First-principles simulations at the nanoscale (and towards the exascale) using QUANTUM ESPRESSO

P. Giannozzi
Università di Udine, Italy

Supermassive Computations in Theoretical Physics
FBK Trento, 2015/02/11

Quantum simulation of matter at the nanoscale

Nanoscale: phenomena happening on a scale of lengths up to a few tens of nm. Basic theoretical tools:

- Density-Functional Theory (DFT) (P. Hohenberg, W. Kohn, and L. Sham, 1964-65)
- Pseudopotentials (J.C. Phillips, M.L. Cohen, M. Schlüter, D. Vanderbilt and many others, 1960-2000)
- Car-Parrinello and other iterative techniques (SISSA 1985, and many other places since)

Sometimes referred to as *The Standard Model* of materials science.

the saga of time and length scales

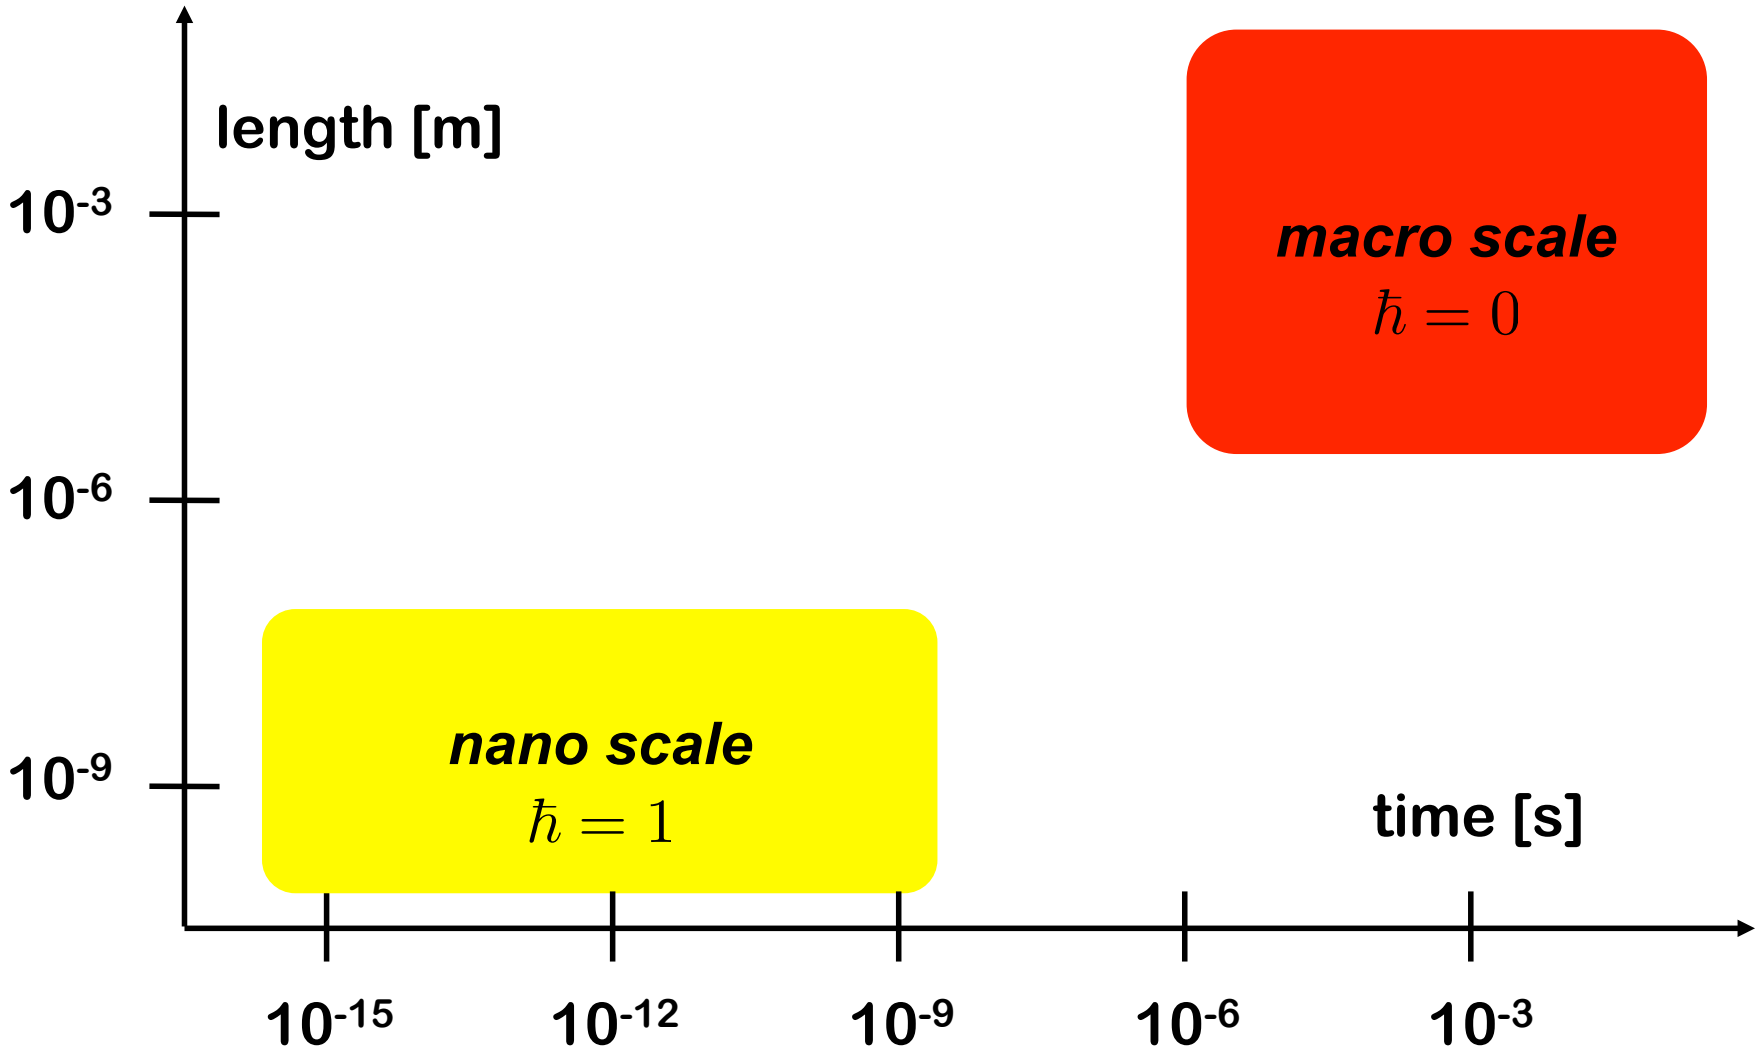

the saga of time and length scales

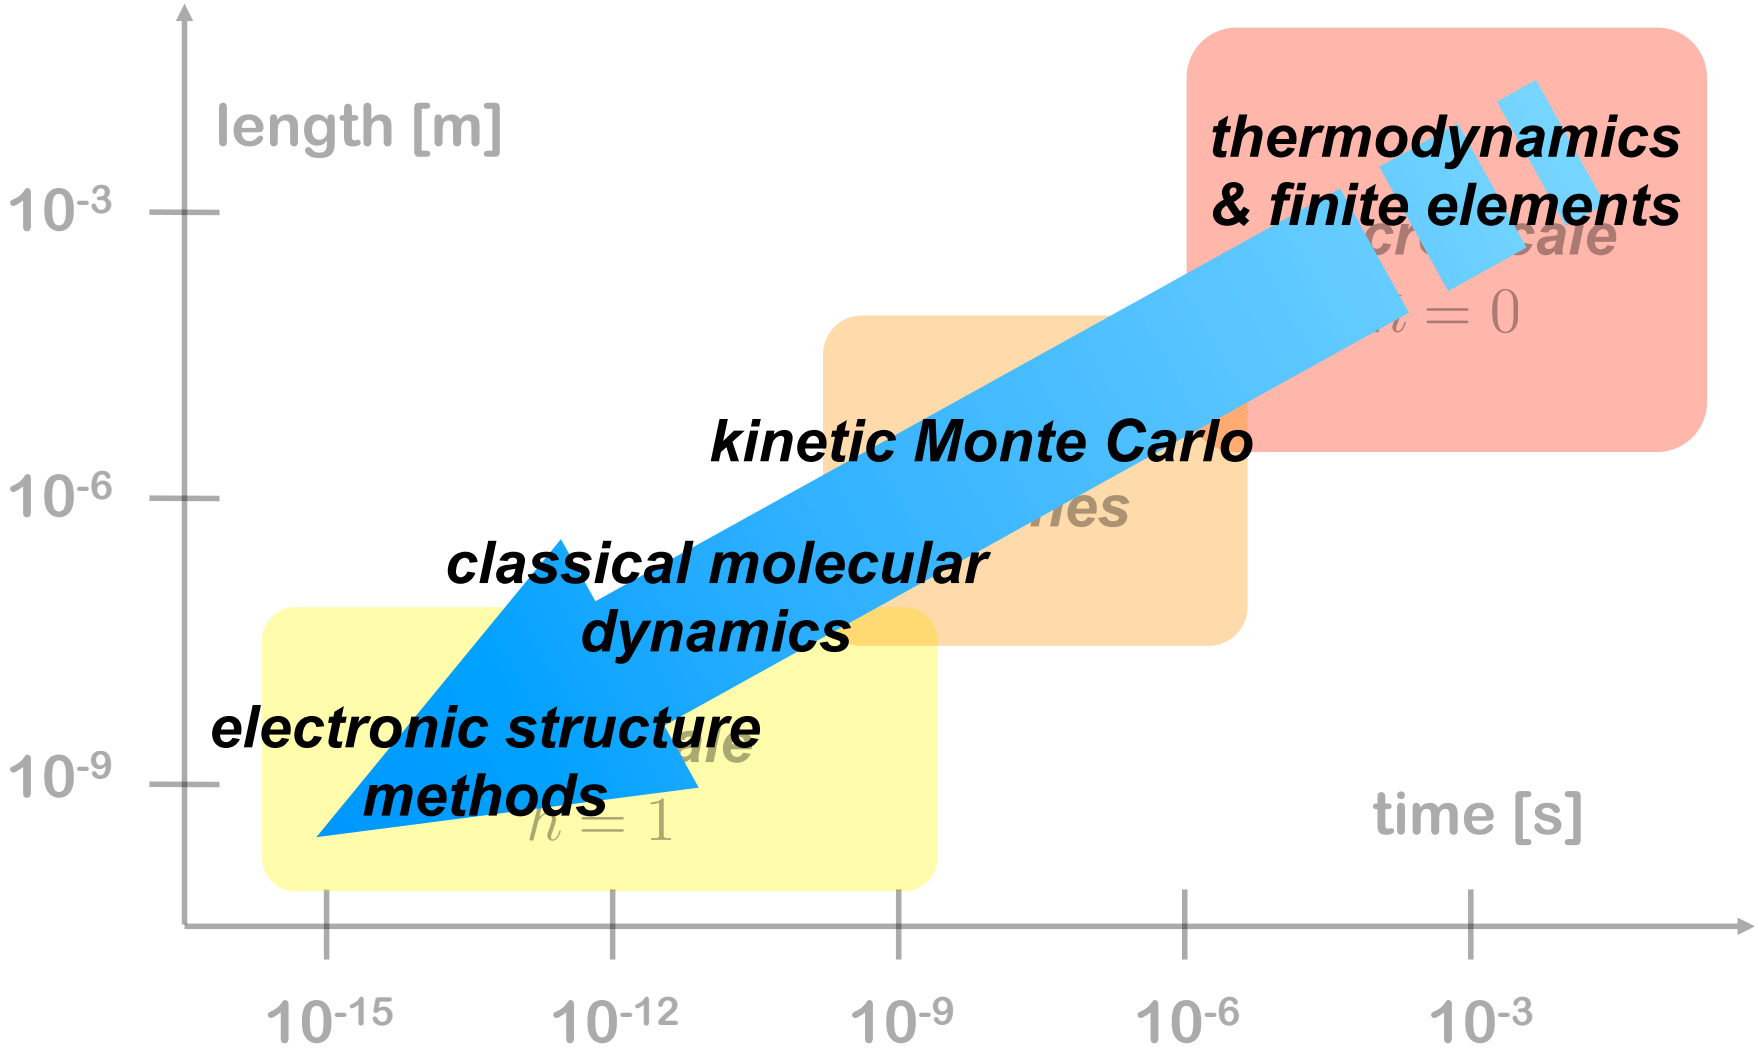

size vs. accuracy

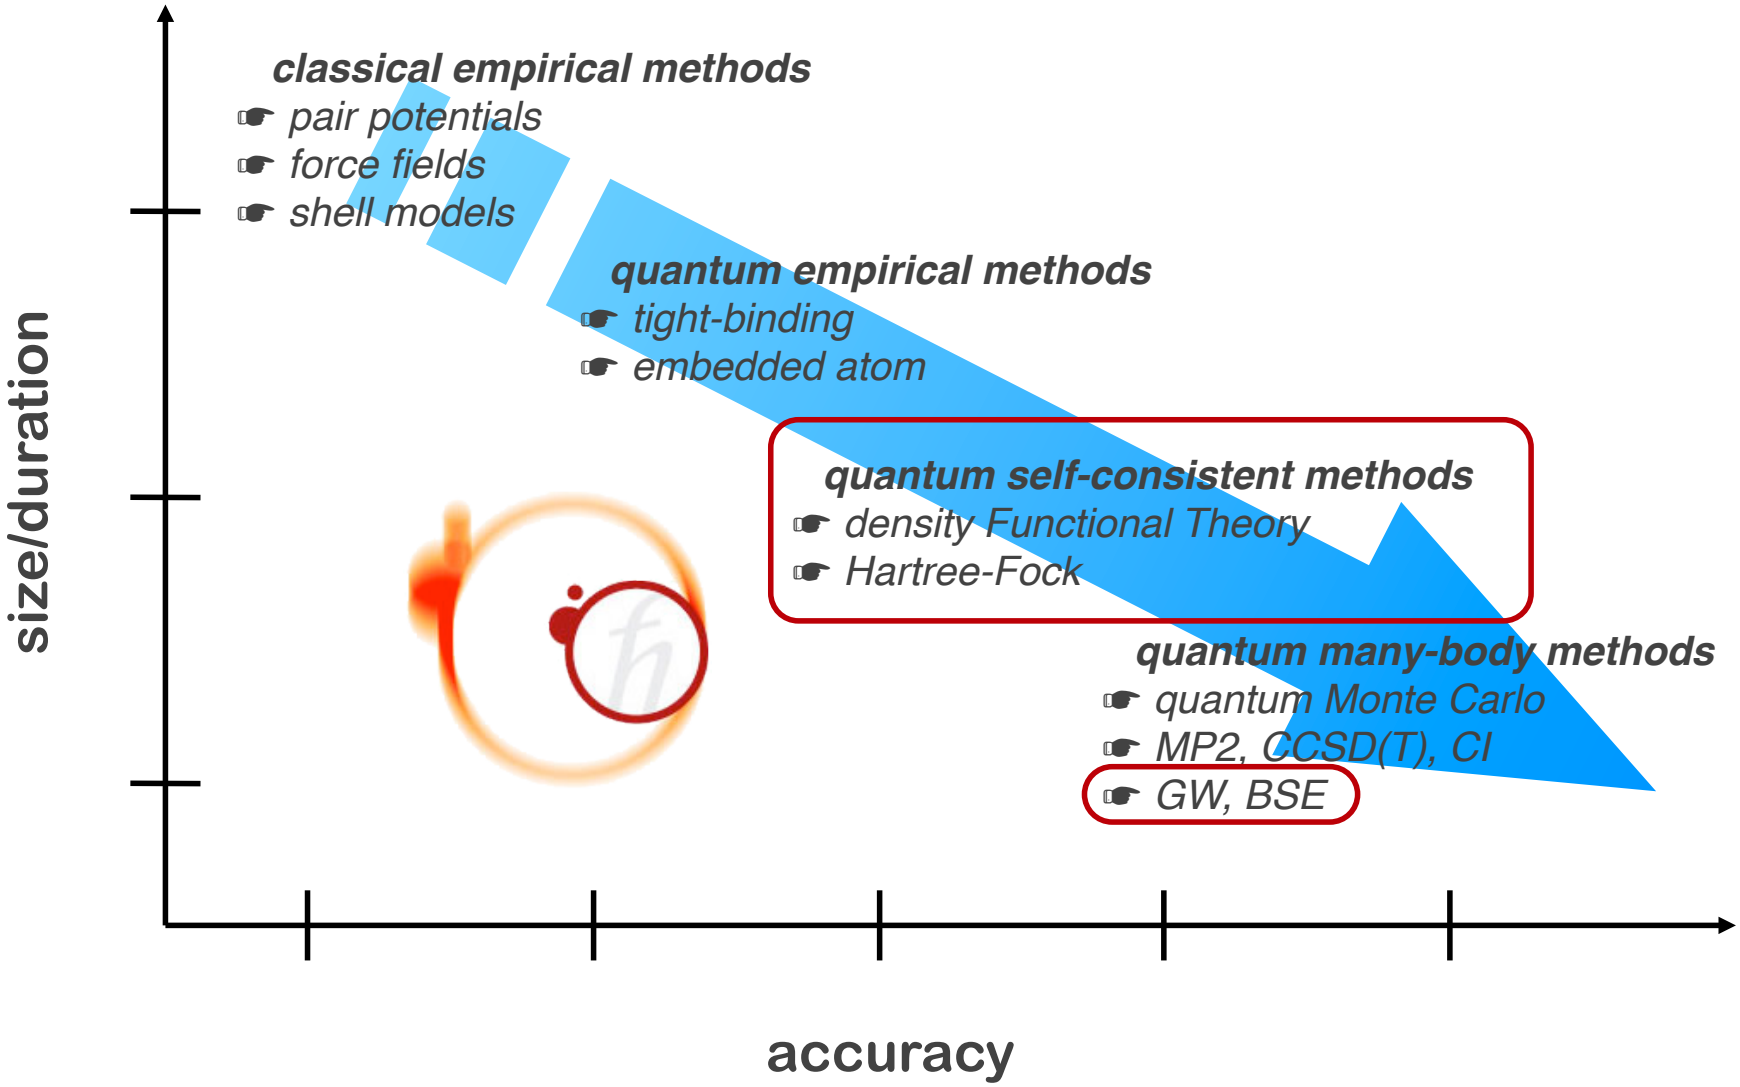

At the nanoscale: new materials

Most common atomic configurations in amorphous $CdTeO_x$, $x = 0.2$; Phys. Rev. **B** 79, 014205 (2009).

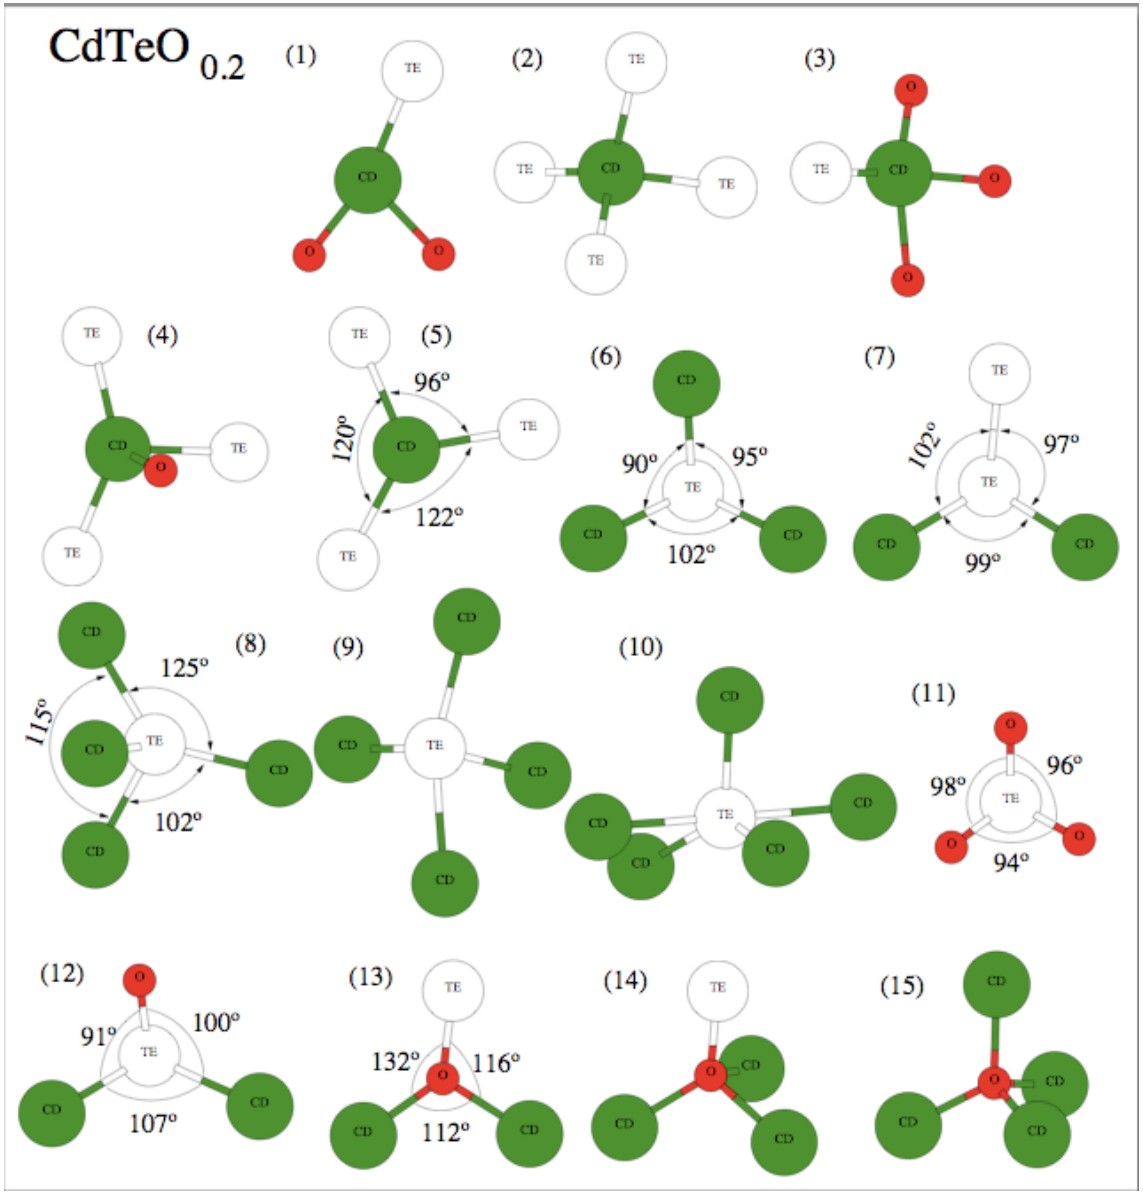

At the nanoscale: new devices

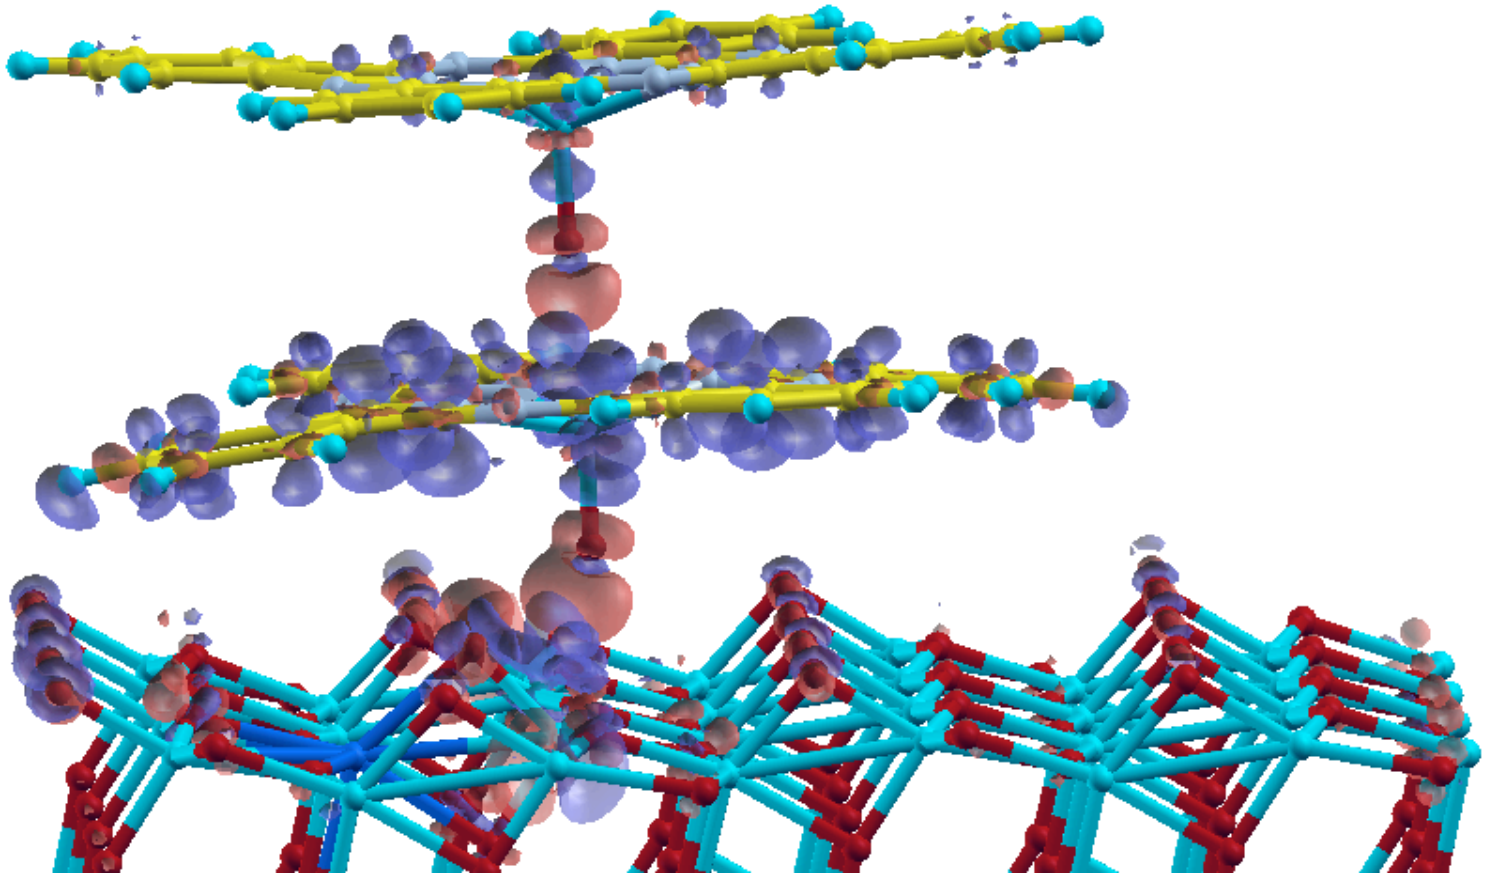

Organic-inorganic semiconductor heterojunction, phtalocyanine over TiO_2 anatase surface; Chem. Mater. **21**, 4555 (2009).

At the nanoscale: nanocatalysis

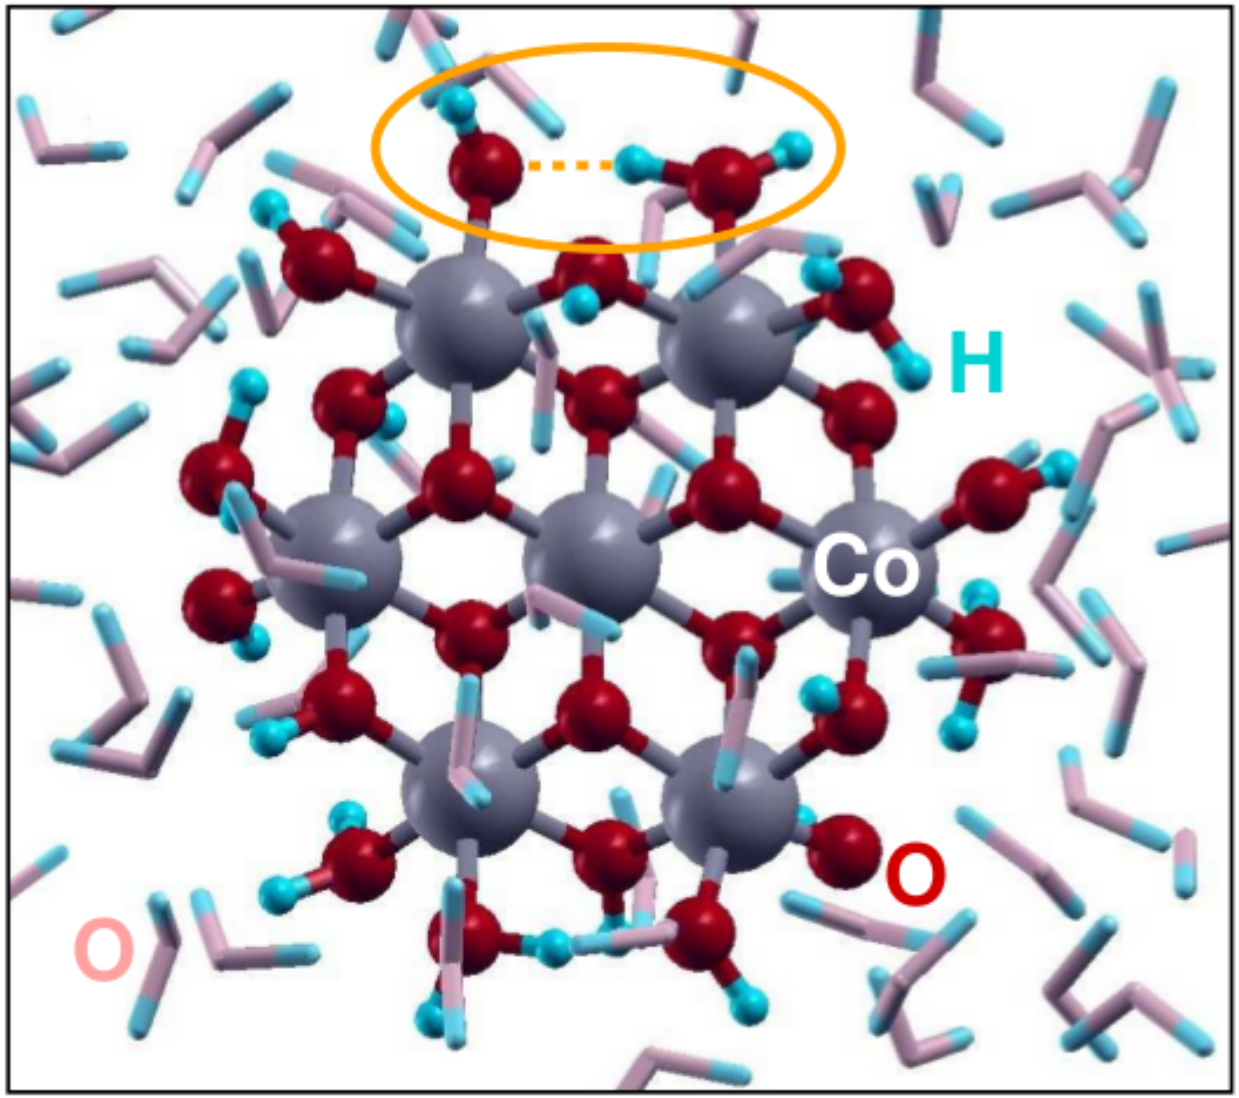

Cobalt-base catalyser for water splitting: J. Am. Chem. Soc. **135**, 15353 (2013)

At the nanoscale: biological systems

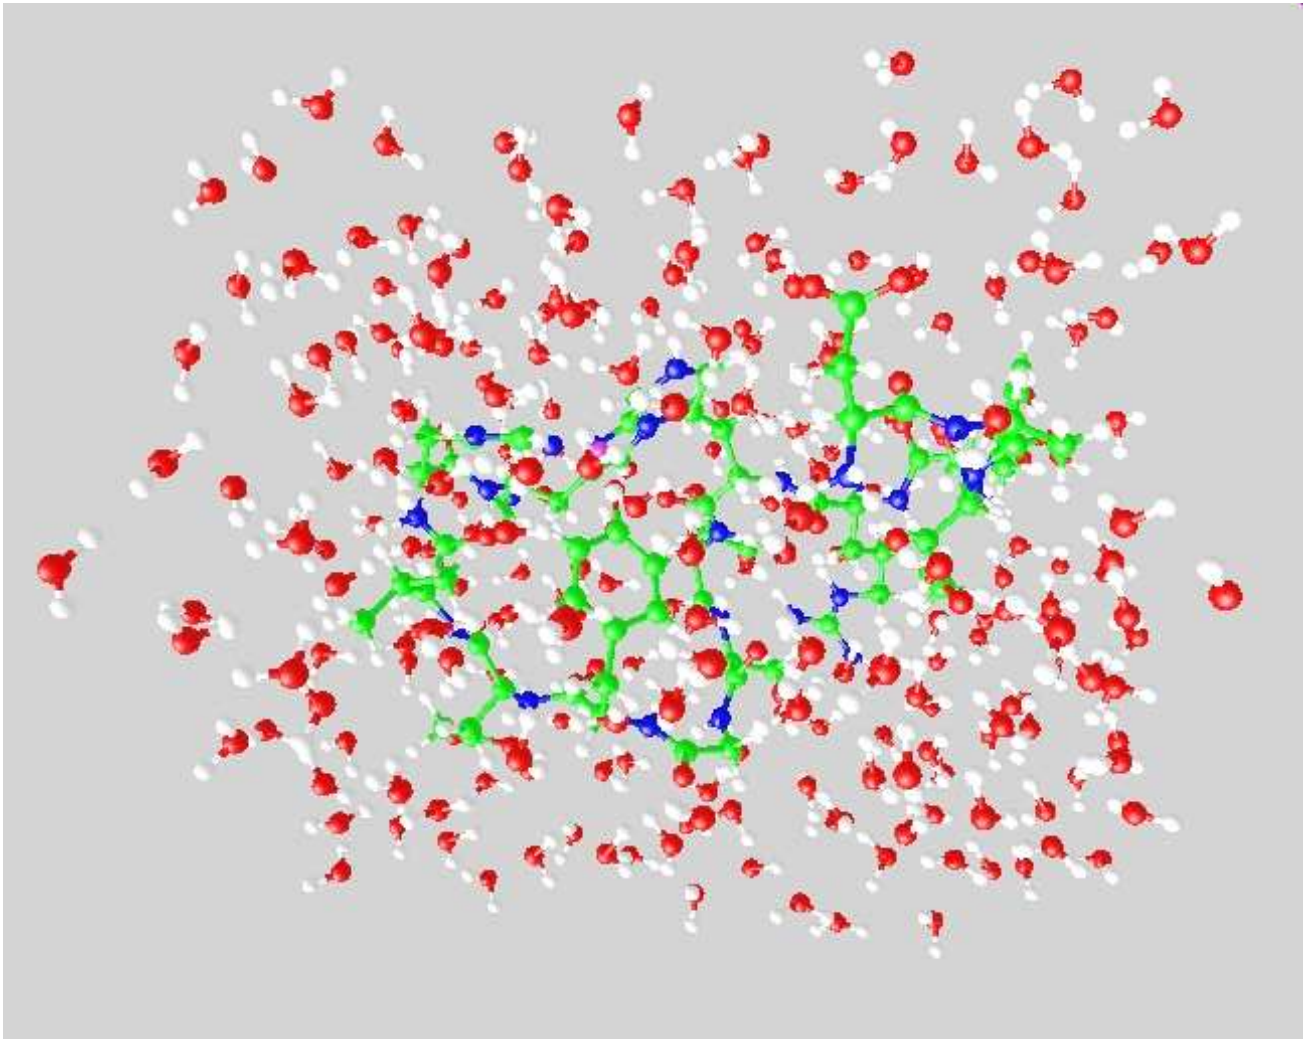

Metal-β-amyloid interactions; Metallomics **4**, 156 (2012).

Towards the exascale: massive parallelization

C@Ir(001)
443 atoms
2987 electrons

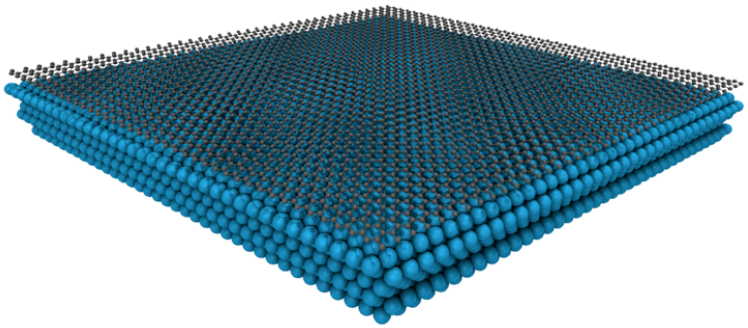

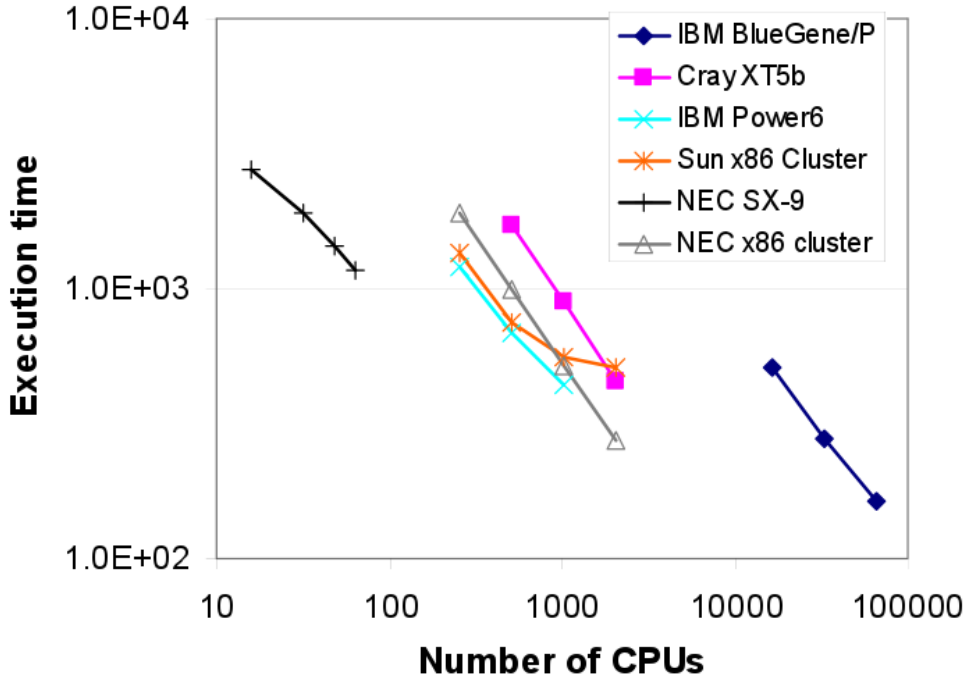

... still not forgetting smaller machines! In the figure, Nicola Marzari's smartphone running QUANTUM ESPRESSO

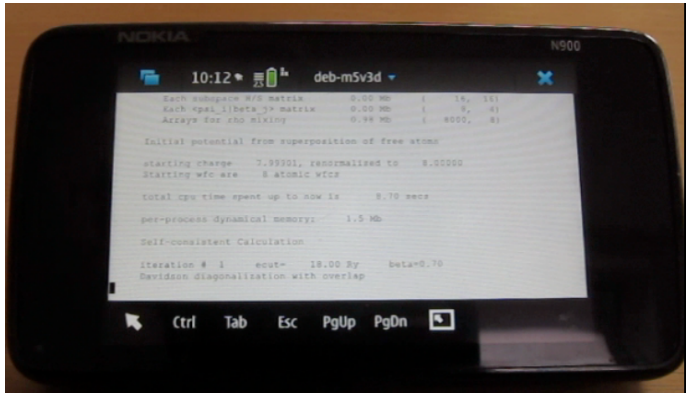

First-principles simulations

Time-dependent Schrödinger equation for nuclei $\mathbf{R} \equiv \{\vec{R}_I\}$ and electrons $\mathbf{r} \equiv \{\vec{r}_i\}$:

$$i\hbar \frac{\partial \hat{\Phi}(\mathbf{r}, \mathbf{R}; t)}{\partial t} = \left(-\sum_I \frac{\hbar^2}{2M_I} \nabla^2_{\vec{R}_I} - \sum_i \frac{\hbar^2}{2m} \nabla^2_{\vec{r}_i} + V(\mathbf{r}, \mathbf{R}) \right) \hat{\Phi}(\mathbf{r}, \mathbf{R}; t)$$

Born-Oppenheimer (or *adiabatic*) approximation, valid for $M_I >> m$:

$$\hat{\Phi}(\mathbf{r}, \mathbf{R}; t) \simeq \Phi(\mathbf{R}) \Psi(\mathbf{r}|\mathbf{R}) e^{-i\hat{E}t/\hbar}$$

Problem splits into an *electronic problem depending upon nuclear positions*:

$$\left(-\sum_i \frac{\hbar^2}{2m} \nabla^2_{\vec{r}_i} + V(\mathbf{r}, \mathbf{R}) \right) \Psi(\mathbf{r}, \mathbf{R}) = E(\mathbf{R}) \Psi(\mathbf{r}, \mathbf{R})$$

and a *nuclear problem under an effective interatomic potential* $E(\mathbf{R})$, typically treated as *classical*, with forces on nuclei: $\boxed{\mathbf{F}_I = -\nabla_{\vec{R}_I} E(\mathbf{R}).}$

Density-Functional Theory

Transforms the many-electron problem into an equivalent problem of (fictitious) non-interacting electrons, the *Kohn-Sham equations*:

$$H\phi_v \equiv \left(-\frac{\hbar^2}{2m}\nabla^2_{\vec{r}} + V_{\mathbf{R}}(\vec{r}) \right) \phi_v(\vec{r}) = \epsilon_v \phi_v(\vec{r})$$

The effective potential is a *functional* of the charge density:

$$V_{\mathbf{R}}(\vec{r}) = -\sum_I \frac{Z_I e^2}{|\vec{r} - \vec{R}_I|} + v[n(\vec{r})], \qquad n(\vec{r}) = \sum_v |\phi_v(\vec{r})|^2$$

(Hohenberg-Kohn 1964, Kohn-Sham 1965). Exact form is unknown, but simple approximate forms yielding very accurate (ground-state) results are known.

Density-Functional Theory II

The total energy is also a functional of the charge density:

$$E \Rightarrow E[\{\phi\}, \mathbf{R}] = -\frac{\hbar^2}{2m}\sum_v \int \phi_v^*(\vec{r})\nabla^2\phi_v(\vec{r})d\vec{r} + \int V_{\mathbf{R}}(\vec{r})n(\vec{r})d\vec{r} +$$

$$+ \frac{e^2}{2}\int \frac{n(\vec{r})n(\vec{r'})}{|\vec{r}-\vec{r'}|}d\vec{r}d\vec{r'} + E_{xc}[n(\vec{r})] + \sum_{I\neq J}\frac{e^2}{2}\frac{Z_I Z_J}{|\vec{R}_I - \vec{R}_J|}$$

Kohn-Sham equations arise from the minimization of the energy functional:

$$E(\mathbf{R}) = \min_{\phi} E[\{\phi\}, \mathbf{R}], \qquad \int \phi_i^*(\vec{r})\phi_j(\vec{r})d\vec{r} = \delta_{ij}$$

Hellmann-Feynman theorem holds. Forces on nuclei:

$$\vec{F}_I = -\nabla_{\vec{R}_I} E(\mathbf{R}) = -\int n(\vec{r})\nabla_{\vec{R}_I} V_{\mathbf{R}}(\vec{r})d\vec{r}$$

The tricks of the trade

- expanding the Kohn-Sham orbitals into a suitable *basis* set turns DFT into a multi-variate minimization problem, and the Kohn-Sham equations into a *non-linear matrix eigenvalue problem*

- the use of *pseudopotentials* allows one to ignore chemically inert core states and to use *plane waves*

- plane waves are *orthogonal* and the *matrix elements* of the Hamiltonian are usually easy to calculate; the *completeness* of the basis is easy to check

- plane waves allow to efficiently calculate *matrix-vector products* and to solve the Poisson equation using *Fast Fourier Transforms* (FFTs)

Accuracy vs. Approximations

Theoretical approximations / limitations of DFT:

- the Born-Oppenheimer approximation
- DFT functionals (LDA, GGA, ...)
- pseudopotentials
- no easy access to excited states and/or quantum dynamics

Numerical approximations / limitations:

- finite/limited size/time
- finite basis set
- differentiation / integration / interpolation

Requirements on effective software for quantum simulations at the nanoscale

- Challenging calculations stress the limits of available computer power: software should be **fast and efficient**
- Diffusion of first-principle techniques among non-specialists requires software that is **easy to use and** (reasonably) **error-proof**
- Introducing innovation requires new ideas to materialize into new algorithms through codes: software should be **easy to extend and to improve**
- Complex problems require a mix of solutions coming from different approaches and methods: software should be **interoperable with other software**
- Finaly, scientific ethics requires that **results should be reproducible** and **algorithms susceptible of validation**

The QUANTUM ESPRESSO distribution

QUANTUM ESPRESSO stands for *Quantum opEn-Source Package for Research in Electronic Structure, Simulation, and Optimization*

QUANTUM ESPRESSO is a *distribution* (an integrated suite) of software for atomistic calculations based on electronic structure, using density-functional theory, a plane-wave basis set, pseudopotentials. Freely available under the terms of the GNU General Public License

The main goals of QUANTUM ESPRESSO are

- *innovation* in methods and algorithms
- *efficiency* on modern computer architectures

A great effort is also devoted to *user friendliness* and to the formation of a *users' and developers' community*

QUANTUM ESPRESSO **contributors**

QUANTUM ESPRESSO receives contributions from many individuals and partner institutions in Europe and worldwide. Who "owns" QUANTUM ESPRESSO?

QUANTUM ESPRESSO **Foundation**

The QUANTUM ESPRESSO Foundation: a non–profit ("limited by guarantee") company, based in London, that

- coordinates and supports research, education, and outreach within the QUANTUM ESPRESSO community
- owns the trademarks and protects the open-source character of QUANTUM ESPRESSO
- raises funds to foster the QUANTUM ESPRESSO project

QUANTUM ESPRESSO **Foundation Members**

Current QEF members:

- Scuola Internazionale Superiore di Studi Avanzati (SISSA), Trieste
- Ecole Polytechnique Fédérale de Lausanne (EPFL)
- International Centre for Theoretical Physics (ICTP), Trieste
- Consiglio Nazionale delle Ricerche (IOM-CNR), Italy
- CINECA supercomputing center, Bologna
- University of North Texas
- Duke University
- ...

Development

The distribution is maintained as a single SVN (Subversion) tree. Available to everyone anytime via anonymous access.

- Web site: `http://www.quantum-espresso.org`
- Developers' portal: `http://www.qe-forge.org`

Mailing list (public):

- `pw_forum@pwscf.org`: for general discussions
- `qe_developers@qe-forge.org`: used by developers for technical discussions
- `qe_commits@qe-forge.org`: used by developers, receives commit messages

Developers' community: qe-forge

FORGE | Home | My Stuff | Projects | Snippets | Log In | Register New Account

Home »

Welcome to QEforge

QEforge is a web portal offering support to researchers active in the field of computer simulation and numerical modeling of matter and materials at the atomic scale. The most popular source code management (CVS, SVN or Git) systems, mailing lists, public forums, download space, wiki pages, and much more are provided through the Gforge engine.

QEforge is a free service offered to researchers from all over the world by the Quantum ESPRESSO Foundation, not restricted by any means to developers/users of the Quantum ESPRESSO project. Anybody is warmly welcome to join the QEforge community. Researchers participating in it will retain complete control over their projects (including the right not to release them), while enjoying the advanced development tools that QEforge provides.

You do not need to register if you just want to download publicly released files from qe-forge. Due to potential problems with spammers, *registration cannot be automatic*. In order to complete the registration procedure, please *send a presentation e-mail* from your institutional address to support (at) qe-forge.org, indicating:

1. institution to which you belong
2. position within the institution
3. purpose of your subscription request.

Please note that access to private repositories may be restricted even for registered users. Please directly contact the administrators of the project you are interested in, if you cannot access a repository.

If you experience any problem not mentioned above, please write to support (at) qe-forge.org.

Disclaimer

QEforge is provided for free by the Quantum ESPRESSO Foundation (QEF) to researchers active in the field of computer simulation of matter at the atomic scale, worldwide. Researchers using services from QEforge will retain all property rights and complete control over their project and related material, including the right not to release anything.

The QEforge is strongly committed to provide a high-quality reliable service. QEforge is continuously monitored, problems are promptly identified and fixed. The contents of QEforge is regularly saved to backup disks.
This being said, the QEforge site and services are provided on a "as is" basis. The QEF cannot be held liable for any damages or losses possibly caused by the unavailability, whether temporary or permanent, of the QEforge site and services, or by any defects or errors in the services.

QEforge can be used only for the development of code related to research activities in the field of computer simulation of matter at the atomic scale. We regret that we cannot grant access to other kinds of research activities given the limits of our resources. Development of commercial software is not allowed on QEforge. Accounts and projects not complying with the above terms may be suspended. Accounts and projects engaged in ranking manipulation, spamming, unlawful activites will be terminated without notice and all related contents destroyed.

Accounts not showing any activity since more than one year may be suspended; if no request for reactivation is received after a further year, they may be cancelled. The same applies to projects that have released no files.

Activity

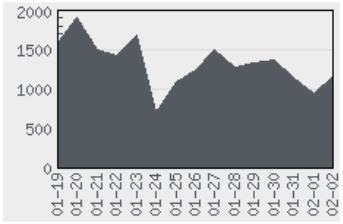

Recently Registered Projects

(2014-12-04) QE-FermiSurfer interface
(2014-11-13) DFPT-tetra
(2014-07-10) Unfold
(2014-01-13) thermo_pw
(2013-10-01) DFPT+U
(2013-04-04) Frozen Density Embedding
(2013-03-14) ASE interface for QE
(2013-02-22) MiniDFT
(2013-01-15) QE-tutorial
(2012-11-22) Third order DFPT

Top Downloads

(810,791) Quantum ESPRESSO
(58,530) WanT
(13,980) Yambo
(4,613) PSlibrary
(3,096) QE-GPU
(1,523) QE-GIPAW
(977) phiGEMM
(838) Quasiharmonic Approximation
(831) DMFT
(822) GRID lattice dynamics

Browse Project Topics

Quantum ESPRESSO related (23)
Classical Simulation (2)
Density-Functional Theory calculation (15)
Many-Body calculation (5)
Libraries (2)
Utilities (3)
Visualization tools (2)

Support and Docs

Local Support Project

Currently 45 public projects, 570 registered users, 66 QE developers registered
(not all of them active, though!)

Users' community: factoids

- About 1800 registered users for the pw_forum mailing list
- An average of $\sim$ 10 messages a days on pw_forum
- latest version (5.1.1) downloaded almost 20000 [*] times

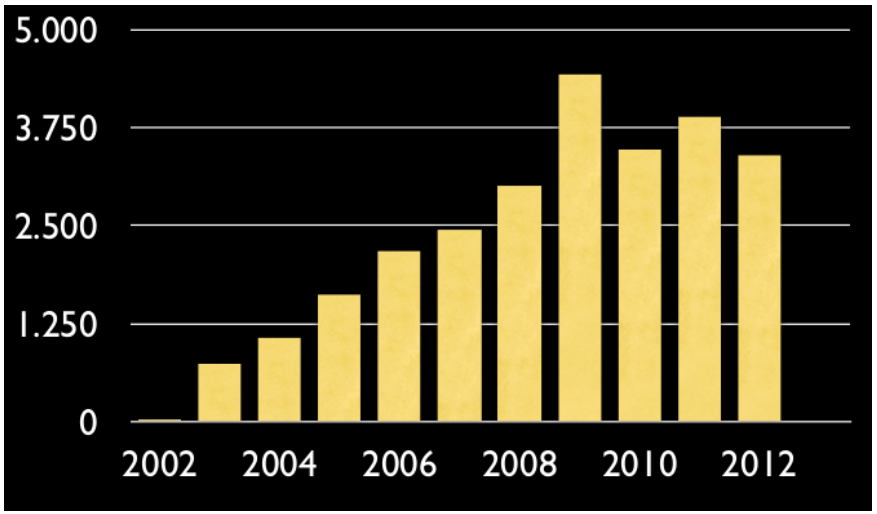

- 30 Schools or tutorials since 2002, attended by $\sim$ 1200 users
- 3 developers' schools since 2013, latest in January 2015

[*] *this number is likely inflated by bots, failed downloads, etc.*

Schools and tutorial using QUANTUM ESPRESSO

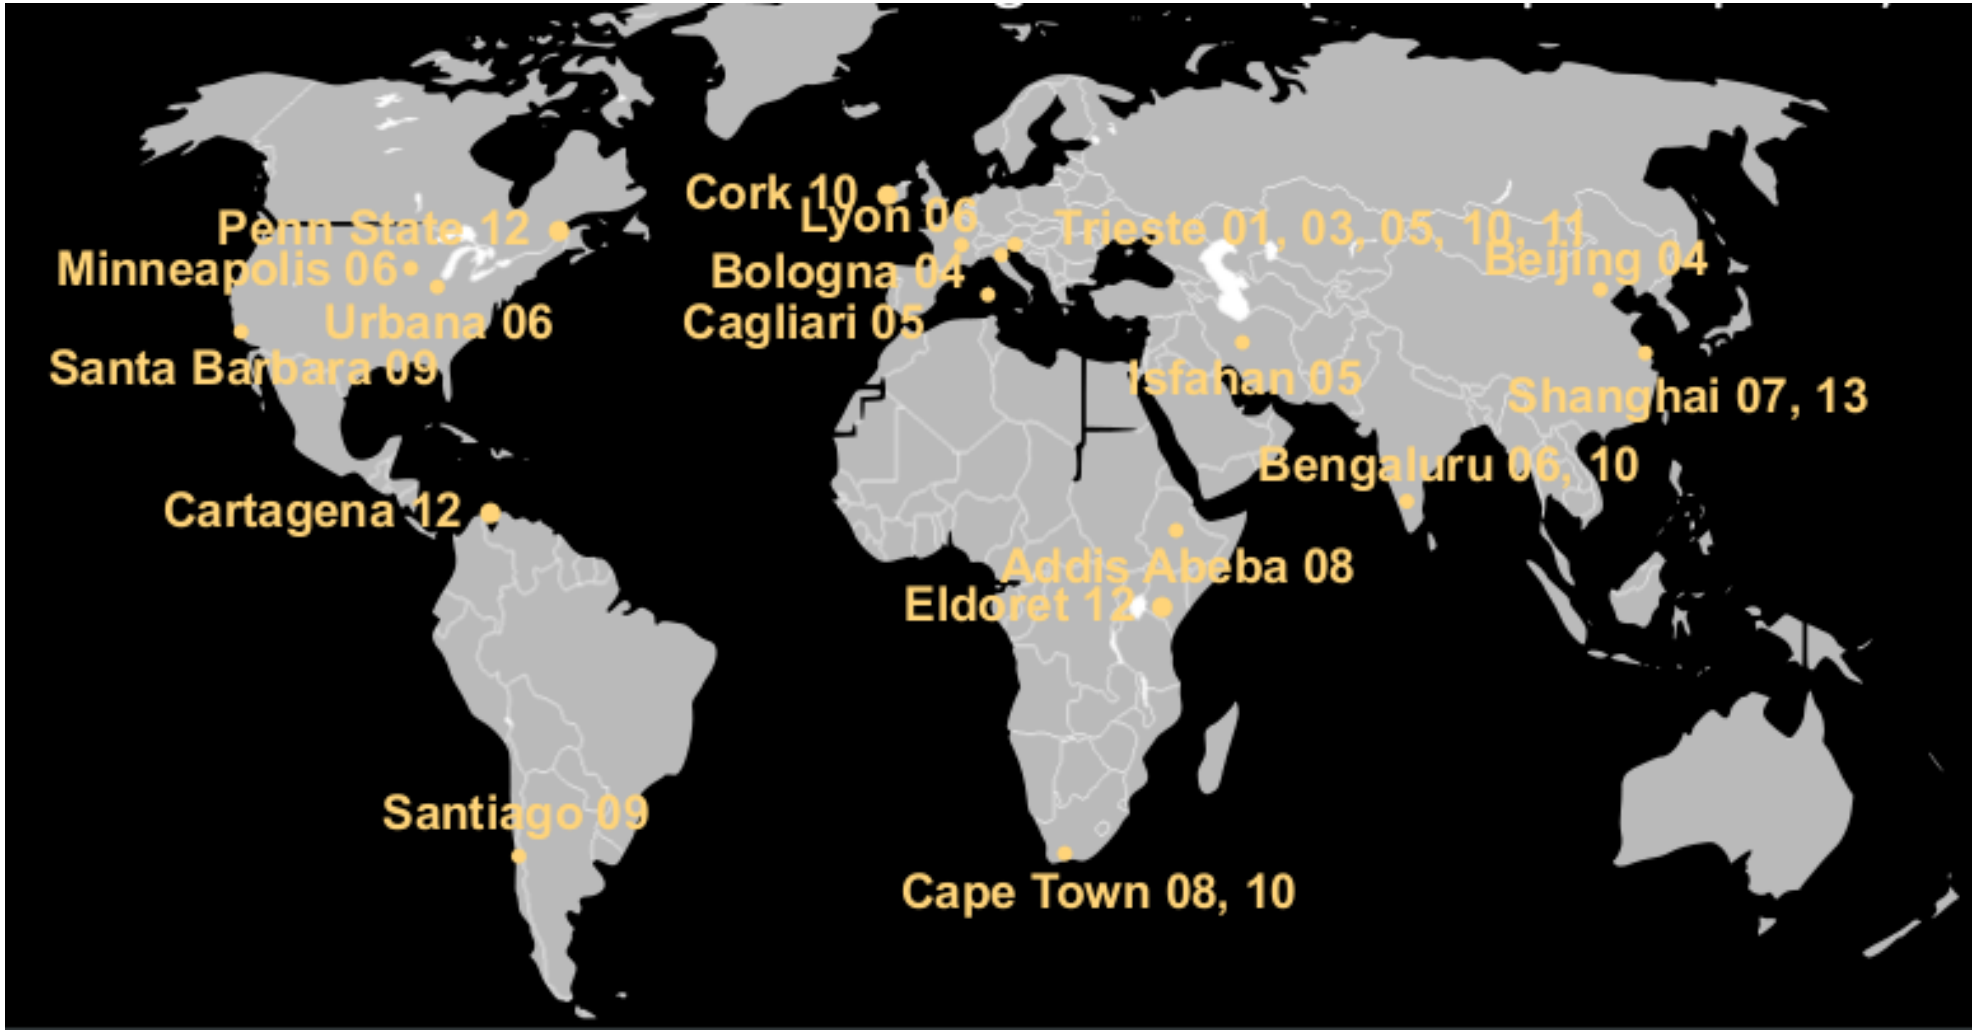

More: Penn State, June 2014; University of Tokyo, April 2014; Pune, July 2014.
Next: Cordoba, September 2015

IOP PUBLISHING JOURNAL OF PHYSICS: CONDENSED MATTER

J. Phys.: Condens. Matter **21** (2009) 395502 (19pp) doi:10.1088/0953-8984/21/39/395502

QUANTUM ESPRESSO: a modular and open-source software project for quantum simulations of materials

Paolo Giannozzi[1,2], Stefano Baroni[1,3], Nicola Bonini[4], Matteo Calandra[5], Roberto Car[6], Carlo Cavazzoni[7,8], Davide Ceresoli[4], Guido L Chiarotti[9], Matteo Cococcioni[10], Ismaila Dabo[11], Andrea Dal Corso[1,3], Stefano de Gironcoli[1,3], Stefano Fabris[1,3], Guido Fratesi[12], Ralph Gebauer[1,13], Uwe Gerstmann[14], Christos Gougoussis[5], Anton Kokalj[1,15], Michele Lazzeri[5], Layla Martin-Samos[1], Nicola Marzari[4], Francesco Mauri[5], Riccardo Mazzarello[16], Stefano Paolini[3,9], Alfredo Pasquarello[17,18], Lorenzo Paulatto[1,3], Carlo Sbraccia[1,†], Sandro Scandolo[1,13], Gabriele Sclauzero[1,3], Ari P Seitsonen[5], Alexander Smogunov[13], Paolo Umari[1] and Renata M Wentzcovitch[10,19]

Cited approx. 3300 times since publication

Structure of the distribution

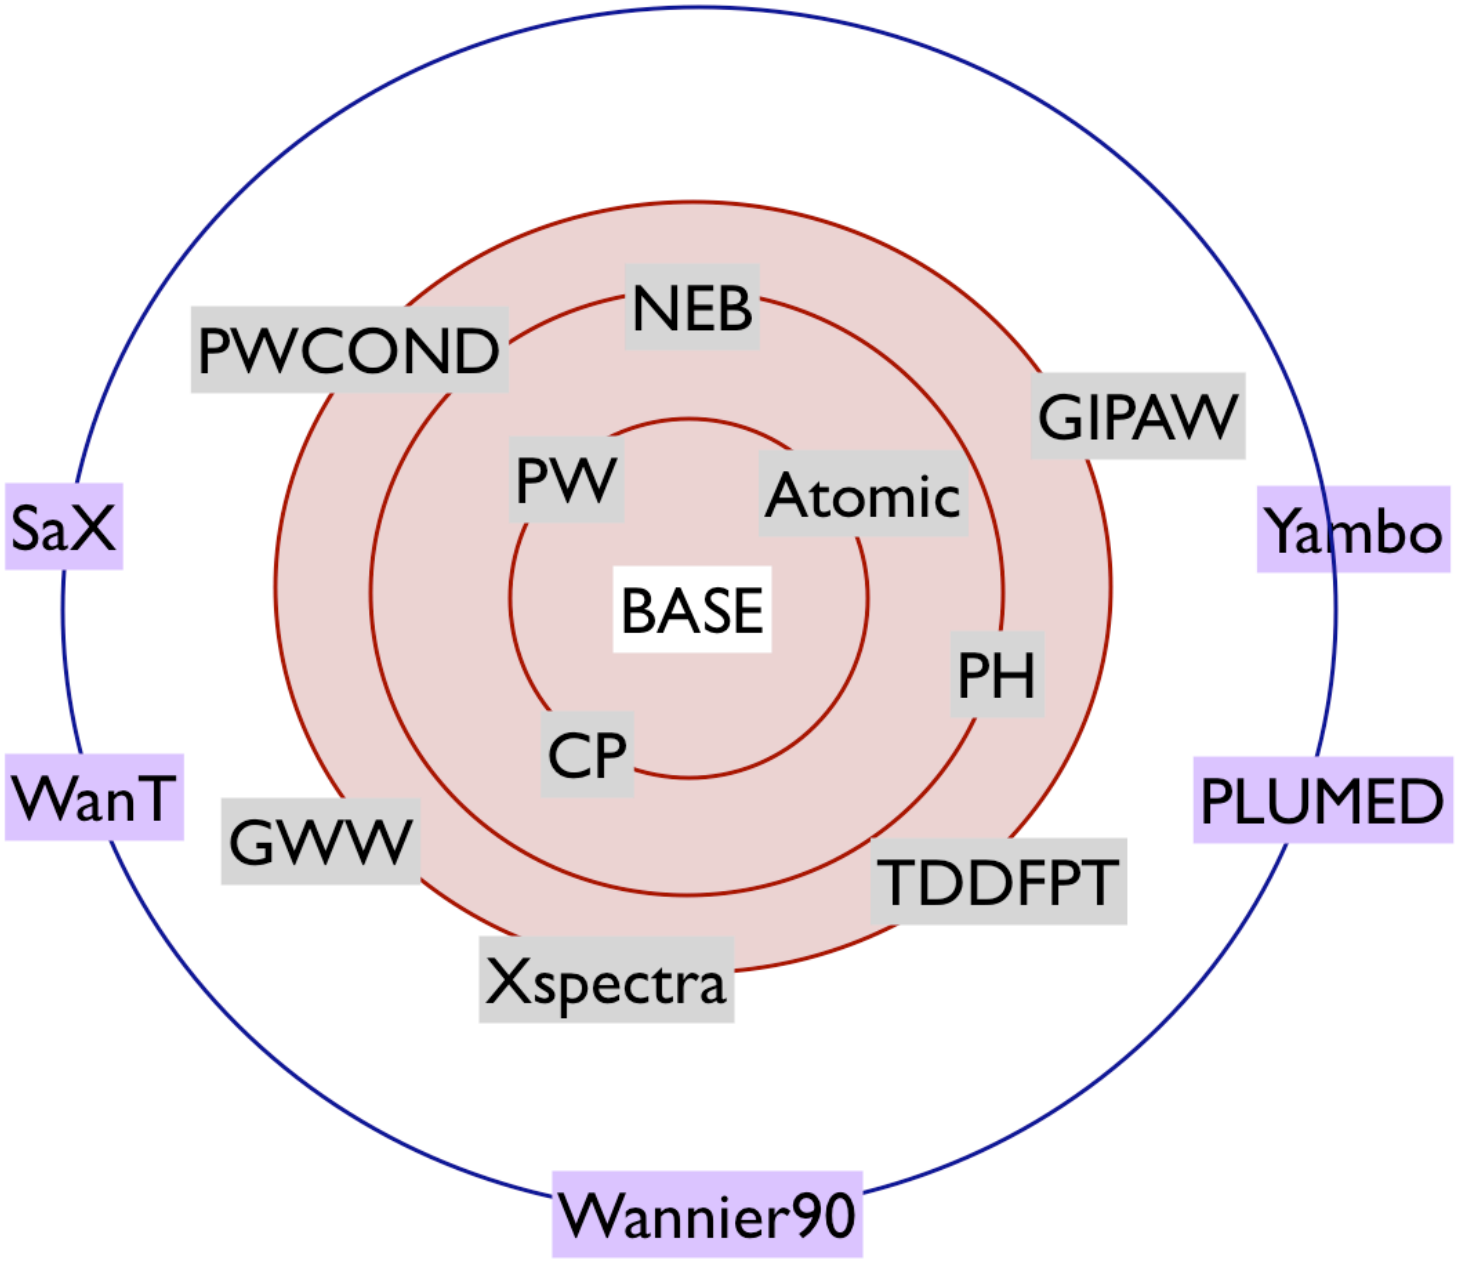

Technical characteristics (coding)

- 380000+ Fortran-95 lines, with various degrees of sophistication (i.e. use of advanced f95 features) – no "dusty decks" any longer
- use of standard library routines (lapack, blas, fftw) to achieve portability – Machine-optimized libraries can (should) (*must!*) be used if available
- C-style preprocessing options allow to keep a single source tree for all architecturesi (GPUs excepted) from PC's to BG's (BlueGene)
- various parallelization levels via MPI calls or OpenMP directives, hidden into calls to a few routines – almost unified serial and parallel versions; parallel code can (usually) be written without knowing the details of how parallelism works.

XML-based data file format

Data format for easy data exchange between different codes:

- a *directory* instead of a single file
- a *formatted* 'head' file contains structural data, computational details, and links to files containing large datasets
- *binary* files for large datasets, one large record per file

Implementation tool: `iotk` toolkit, a lightweight library. Advantages:

- *efficient*: exploits the file system and binary I/O
- *extensible*: based on "fields" introduced by XML syntax
  `<field> ... </field>`
- *easy* to read, write, and understand

What can QUANTUM ESPRESSO do?

- Structural modeling (equilibrium structures of molecules, crystal, surfaces)
- Linear response functions (vibrational and dielectric properties); some non-linear ones (third-order force constants and dielectric response, non-resonant Raman)
- Chemical reactivity and transition-path sampling (Nudged Elastic Band, NEB)
- Dynamical modeling (ab-initio molecular dynamics, Car-Parrinello MD)
- Computational microscopy (simulation of STM images)
- Quantum (ballistic) transport

Advanced QUANTUM ESPRESSO **capabilities**

- several "beyond-DFT" methods: DFT+U, meta-GGA, hybrid functionals, nonlocal van-der-Waals functionals

- free-energy sampling (metadynamics, with PLUMED plugin)

- computational spectroscopy
  - lattice and molecular vibrations: Raman, Infrared, Neutrons
  - magnons and spin excitations
  - photoemission (with Many-Body Perturbation Theory, MBPT)
  - optical/UV absorption (Time-dependent DFT, MBPT)
  - NMR chemical shifts
  - X-ray spectra, core level shifts

Computer requirements of quantum simulations

Quantum simulations are both CPU *and* RAM-intensive.
Actual CPU time and RAM requirements depend upon:

- *size of the system under examination:* As a rule of thumb, CPU $\propto N^{2\div 3}$, RAM $\propto N^2$, where $N =$ number of atoms in the unit cell (or supercell)

- *kind of system:* type and arrangement of atoms, influencing the number of plane waves, of Kohn-Sham orbitals, of **k**-points (in periodic systems) needed...

- *desired results:* computational effort increases from simple self-consistent (single-point) calculation to structural optimization to reaction pathways, molecular-dynamics simulations, ...

CPU time mostly spent in FFT and linear algebra
RAM mostly needed to store Kohn-Sham states

Typical computational requirements

Basic step: self-consistent ground-state DFT electronic structure.

- Simple crystals, small molecules, up to $\sim$ 50 atoms – CPU seconds to hours, RAM up to 1-2 Gb: *may run on single PC*

- Surfaces, larger molecules, complex or defective crystals, up to a few hundreds atoms – CPU hours to days, RAM up to 10-20 Gb: *requires PC clusters or conventional parallel machines*

- Complex nanostructures or biological systems – CPU days to weeks or more, RAM tens to hundreds Gb: *massively parallel machines*

Main factor pushing towards parallel machines is the large CPU requirements — but the need to distribute RAM may also be a strong driving factor.

Parallelization of QUANTUM ESPRESSO

Several *parallelization levels* are implemented; most of them require *fast* interprocess communications.

Scalability of realistic calculations on up to tens of thousands cores, using mixed MPI-OpenMP parallelization, has been demonstrated.

Careful optimization of nonscalable RAM and computations required! Scalability *strongly depends upon the kind and size of system*!

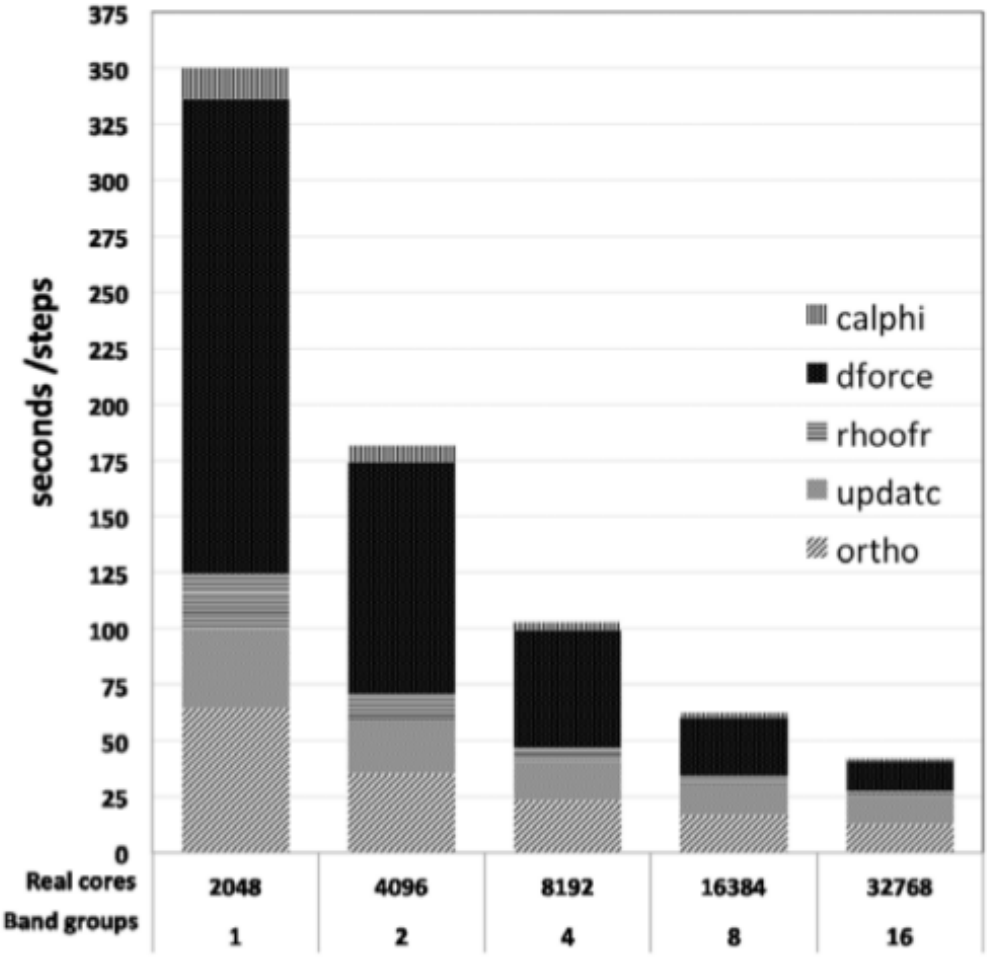

CP Scalability on BG/Q, 1532-atom porphyrin-functionalized carbon nanotube (data from paper appearing in next slide)

Summary of parallelization levels

N. Varini et al. / Computer Physics Communications 184 (2013) 1827–1833

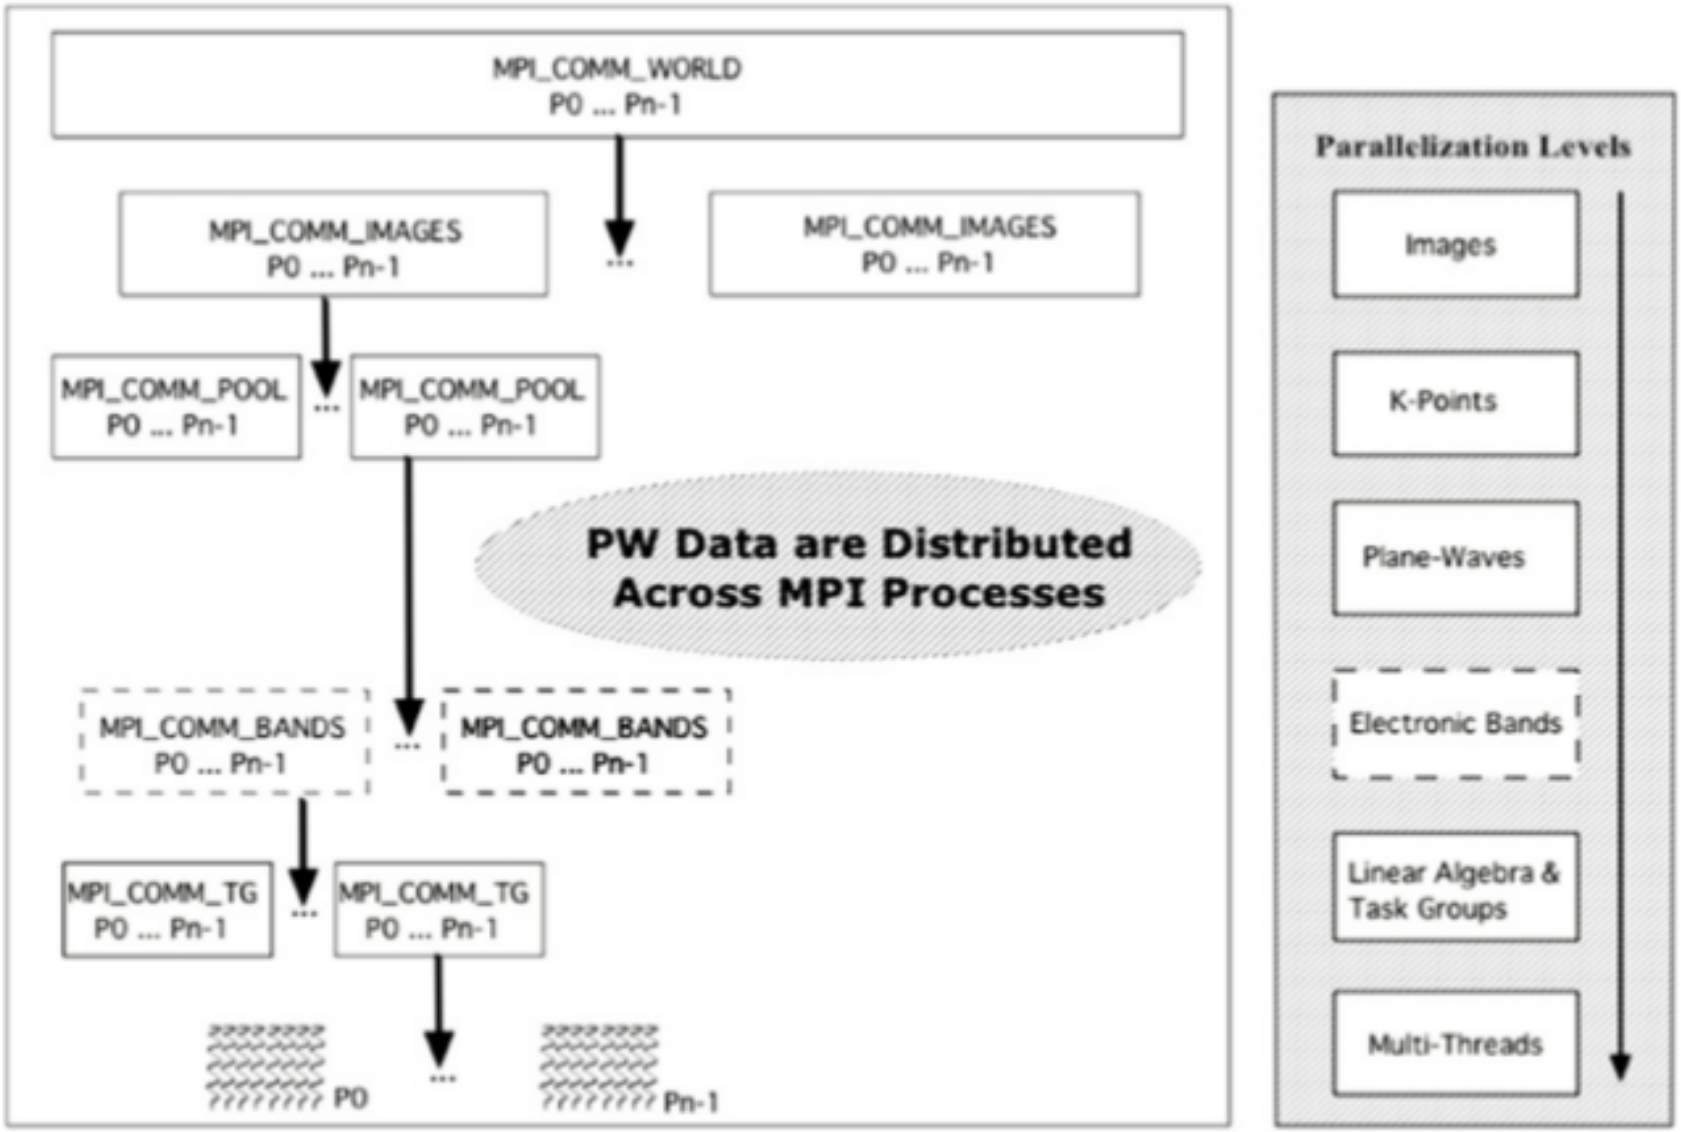

MPI Communicators Hierarchy

Summary of parallelization levels (2)

group	distributed quantities	communications	performances
image	NEB images, phonon modes	very low	linear CPU scaling, fair to good load balancing; does not distribute RAM
pool	$\mathbf{k}$-points	low	almost linear CPU scaling, fair to good load balancing; may distribute some RAM
bands	Kohn-Sham orbitals	high	improves scaling
plane-wave	PW, **G**-vector coefficients, **R**-space FFT arrays	high	good CPU scaling, good load balancing, distributes most RAM
task	FFT on electron states	high	improves load balancing
linear-algebra	subspace hamiltonians and constraints matrices	very high	improves scaling, distributes more RAM
OpenMP	FFT, libraries	intra-node	extends scaling on multicore machines

Importance of collaboration with computing centers

DEISA EXTREME COMPUTING INITIATIVE

Distributed European Infrastructure for Supercomputing Applications

Ab-initio simulations of Protein-Surface Interactions mediated by WATer

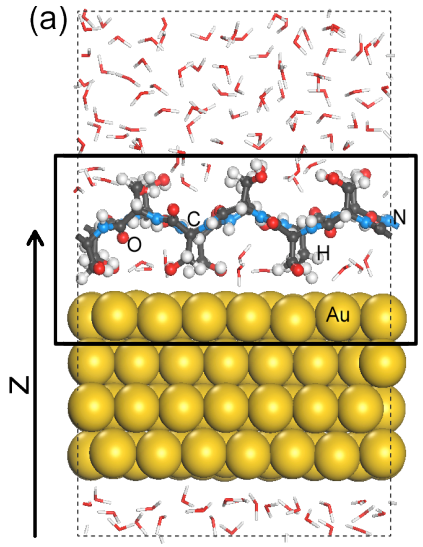

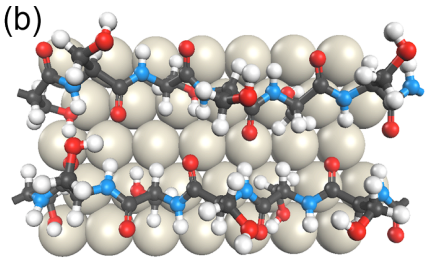

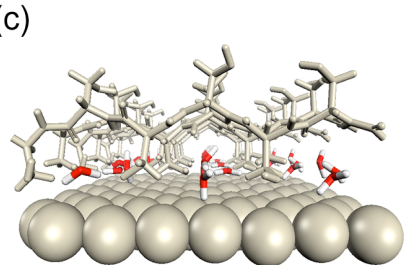

Density map of Oxygen in the hydration layers

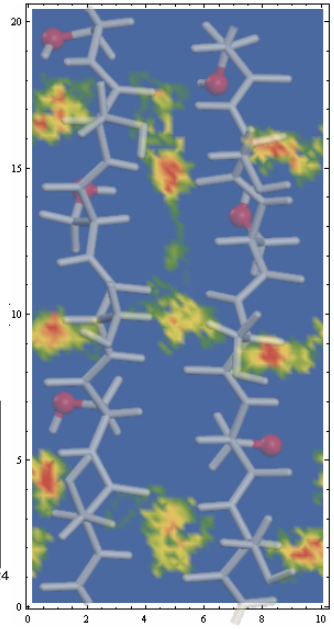

(a-c) Löwdin charges (dq) for selected atoms as a function of the positions, calculated with respect to the corresponding formal atomic values

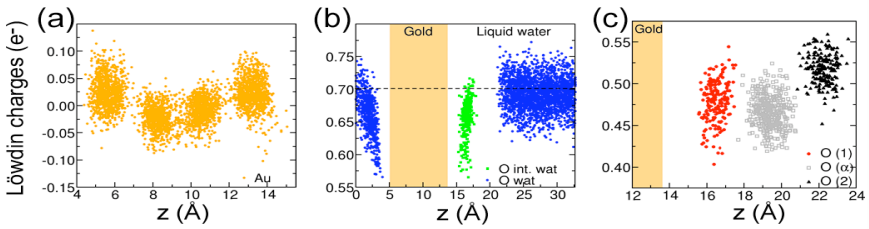

S. Corni, A. Calzolari, G. Cicero, C. Cavazzoni, A. Catellani and R. Di Felice

QUANTUM ESPRESSO on GPU's

fspiga / QE-GPU

Watch 5 | Star 11 | Fork 8

GPU-accelerated Quantum ESPRESSO

323 commits | 2 branches | 4 releases | 2 contributors

branch: master | QE-GPU / +

Tested yet but I wuold not suggest to use it because of the old-fashi...

fspiga authored on Oct 27, 2014 latest commit 5d5c24bf33

GPU	Fix for NVIDIA cuda-proxy	4 months ago
qe-patches	Patches generated for 14.03.0	10 months ago
DISCLAMER	Ready for v14.06.0	8 months ago
License	Uniformed name convention across the entire package	a year ago
README	Updated README	4 months ago
RELEASE_NOTES	Tested yet but I wuold not suggest to use it because of the old-fashi...	4 months ago

Code

Pull Requests 0

Pulse

Graphs

HTTPS clone URL

https://github.com/fsp

You can clone with HTTPS or Subversion.

Download ZIP

README

```
QE-GPU: GPU-Accelerated Quantum ESPRESSO
----------------------------------------

Quantum ESPRESSO is an integrated suite of Open-Source computer codes for
electronic-structure calculations and materials modeling at the nanoscale.

The aim of QE-GPU is to create a "plugin-like" component for the standard
Quantum ESPRESSO package that allows to exploit the capabilities of NVIDIA
GPU graphics cards in order to allow materials scientists to do better and
fast science. GPU acceleration is currently available for the Plane-Wave
Self-Consistent Field (PWscf) code and the energy barriers and reaction
```

Perspectives and Outlook

- More packages for advanced methodologies
- Better-structured distribution, with interfaces to external codes and to python scripting
- Porting to new architectures: hybrid CPU-GPUs, Intel Xeon Phi
- Towards the exascale (really!): communication-reducing and latency-hiding algorithms, parallelization everywhere
- ...

Credits

- Thanks to all people whose slides and pictures I borrowed
- Thanks to all people who contributed to QUANTUM ESPRESSO
- ...and thanks to you all