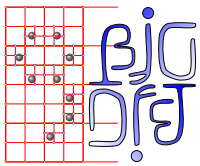

www.bigdft.org

Journées SUCCES – Grenoble – 2017

INSTITUT DE NANOSCIENCES ET CRYOGÉNIE
CEA GRENOBLE, FRANCE

Calcul intensif dans les matériaux : exemple du code ab initio BigDFT

L. Genovese, B; Videau, T. Deutsch

L_Sim – CEA Grenoble

October 17, 2017

Atomistic Simulations

www.bigdft.org

Material Science – Chemistry

- Theory – Experiment – Simulation
- Hardware – Computers
- Algorithms

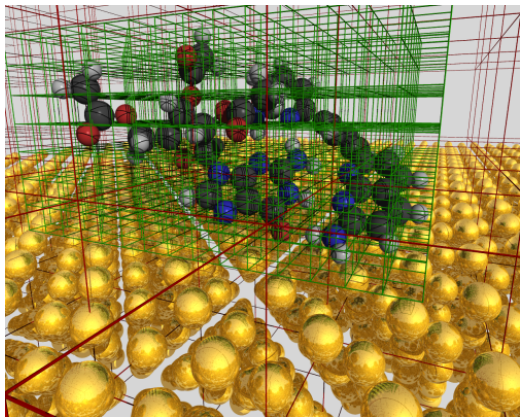

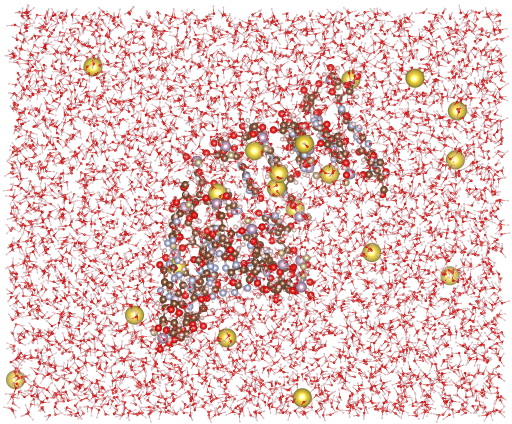

- Atomistic Simulations
  - Force fields (Molecular Dynamics)
  - Semi-Empirical (tight-binding) methods
  - Density Functional Theory
  - Quantum Chemistry (wavefunction-based methods)
  - Quantum Monte-Carlo

- More than 40,000 DFT users

Atomistic Simulations

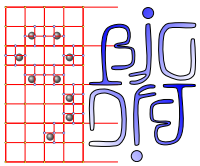

www.bigdft.org

Material Science – Chemistry

- Theory – Experiment – Simulation
- Hardware – Computers
- Algorithms

- Atomistic Simulations
  - Force fields (Molecular Dynamics)
  - Semi-Empirical (tight-binding) methods
  - Density Functional Theory
  - Quantum Chemistry (wavefunction-based methods)
  - Quantum Monte-Carlo

- More than 40,000 DFT users

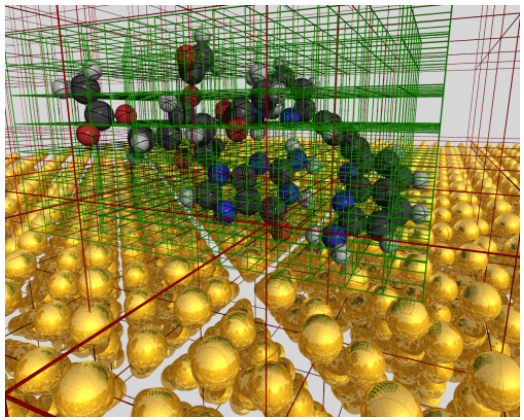

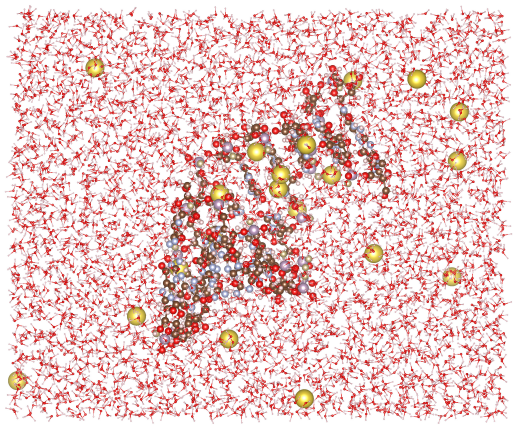

Large use of the Density Functional Theory

www.bigdft.org

- Introduction
- BigDFT code
- MPI/OpenMP
- GPU
- Linear
- Perspectives

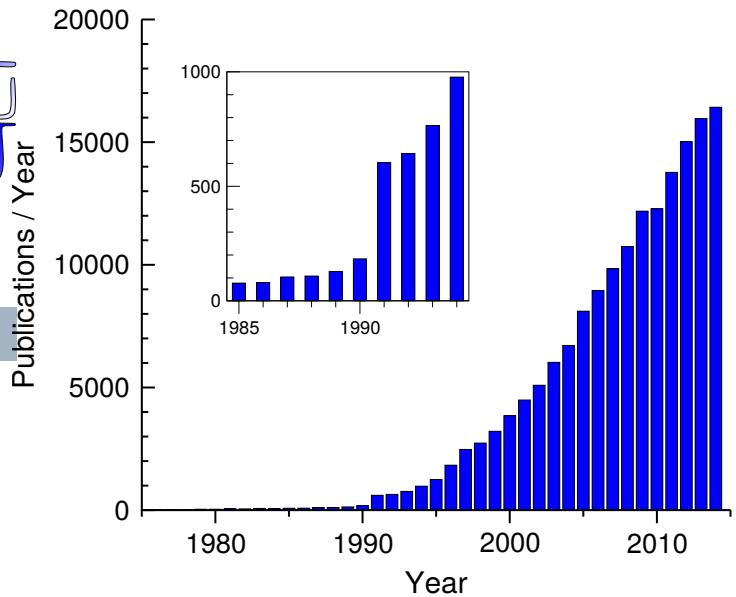

Number of publications per year (1975-2014) on topics "density functional" or "DFT". From Mavropoulos, 2015.

RESEARCH ARTICLE SUMMARY

DFT METHODS

Reproducibility in density functional theory calculations of solids

Kurt Lejaeghere,* Gustav Bihlmayer, Torbjörn Björkman, Peter Blaha, Stefan Blügel, Volker Blum, Damien Caliste, Ivano E. Castelli, Stewart J. Clark, Andrea Dal Corso, Stefano de Gironcoli, Thierry Deutsch, John Kay Dewhurst, Igor Di Marco, Claudia Draxl, Marcin Dułak, Olle Eriksson, José A. Flores-Livas, Kevin F. Garrity, Luigi Genovese, Paolo Giannozzi, Matteo Giantomassi, Stefan Goedecker, Xavier Gonze, Oscar Grånäs, E. K. U. Gross, Andris Gulans, François Gygi, D. R. Hamann, Phil J. Hasnip, N. A. W. Holzwarth, Diana Iuşan, Dominik B. Jochym, François Jollet, Daniel Jones, Georg Kresse, Klaus Koepernik, Emine Küçükbenli, Yaroslav O. Kvashnin, Inka L. M. Locht, Sven Lubeck, Martijn Marsman, Nicola Marzari, Ulrike Nitzsche, Lars Nordström, Taisuke Ozaki, Lorenzo Paulatto, Chris J. Pickard, Ward Poelmans, Matt I. J. Probert, Keith Refson, Manuel Richter, Gian-Marco Rignanese, Santanu Saha, Matthias Scheffler, Martin Schlipf, Karlheinz Schwarz, Sangeeta Sharma, Francesca Tavazza, Patrik Thunström, Alexandre Tkatchenko, Marc Torrent, David Vanderbilt, Michiel J. van Setten, Veronique Van Speybroeck, John M. Wills, Jonathan R. Yates, Guo-Xu Zhang, Stefaan Cottenier*

INTRODUCTION: The reproducibility of results is one of the underlying principles of science. An observation can only be accepted by the scientific community when it can be confirmed by independent studies. However, reproducibility does not come easily. Recent works have painfully exposed cases where previous conclusions were not upheld. The scrutiny of the scientific community has also turned to research involving computer programs, finding that reproducibility depends more strongly on implementation than commonly thought. These problems are especially relevant for property predictions of crystals and molecules, which hinge on precise computer implementations of the governing equation of quantum physics.

RATIONALE: This work focuses on density functional theory (DFT), a particularly popular quantum method for both academic and industrial applications. More than 15,000 DFT papers are published each year, and DFT is now increasingly used in an automated fashion to build large databases or apply multiscale techniques with limited human supervision. Therefore, the reproducibility of DFT results underlies the scientific credibility of a substantial fraction of current work in the natural and engineering sciences. A plethora of DFT computer codes are available, many of them differing considerably in their details of implementation, and each yielding a certain "precision" relative to other codes. How is one to decide for more than a few simple cases which code predicts the correct result, and which does not? We devised a procedure to assess the precision of DFT methods and used this to demonstrate reproducibility among many of the most widely used DFT codes. The essential part of this assessment is a pairwise comparison of a wide range of methods with respect to their predictions of the equations of state of the elemental crystals. This effort required the combined expertise of a large group of code developers and expert users.

RESULTS: We calculated equation-of-state data for four classes of DFT implementations, totaling 40 methods. Most codes agree very well, with pairwise differences that are comparable to those between different high-precision experiments. Even in the case of pseudization approaches, which largely depend on the atomic potentials used, a similar precision can be obtained as when using the full potential. The remaining deviations are due to subtle effects, such as specific numerical implementations or the treatment of relativistic terms.

CONCLUSION: Our work demonstrates that the precision of DFT implementations can be determined, even in the absence of one absolute reference code. Although this was not the case 5 to 10 years ago, most of the commonly used codes and methods are now found to predict essentially identical results. The established precision of DFT codes not only ensures the reproducibility of DFT predictions but also puts several past and future developments on a firmer footing. Any newly developed methodology can now be tested against the benchmark to verify whether it reaches the same level of precision. New DFT applications can be shown to have used a sufficiently precise method. Moreover, high-precision DFT calculations are essential for developing improvements to DFT methodology, such as new density functionals, which may further increase the predictive power of the simulations.

ON OUR WEB SITE

Read the full article at http://dx.doi.org/10.1126/science.aad3000

The list of author affiliations is available in the full article online.
*Corresponding author. E-mail: kurt.lejaeghere@ugent.be (K.L.); stefaan.cottenier@ugent.be (S.C.)
Cite this article as K. Lejaeghere et al., Science 351, aad3000 (2016). DOI: 10.1126/science.aad3000

Recent DFT methods yield reproducible results. Whereas older DFT implementations predict different values (red darts), codes have now evolved to mutual agreement (green darts). The scoreboard illustrates the good pairwise agreement of four classes of DFT implementations (horizontal direction) with all-electron results (vertical direction). Each number reflects the average difference between the equations of state for a given pair of methods, with the green-to-red color scheme showing the range from the best to the poorest agreement.

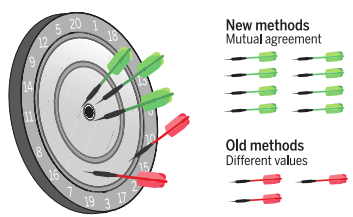

Scorecard

	0.3	0.3	0.6	1.0	0.9	0.3	1.5	0.6	0.9	0.4	1.0	0.4	0.4	6.3	13.5	1.1	2.1	0.7	1.4
0.3		0.1	0.5	0.9	0.8	0.2	1.5	0.6	0.8	0.4	1.0	0.5	0.3	6.3	13.4	1.1	2.1	0.7	1.4
0.3	0.1		0.5	0.9	0.8	0.2	1.5	0.6	0.8	0.4	0.9	0.5	0.3	6.3	13.4	1.1	2.1	0.7	1.4
0.6	0.5	0.5		0.8	0.6	0.4	1.5	0.6	0.8	0.6	1.0	0.7	0.5	6.3	13.2	1.0	1.9	0.6	1.3
1.0	0.9	0.9	0.8		0.9	0.9	1.8	0.9	1.3	1.0	1.4	1.0	0.9	6.4	13.0	1.2	1.8	1.0	1.6
0.9	0.8	0.8	0.6	0.9		0.8	1.7	0.7	1.1	0.8	1.3	1.0	0.8	6.5	13.2	1.1	1.8	0.8	1.5
0.3	0.2	0.2	0.4	0.9	0.8		1.5	0.5	0.8	0.3	1.0	0.5	0.3	6.2	13.4	1.0	2.0	0.6	1.4

Outline

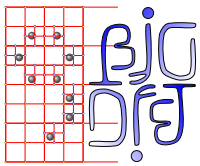

www.bigdft.org

Ab initio calculations with DFT

Several advantages

- ✔ Ab initio: No adjustable parameters
- ✔ DFT: Quantum mechanical (fundamental) treatment

Main limitations

- ✘ Approximated approach (electron correlations)
- ✘ Requires high computer power, limited to few hundreds atoms in most cases

Wide range of applications

Nanoscience, biology, materials

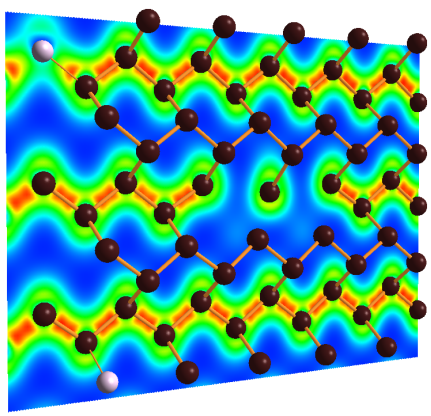

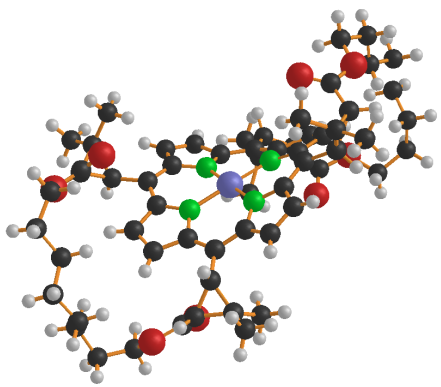

Finding the most stable atomic configuration

Need a large number of calculations

Different methods linked with BigDFT:

- Minima Hopping (S. Goedecker, Basel)
- Activation-Relaxation Technique (N. Mousseau, Montreal)

Applied on different systems

Benefit from high flexibility and performances

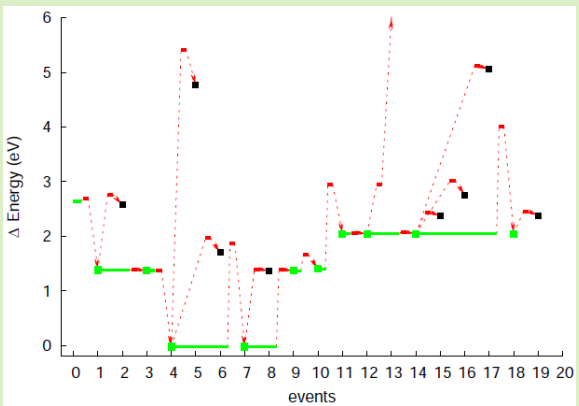

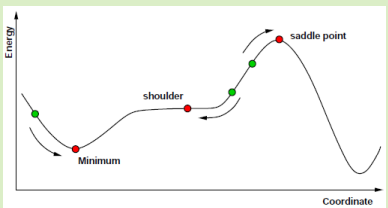

Interesting for potential synthesis pathways

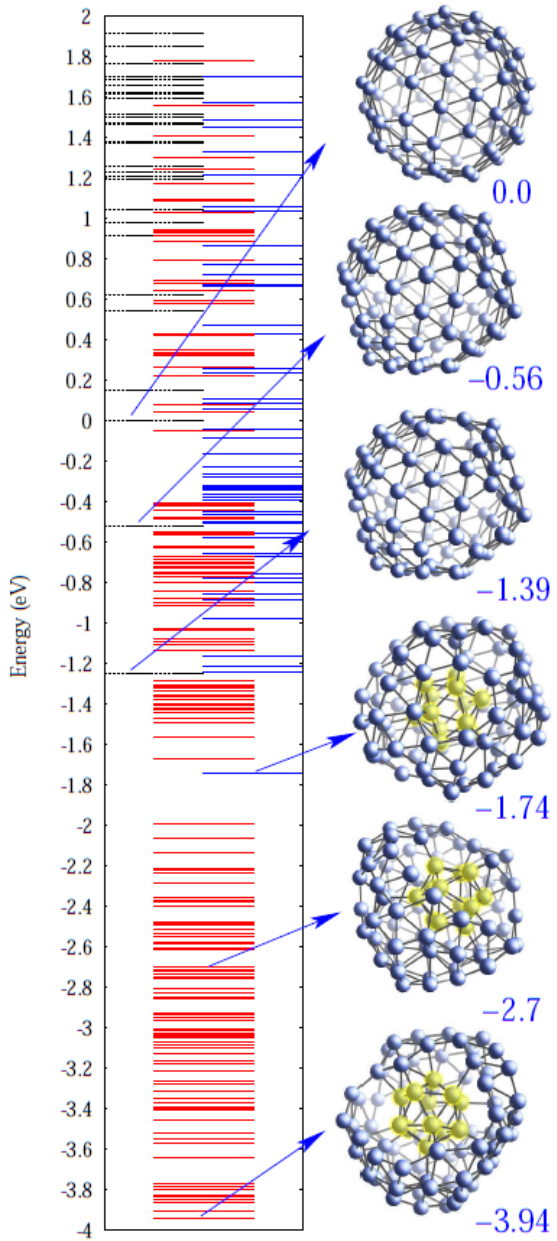

A basis for nanosciences: the BigDFT project

www.bigdft.org

STREP European project: BigDFT(2005-2008)

Four partners, 15 contributors:
CEA-INAC Grenoble, U. Basel, U. Louvain-la-Neuve, U. Kiel

Aim: To develop an ab-initio DFT code based on Daubechies Wavelets for large system calculations,
distributed freely (GNU-GPL license)

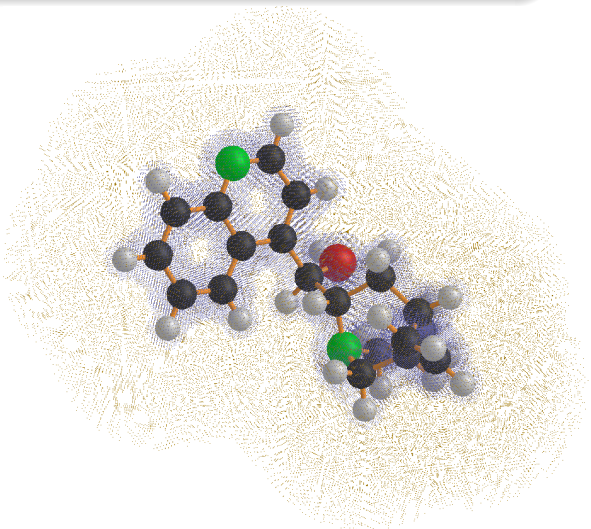

References

"Daubechies wavelets as a basis set for density functional pseudopotential calculations",

L. Genovese, A. Neelov, S. Goedecker, T. Deutsch, *et al.*, J. Chem. Phys. **129**, 014109 (2008)

"Daubechies wavelets for linear scaling density functional theory", S. Mohr, L. Genovese, T. Deutsch,

S. Goedecker, *et al.*, J. Chem. Phys. **140**, 204110 (2014)

Goal

www.bigdft.org

Wavelets

an ideal basis for electronic structure calculations – flexible, systematic etc.

(Linear-scaling) DFT

allows us to access very large system sizes via the use of a localized minimal basis set

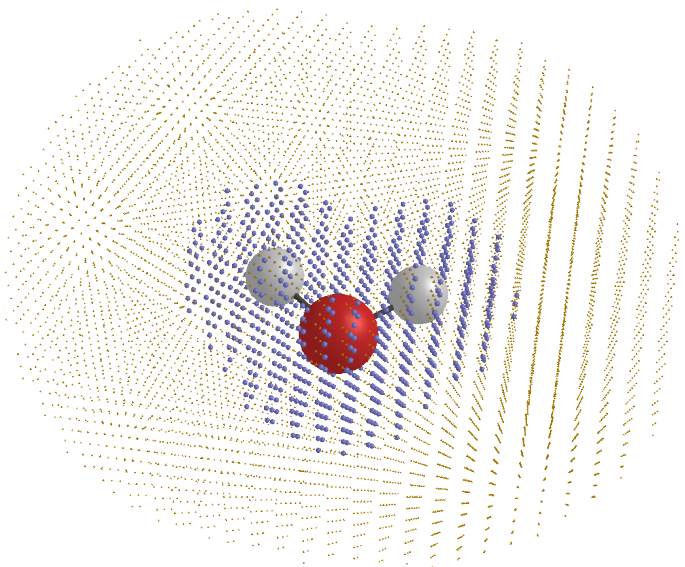

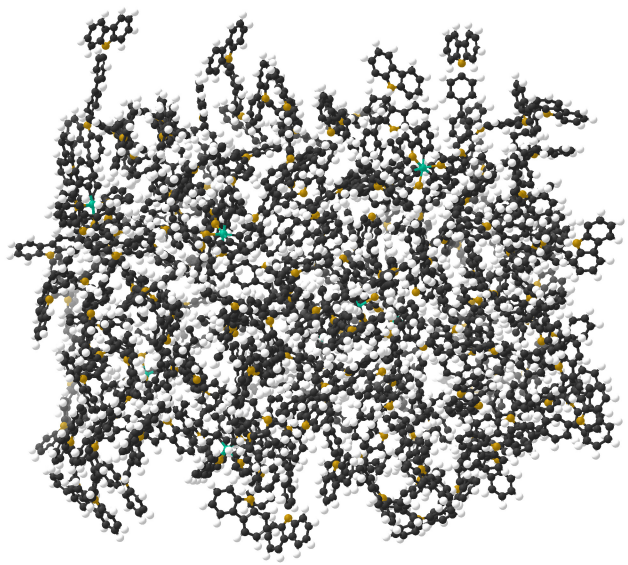

we want to combine the two...

Goal

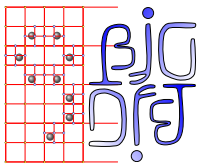

www.bigdft.org

Massively parallel architectures

... and run calculations on large, *realistic* systems, using massively parallel machines

Why do we use wavelets in BigDFT?

Adaptivity

One grid, two resolution levels in BigDFT:

- 1 scaling function ("coarse region")
- 1 scaling function and 7 wavelets ("fine region")

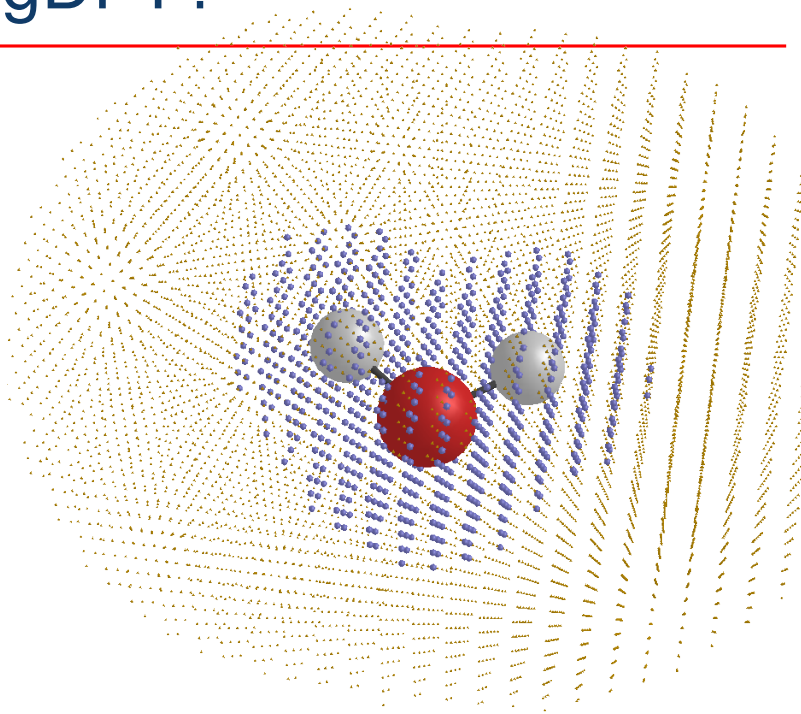

- Ideal for big inhomogeneous systems
- Efficient Poisson solver, capable of handling different boundary conditions – free, wire, surface, periodic
- Established code with many capabilites

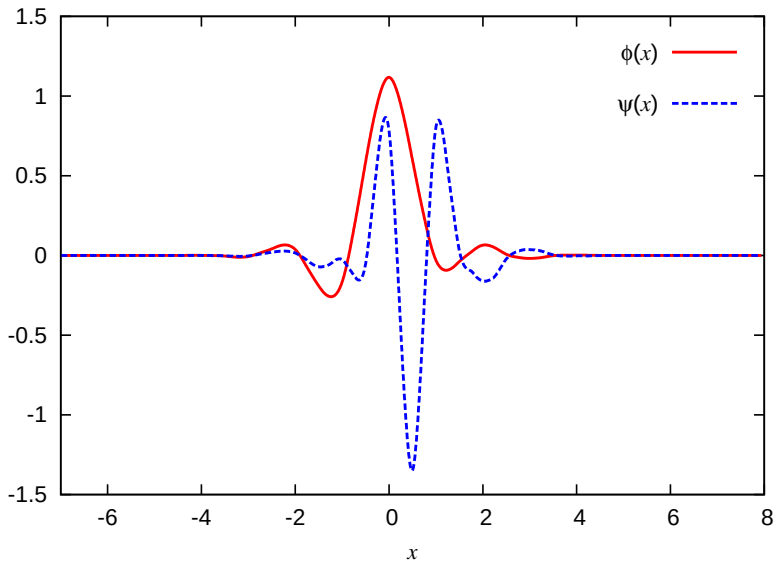

A brief description of wavelet theory

Two kind of basis functions

A Multi-Resolution real space basis

The functions can be classified following the resolution level they span.

Scaling Functions

The functions of low resolution level are a linear combination of high-resolution functions

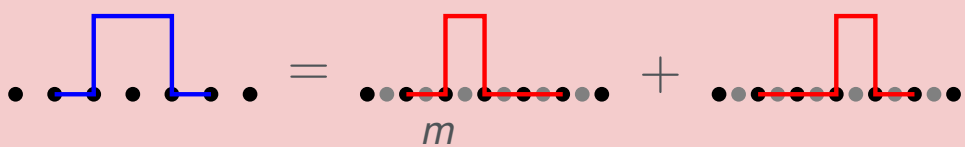

$$\phi(x) = \sum_{j=-m}^{m} h_j \phi(2x-j)$$

Centered on a resolution-dependent grid: $\phi_j = \phi_0(x-j)$.

A brief description of wavelet theory

Wavelets

They contain the DoF needed to complete the information which is lacking due to the coarseness of the resolution.

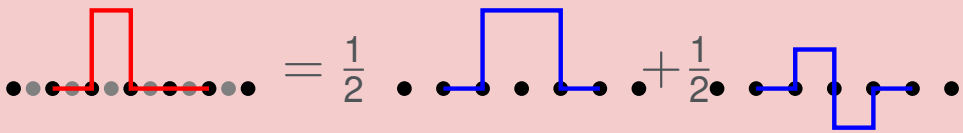

$$\phi(2x) = \sum_{j=-m}^{m} \tilde{h}_j \phi(x-j) + \sum_{j=-m}^{m} \tilde{g}_j \psi(x-j)$$

Increase the resolution without modifying grid space

SF + W = same DoF of SF of higher resolution

$$\psi(x) = \sum_{j=-m}^{m} g_j \phi(2x-j)$$

All functions have compact support, centered on grid points.

Adaptivity of the mesh

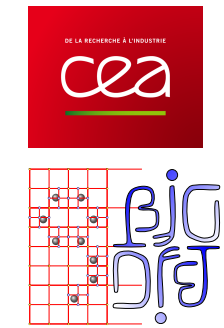

www.bigdft.org

Atomic positions (H_2O)

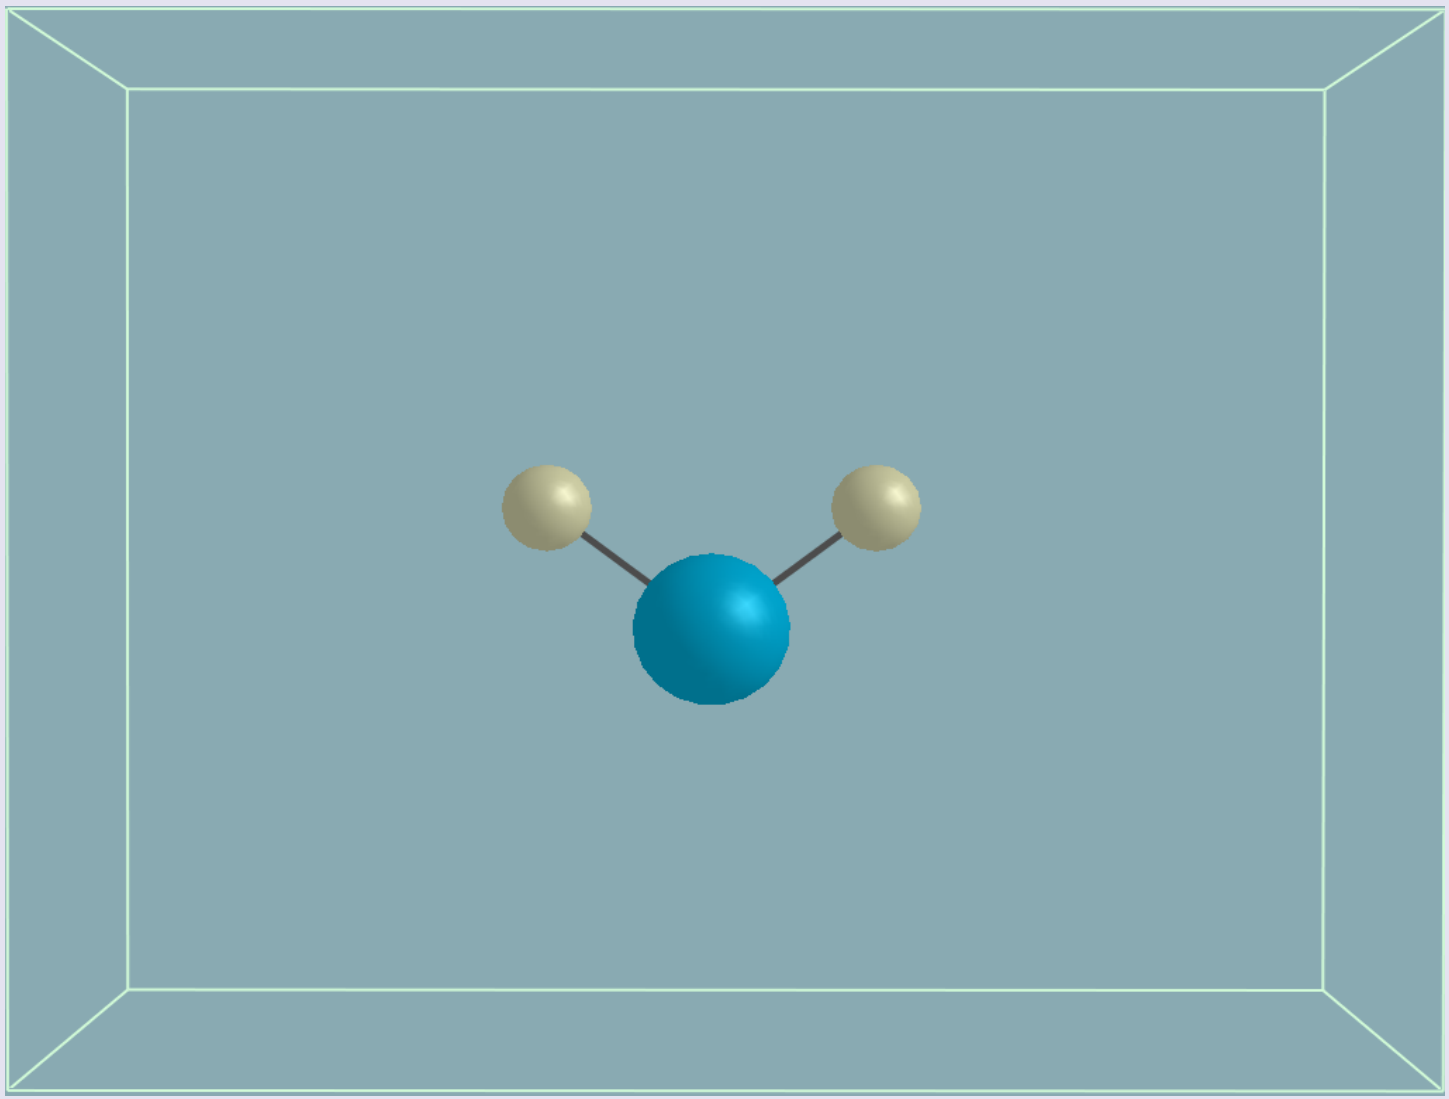

Adaptivity of the mesh

Fine grid (high resolution)

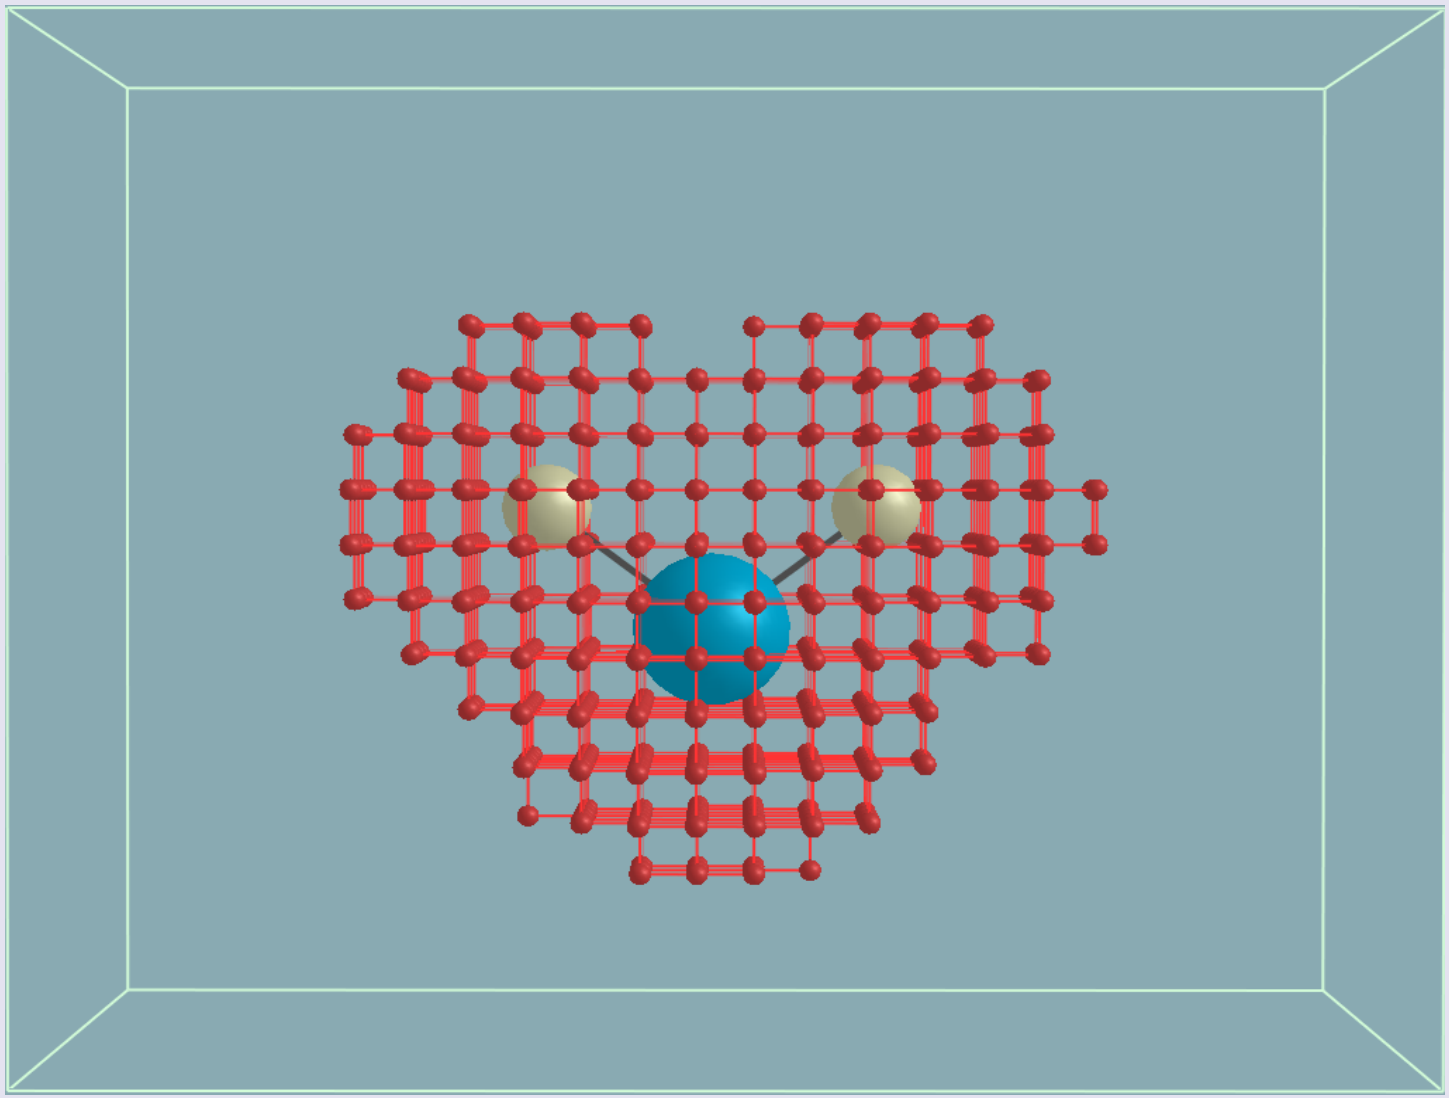

Adaptivity of the mesh

www.bigdft.org

Coarse grid (low resolution)

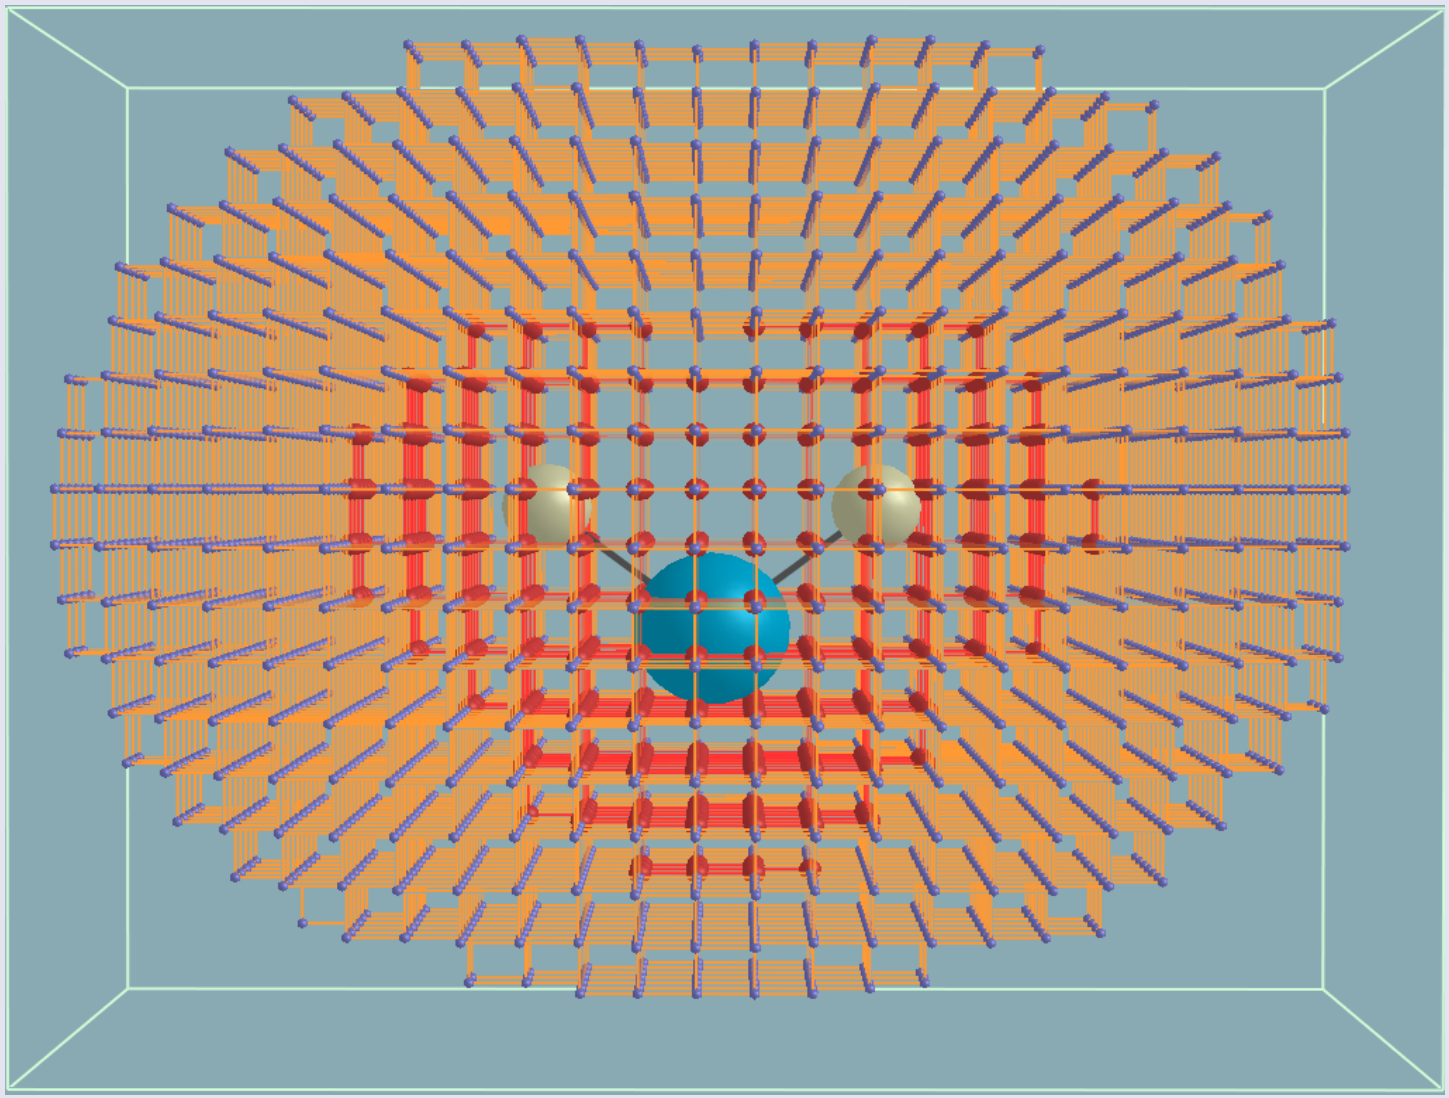

Systematic basis set

Two parameters for tuning the basis

- The grid spacing
- The extension of the grid

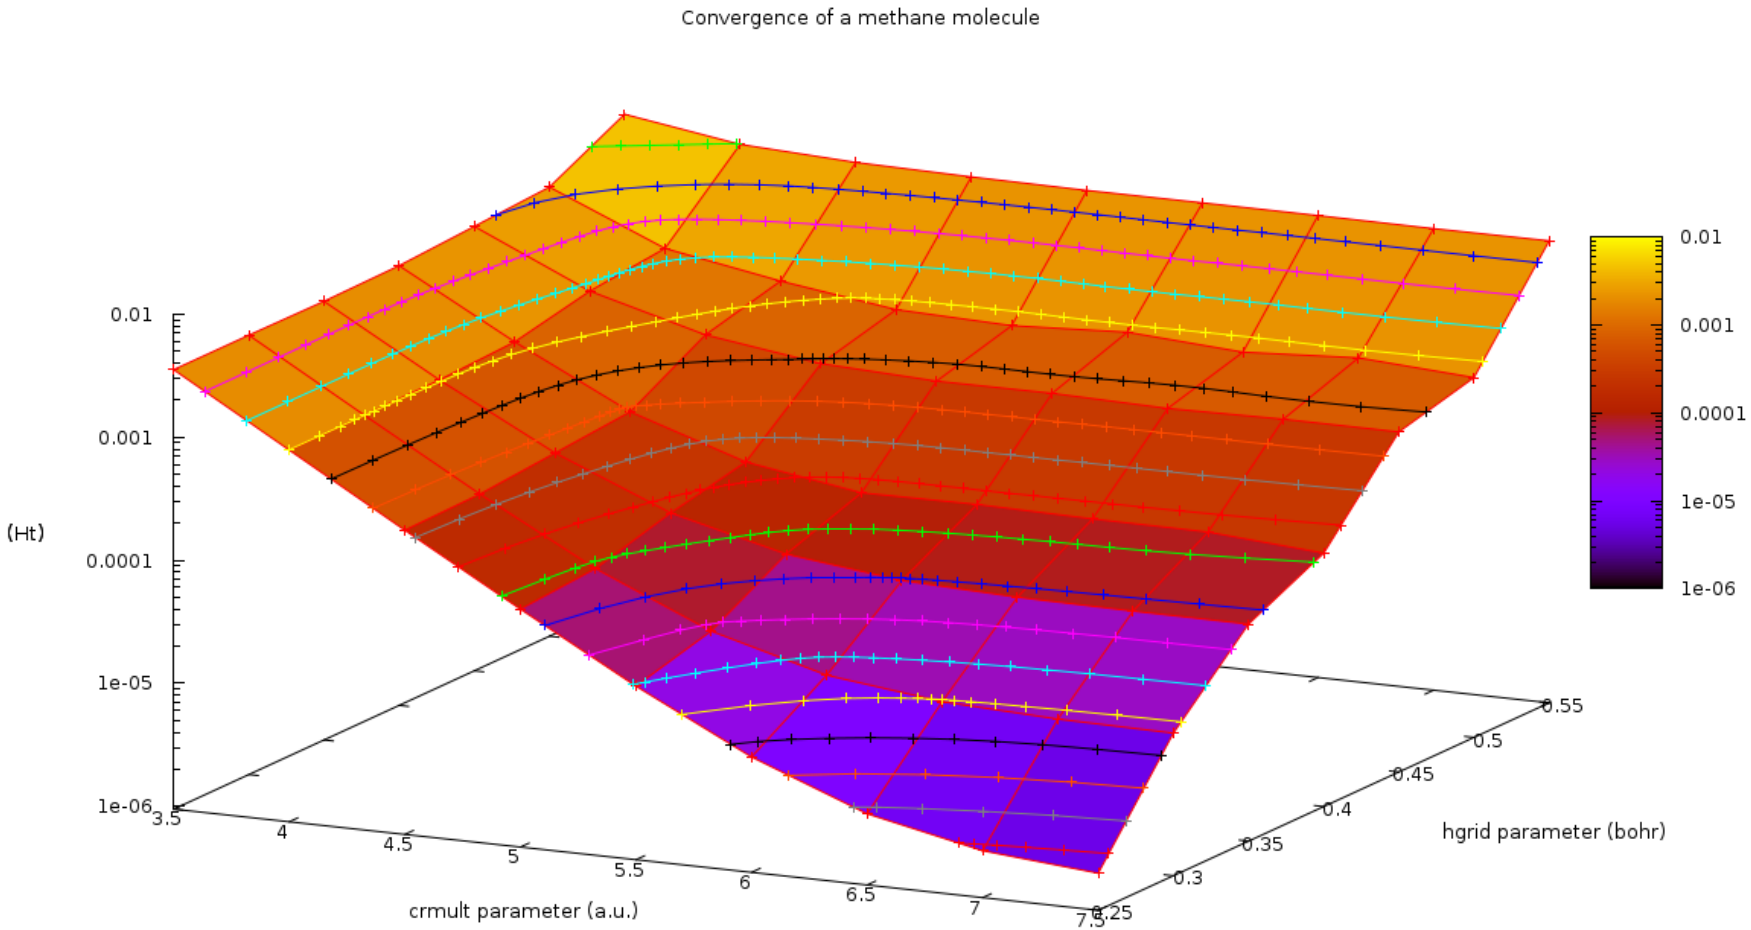

Deltatest benchmark: $\Delta = 1.0$

Science 351, 6280 (2016)

www.bigdft.org

For first three rows

Code	Version	Basis	Electron treatment	Δ-value	Authors
WIEN2k	13.1	LAPW/APW+lo	all-electron	0 meV/atom	S. Cottenier
BigDFT	1.7.6	Daubechies wavelets	HGHk-semicore and NLCC 2015 norm-conserving	0.1 meV/atom	BigDFT [11]
Elk	3.1.5	APW+lo	all-electron	0.2 meV/atom	Elk [14]
VASP	5.2.12	plane waves	PAW 2015 GW-ready (5.4)	0.2 meV/atom	K. Lejaeghere
Quantum ESPRESSO	5.1	plane waves	SSSP Accuracy (mixed NC/US/PAW potential library)	0.2 meV/atom	QuantumESPRESSO [12]
FLEUR	0.26	LAPW (+lo)	all-electron	0.2 meV/atom	FLEUR [9]
FHI-aims	081213	tier2 numerical orbitals	all-electron (relativistic atomic_zora scalar)	0.3 meV/atom	ASE [2]

Screenshot of DeltaTest webpage as of 24/02/16, elements up to Ar, new NLCC - HGH - NC - PSP (S. Saha)

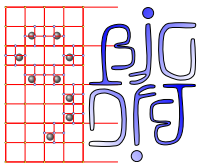

www.bigdft.org

A code both for Solid-State and Quantum Chemistry

- 3D periodic, Surfaces and Free BC (← Poisson Solver)
- Very high precision (analytic KS operators)
- Usage of analytic HGH pseudopotentials
- AE accuracy, benchmarked in G2-1, S22, DeltaTest

Present functionalities

Traditional functionalities for GS Kohn-Sham DFT (including metals, Hybrid Functionals), LR-TDDFT, empirical VdW
Exhaustive library of Structural Prediction, $O(N)$ calculations

Under implementation

Non orthorhombic cells, Systems embedded in electrostatic environments JCP 144, 014103 (2016),
Resonant States extraction

Open Source (GPL) Launchpad site

www.bigdft.org

- Introduction
- BigDFT code
- MPI/OpenMP
- GPU
- Linear
- Perspectives

BigDFT

Overview | Code | Bugs | Blueprints | Translations | Answers

Registered 2013-02-13 by Luigi Genovese

BigDFT is an ab initio code based on Daubechies wavelets.

Such functions have features which make them a powerful and promising basis set for application in materials science. These are a compact support multiresolution basis, and form one of the few examples of systematic real space basis sets. They are an optimal basis for expanding localised information. The real space description they provide allows to build an efficient, clean method to treat systems in complex environments, like surfaces geometries or system with a net charge. The mathematical properties of the formalism are optimal to build a robust, highly optimised code, conceived for systems of few hundred atoms, with excellent efficiency on parallel computers

BigDFT is a DFT massively parallel electronic structure code using a wavelet basis set with the capability to use a linear scaling method.
Wavelets form a real space basis set distributed on an adaptive mesh (two levels of resolution in our implementation).
GTH or HGH pseudopotentials are used to remove the core electrons.
The code BigDFT is available in ABINIT v5.5 and higher but can also be downloaded in a standalone version from the website.
Thanks to our Poisson solver based on a Green function formalism, periodic systems, surfaces and isolated systems can be simulated with explicit boundary conditions.
The Poisson solver can also be downloaded and used independently and is integrated in ABINIT, Octopus and CP2K.
The code is free software, available under GNU-GPL license and the BigDFT developer community encourages anyone willing to contribute to join the team.

Change branding

Home page Wiki

BigDFT breakdown process (1.8.0)

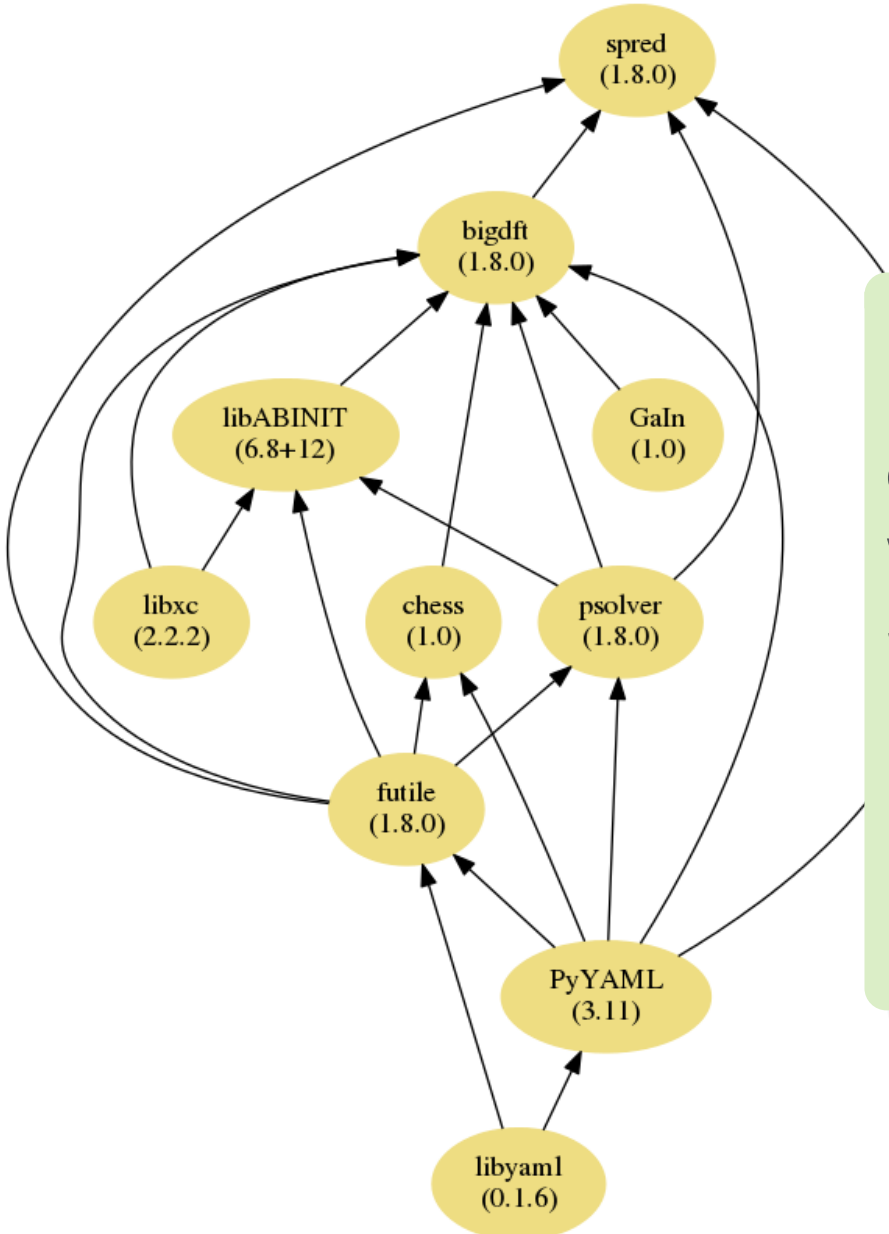

Each section of BigDFT is, when appropriate, defined as a module with its own build system and compilation instructions. At present:

- FUTILE 1.0
- CheSS 1.0
- PSolver 1.8

Completely separated from BigDFT code

www.bigdft.org

Chebyshev **S**parse **S**olvers — Fundamental for $O(N)$ BigDFT

- Can be built completely independently from BigDFT
- Comes with a set of tests and examples
- Integration in ELSI package (PEXSI, OMM, ELPA,. . .)

Comparison with PEXSI in BigDFT

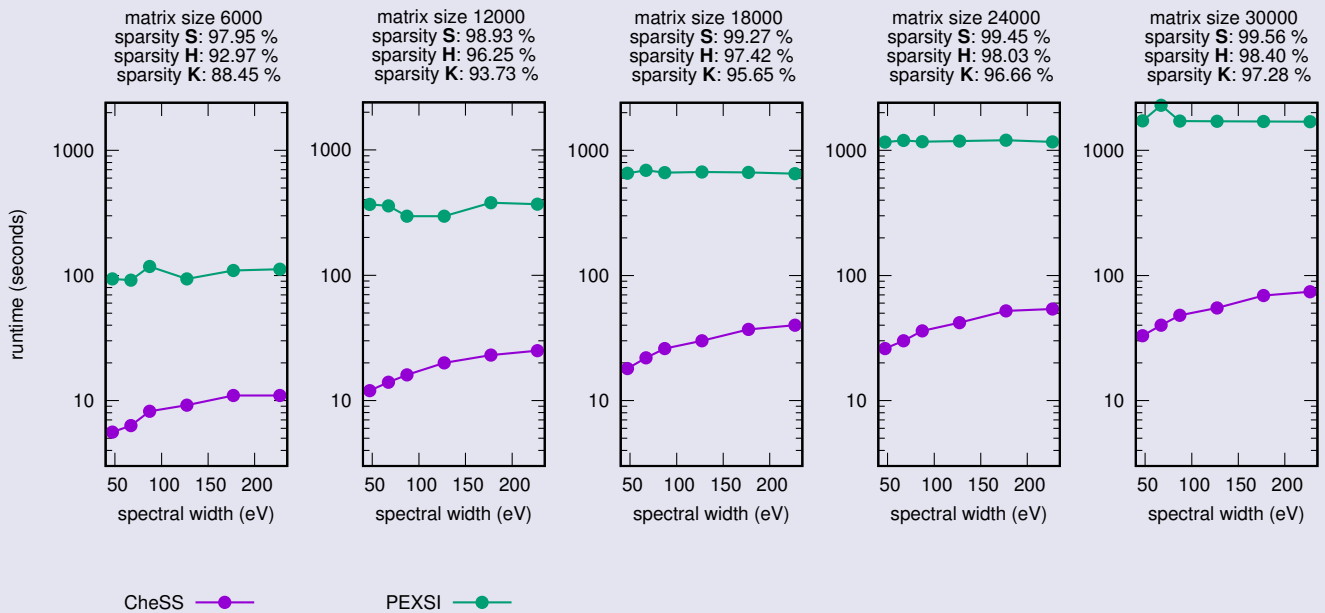

Outline

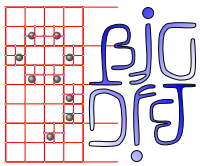

www.bigdft.org

Massively parallel (MPI + OpenMP)

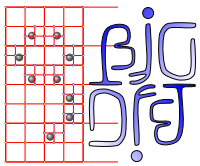

www.bigdft.org

Two kinds of parallelisation

- By orbitals (Hamiltonian application, preconditioning)
- By components (overlap matrices, orthogonalisation)

A few (but large) packets of data

More demanding in bandwidth than in latency

- Better data locality (Hamiltonian application and orthonormality)
- Optimal speedup (eff. $\sim$ 85%), also for big systems

Cubic scaling code

For systems bigger than 500 atomes (1500 orbitals) : orthonormalisation operation is predominant (N^3)

Orbital distribution scheme

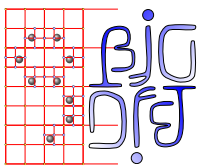

www.bigdft.org

Used for the application of the hamiltonian

The hamiltonian (convolutions) is applied separately onto each wavefunction

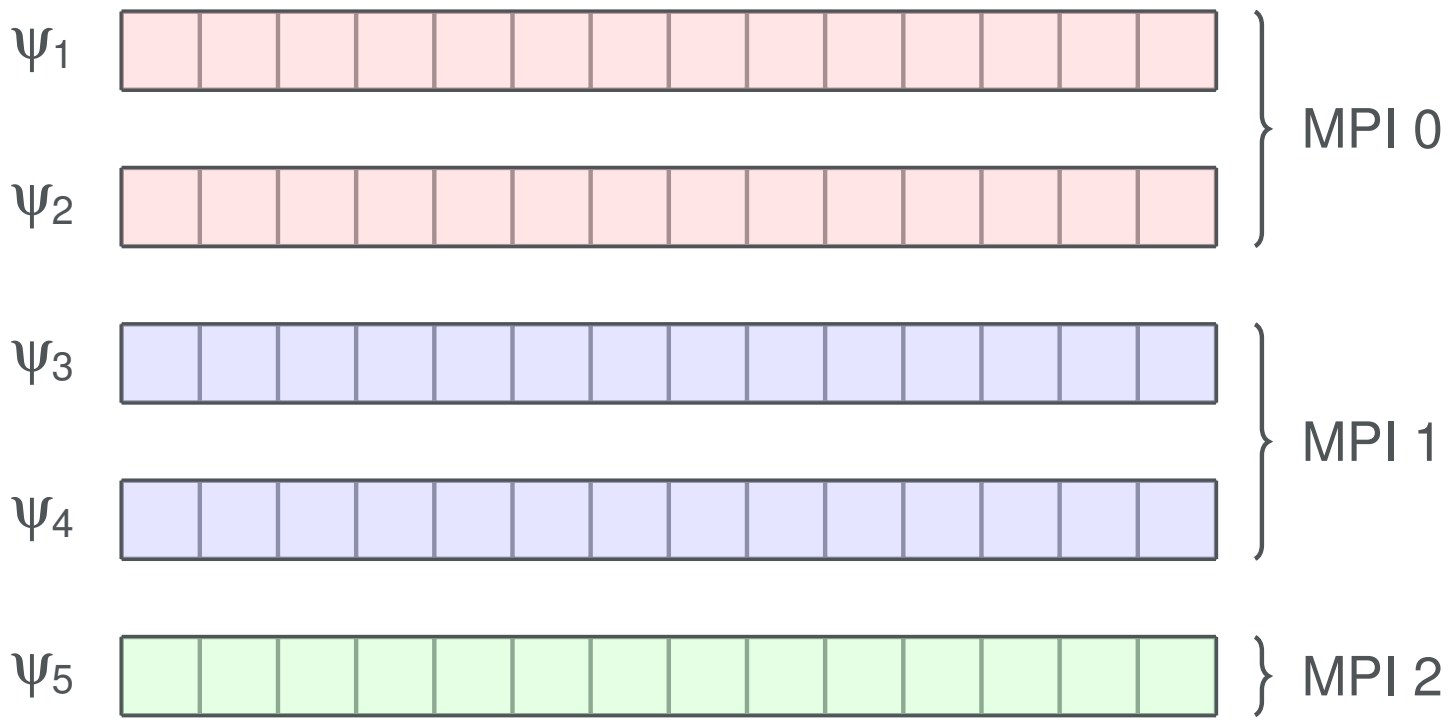

Coefficient distribution scheme

Used for scalar product & orthonormalisation

BLAS routines (level 3) are called, then result is reduced

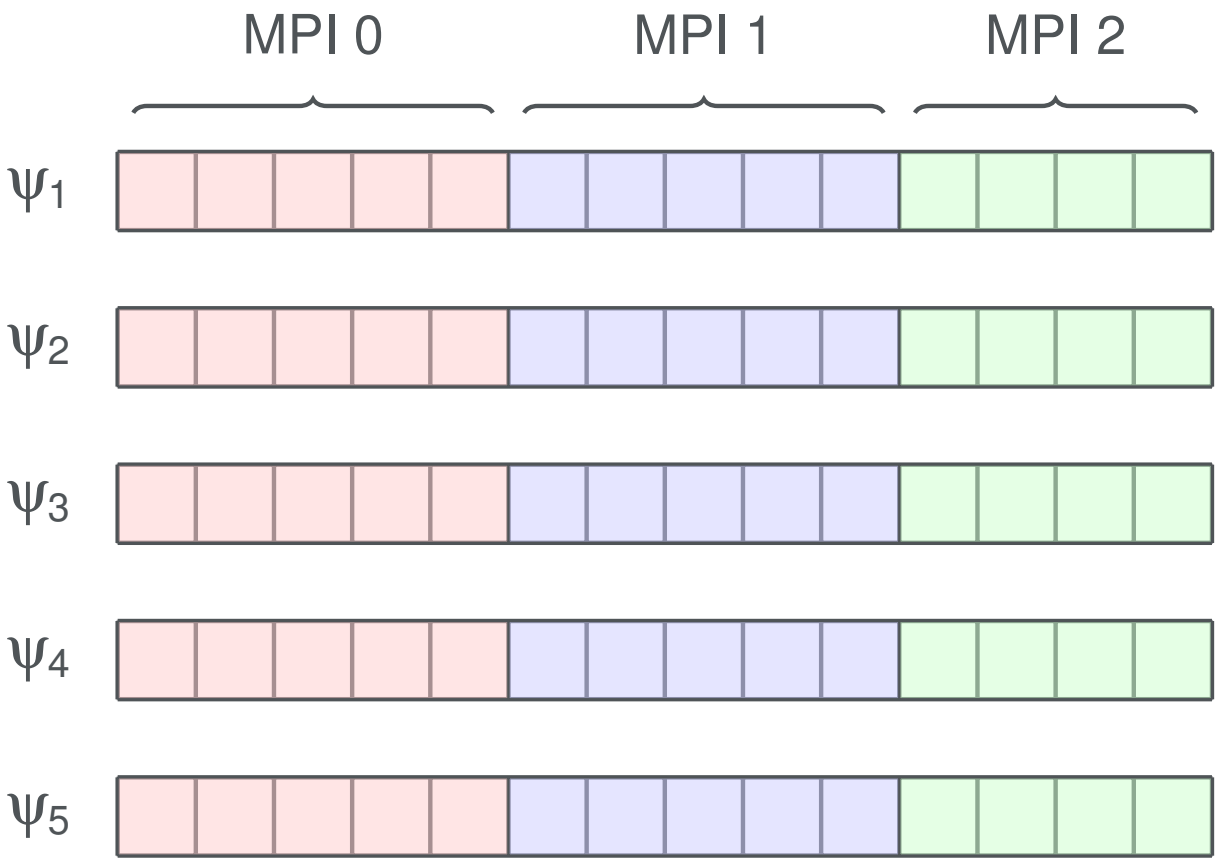

Communications are performed via MPI_ALLTOALLV

Using GPUs in a Big systematic DFT code

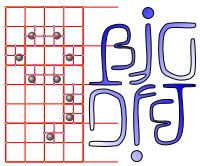

www.bigdft.org

Nature of the operations

- Operators approach via convolutions
- Linear Algebra due to orthogonality of the basis
- Communications and calculations do not interfere

☛ A number of operations which can be accelerated

Evaluating GPU convenience

Three levels of evaluation

1. Bare speedups: GPU kernels *vs.* CPU routines
   Does the operations are suitable for GPU?
2. Full code speedup on one process
   Amdahl's law: are there hot-spot operations?
3. Speedup in a (massively?) parallel environment
   The MPI layer adds an extra level of complexity

BigDFT code on Hybrid architectures (2009)

No Hot-spot operations

Different code sections can be ported on GPU
up to 20x speedup for some operations,
7x for the full parallel code

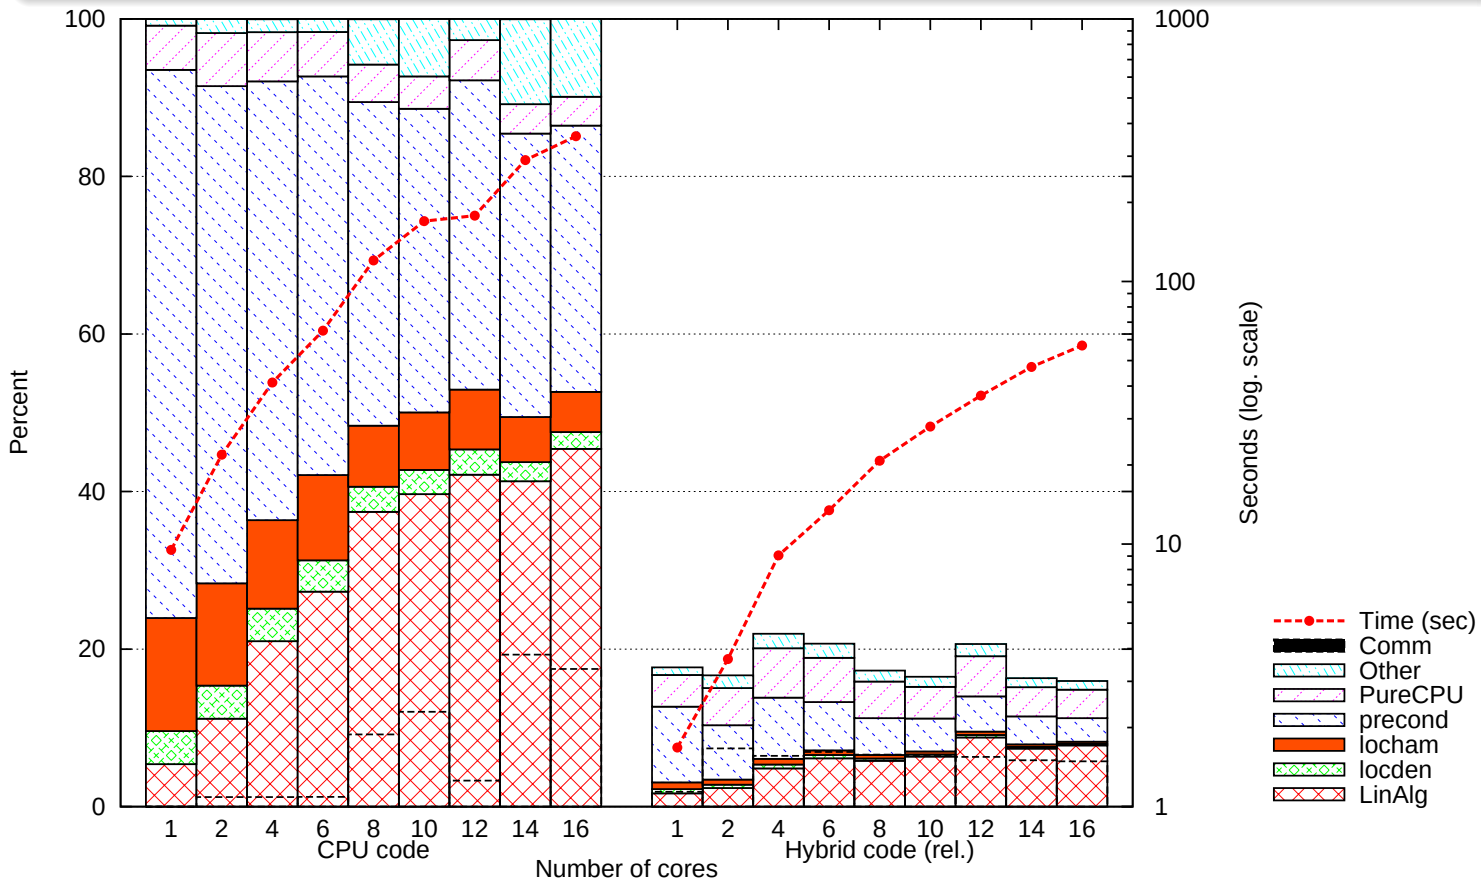

BigDFT code on Hybrid architectures (2009)

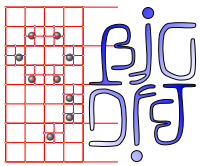

www.bigdft.org

No Hot-spot operations

Different code sections can be ported on GPU

up to 20x speedup for some operations,

7x for the full parallel code

Bull-Fourier 2009 Prize

Reference Paper: L. Genovese *et al.*, J. Chem. Phys. **131**, 034103 (2009)

Use of Titan (Oak Ridge) with 18,000 GPUs

Hybrid and Heterogeneous runs with OpenCL (2011)

NVidia S2070

Connected each in a workstation

BigDFT may run on both

ATI HD 6970

Sample BigDFT run: Graphene, 4 C atoms, 52 kpts

No. of Flop: $8.053 \cdot 10^{12}$

MPI	1	1	4	1	4	8
GPU	NO	NV	NV	ATI	ATI	NV + ATI
Time (s)	6020	300	160	347	197	109
Speedup	1	20.07	37.62	17.35	30.55	55.23
GFlop/s	1.34	26.84	50.33	23.2	40.87	73.87

Supercomputer in a workstation

Exact Exchange operator acceleration (2016)

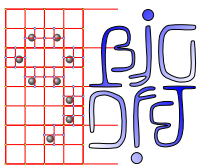

www.bigdft.org

PS is used for the Exact Exchange operator

Recently accelerated via GPU

conjunction with GPU-Direct (improves communications)

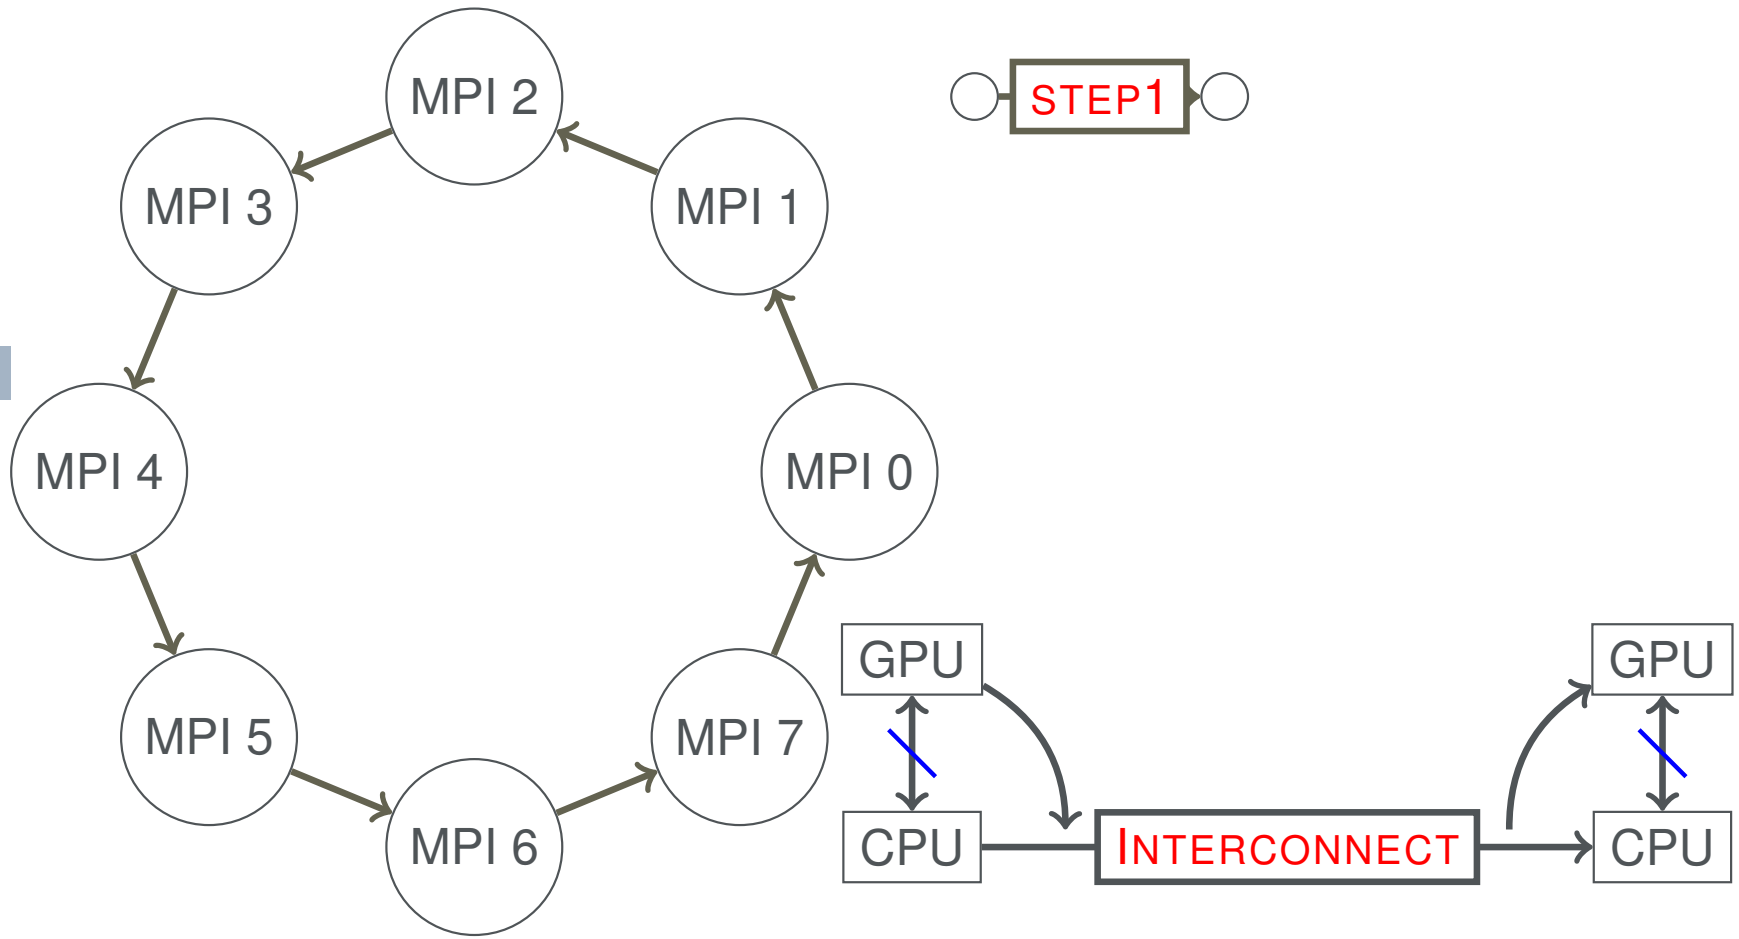

Exact Exchange operator acceleration (2016)

PS is used for the Exact Exchange operator

Recently accelerated via GPU
conjunction with GPU-Direct (improves communications)

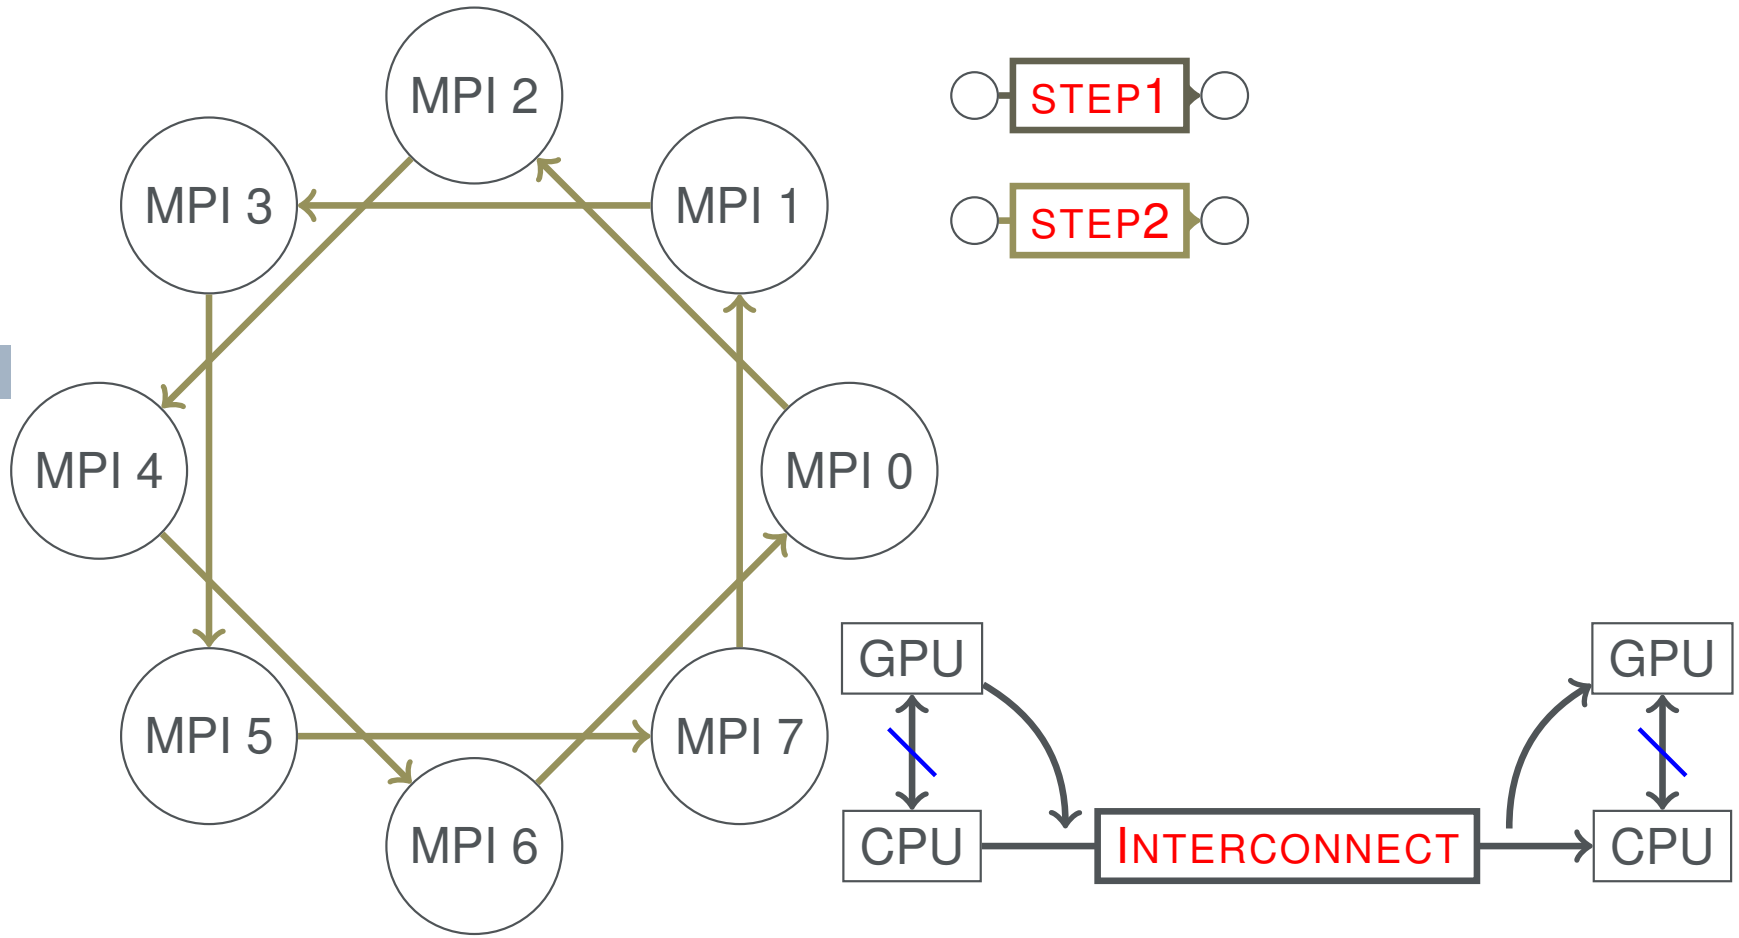

Exact Exchange operator acceleration (2016)

PS is used for the Exact Exchange operator

Recently accelerated via GPU
conjunction with GPU-Direct (improves communications)

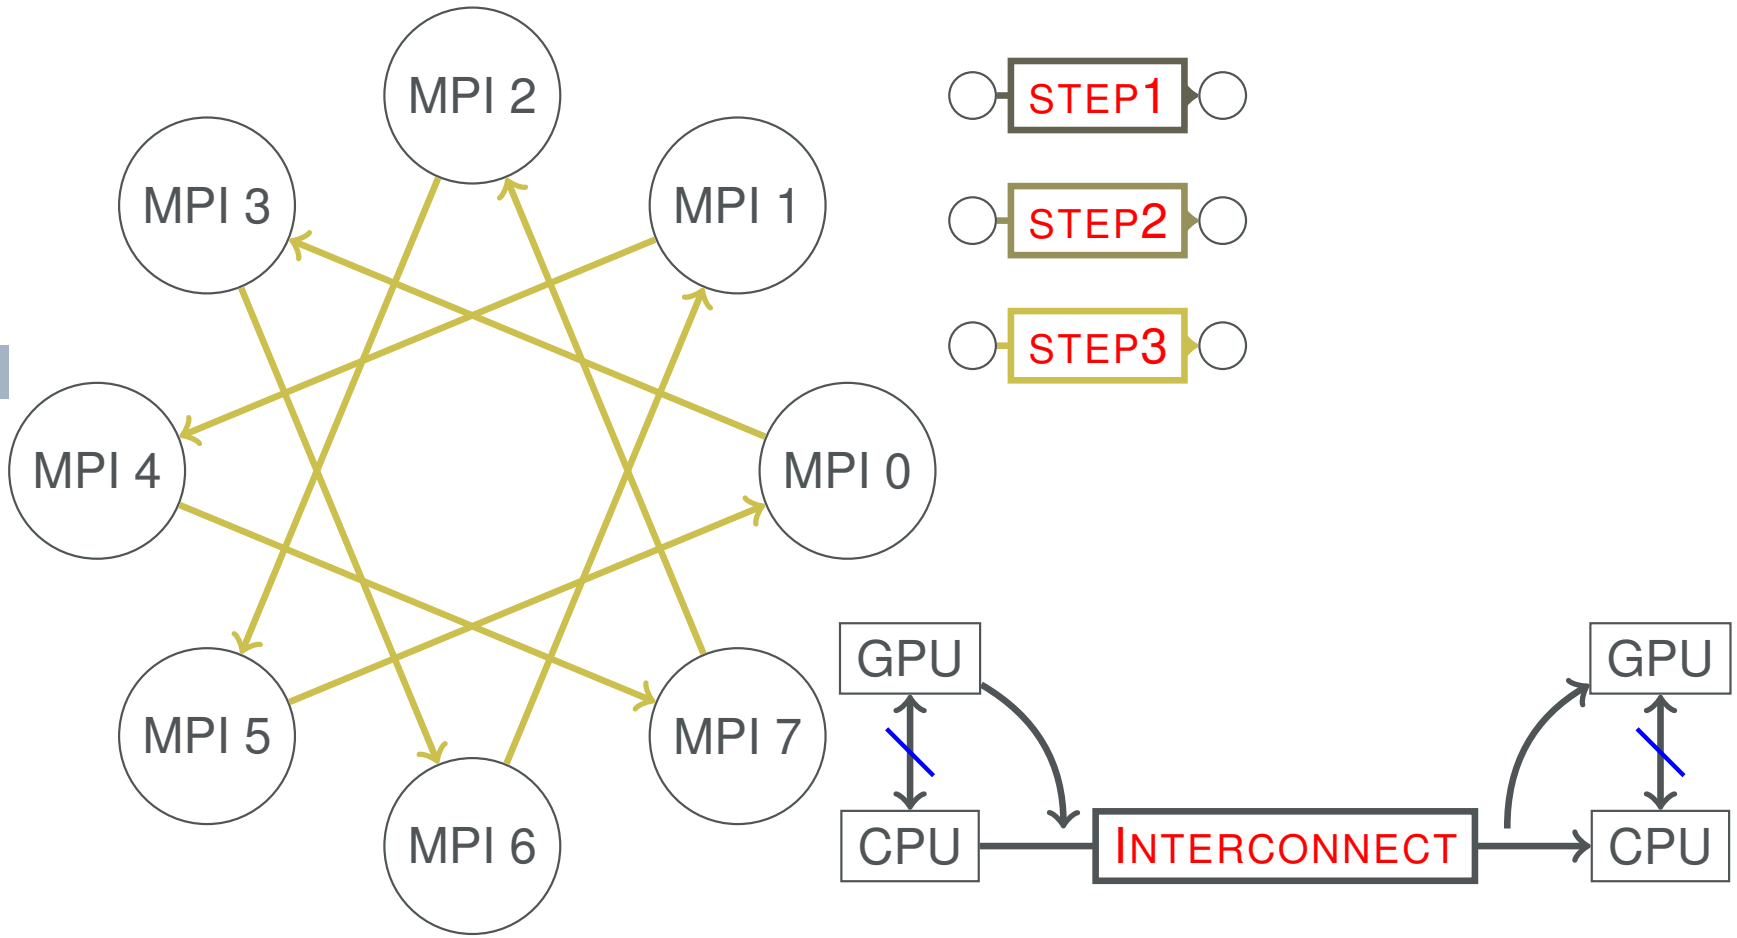

Exact Exchange operator acceleration (2016)

PS is used for the Exact Exchange operator

Recently accelerated via GPU
conjunction with GPU-Direct (improves communications)

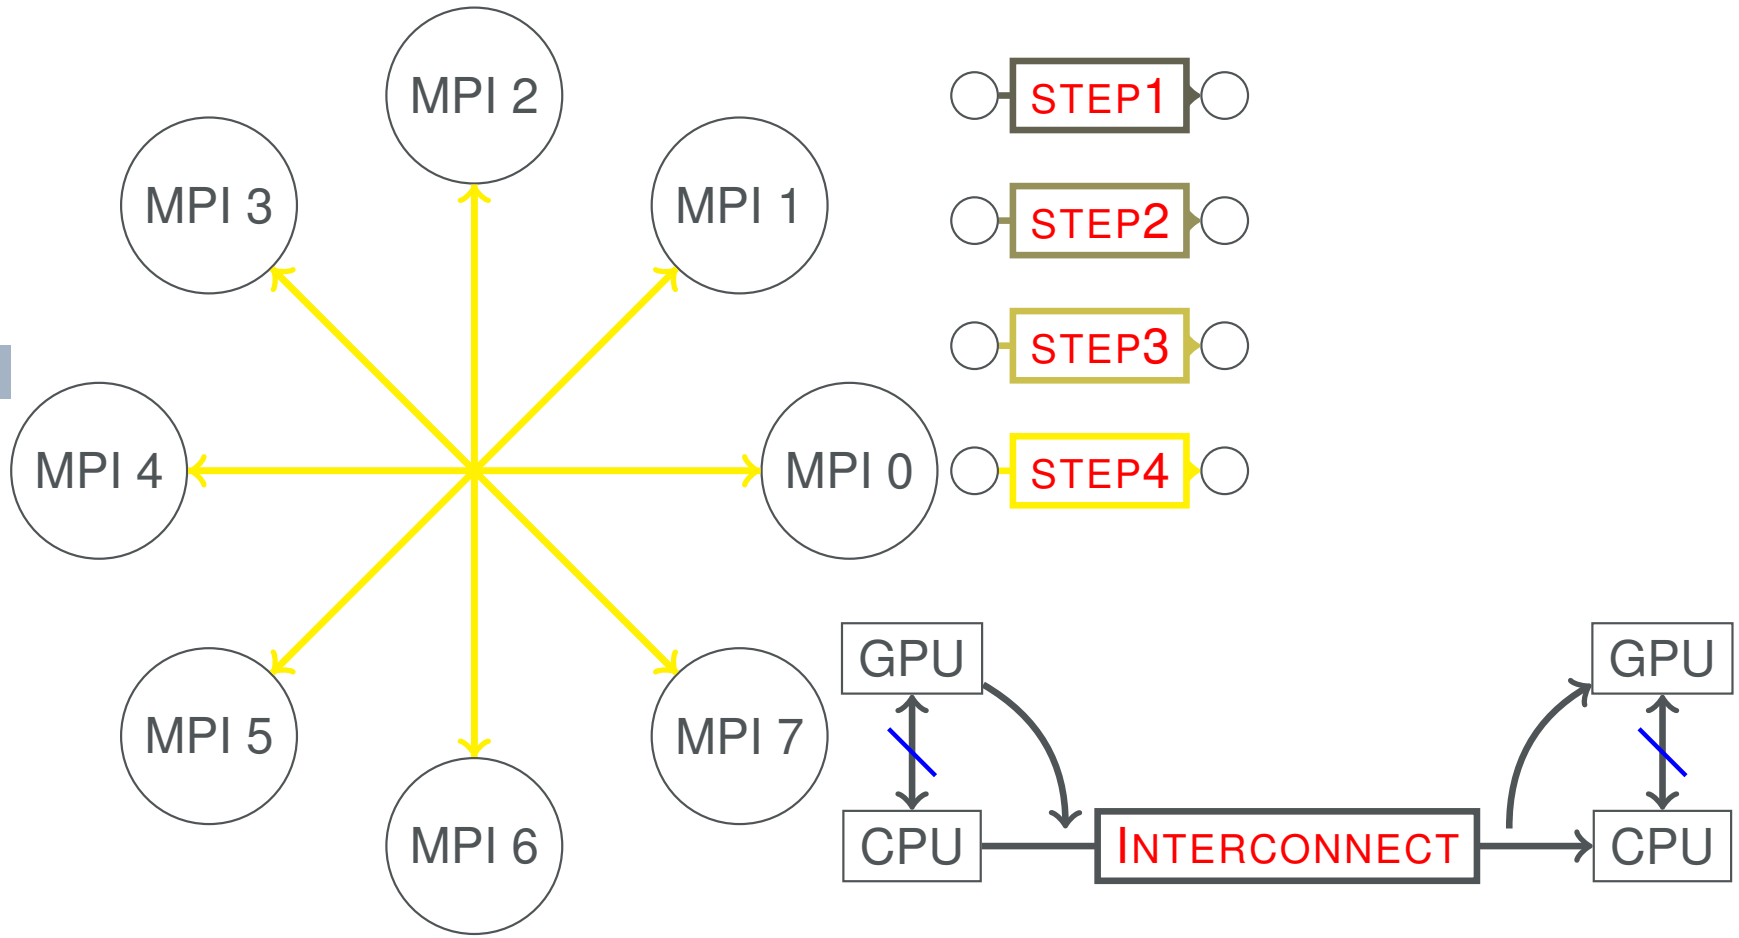

Hybrid Functionals for large systems (2016) γ=PBE0/PBE

www.bigdft.org

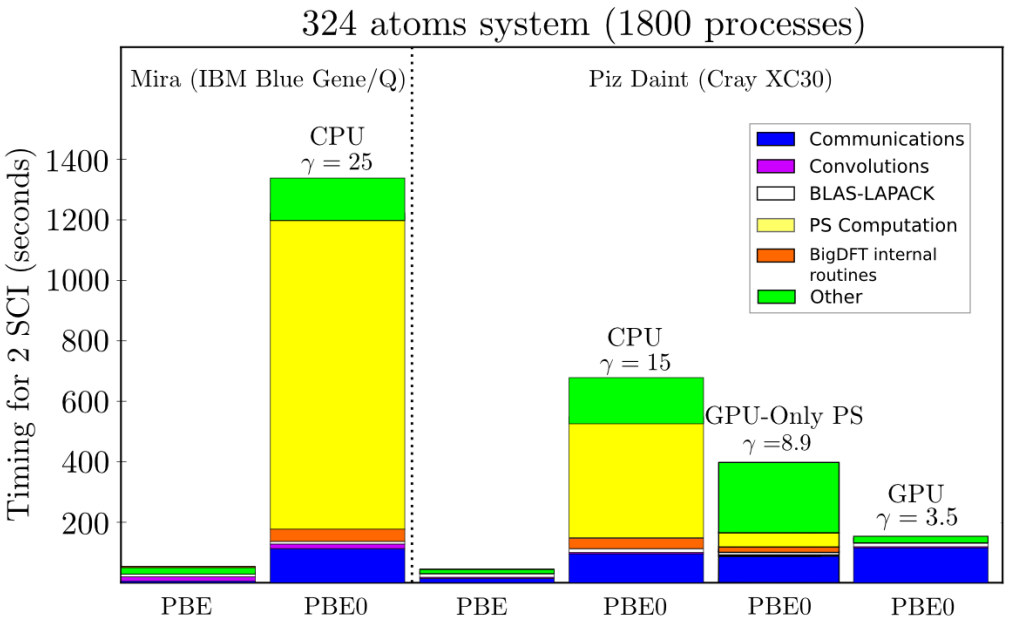

UO_2 systems:

Atoms	Orbitals
12	200
96	1432
324	5400
768	12800
1029	17150

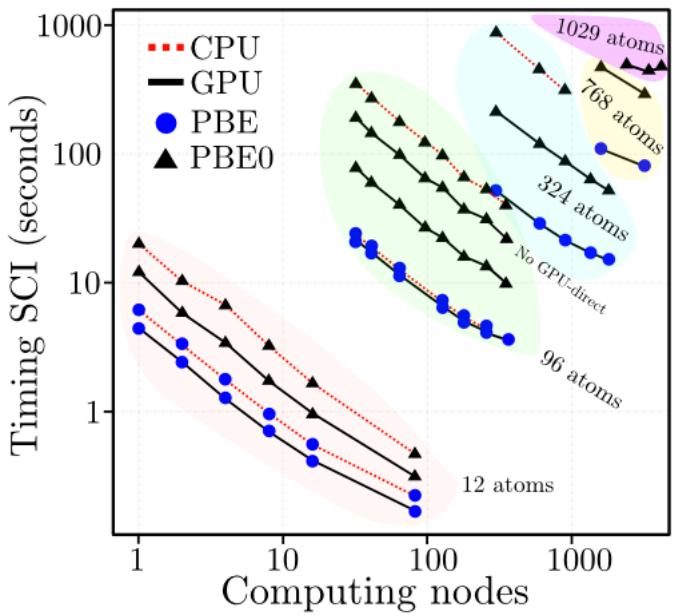

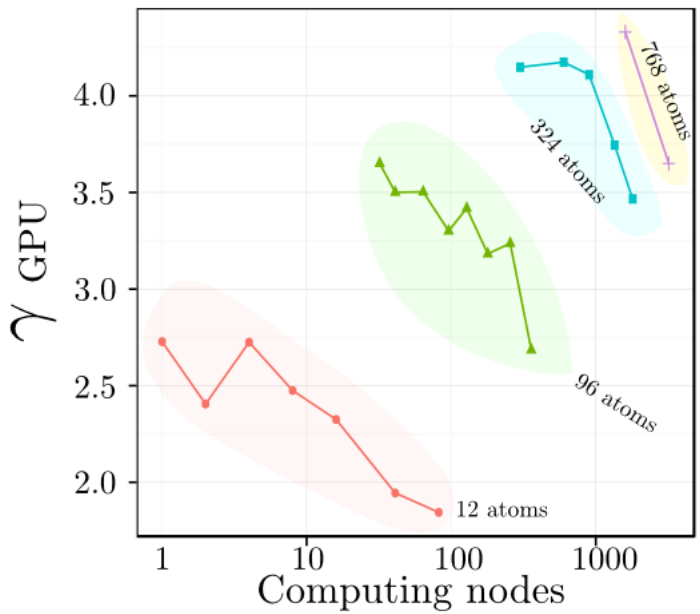

Scaling of BigDFT

www.bigdft.org

"Traditional" BigDFT code

We can reach systems containing up to a few hundred electrons thanks to wavelet properties and efficient parallelization: (MPI + OpenMP + GPU)

Varying the number of atoms N

DFT operations scale differently:

- $O(N \log N)$: Poisson solver
- $O(N^2)$: convolutions
- $O(N^3)$: linear algebra

and have different prefactors:

- $c_{O(N^3)} \ll c_{O(N^2)} \ll c_{O(N \log N)}$

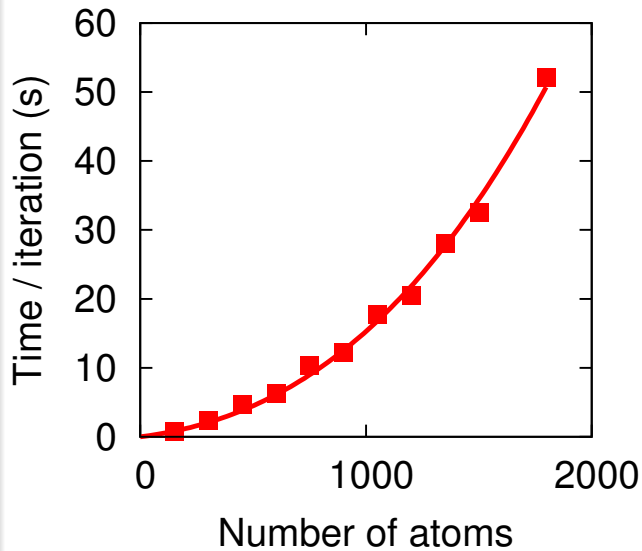

For bigger systems the $O(N^3)$ will dominate

☛ (first) motivation for a new approach

Local orbitals and linear scaling

KS orbitals

Linear combinations of support functions $\phi_\alpha(\mathbf{r})$:

$$\Psi_i(\mathbf{r}) = \sum_\alpha c_i^\alpha \phi_\alpha(\mathbf{r})$$

- localized around atoms
- expanded in wavelets
- optimized in-situ

Density Matrix

Defined via the kernel $K^{\alpha\beta}$ in the $\phi_\alpha(\mathbf{r})$ basis:

$$\rho(\mathbf{r},\mathbf{r}') = \sum_i f_i \Psi_i(\mathbf{r})\Psi_i(\mathbf{r}') = \sum_{\alpha,\beta} \phi_\alpha(\mathbf{r}) K^{\alpha\beta} \phi_\beta(\mathbf{r}')$$

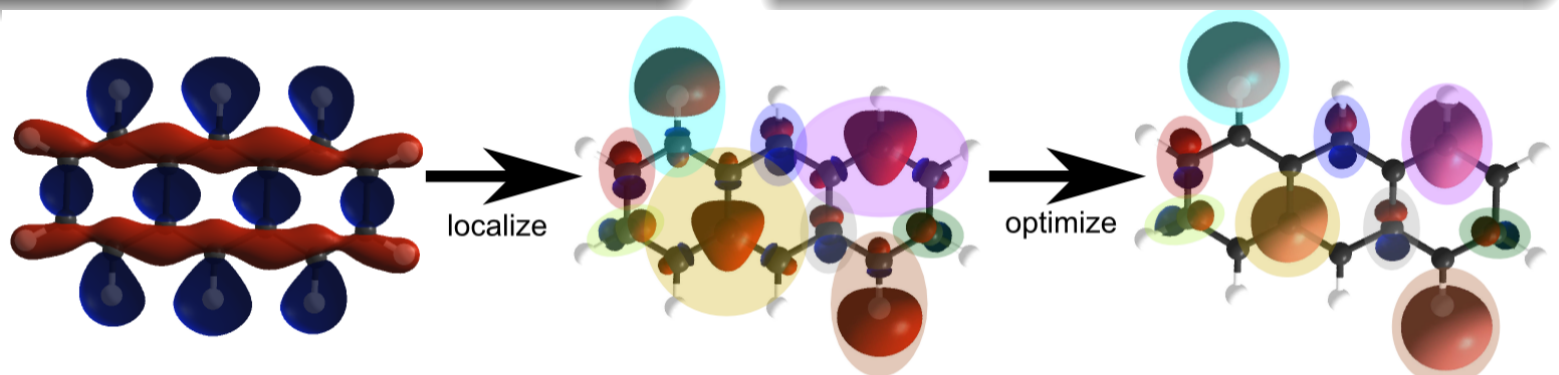

extended Kohn-Sham orbitals
cubic scaling, high accuracy

localized support functions (LCAO)
linear scaling, low accuracy

localized adaptive support functions
linear scaling, high accuracy

Localized form for the Density and the Hamiltonian

Localization $\rightarrow$ sparse matrices $\rightarrow O(N)$

Comparison with the cubic version

Energy and forces with accuracy of a systematic approach

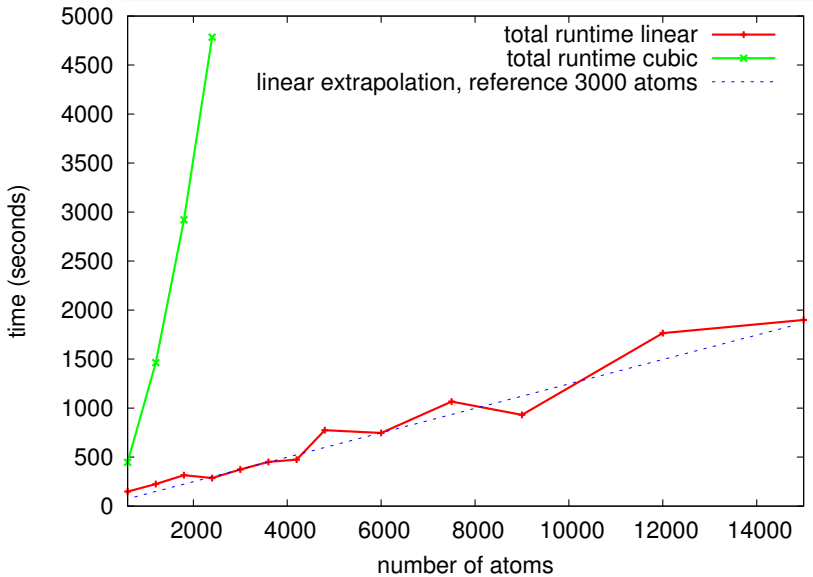

- 20 min for 18 000 atoms
- CPU Time and memory $\propto$ number of atoms
- Precise DFT computing for thousands atoms at Institute-Scale ($10^2 - 10^3$ CPU cores)

Different levels of precision ☛ cutoff radii

Without fine-tuning converges to absolute energy differences of the order of 10 meV/atom, and almost exact forces.

High flexibility, like the cubic code

- Charged systems, various BC (free, surfaces, periodic)
- System sizes: 100 - 30K atoms $\rightsquigarrow$ 100 k Basis functions

Recent improvements

Algorithm is robust and reliable on a variety of systems

Accurate and efficient linear scaling DFT calculations with universal applicability

S. Mohr, L. E. Ratcliff, L. Genovese, D. Caliste, P. Boulanger, S. Goedecker and T. Deutsch

Phys. Chem. Chem. Phys., 2015, **17**, 47, 31360-31370.

DOI: 10.1039/c5cp00437c

Included in the **Real-space numerical grid methods in quantum chemistry** themed issue of *PCCP*

Guest-edited by Luca Frediani (The Arctic University of Norway) and Dage Sundholm (University of Helsinki)

Why Large Scale DFT?

Present-day situation

www.bigdft.org

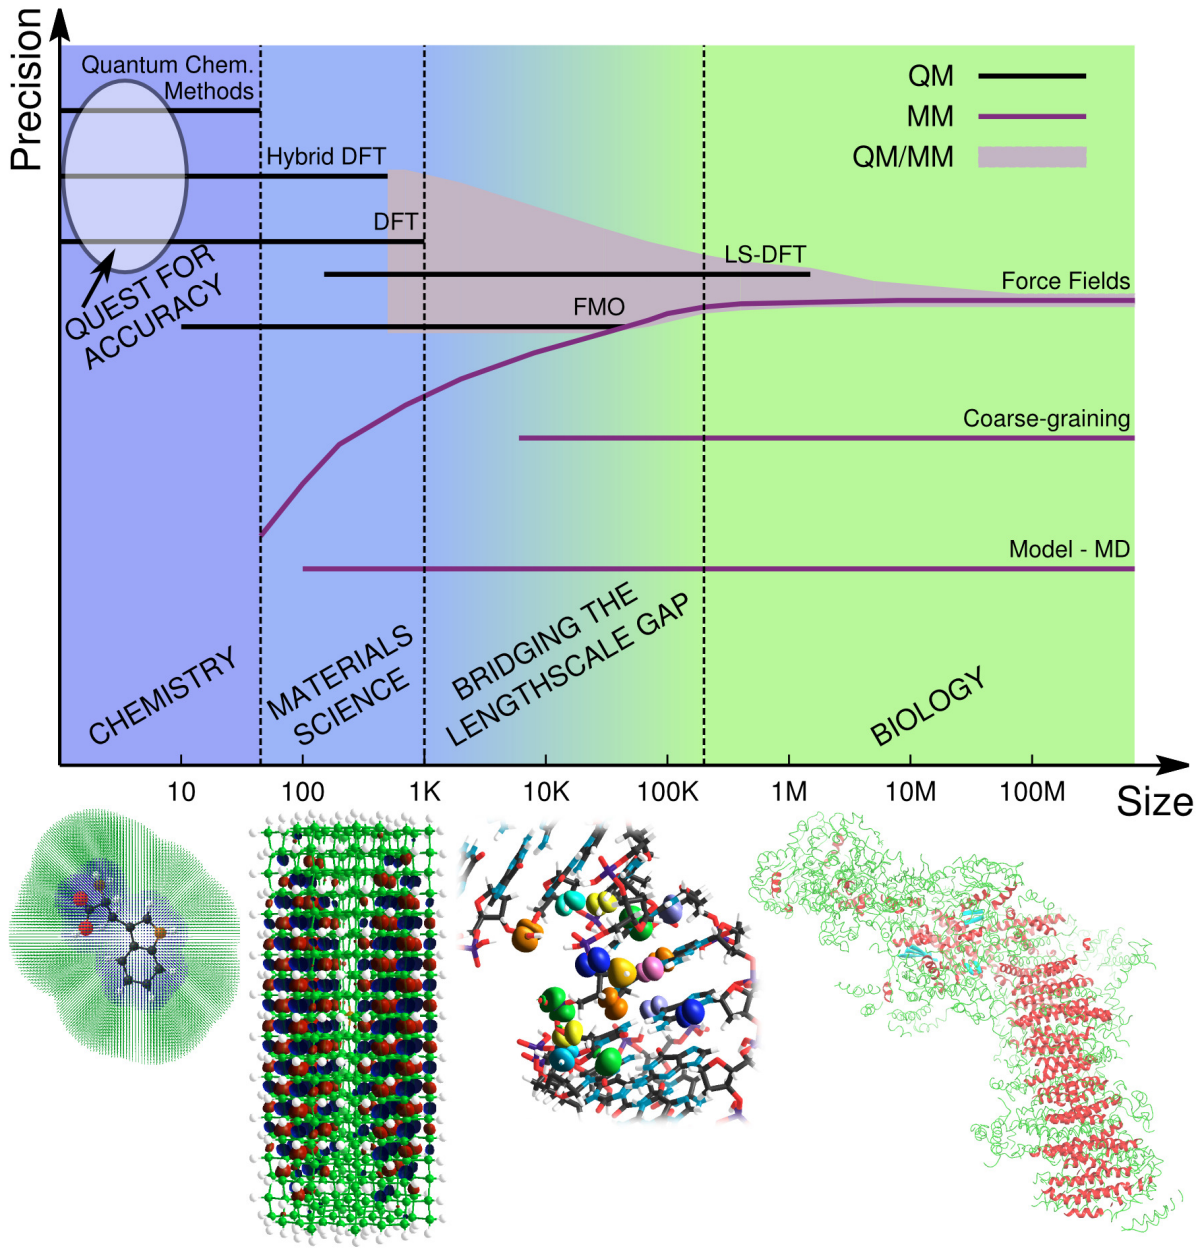

Summary and future directions

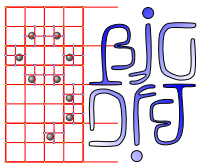

www.bigdft.org

Linear-Scaling DFT calculations based on wavelets

- Robust convergence, high accuracy and flexibility (BC)
- Reduction in degrees of freedom $\rightarrow$ large systems via moderate sized machines ($\sim$ TFlop/s) Lab-scale
- Optimal mapping between KS DoF and atoms
- Different level of descriptions (controlling the precision) QM $\supset$ Fragments $\supset$ Atomic charges
- Opens up new possibilities

Challenges and future directions

- Explore interplay environment $\leftrightarrow$ electronic excitations (CDFT, QM/MM, statistics. . .)
- Provide high quality back end for extraction of atomic multipoles from QM calculations
- Towards a control of the level of theory (QM/QM)

Acknowledgments

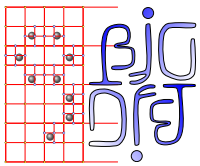

www.bigdft.org

Order N methods Stephan Mohr, Laura E. Ratcliff, Paul Boulanger

Group of Stefan Goedecker B. Schaefer, A. Ghazemi, S. Saha, G. Fisicaro, A. Degomme, J. Flores-Livas(Basel University)

Link with ABINIT and python bindings Damien Caliste (CEA)

Resonant states M. Morinière, I. Duchemin (CEA), B. Nectoux, E. Cancés (CERMICS)

Optimized convolutions B. Videau, J.-F. Méhaut (LIG, computer scientists, Grenoble)