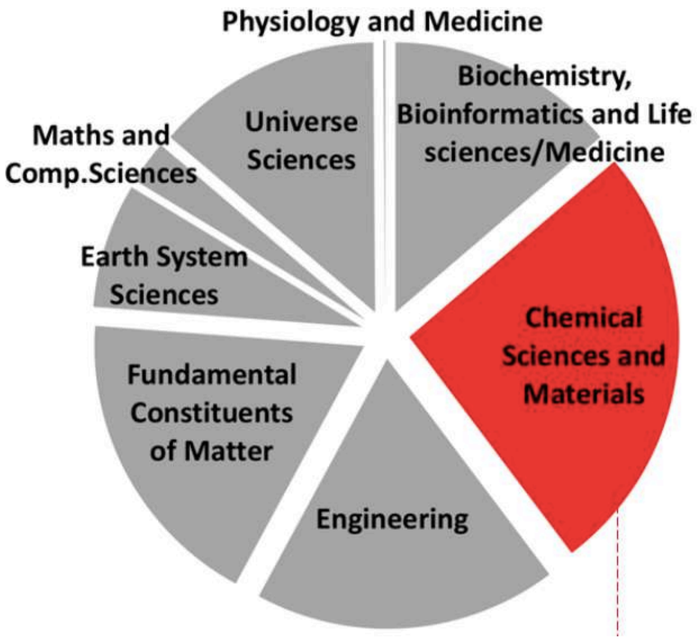

SELECTED ACTIVITIES

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

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

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

MaX flagship codes:

- electronic structure methods
- implement diverse computational approaches
- open source community codes
- large in terms of user base (~ 5000 citations/year)

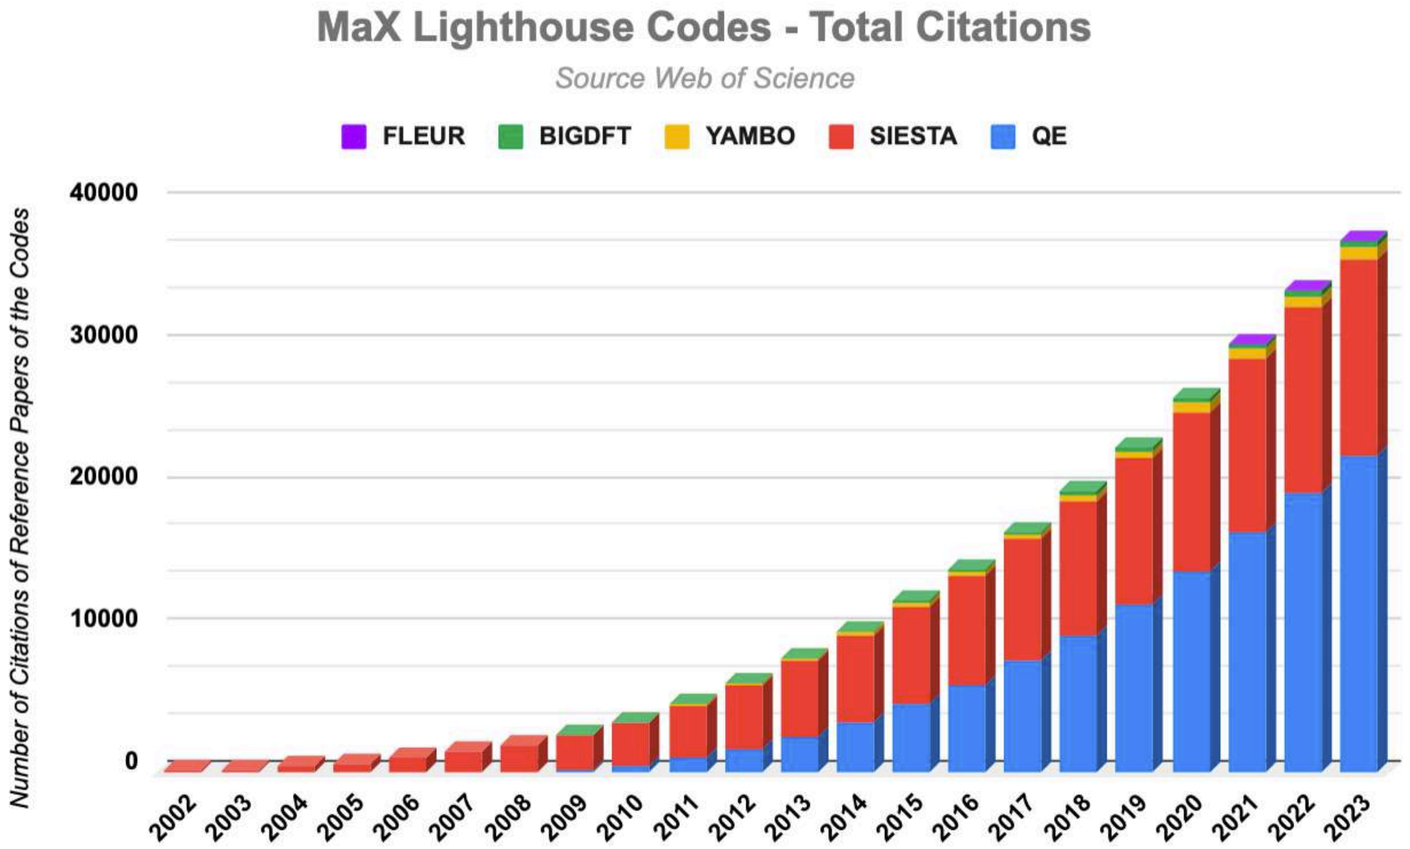

MaX flagship codes:

- electronic structure methods
- implement diverse computational approaches
- open source community codes
- large in terms of user base (~ 5000 citations/year)

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

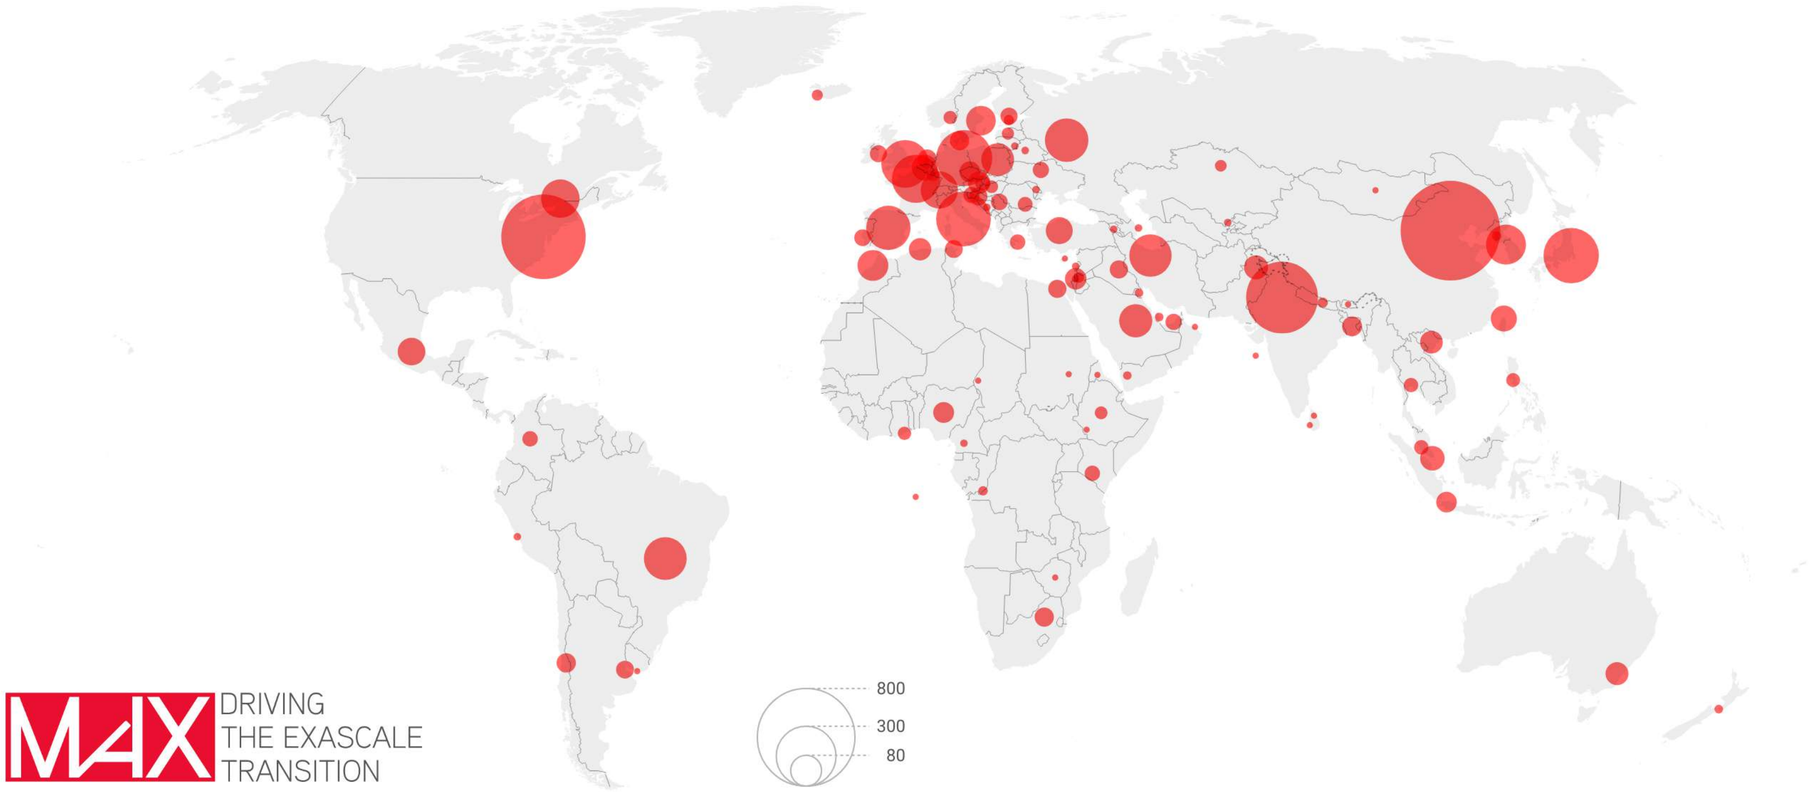

MAX DRIVING THE EXASCALE TRANSITION

Source: Web of Science • Created with Datawrapper

A partnership with the required skills

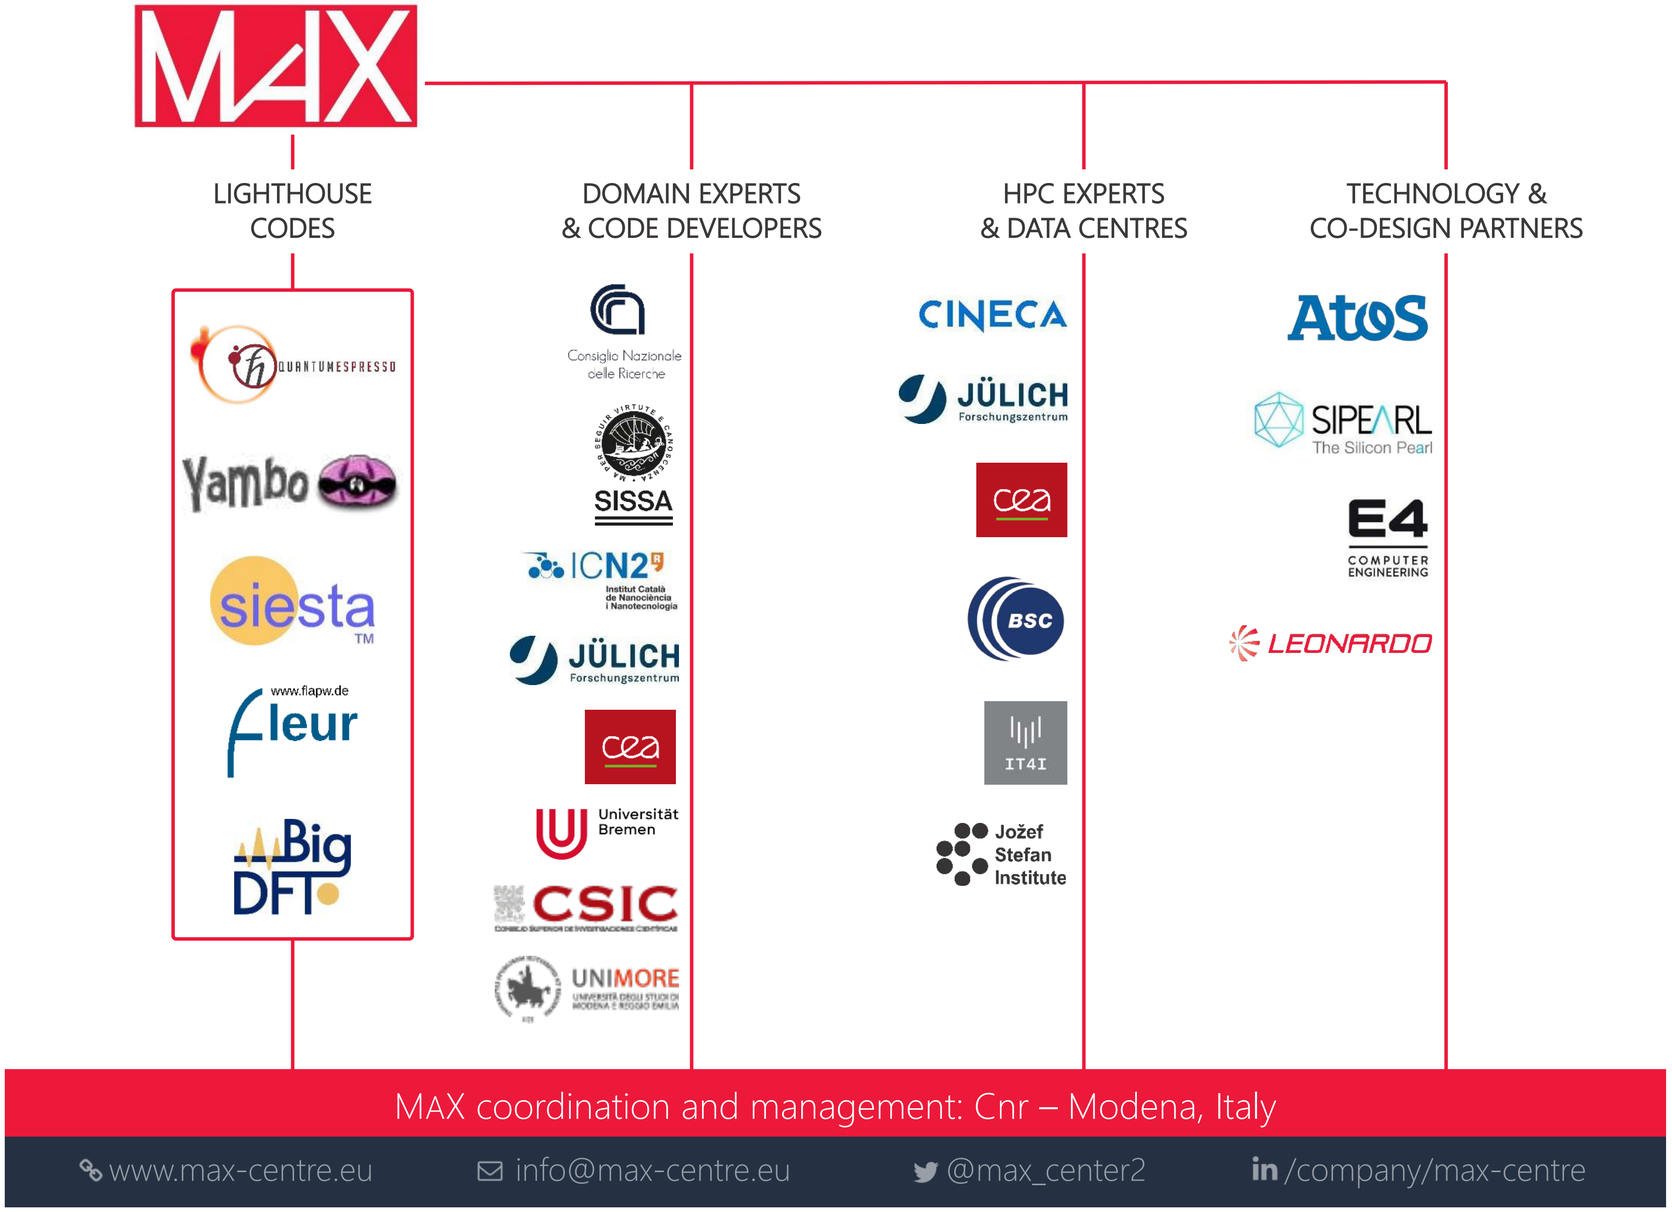

MaX work-plan

Revolves around:

- **Codes** (Lighthouse applications)
- **Exascale workflows, Data, &** Scientific grand challenges
- **Technical challenges** and deployment
- **Co-design** and technology exploitation
- **Training**

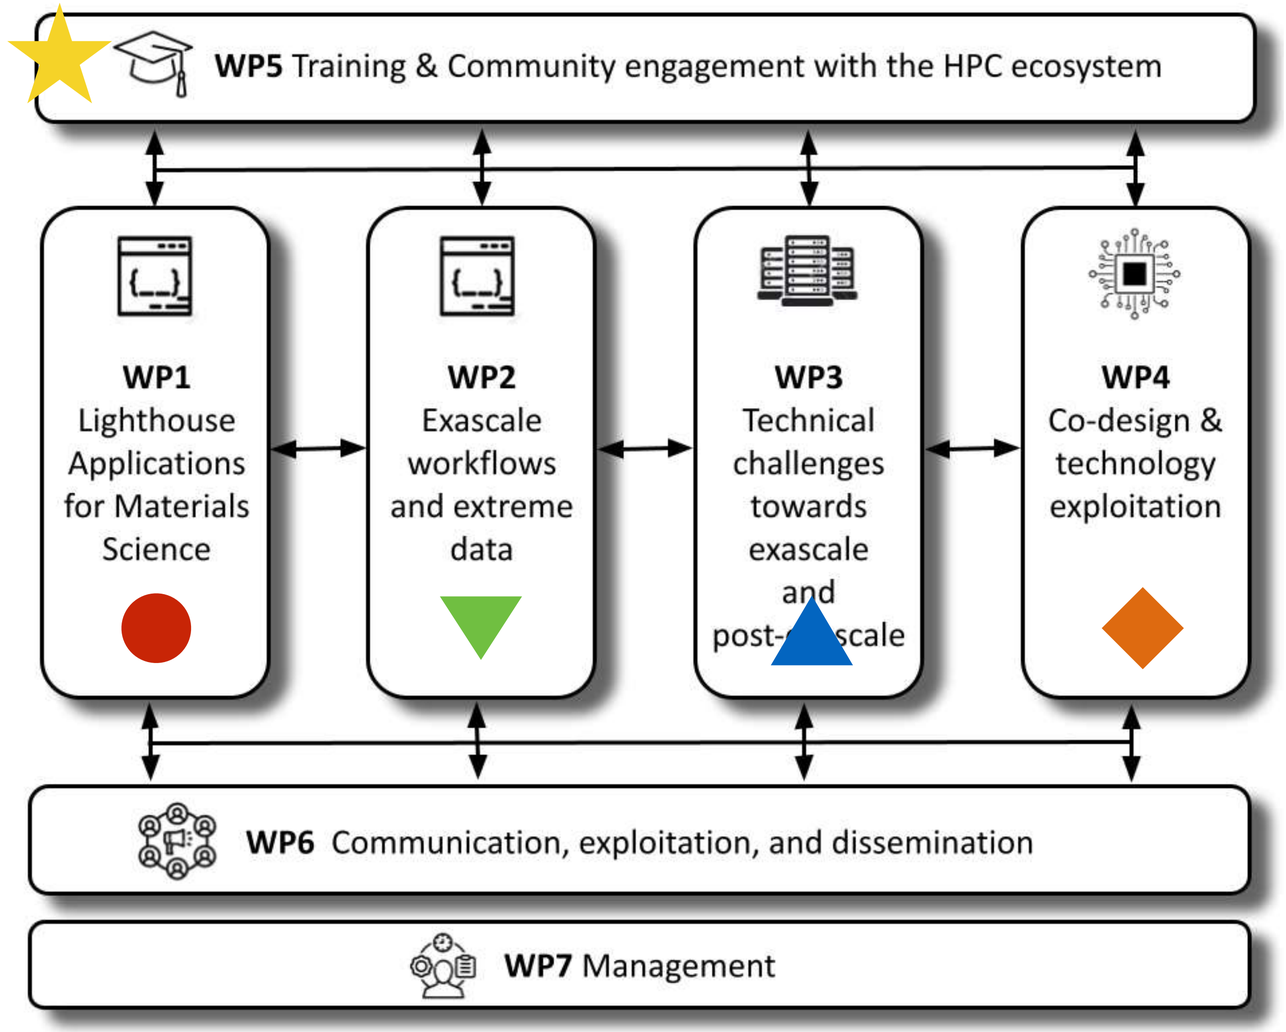

MaX work-plan

Revolves around:

- **Codes** (Lighthouse applications)
- **Exascale workflows, Data, &** Scientific grand challenges
- **Technical challenges** and deployment
- **Co-design** and technology exploitation
- **Training**

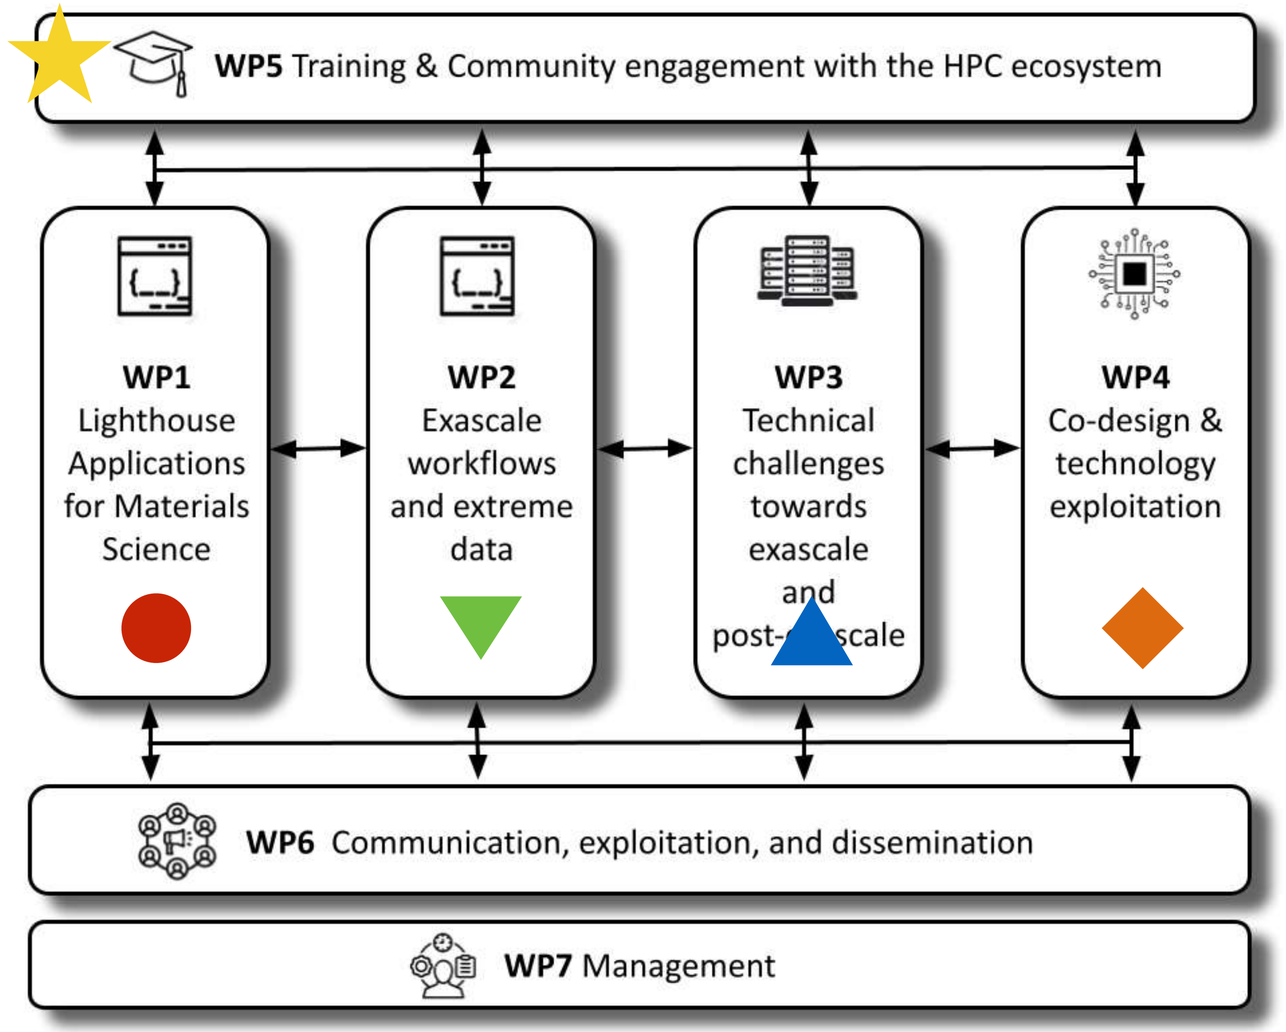

Codes

GOAL: turn MaX flagship codes into exascale-enabled applications

- **large scale MPI** parallelism (order of 10000 tasks)
- combined with **GPU awareness**

(T1.1) Single-node optimisation

- make sure **MaX codes** can exploit **accelerated nodes** featuring **multiple GPU brands**

(T1.2) Multi-node parallel efficiency

- make the codes scalable in the presence of GPUs

(T1.3) Scientific software engineering

- support long-term maintainability and community contributions

(T1.4) New Scientific features

```
229       do jb=Sx_lower_band,Sx_upper_band
230         !
231         if (.not.PAR_IND_G_b%element_1D(jb)) cycle
232         !
233         isc%os(1)=jb
234         iscp%os=isc%os
235         !
236         call DEV_SUB(scatter_Bamp)(isc)
237         !
238         ! Normal case, the density matrix is diagonal
239         !
240         if (isc%is(1)/=iscp%is(1)) then
241           call DEV_SUB(scatter_Bamp)(iscp)
242         else
243           ! dev2dev, iscp%rhotw = isc%rhotw
244           call dev_memcpy(DEV_VAR(iscp%rhotw),DEV_VAR(isc
245         endif
246         !
247         DP_Sx_l=DEV_SUB(Vstar_dot_VV)(isc%ngrho,DEV_VAR(i
248     &                                 DEV_VAR(isc%rhotw),
249         DP_Sx=DP_Sx + DP_Sx_l * ( -4._SP/spin_occ*pi*E%f(
250         !
251         if (master_thread.and.is_ibz==1.and.n_lt_steps>0)
252         !
253       enddo
```

Codes

GOAL: turn MaX flagship codes into exascale-enabled applications

- **large scale MPI** parallelism (order of 10,000 tasks)
- combined with **GPU awareness**

Jupiter (@FZJ-JSC): > 1 ExaFlops
~ 6000 nodes with 4 Nvidia GH200/node
~ 24,000 GPUs in total

Codes

GOAL: turn MaX flagship codes into exascale-enabled applications

- **large scale MPI** parallelism (order of 10,000 tasks)
- combined with **GPU awareness**

Jupiter (@FZJ-JSC): > 1 ExaFlops
~ 6000 nodes with 4 Nvidia GH200/node
~ 24,000 GPUs in total

MaX-2

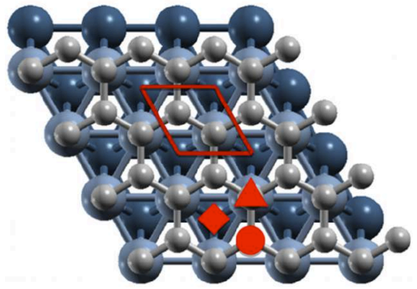

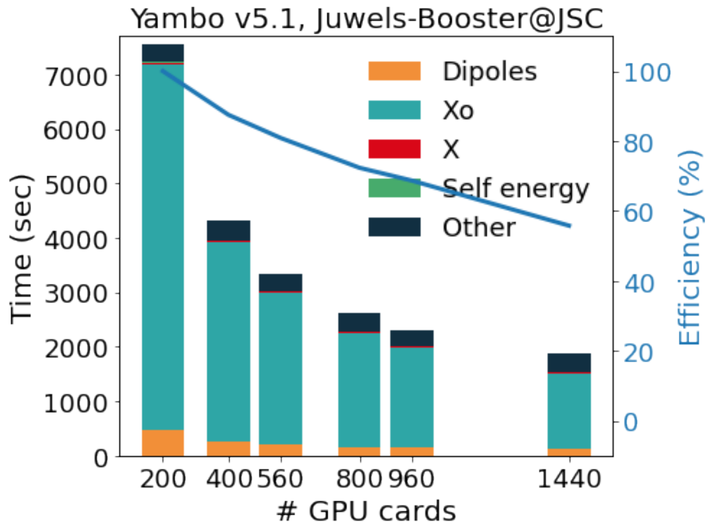

Juwels-Booster:
4 Nvidia A100 / node

runs: up to 360 J-B nodes about 40% of the whole machine (960 nodes)

Codes

GOAL: turn MaX flagship codes into exascale-enabled applications

- **large scale MPI** parallelism (order of 10,000 tasks)
- combined with **GPU awareness**

Jupiter (@FZJ-JSC): > 1 ExaFlops
~ 6000 nodes with 4 Nvidia GH200/node
~ 24,000 GPUs in total

MaX-3

Leonardo @ CINECA:
4 Nvidia A100 next / node

runs: up to 3000 nodes about 87% of the whole machine (3456 nodes)

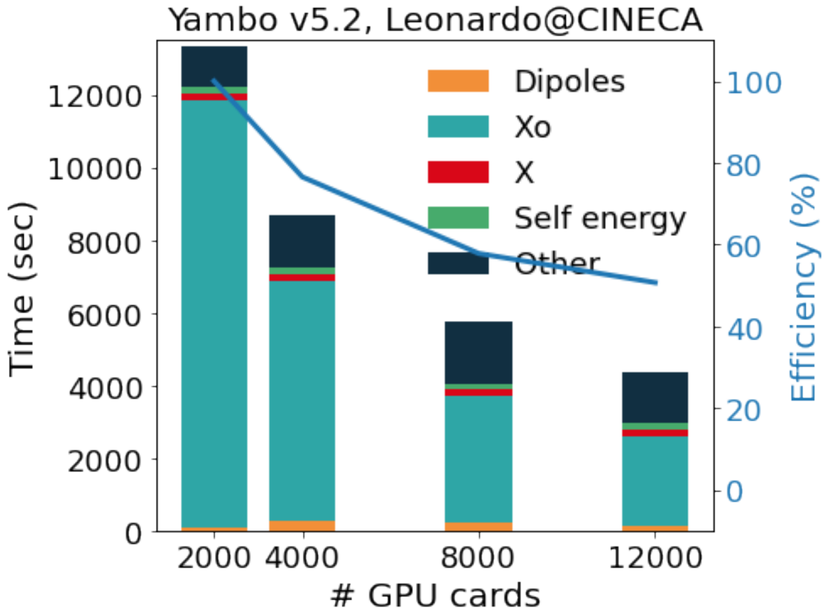

MAX DRIVING THE EXASCALE TRANSITION

Yambo

Contributors: AF, N. Spallanzani, D. Sangalli, A. Marini, D. Varsano

Support from vendors: Nvidia, AMD, Intel

HPC Centres: Cineca, IT4I

Ab initio many-body perturbation theory (MBPT)

Yambo

- A. Marini, C. Hogan, M. Gruning, D. Varsano, Comp.Phys.Comm. **180**, 1392 (2009)
- D. Sangalli, et al, J. Phys.: Condens. Matter. **31**, 325902 (2019)

https://github.com/yambo-code

https://www.yambo-code.eu

- Implements **many-body methods** (post-DFT) for finite and extended systems
- eg: quasiparticle energies (GW), optical absorption, spectroscopies
- **plane wave basis set** and pseudo-potentials
- Natively interfaced to Quantum ESPRESSO (MaX) and Abinit

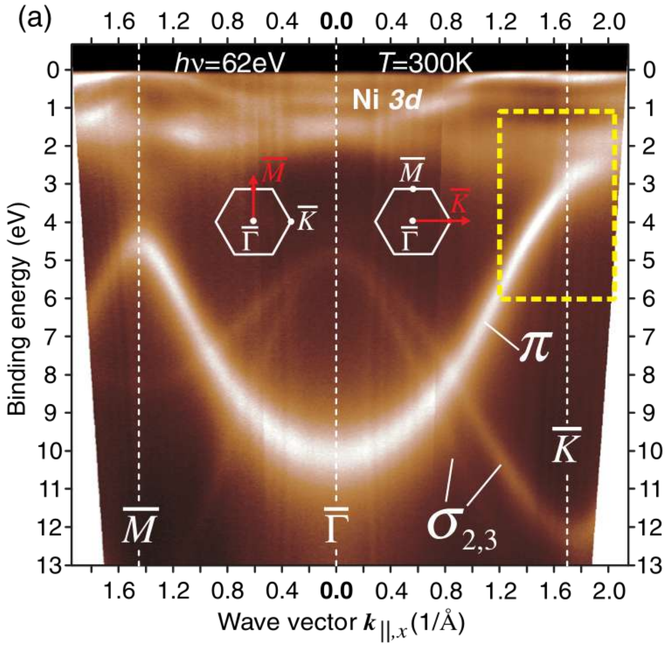

Varykhalov et al, PRX **2**, 041017 (2012)

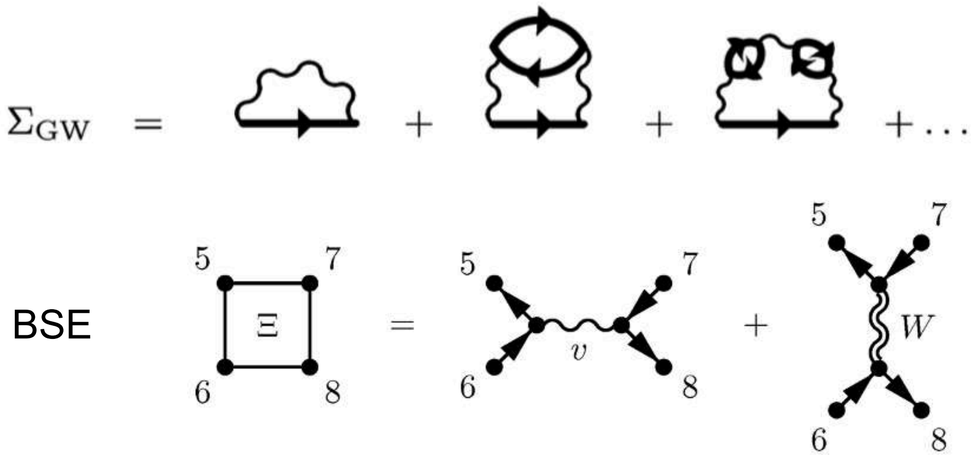

Ab initio many-body perturbation theory (MBPT)

$$\left[-\frac{1}{2}\nabla^2 + v^{\mathrm{KS}}(\mathbf{x})\right]\psi_{n\mathbf{k}}(\mathbf{x}) = \epsilon_{n\mathbf{k}}\psi_{n\mathbf{k}}(\mathbf{x})$$

- **DFT** simulation as **input**
- interfaced to QE (MaX), Abinit

$$G(\mathbf{x},\mathbf{x}',\omega) = \sum_{n\mathbf{k}} \frac{\psi_{n\mathbf{k}}(\mathbf{x})\psi^*_{n\mathbf{k}}(\mathbf{x}')}{\omega - \epsilon_{n\mathbf{k}} \pm i0^+}$$

- the **DFT Green's function** used to compute **MBPT** quantities

- microscopic screening

$$\chi_{\mathbf{q}}(\mathbf{G},\mathbf{G}',\omega)$$

$\Sigma_{\mathrm{GW}} =$ 

BSE 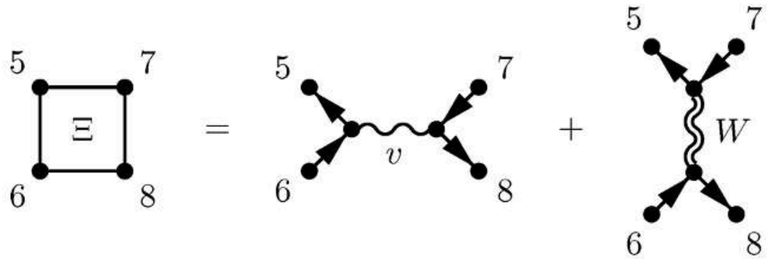

- A. Marini, C. Hogan, M. Gruning, D. Varsano, Comp. Phys. Comm. **180**, 1392 (2009)
- D. Sangalli, et al, J. Phys.: Condens. Matter. **31**, 325902 (2019)

Ab initio many-body perturbation theory (MBPT)

MAX DRIVING THE EXASCALE TRANSITION

Optical properties of materials

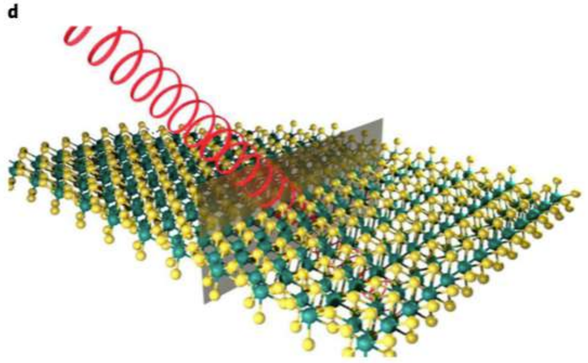

Varsano, Molinari et al,
Nature Nanotech 2020

Interfaces

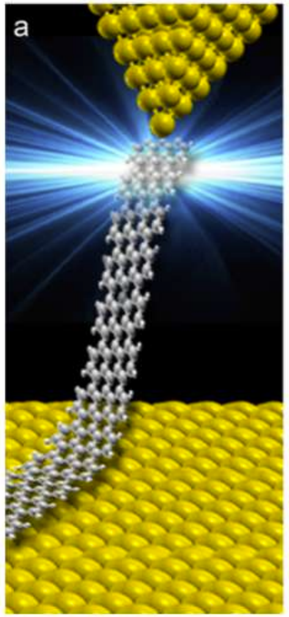

Cardoso, Ferretti, Prezzi et al,
Nano Lett 2018

Ultrafast spectroscopies

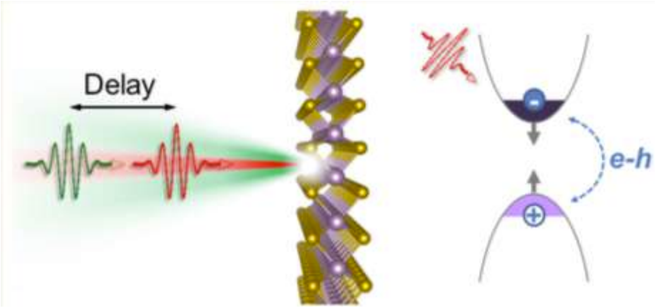

Pogna, Sangalli, Marini, Prezzi et al,
ACS Nano 2016

Photovoltaics, photocatalysis

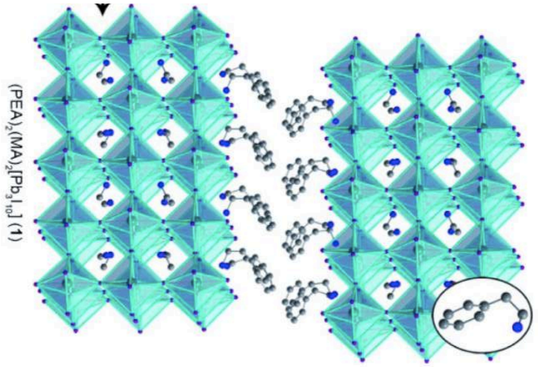

Palummo, Varsano et al,
Adv. En. Mater. 2020

Electron-phonon coupling

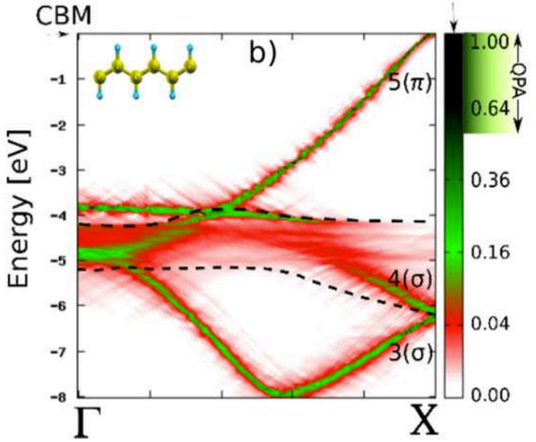

Cannuccia, Marini,
Eur. Phys. J. B 2012

Excited states and Theoretical spectroscopy

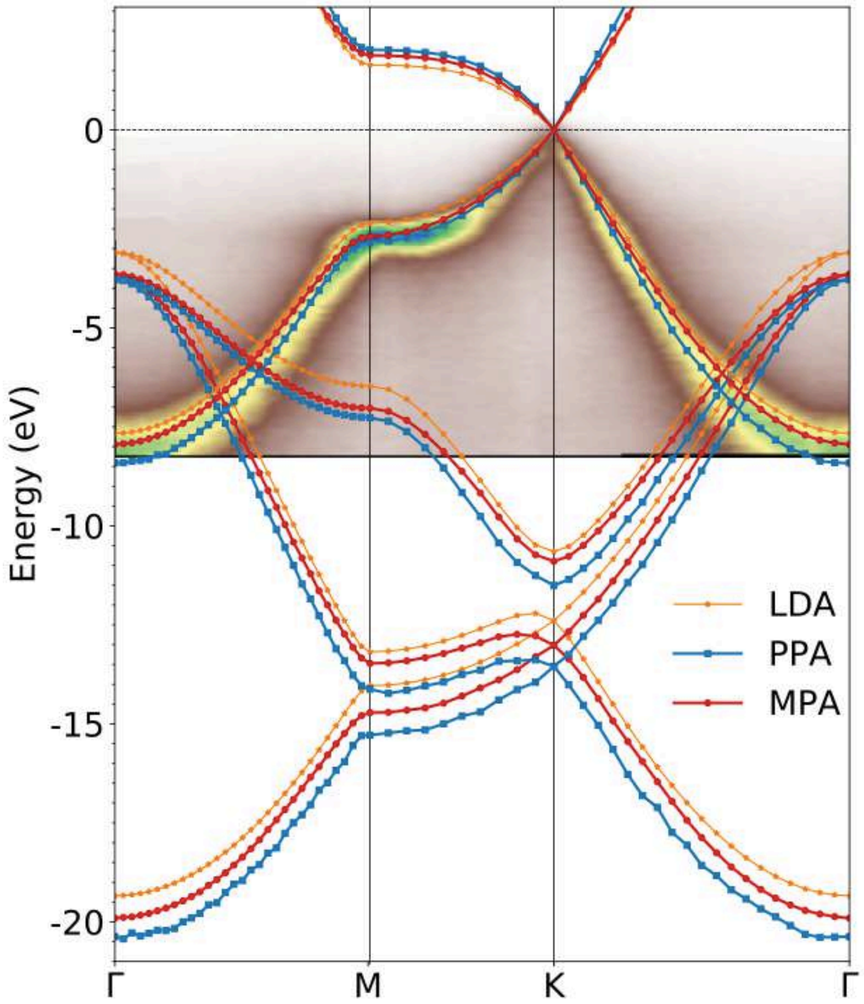

A. Guandalini, D.A. Leon, P. D'Amico, C. Cardoso, AF, M. Rontani, D. Varsano, PRB **109**, 075120 (2024)

ARPES and EELS of free standing **Graphene**

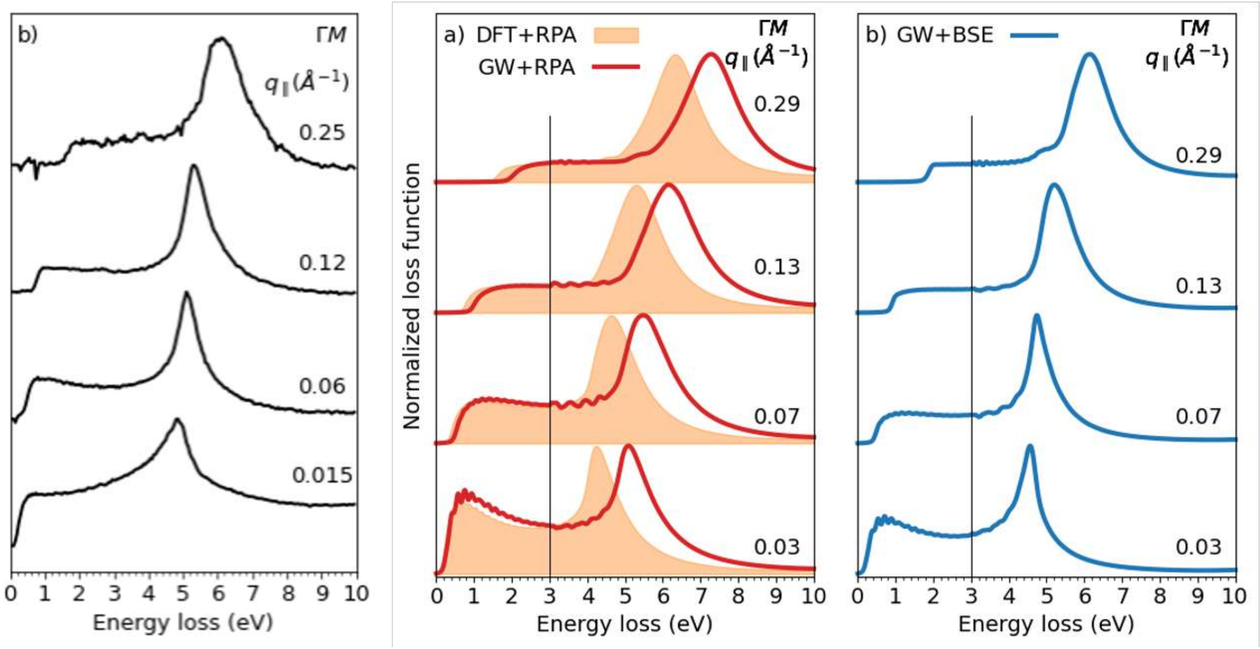

A. Guandalini, R. Senga, Y. Lin, K. Suenaga, AF, D. Varsano, A. Recchia, P. Barone, F. Mauri, T. Pichler, C. Kramberger, NanoLett **23**, 11835 (2023)

Excited states and Theoretical spectroscopy

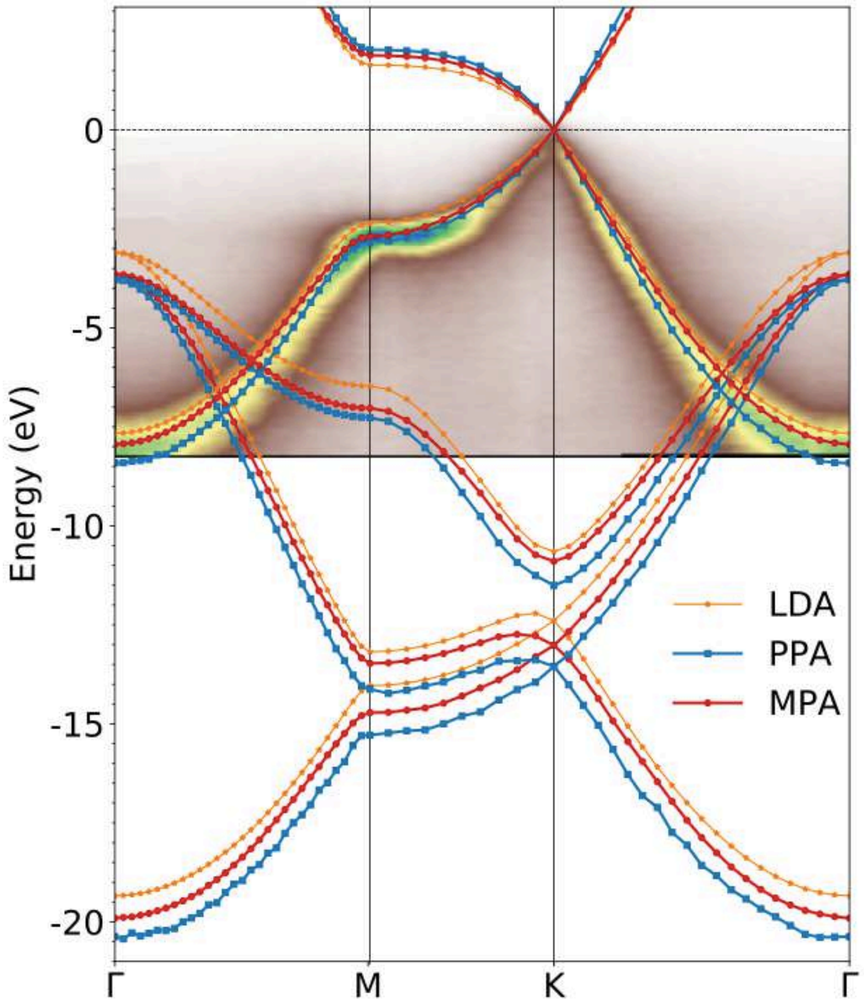

A. Guandalini, D.A. Leon, P. D'Amico, C. Cardoso, AF, M. Rontani, D. Varsano, PRB **109**, 075120 (2024)

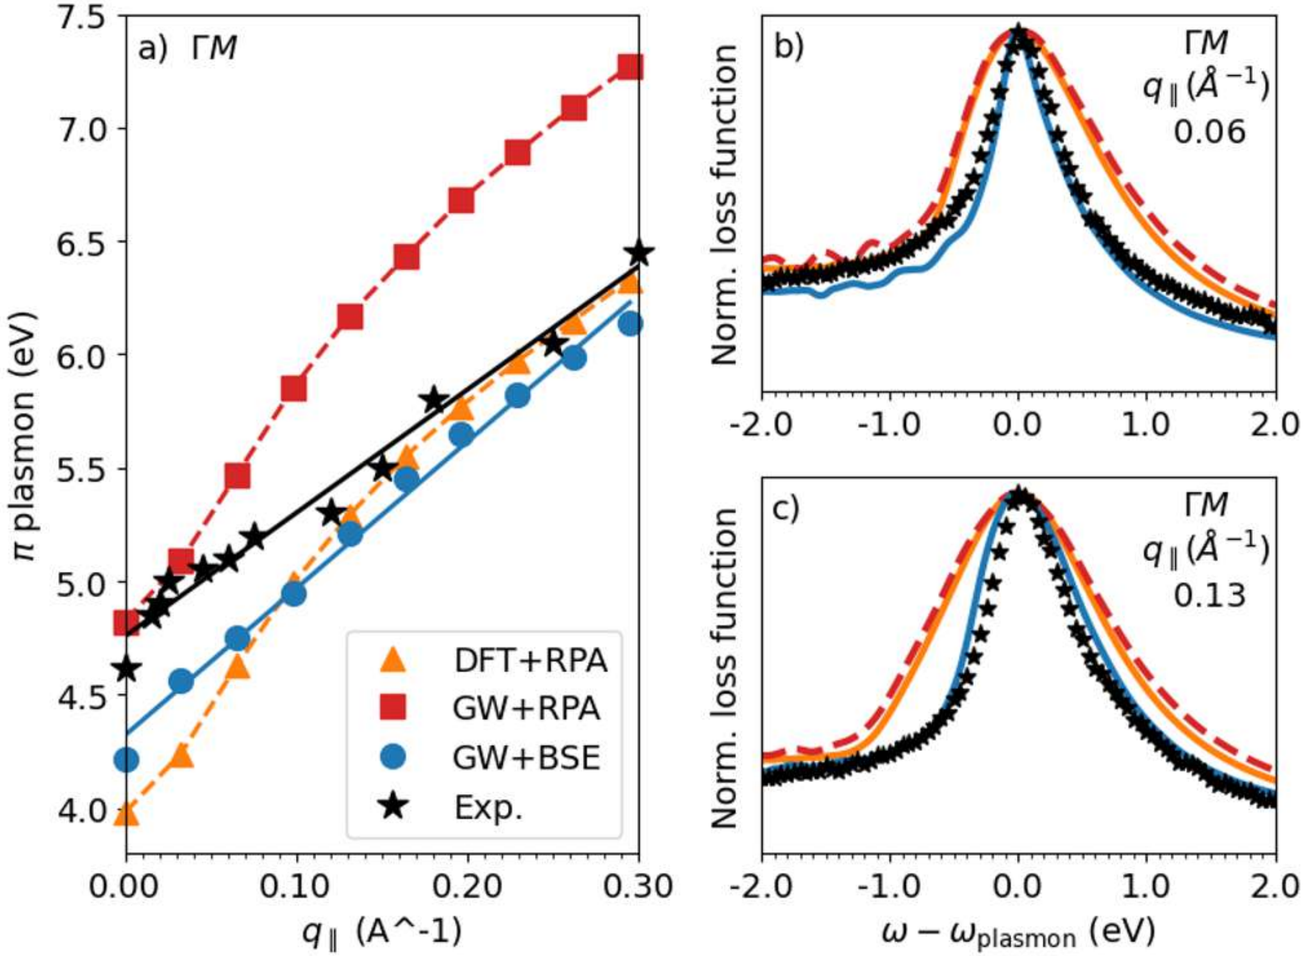

A. Guandalini, R. Senga, Y. Lin, K. Suenaga, AF, D. Varsano, A. Recchia, P. Barone, F. Mauri, T. Pichler, C. Kramberger, NanoLett **23**, 11835 (2023)

performance portability

https://github.com/yambo-code

https://www.yambo-code.eu

DeviceXlib

- deviceXlib is a **library** that wraps **device-oriented routines** and utilities, such as device data allocation, initialization, host-device data transfers, linear algebra…
- deviceXlib supports **CUDA-Fortran, OpenACC and OpenMP-GPU** programming models.
- Includes support for **CUDA, ROCm/HIP, OneAPI** environments**.**
- deviceXlib wraps a subset of functions from NVIDIA cuBLAS, INTEL oneMKL BLAS, and AMD rocBLAS libraries.

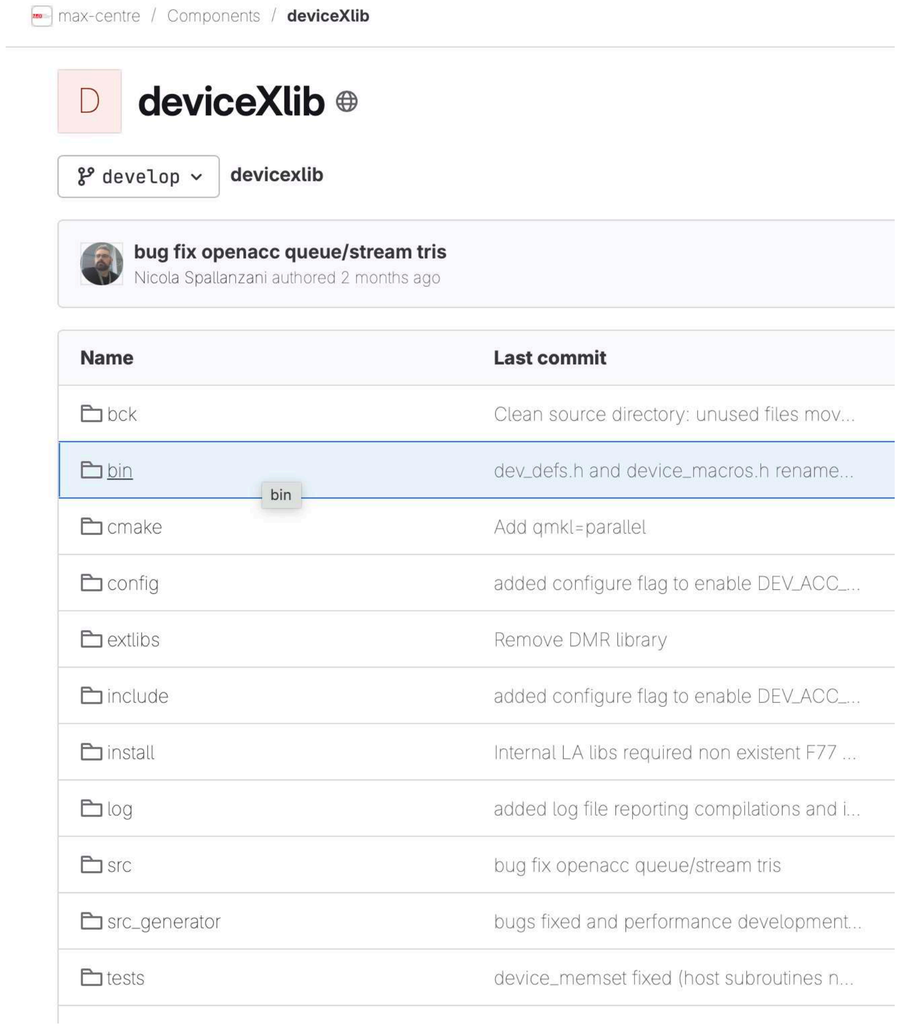

http://gitlab.com/max-centre/components/devicexlib/

Performance portability

- Released for **production on NVIDIA GPUs** using CUDA-Fortran (ref production-version)

- Implementation of **OpenACC (NVIDIA) and OpenMP-GPU (AMD and INTEL)** refined and optimized
- Done by exploiting **DeviceXlib 0.9.0**

- Yambo supports **multiple GPU backends**
- Source base significantly streamlined
- **Yambo running** on
  - NVIDIA-GPUs (eg Leonardo,MN-V)
  - AMD-GPUs (eg LUMI)
  - ARM-chips (eg Deucalion, Fugaku)

```
 rho_tw_rs_p => DEV_VAR(isc%rho_tw_rs)
 rhotw_p     => DEV_VAR(isc%rhotw)
 [...]
 !
 ! ordinary implementation
 !
 !DEV_ACC data present(rho_tw_rs_p,WF_symm_i_p,WF_symm_o_p)
 !DEV_ACC parallel loop async
 !DEV_CUF kernel do(1) <<<*,*>>>
 !DEV_OMPGPU target map(present,alloc:rho_tw_rs_p,WF_symm_i_p,WF_symm_o_p)
 !DEV_OMPGPU teams loop
 !DEV_OMP parallel default(shared), private(ir)
 !DEV_OMP do
 do ir = 1, fft_size
   rho_tw_rs_p(ir) = cmplx(conjg(WF_symm_i_p(ir,1))*WF_symm_o_p(ir,1),kind=DP)
 enddo
  !
 if (n_spinor==2) then
   !DEV_ACC parallel loop async
   !DEV_CUF kernel do(1) <<<*,*>>>
   !DEV_OMPGPU teams loop
   !DEV_OMP do
   do ir = 1, fft_size
     rho_tw_rs_p(ir) = rho_tw_rs_p(ir)+cmplx(conjg(WF_symm_i_p(ir,2))* &
                     & WF_symm_o_p(ir,2),kind=DP)
   enddo
   !
 endif
 !DEV_OMP end parallel
 !DEV_OMPGPU end target
 !DEV_ACC end data
 [...]
 !
 ! perform the actual FFT
 !
#if defined _GPU_LOC
#  if defined _CUDA
 call fft_3d_cuda(rho_tw_rs_p,fft_dim,+1,cufft_plan)
#  elif defined _HIP
 call fft_3d_hip(rho_tw_rs_p,fft_dim,+1,hipfft_plan)
#  endif
```

Yambo @ Leonardo

GOAL: turn MaX flagship codes into exascale-enabled applications

- **large scale MPI** parallelism (order of 10,000 tasks)
- combined with **GPU awareness**

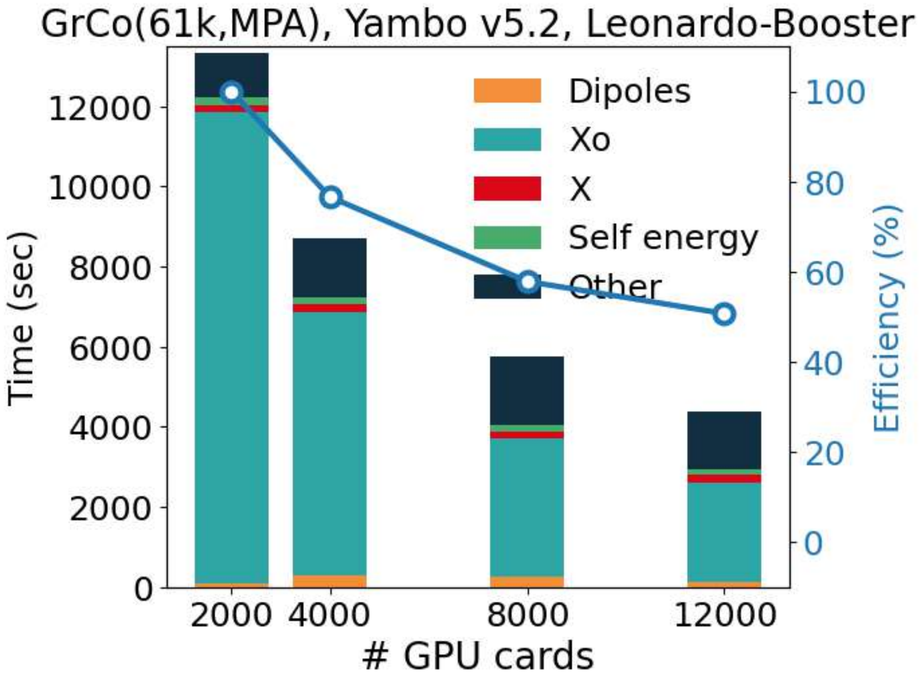

Leonardo @ CINECA:
4 Nvidia A100 next / node

runs: up to 3000 nodes about 87% of the whole machine (3456 nodes)

Yambo @ LUMI

- Featuring **AMD MI250**
- Comparison with Nvidia A100

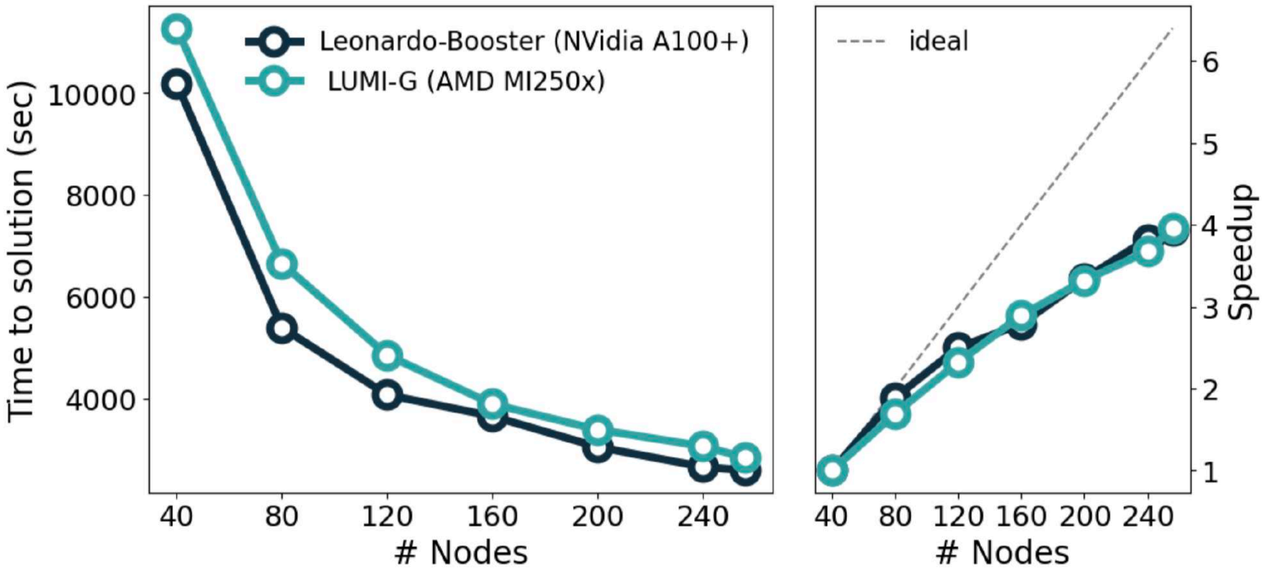

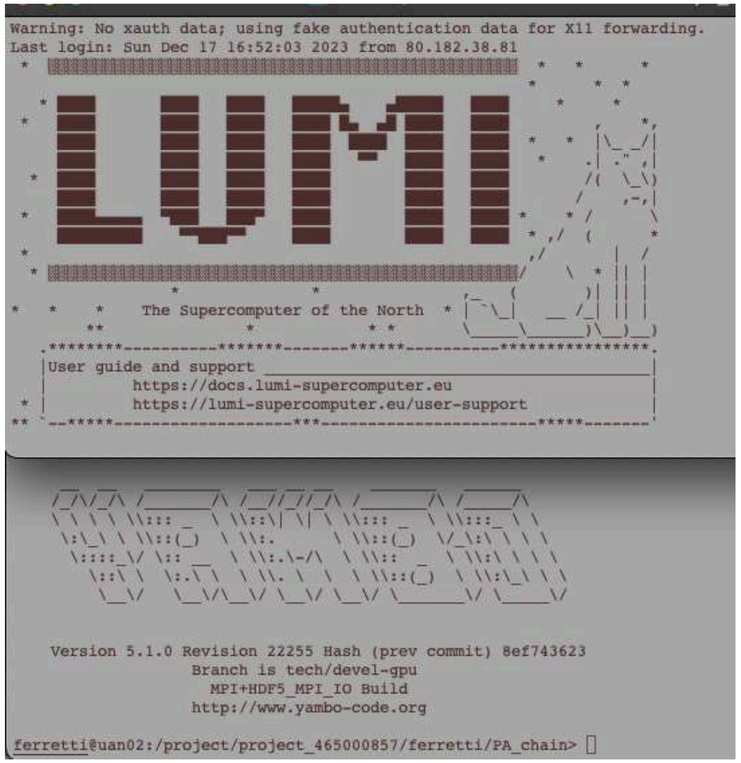

- Initial support of Cray v15 extended to Cray v17

Yambo @ MareNostrum-V

- Featuring **Nvidia H100**
- Use of Nvidia **MPS tested**
- Test relevant also for **porting on Jupiter**

Data: courtesy of J. Gutierrez and R. Grima (BSC)

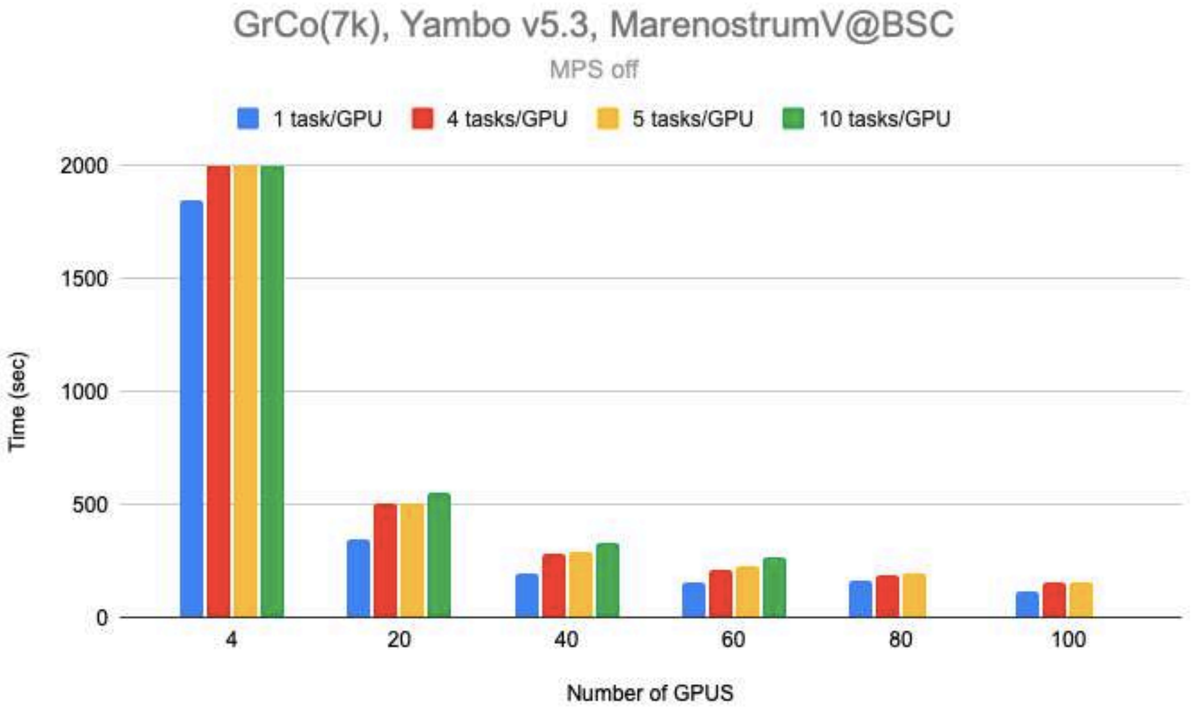

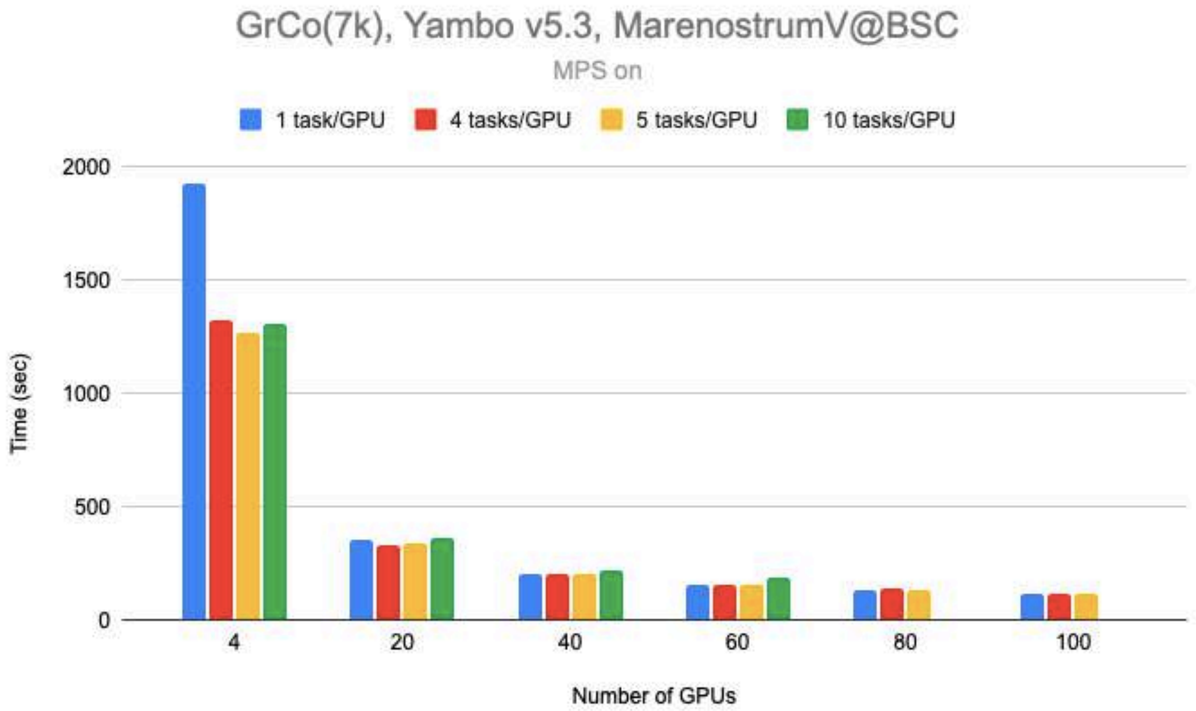

MPS off

MPS on

Yambo @ ARM-based chips

- Runs on **Deucalion featuring Fujitsu A64FX**
- Similar runs also on **Fugaku**
- Comparison with x86 chips (**Leonardo DCGP, LUMI-C, Deucalion-x86**)

HANAMI

(Courtesy of Costanza Borghesi and Nicola Spallanzani)

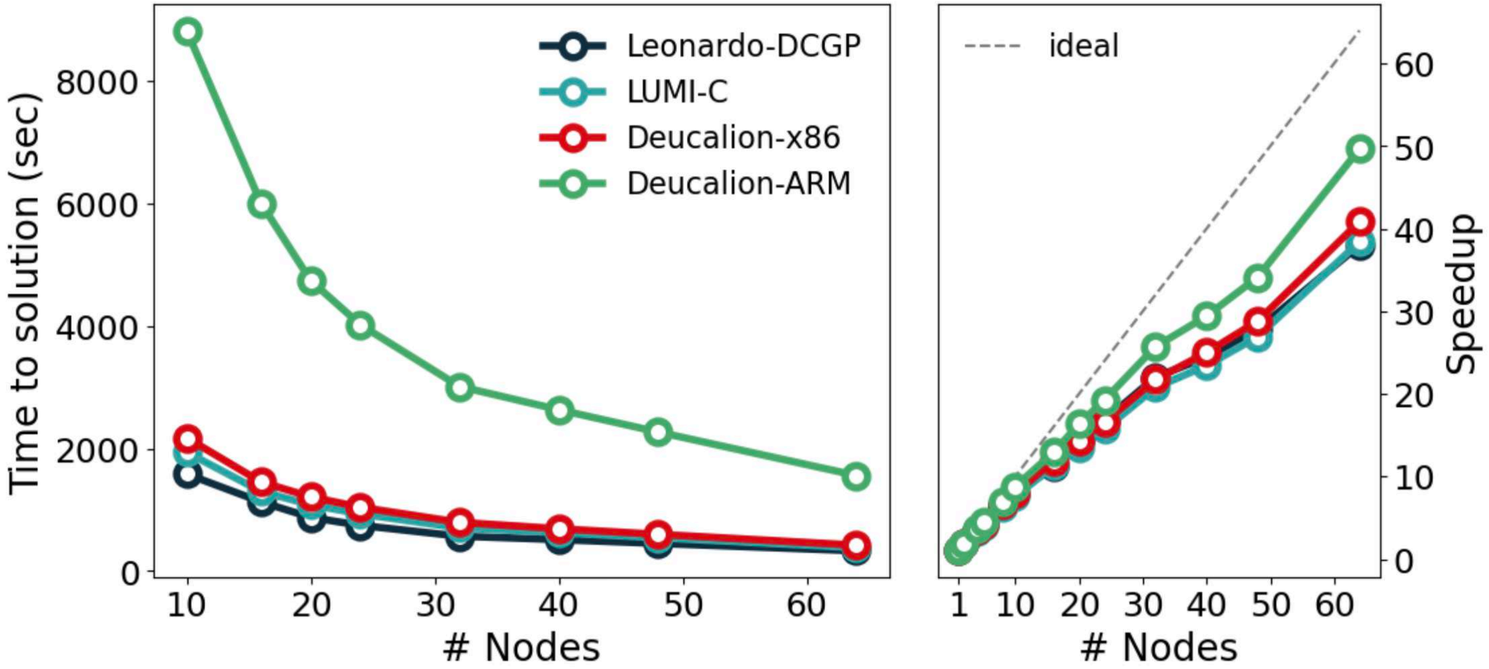

Algorithmic and parallel optimisation

https://github.com/yambo-code

https://www.yambo-code.eu

Optimisation: Response function and MPA

- The **Yambo mini-app** has been used to run extended **profiling** on **different architectures** (both CPU and GPU)
- Bottlenecks identified and **optimisation strategies** proposed
- multiple solutions implemented and validated **in the mini-app**
- Eventually, the most efficient **solution was ported to Yambo**

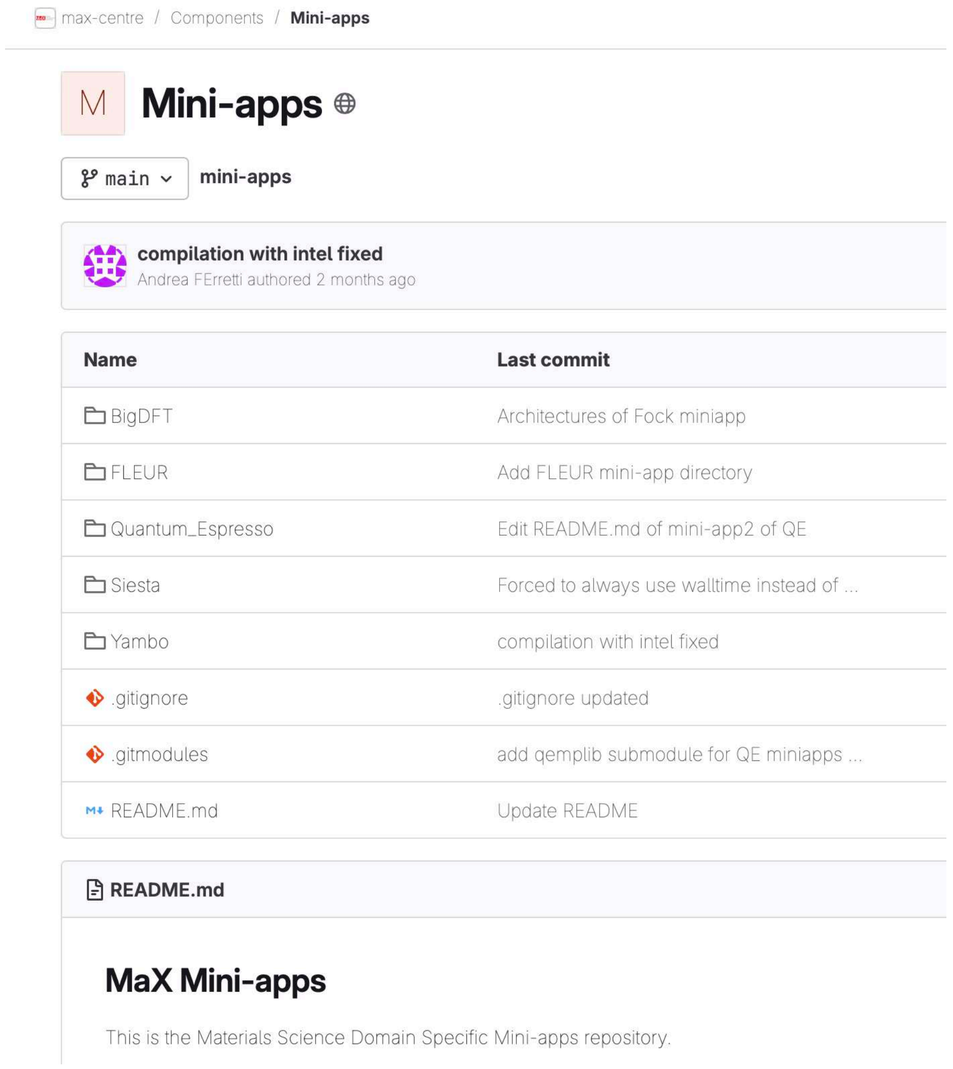

http://gitlab.com/max-centre/components/mini-apps

Optimisation: Response function and MPA

- The **Yambo mini-app** has been used to run extended **profiling** on **different architectures** (both CPU and GPU)
- Bottlenecks identified and **optimisation strategies** proposed
- multiple solutions implemented and validated **in the mini-app**
- Eventually, the most efficient **solution was ported to Yambo**

- Contributions from a large number of MaX partners: CNR, EVIDEN, SIPEARL, IT4I, CINECA
- **Overall**: **massive gain on CPU (50%), smaller though sensible gain on GPUs (20%)**
- Xo **algorithm refactored, additional parallelism** in MPA
- Allowed also to expose additional CGEMM calls on GPUs, suitable for **precision emulation on AI-oriented cards**

Performance progress on Xo driver and MPA self-energy

Yambo version	Nodes	Tasks	Time-Profile	Xo	Self Energy
5.3.0	4	16	2327.00	1789.00	462.63
5.4-beta	4	16	1080.00	656.00	348.40
		speedup	2.1	2.7	1.3

Optimisation: Response function and MPA

- The **Yambo mini-app** has been used to run extended **profiling** on **different architectures** (both CPU and GPU)
- Bottlenecks identified and **optimisation strategies** proposed
- multiple solutions implemented and validated **in the mini-app**
- Eventually, the most efficient **solution was ported to Yambo**

- Contributions from a large number of MaX partners: CNR, EVIDEN, SIPEARL, IT4I, CINECA
- **Overall**: **massive gain on CPU (50%), smaller though sensible gain on GPUs (20%)**
- Xo **algorithm refactored, additional parallelism** in MPA
- Allowed also to expose additional CGEMM calls on GPUs, suitable for **precision emulation on AI-oriented cards**

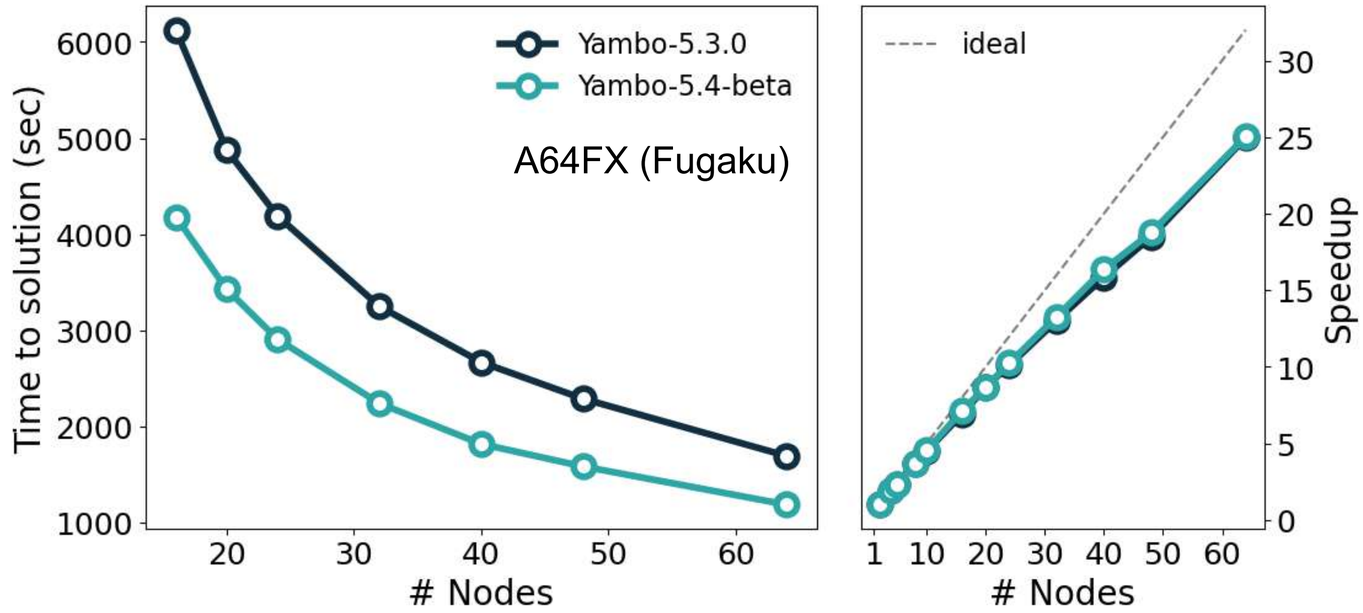

Linear algebra

Parallel efficiency

- **GW runs**: dominated by **load imbalance** and **linear algebra** kernels (dense linear systems)
- **BSE runs**: important role of **linear algebra** solvers (diagonalization)

- **Assessment** of parallel performance and bottleneck identification
- **Linear algebra bottlenecks addressed**

Dense & iterative solvers

- GPU acceleration via **MAGMA (serial) and ELPA (parallel) support** implemented, for both hermitian and non-hermitian diagonalizations
- **GPU** acceleration via **SLEPc enabled**
- **Advanced diagonalisation** algorithms (Krylov-Schur based) enabled via **SLEPc**
- **Nvidia cuSOLVERMp** interfaced and validated (thanks to CINECA)

HANAMI

- Support of **ChASE library** under testing (in collaboration with E. Di Napoli, X. Wu, C. Richefort, Juelich)
- **Overall**: Restructuring of the main interface layer to LA libs ongoing

Yambo-cuSOLVERMp interface

- **Nvidia cuSOLVERMp** interfaced and validated
  (Courtesy of S. Orlandini and M. Montagna — CINECA)

		cuSOLVERMp		ScaLAPACK		
Nodes	Tasks/Node	Linear Algebra	getrf/getrs	Linear Algebra	SU LA	SU func
1	4	118.16	73.58	14760.00	124.9x	200.6x
4	4	116.46	61.05	4080.00	35.0x	66.9x
16	4	119.72	60.41	1063.00	8.9x	17.6x

Figure 12: Comparison of `cusolverMP` and `ScaLapack` in solving a dense linear system of size 21453×21453, as required by the calculation of the response function χ in YAMBO.

Maintenance, releases, deployment

https://github.com/yambo-code

https://www.yambo-code.eu

Software releases:

- Yambo 5.4 (beta)
- Yambo 5.3.0 (current stable)
- Yambo 5.2.1 - 5.2.4

https://github.com/yambo-code/yambo/tags

Maintenance, sustainability, deployment

Software releases:

- Yambo 5.4 (beta)
- Yambo 5.3.0 (current stable)
- Yambo 5.2.1 - 5.2.4

https://github.com/yambo-code/yambo/tags

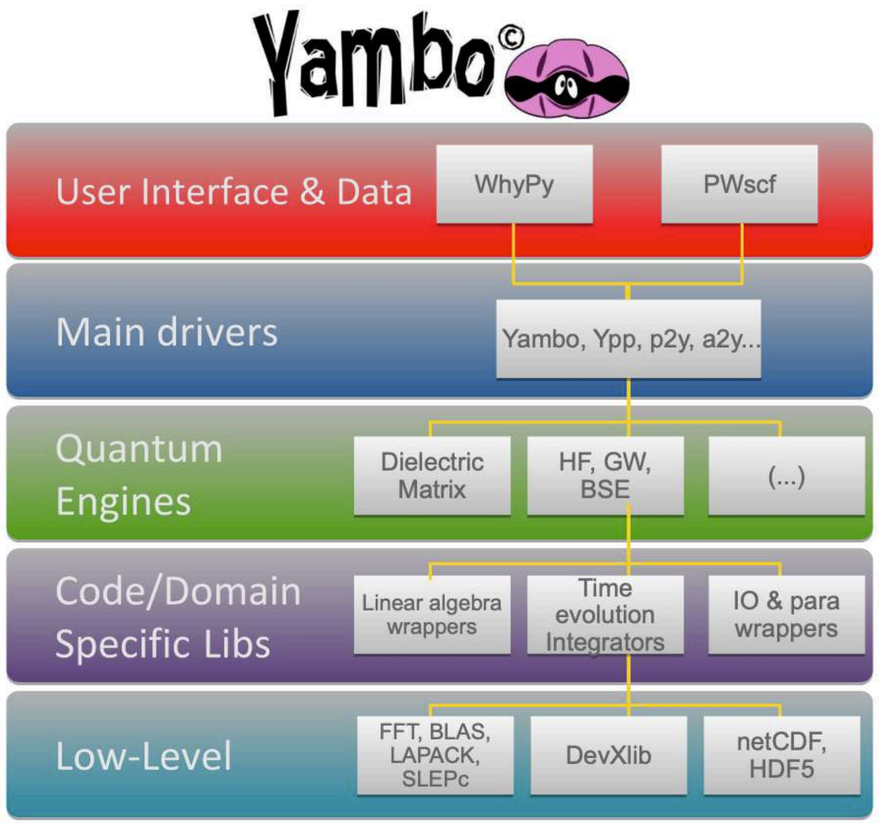

- Based on the principles of **HW/SW "separation of concerns", modularity**, data encapsulation

- Aims at **enabling maintainability** in the long term
- While **keeping the scientific community** engaged

- New main restructuring under design and testing (Yambo 6.0)

MaX-3 achievements (@RP2)

- Automatic build system updated (Cray LUMI env; AMD and INTEL GPU env, ARM)
- **3 different programming models** for GPUs supported (CUDA-F, OpenACC, OpenMP-GPU) **within a single source**
- Source **code streamlined**, with improved readability (=> communities)
- support of **Spack** and **EasyBuild recipes** further developed
- **Mini-app** for relevant **kernel dominating performance** (Xo) extracted and used for blue printing

Benchmarking and profiling

https://gitlab.com/max-centre/JUBE4MaX

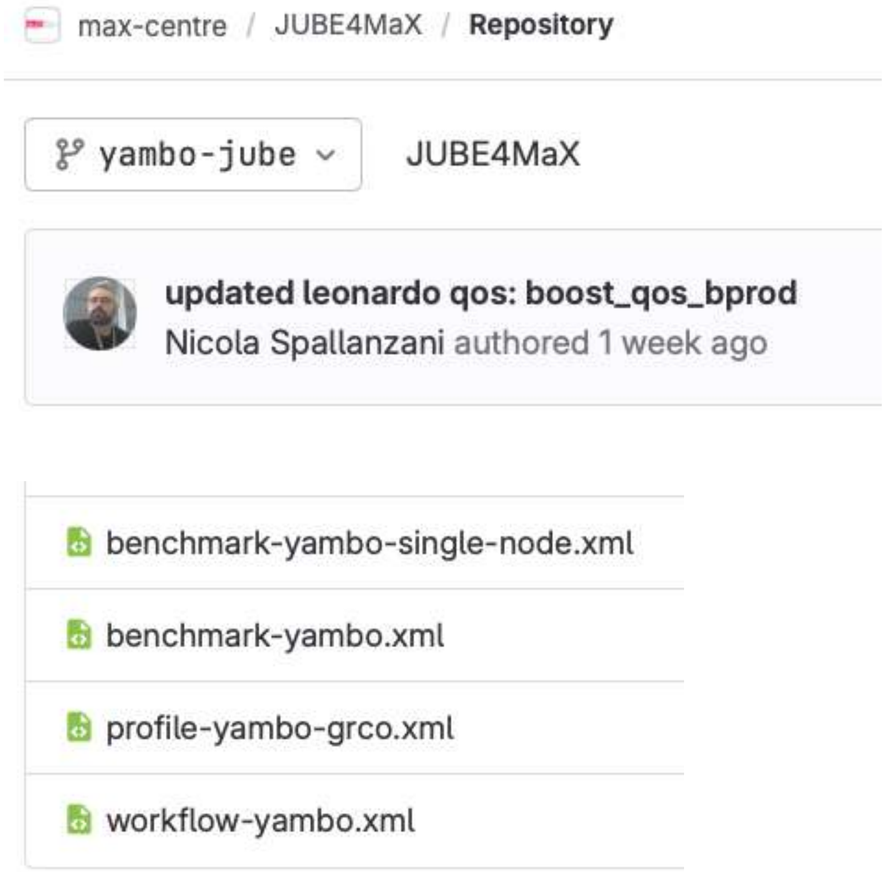

- benchmarks are collected in a **dedicated repo**
- **work done in collaboration with CINECA**

- Using the **JUBE benchmarking environment, to**
  - create benchmark sets
  - run benchmarks on different computer systems
  - evaluate and display the results

```
jube run benchmark-yambo.xml —tag leonardo yambo grco
```

```
jube run workflow-yambo.xml —tag leonardo yambo grco
```

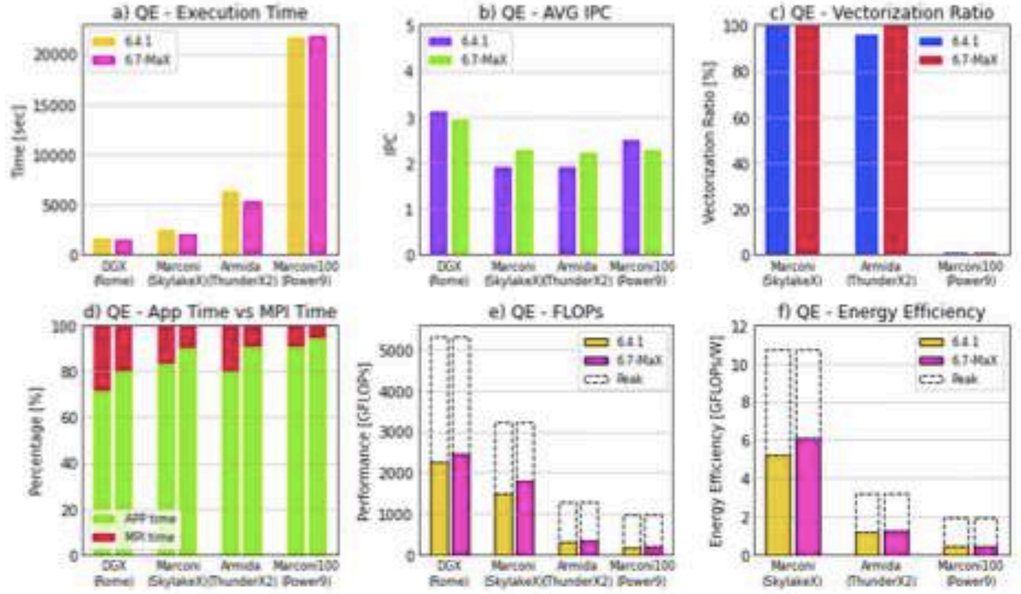

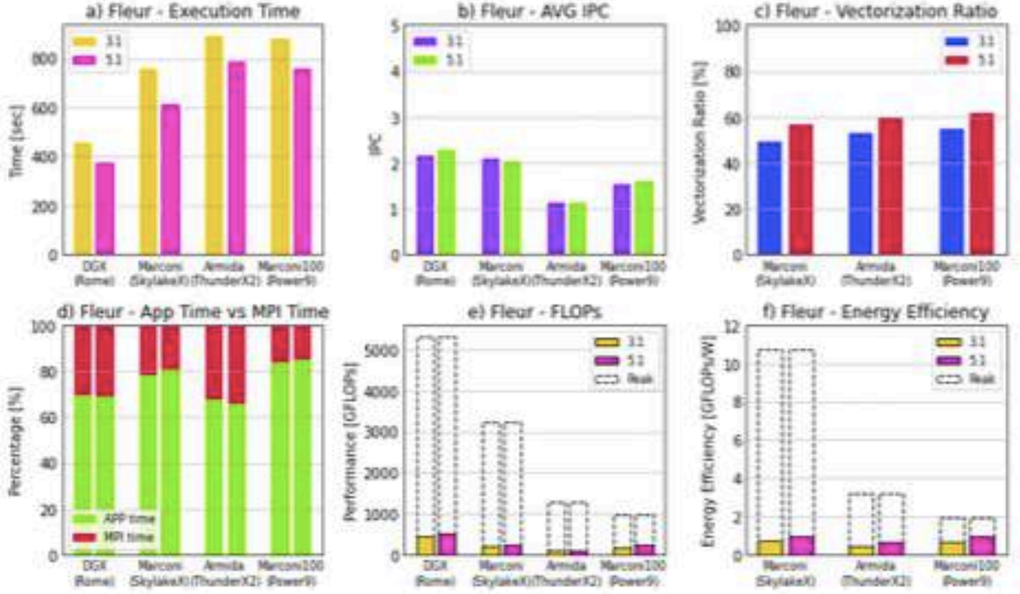

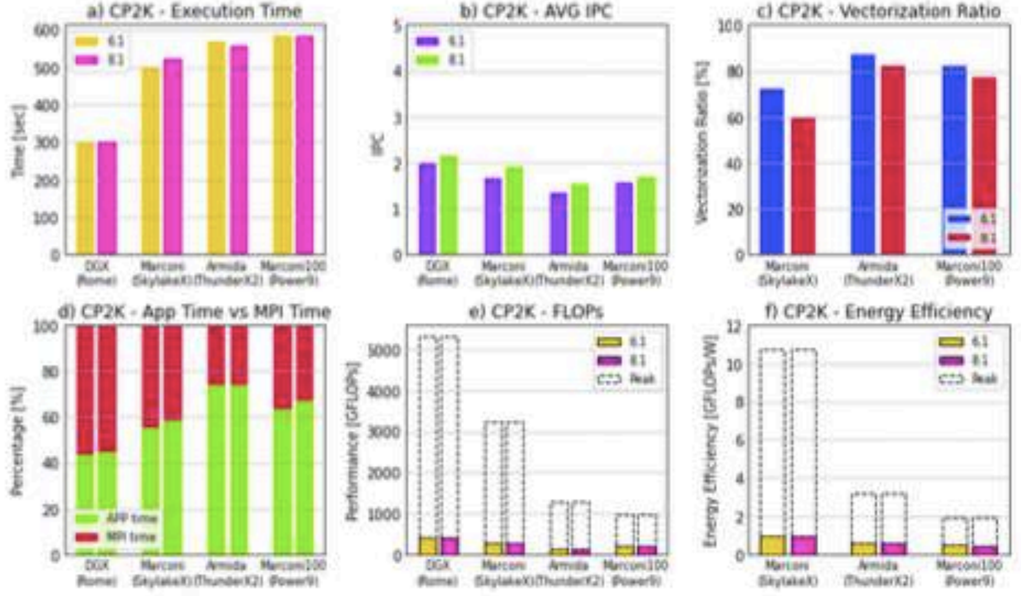

Multiple Metrics

For each code:

- execution time
- instruction per cycle
- vectorization ratio
- app time vs MPI time
- FLOP/s
- energy efficiency

Yambo deployment on EuroHPC machines

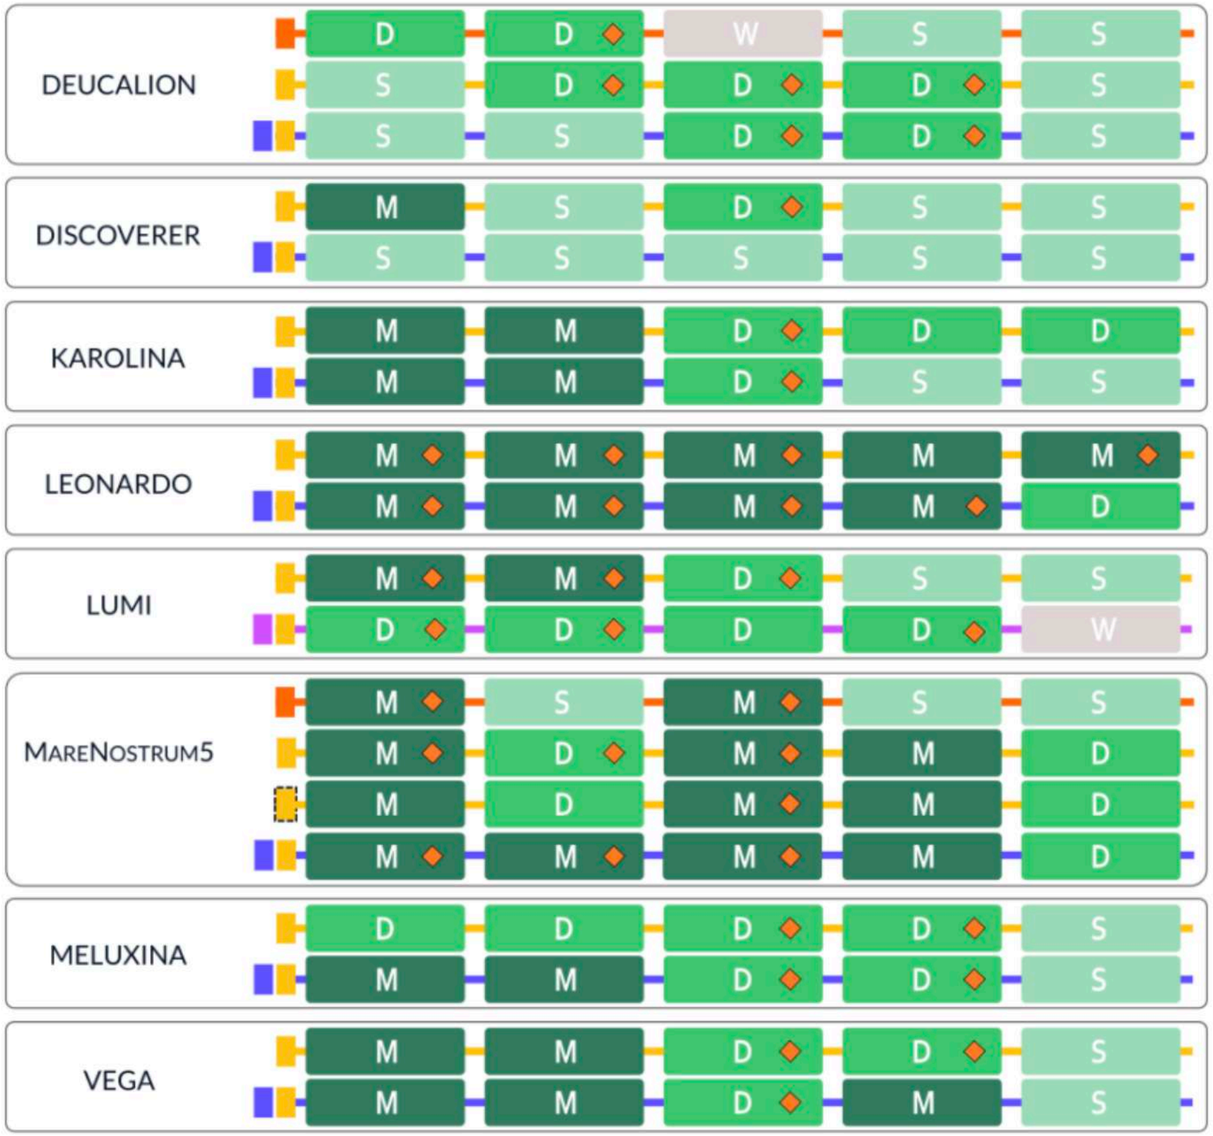

M
Deployed
module available

D
Demonstrated
developers installation

S
Supported architectures
tested on similar architectures

W
Work in progress

◆
Benchmarks
configured and executed

CPU ARM
CPU x86
··· + HBM

GPU AMD
GPU NVIDIA

Yambo deployment on EuroHPC machines

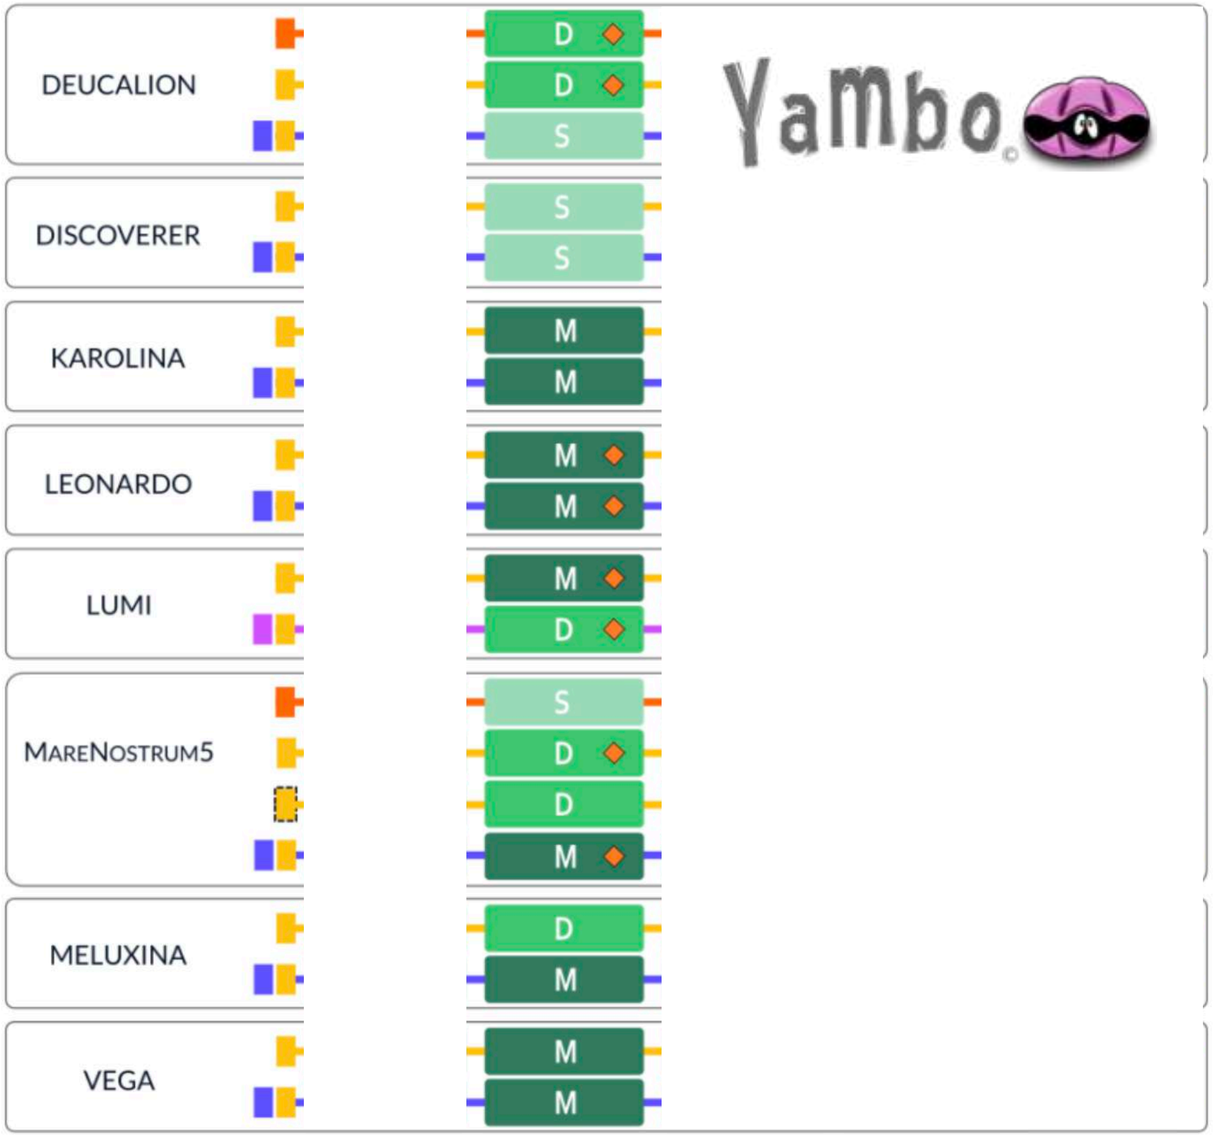

M
Deployed
module available

D
Demonstrated
developers installation

S
Supported architectures
tested on similar architectures

W
Work in progress

◆
Benchmarks
configured and executed

CPU ARM
CPU x86
--- + HBM

GPU AMD
GPU NVIDIA

new features and algorithms

https://github.com/yambo-code

https://www.yambo-code.eu

Adv methods in electronic structure

nature materials

INSIGHT | REVIEW ARTICLE

https://doi.org/10.1038/s41563-021-01013-3

Check for updates

Electronic-structure methods for materials design

Nicola Marzari [1]✉, Andrea Ferretti [2] and Chris Wolverton [3]

Density functional theory (DFT):

- applications ranging from **materials modelling, to quantum chemistry and drug design**
- compatible with **high performance computing** and **high-throughput screening**

beyond DFT:

- **multiple hierarchies** can be climbed
  - wavefunction-based methods
  - **many-body perturbation theory and spectral methods**
  - time-dependent and non-equilibrium methods
  - **ensembles**

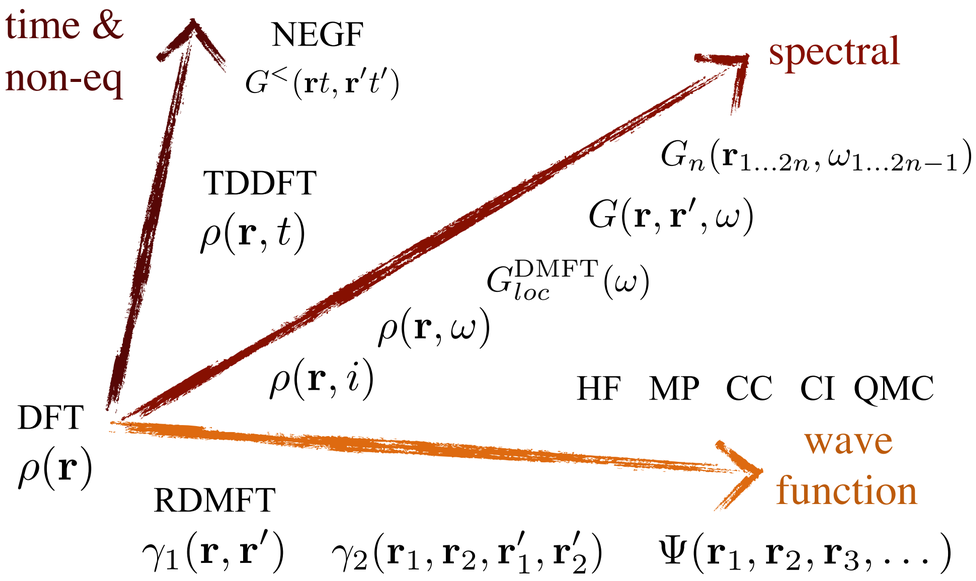

New Features

Coupled electron-ion dynamics

- **Reconstruction of the ionic potential** defined by Quantum ESPRESSO, as required by WP2 workflows
- Also, enables the possibility to compute **GW on top of DFT+U**

- **El-Ph module revamped**: Interface with the LetzElPhC code => exciton-phonon coupling

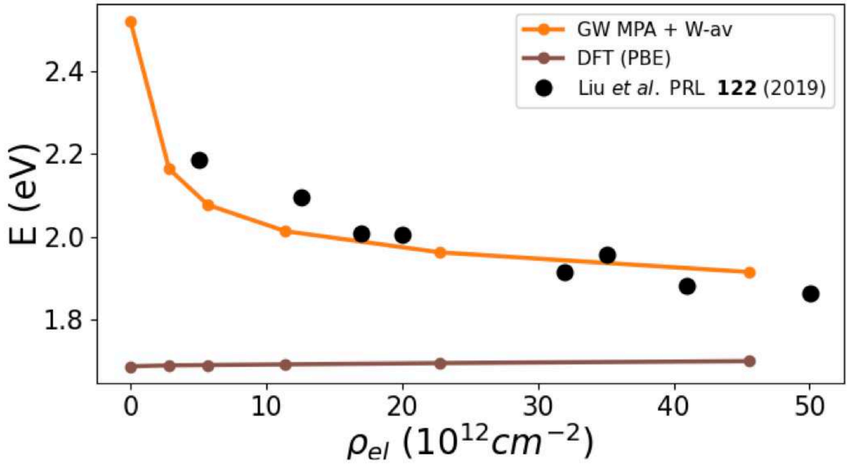

GW self-consistency

- implemented the procedure to perform **quasi-particle self-consistent GW (qsGW)**, as from [1]. In collaboration with C. Morice (Uni Paris Saclay)

- Aimed at reducing the initial state dependence of GW

- Development completed; preliminary testing passed; **undergoing validation**

Convergence accelerators for materials in reduced dimensionality

- The **RIM-W approach** (stochastic accelerator for 2D materials) has been ended to deal with metals
- Made compatible with **MPA full-frequency GW**

- Implemented, demonstrated, and published [2]

[1] T. Kotani, M. van Schilfgaarde, S.V. Faleev, PRB **76**, 165106 (2007)

[2] G. Sesti et al, arXiv 2508:06930 (2025)

STRONG SCALING AND SPEEDUP

Calculation of quasi-particle corrections on a graphene/Co interface (GrCo). The test involves the evaluation of the response function, Hartree-Fock, and correlation self-energy, treated at the GW level. On the left, the scalability test on LEONARDO-BOOSTER (CINECA), while on the right a speedup comparison between LEONARDO-DCGP, LEONARDO-BOOSTER and LUMI-G (CSC) for the same system but with reduced parameters.

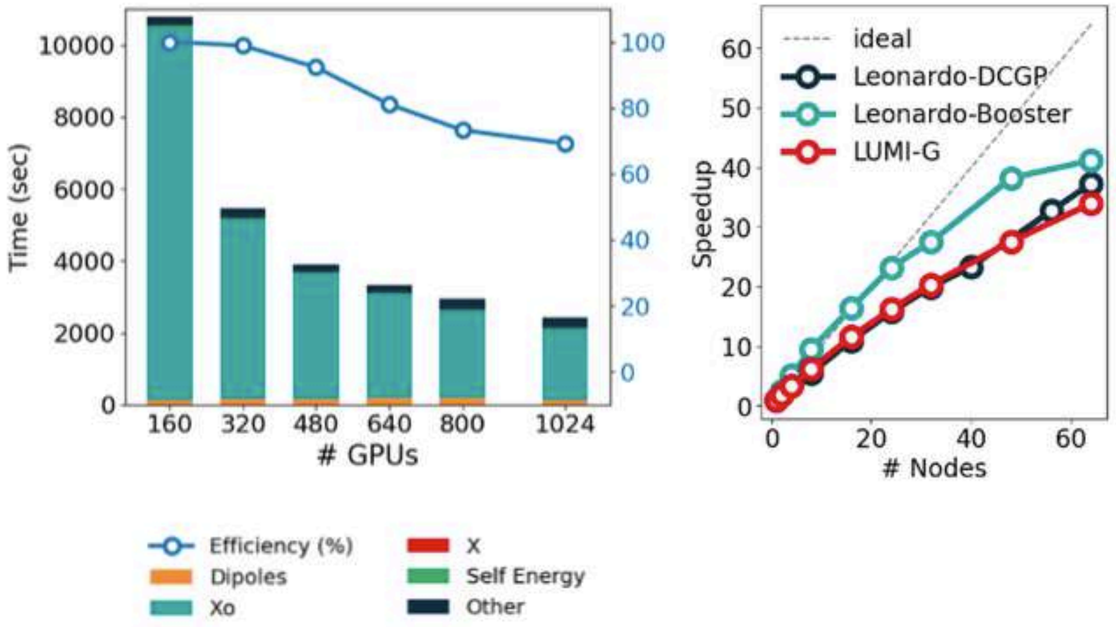

OPTICAL ABSORPTION SPECTRA

Optical absorption spectra computed at different levels of theory using YAMBO. IP@dft, IP@gw: independent particle response, evaluated on top of Kohn-Sham DFT and GW electronic structure. BSE: response function from the Bethe-Salpeter equation. The system studied is a defected (hydroxylated) TiO2 slab, relevant for photocatalysis, counting 210 atoms in the unit cell.

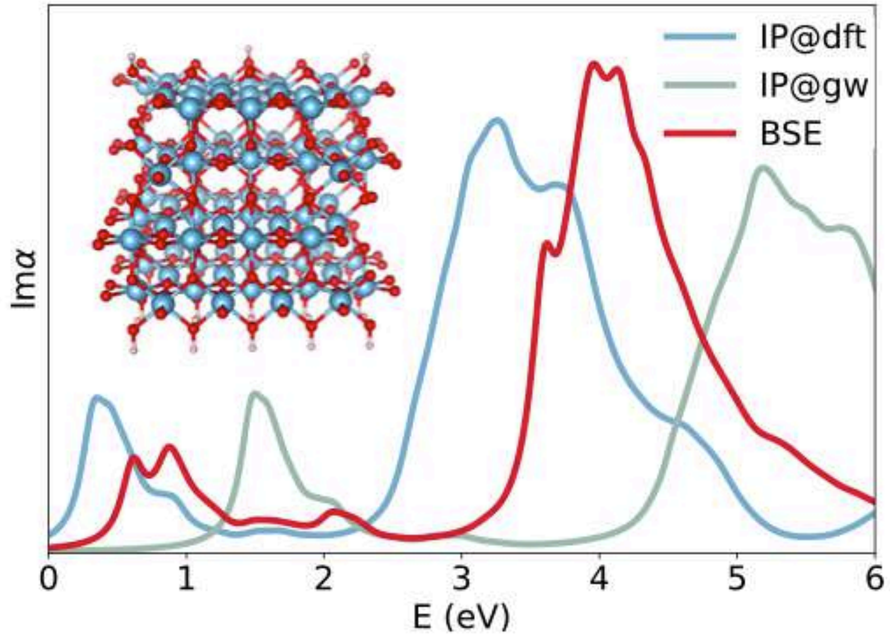

Workflows and data

Exascale Workflows and Data

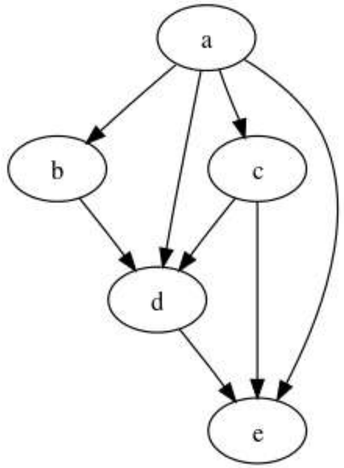

Direct acyclic graph (DAG)
Encodes **complex dependencies**

- using first principles **electronic structure methods**
- materials properties are usually not obtained as single calculations,
- but rather as **chains of interdependent calculations**

=> **workflows** (encoded as DAG)

Exascale Workflows and Data

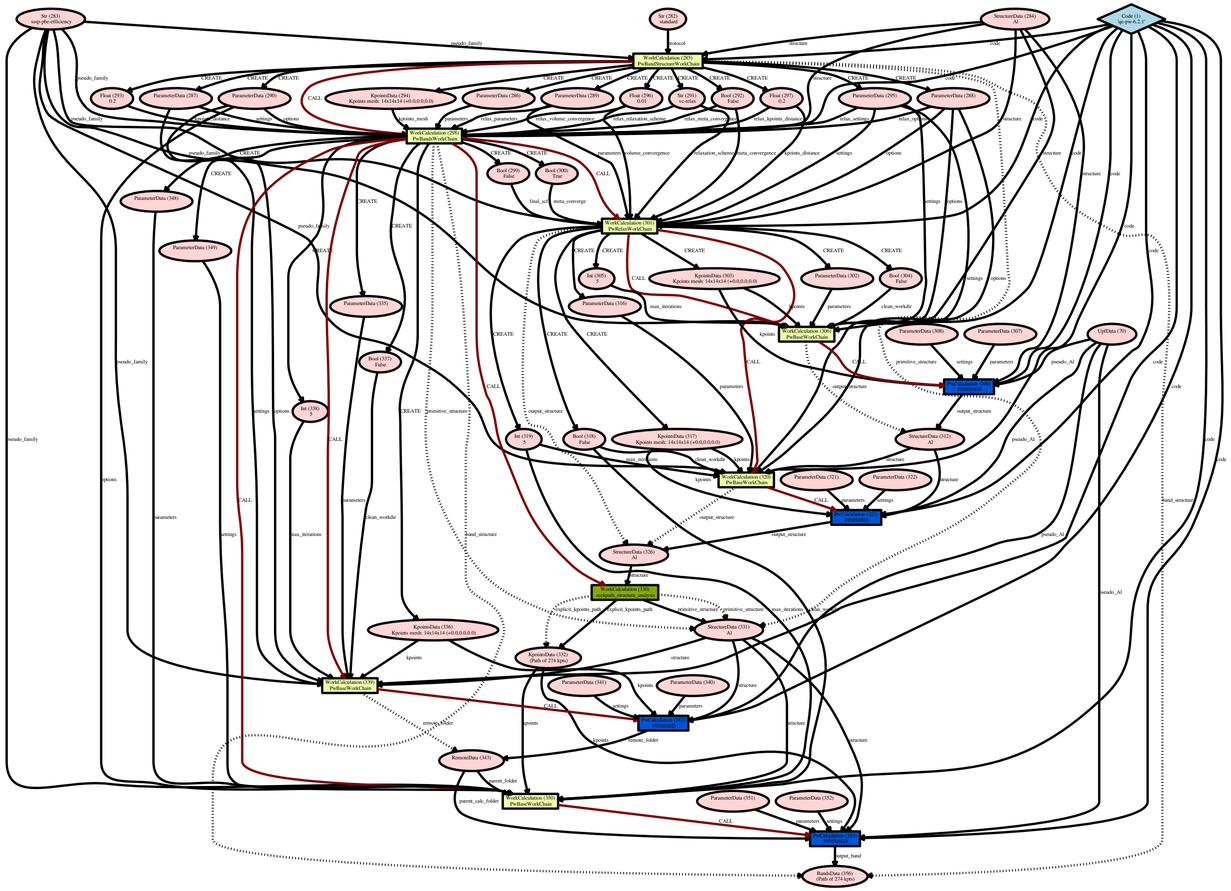

AiiDA

HyperQueue

- a **real-life AiiDA graph** encoding a materials science **workflow** via a DAG

taking workflows to exascale machines:
=> exascale workflow

- In an exascale workflow **each node will be a large scale calculation**
- MaX will use **AiiDA** as an orchestrator combined what the **HyperQueue** metacheduler
- exascale workflows encode **complex logic** and **improve resilience and parallelism**
- and preserve **full provenance and reproducibility**

Exascale Workflows and Data

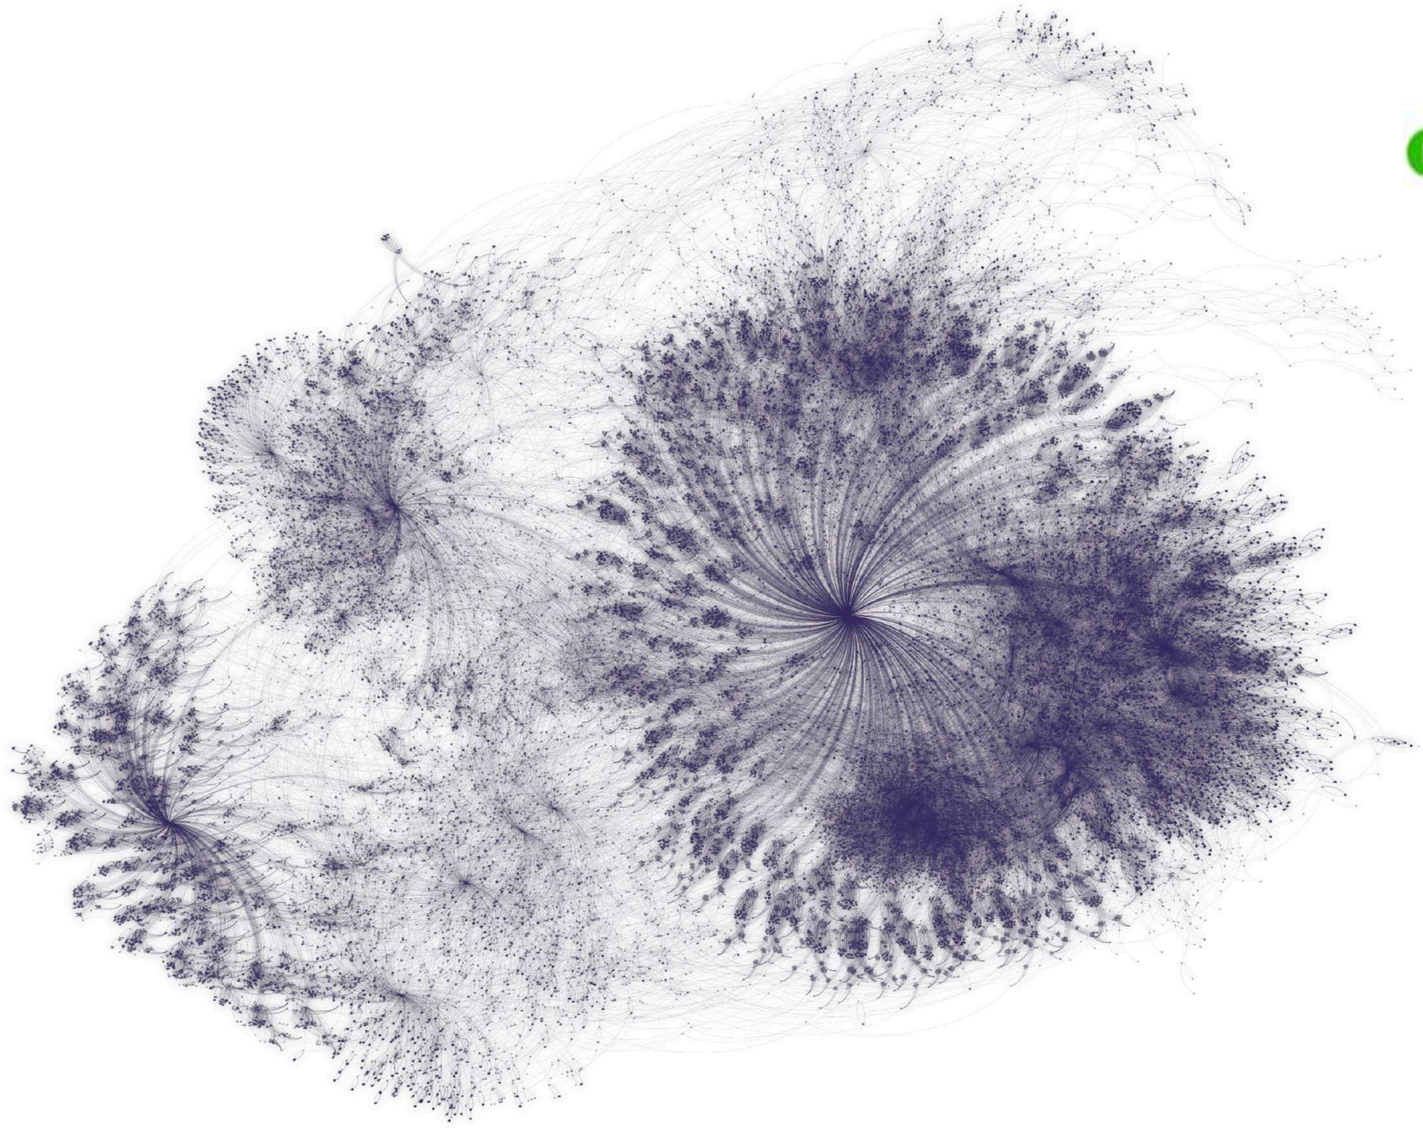

AiiDA

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

Scientific showcases

Thermal conductivity

$$\kappa = \frac{1}{3k_B T^2 V}\int_0^\infty \langle J_\alpha(t) J_\beta(0)\rangle \mathrm{d}t$$

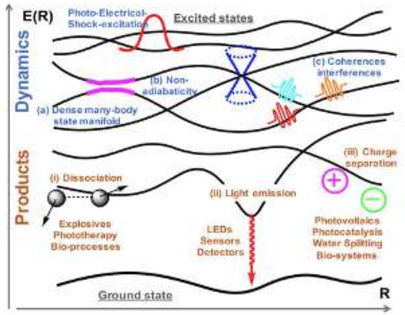

Non-adiabatic MD

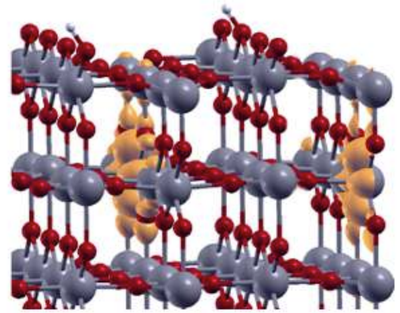

Polaron photocatalysis

Battery materials

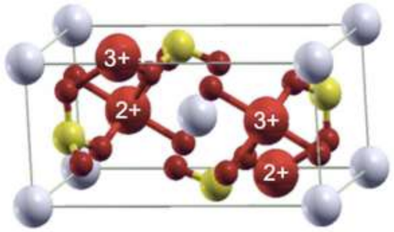

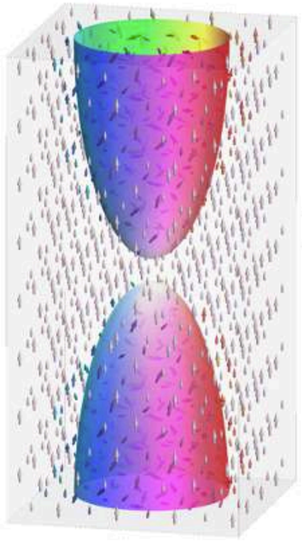

Micromagnetism

Coherent quantum states

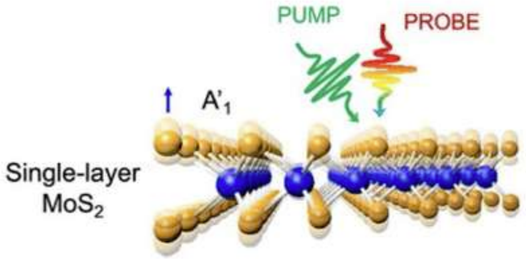

Protein interactions

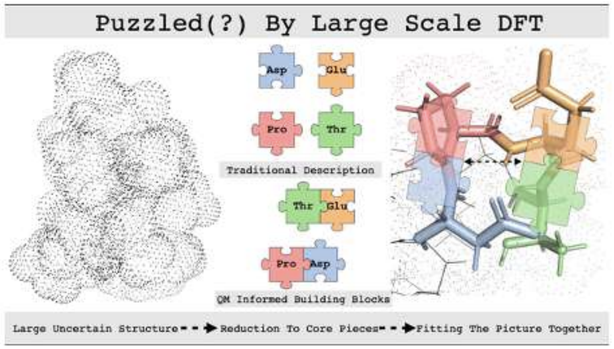

High Throughput explorations

- Li-ion conductors

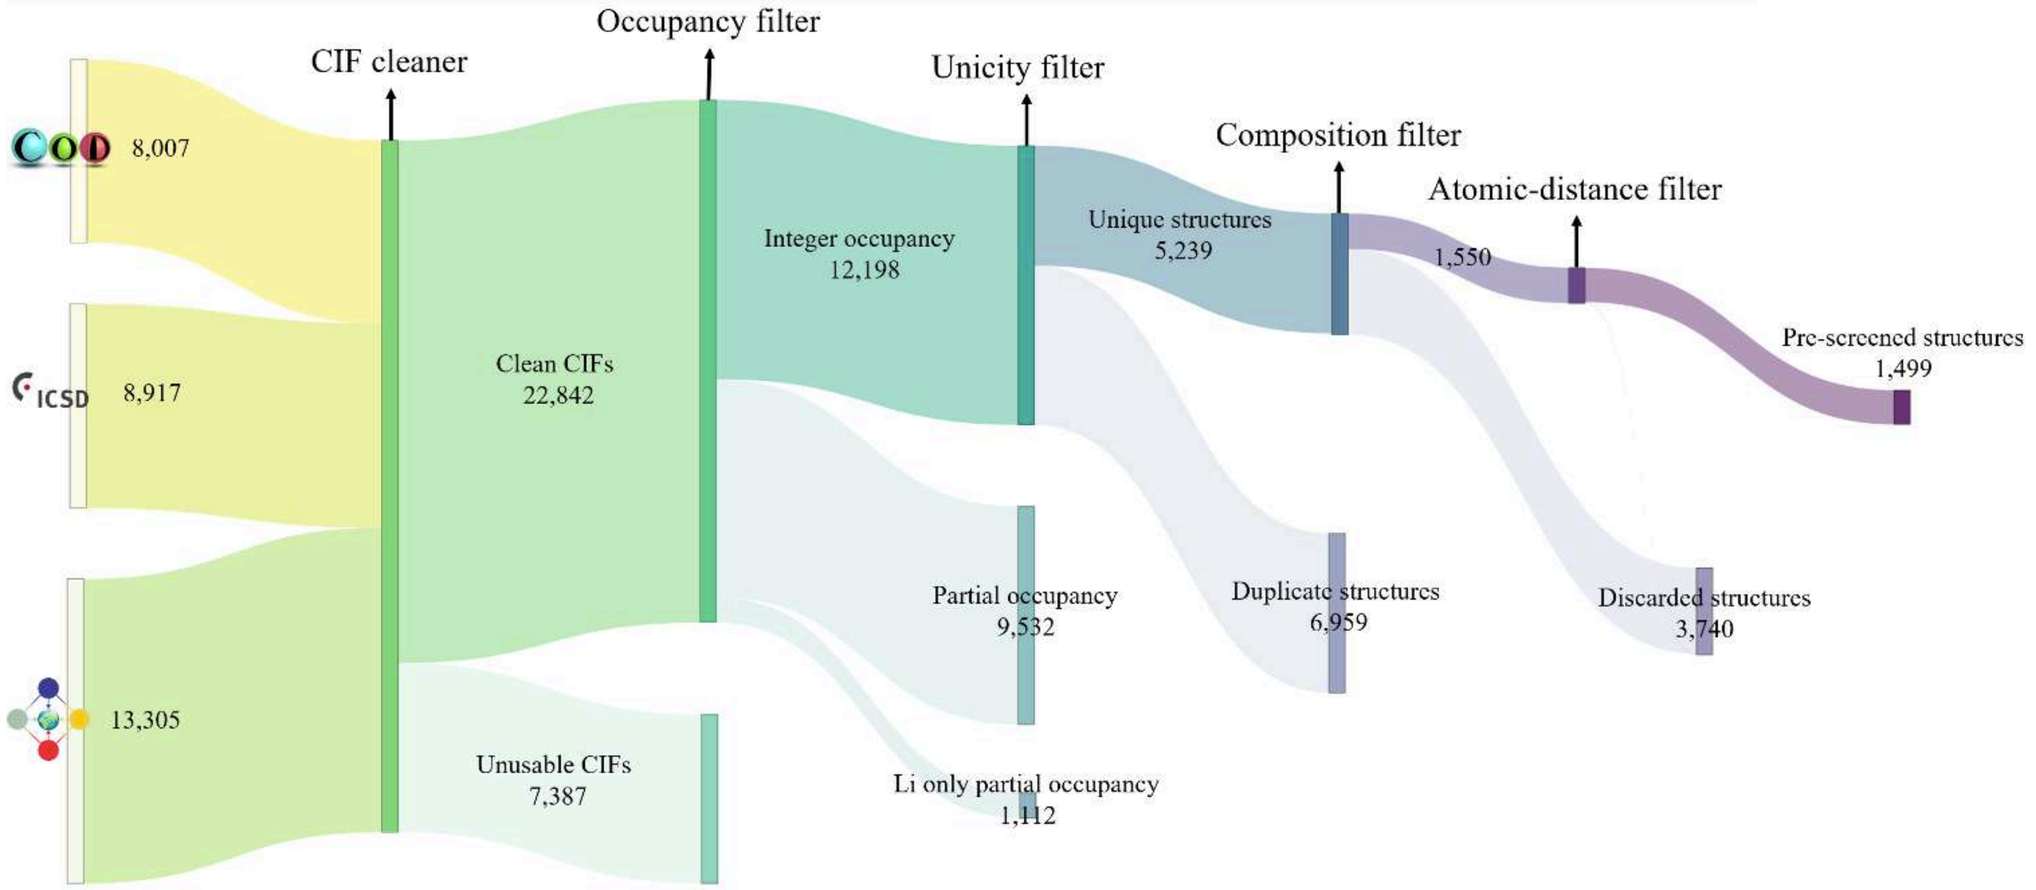

- Courtesy of N. Marzari (EPFL, U. Bremen)

High Throughput explorations

- Li-ion conductors

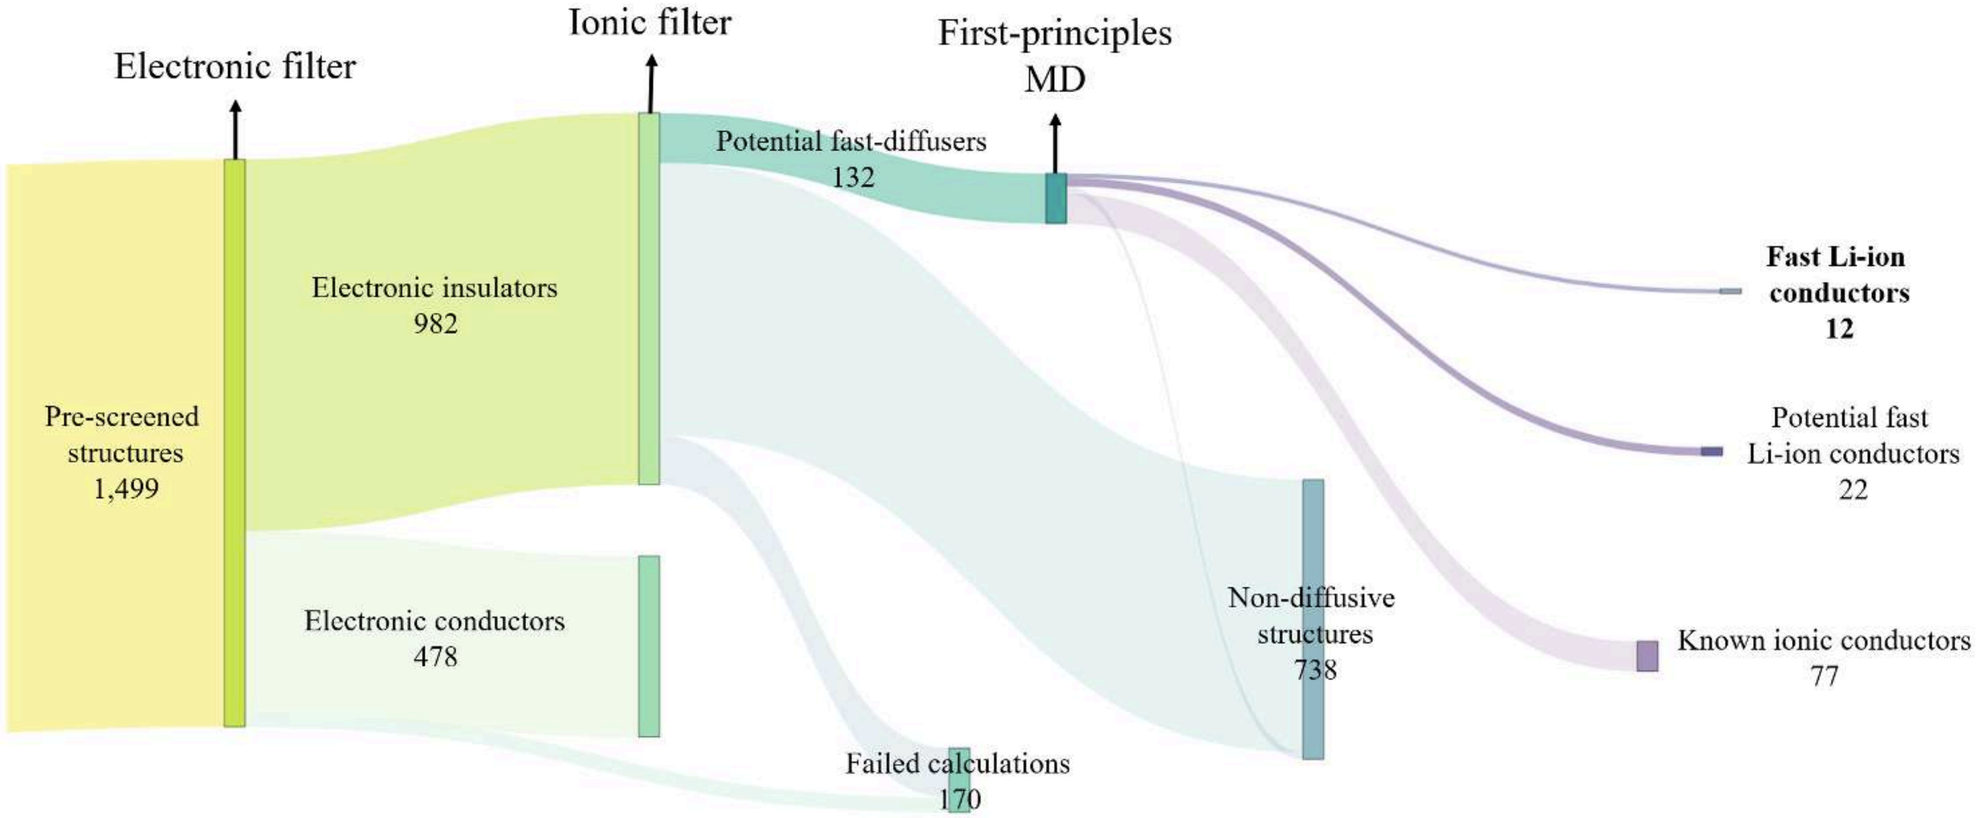

- Courtesy of N. Marzari (EPFL, U. Bremen)

With full workflow provenance

- Li-ion conductors

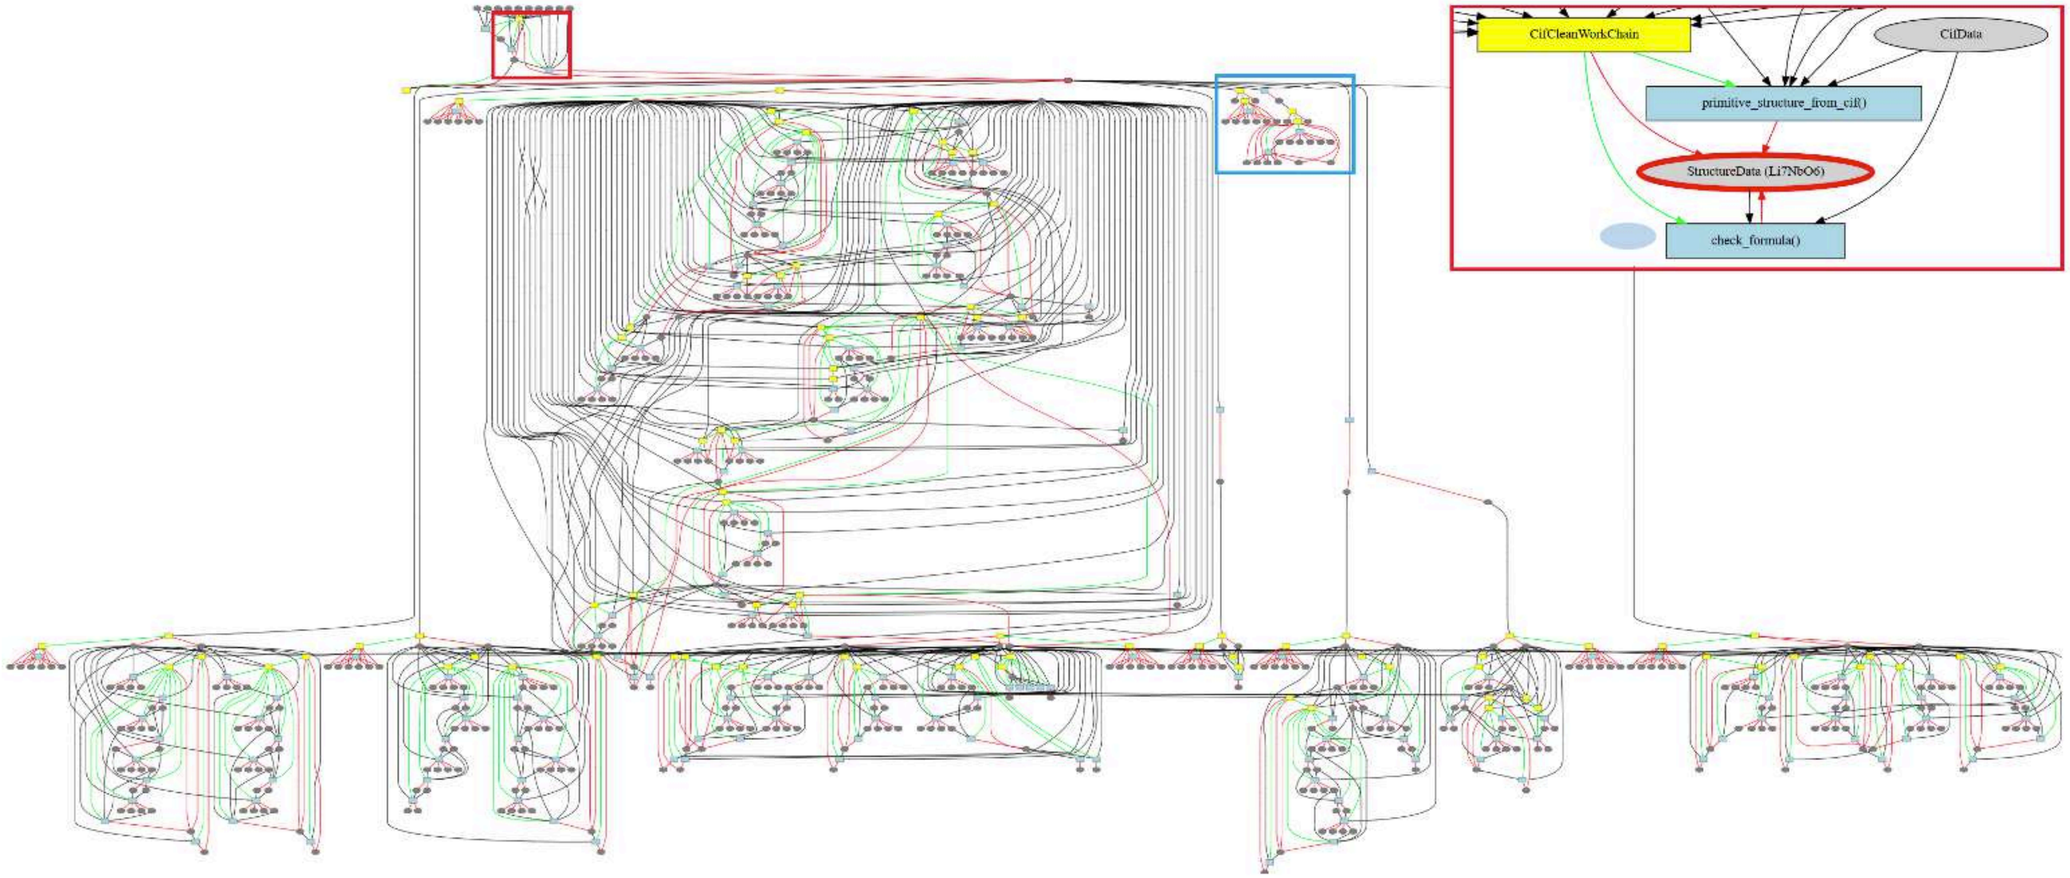

- Courtesy of N. Marzari (EPFL, U. Bremen)

automated workflows for MBPT

npj Comput Materials **9**, 74 (2023)

Towards high-throughput many-body perturbation theory: efficient algorithms and automated workflows

Miki Bonacci [1,2]✉, Junfeng Qiao [3], Nicola Spallanzani[2], Antimo Marrazzo [4], Giovanni Pizzi [3,5], Elisa Molinari [1,2], Daniele Varsano [2], Andrea Ferretti [2] and Deborah Prezzi [2]

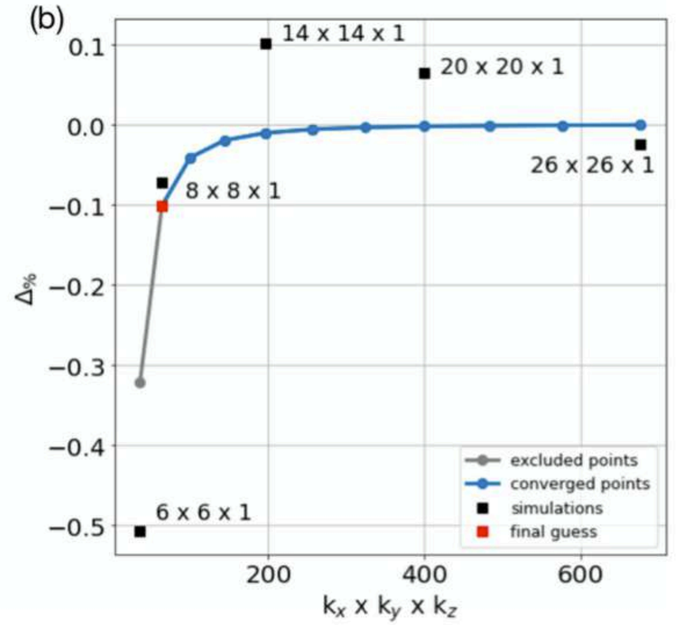

- Algorithms for **automatic convergence** of many-body perturbation theory methods
- **efficient sampling** in multi-dimensional parameter space
- combines yambo workflows with the AiiDA automation engine

Miki Bonacci
now: PSI Zurich, CH

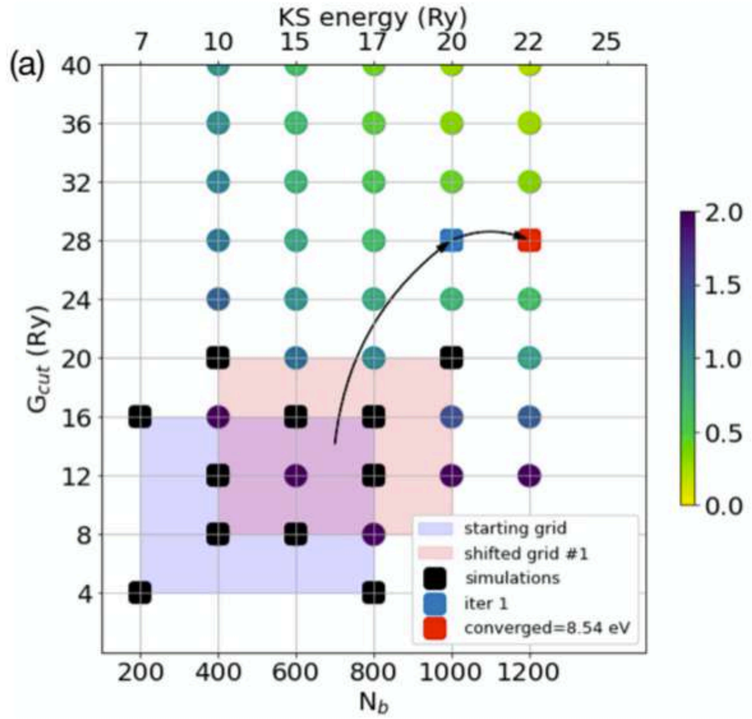

Polarons in TiO2

Defect projected BSE spectra

Boltzmann averaging is performed

Polarons optical responses are decoupled

Surface polaron sensitive to OH_b group

Sub-surface polarons behave as isolated

Multiple polaronic (bulk vs surface), and OH_b configurations.

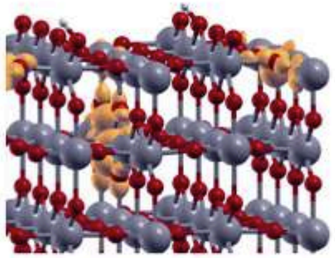

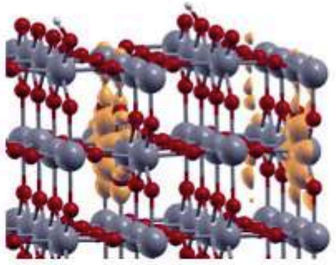

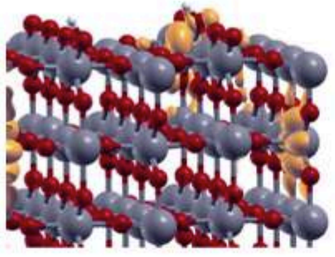

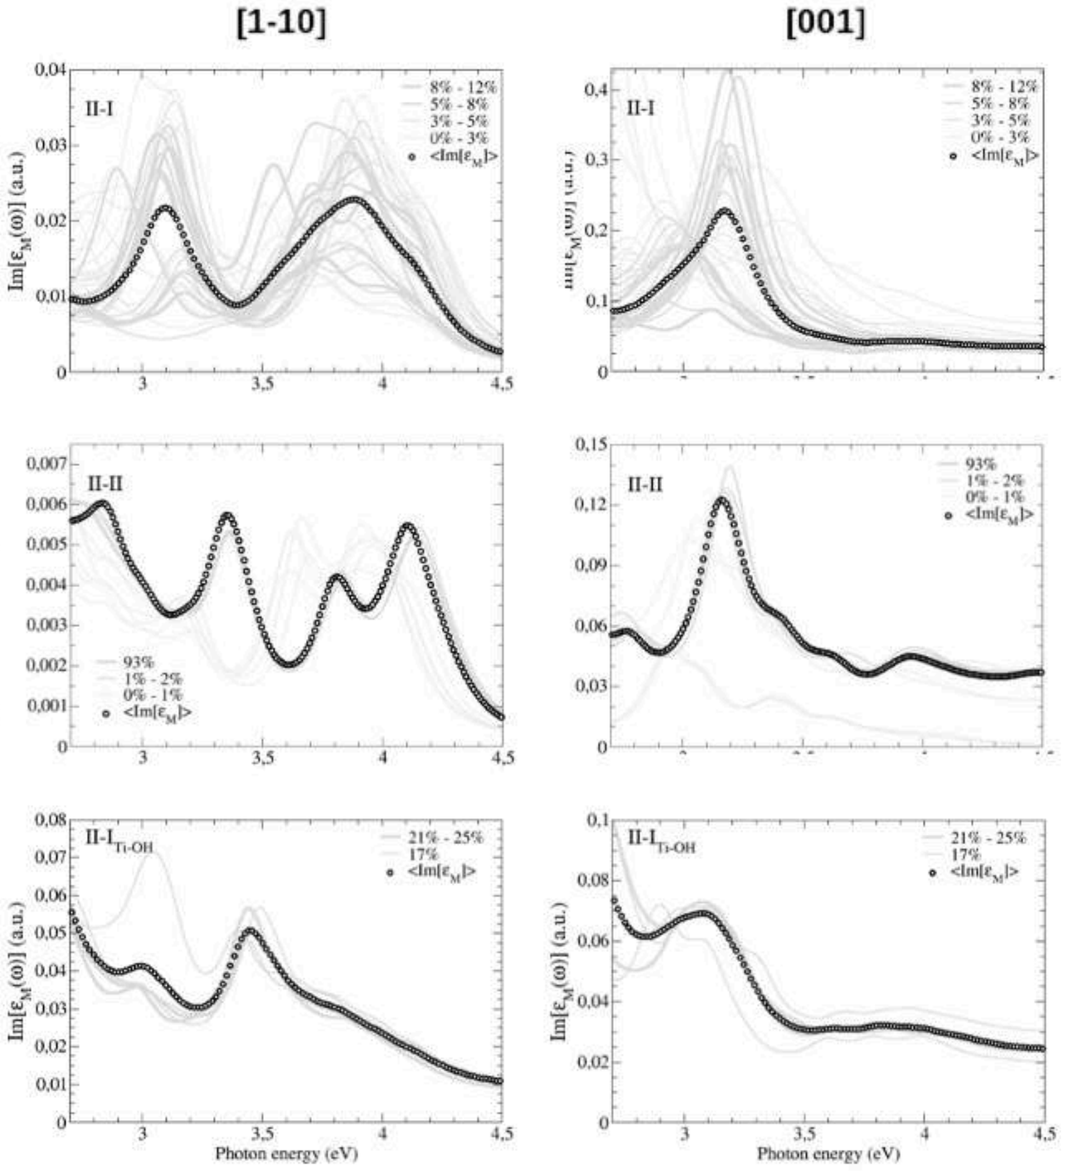

MaterialsCloud as a data infrastructure

https://www.materialscloud.org

Nicola Marzari
EPFL, U.Bremen

Giovanni Pizzi
PSI

MATERIALSCLOUD

Built for seamless sharing of resources
in computational materials science.

LEARN Lectures and tutorials in computational materials science

WORK Simulation tools and services - in the cloud or on your computer

DISCOVER Curated research data with tailored visualizations

EXPLORE Interactive browser for AiiDA provenance graphs

ARCHIVE An open-access, moderated repository for research data in computational materials science

Please cite L. Talirz et al., Sci Data 7, 299 (2020), if you use Materials Cloud in your research.

Vision | News | Open Science and DMPs | Statistics | Partners

https://www.materialscloud.org

MATERIALSCLOUD

https://www.materialscloud.org

Supported by:

Supported by

MARVEL NCCR
nccr-marvel.ch

MaX Centre
max-centre.eu

SNSF
snf.ch

EPFL
epfl.ch

swissuniversities
swissuniversities.ch
Go to project page

Paul Scherrer Institut
psi.ch/en/lms

BIG-MAP
big-map.eu

University of Bremen
Excellence Chair

H2020 MarketPlace
the-marketplace-project.eu

H2020 Intersect
intersect-project.eu

OSSCAR
osscar.org, EPFL Open Science Fund

ERC
erc.europa.eu

PASC
pasc-ch.org

PRACE
prace-ri.eu

NFFA and NEP
nffa.eu

EMMC
emmc.info

DOME 4.0 - EU Project
dome40.eu

H2020 OpenModel
open-model.eu

Materials Cloud Archive

https://archive.materialscloud.org

Enjoy !

Follow us on:

@max_center2

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

company/max-centre/

youtube/channel/MaX Centre eXascale