Large-Scale Electronic Structure Theory in FHI-aims and ELSI

Volker Blum

MEMS Department, Duke University, Durham, NC

http://aims.pratt.duke.edu

ACI-1450280

DMR-1729297

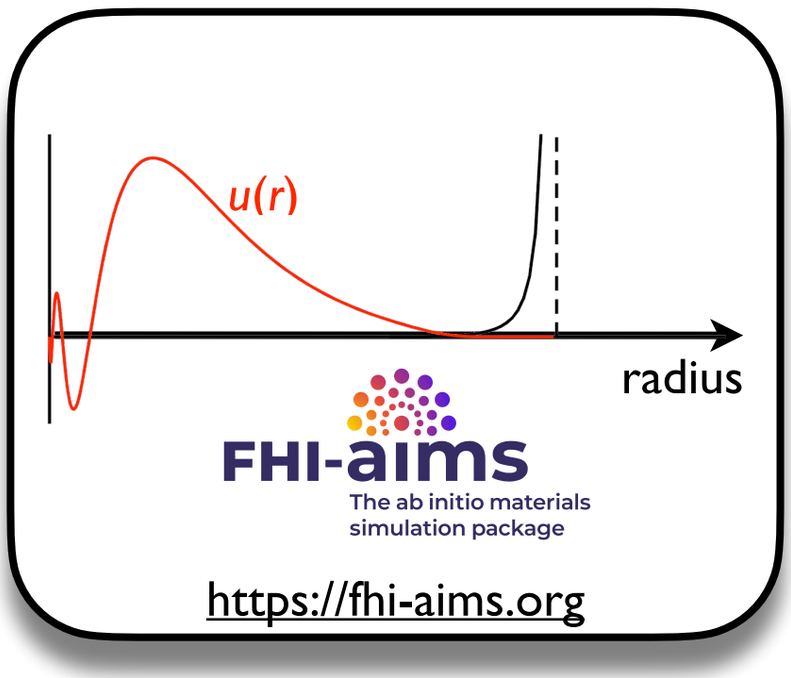

ELSI ELPA

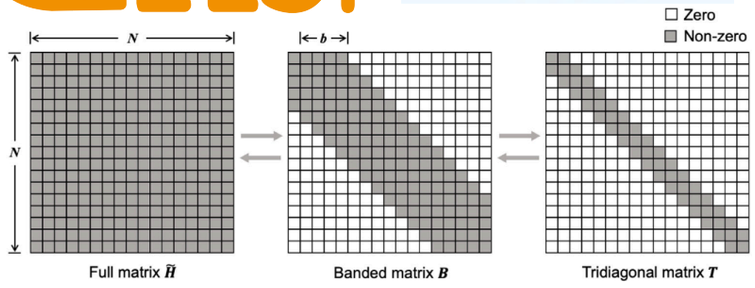

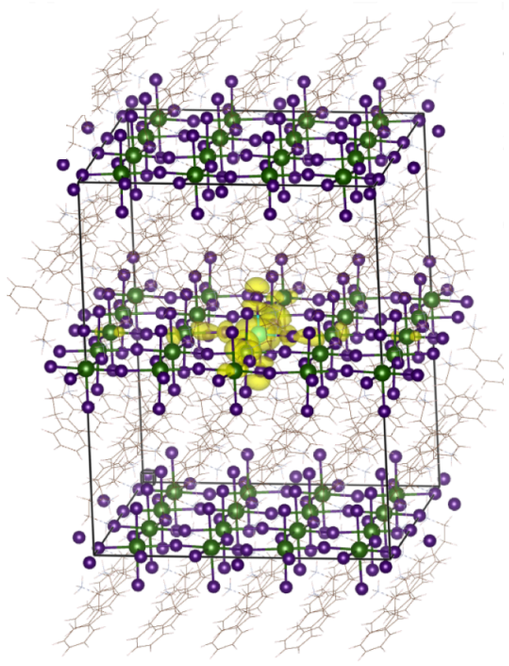

Large-Scale Electronic Structure Theory in FHI-aims and ELSI

Volker Blum

MEMS Department, Duke University, Durham, NC

http://aims.pratt.duke.edu

ACI-1450280

DMR-1729297

CHOISE
Center for Hybrid Organic Inorganic Semiconductors for Energy

U.S. DEPARTMENT OF ENERGY | Office of Science

Numeric Atom-Centered Basis Sets

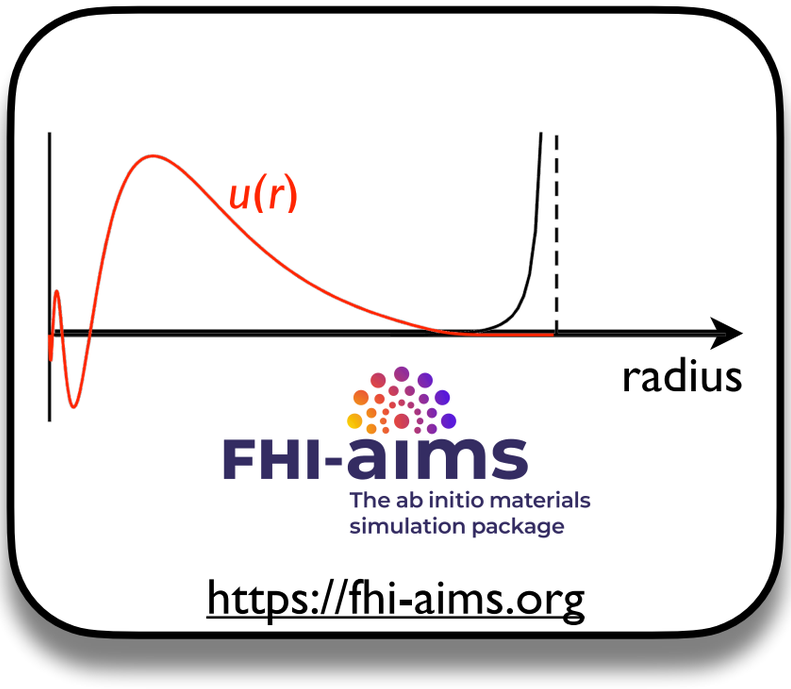

Eigenvalue and Density Matrix Solutions

ELSI ELPA

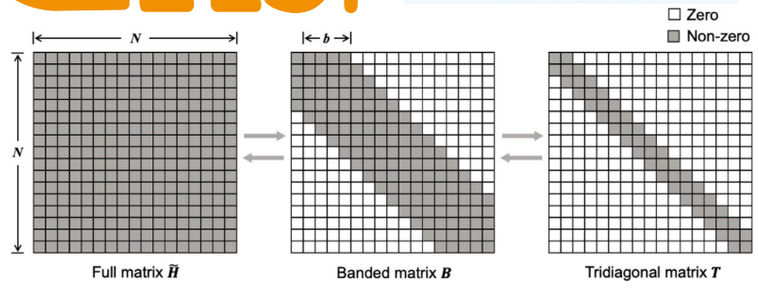

All-Electron Hybrid DFT for Large Systems

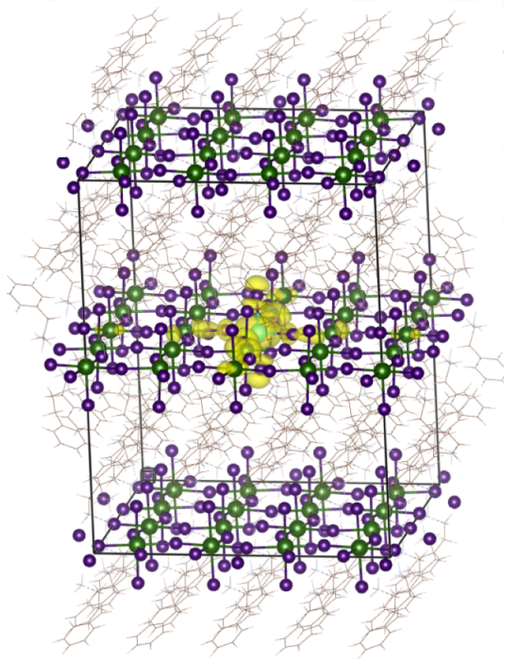

Group (Present and Past)

Ab Initio Materials Simulations Group (Duke Univ.)

Yi Yao

Ruyi Song

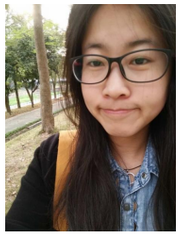

Xixi Qin

Gabrielle Koknat

Tianlin Wang

Uthpala Herath

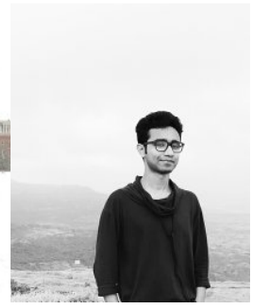

Rayan Chakraborty

Naidel Caturello

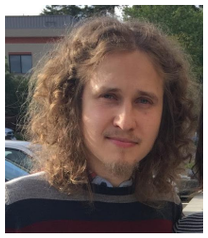

Raul Laasner

Rundong Zhao (now Beihang)

William Huhn (now Intel)

Victor Wen-zhe Yu

Chi Liu

Jack Morgenstein
Jannik Eisenlohr
Harrison York
Reyna Vrbensky
Douglas Heine
Shreyas Karadi
Daniel Duarte
Wentao Zhang

Duke
PRATT SCHOOL of ENGINEERING

(Duke University, Durham, NC, USA)

Our (Long-Term) Approach: FHI-aims

FHI-aims
The ab initio materials simulation package

https://fhi-aims.org

Vision:
Quantum mechanics based simulations of real, complex molecules, materials and their properties without a priori precision and accuracy limitations.

→ Algorithmic choices and priorities:

- All-electron
- Non-periodic and periodic systems on equal footing
- Scalability to large systems (thousands of atoms) without precision limitations
- Seamless scalability from laptop to massively parallel and/or new HPC architectures
- Density functional theory and correlated methods (RPA, GW, CC via CC4S…)

Critical choice

- Numeric atom-centered basis functions (accurate representation of occupied orbitals and densities)

Stewardship: MS1P e.V. (non-profit organization, purpose: basic science)
Advisory board, very active community

FHI-aims Code

https://fhi-aims.org

FHI-aims
The ab initio materials simulation package

FHI-aims team and collaborators: Matthias Scheffler (Berlin), Sebastian Kokott (Berlin), Mariana Rossi (Hamburg), Xinguo Ren (Beijing), Karsten Reuter (Berlin), over 150 individuals with contributions to the project. Development in Berlin, Duke, UNC, Beijing, Helsinki, Warwick, Cardiff, Dresden, Dalhousie, etc.

Coordinators: Blum, Kokott, Rossi, Scheffler

Alaa Akkoush, Heiko Appel, Victor Atalla, Kurt Baarmann, Carsten Baldauf, Alexej Bagrets, Jörg Behler, Daniel Berger, Josh Berryman, Sheng Bi, Benedikt Biedermann, Bjoern Bieniek, Volker Blum, Saeed Bohloul, Tiago Botari, Connor Box, Danilo Brambila, Gabriel Bramley, Daniel Bultrini, Christian Carbogno, Fabio Caruso, Marco Casadei, Michele Ceriotti, Wael Chibani, Sucismita Chutia, Francisco Antonio Delesma, Fabio Della Sala, Maria Dragoumi, Andreas Dolfen, Marc Dvorak, Simon Erker, Ferdinand Evers, Eduardo Fabiano, Matt Farrow, Nicola Ferri, Karen Fidanyan, Jakob Filser, Lukas Gallandi, Ralf Gehrke, Luca Ghiringhelli, Mark Glass, Vivekanand Gobre, Dorothea Golze, Matthias Gramzow, Patrick Guetlein, Stefan Gutzeit, Volker Haigis, Felix Hanke, Paula Havu, Ville Havu, Joscha Hekele, Olle Hellman, Jan Hermann, Oliver Hofmann, Johannes Hoja, Xiaojuan Hu, William Huhn, Lukas Hörmann, Arvid Ihrig, Timo Jacob, Adam Jackson, Svenja Janke, Ran Jia, Rainer Johanni, Erin Johnson, Werner Jürgens, Matthias Kahk, Yosuke Kanai, Levi Keller, Matthias Kick, Woo Youn Kim, Jan Kloppenburg, Alexander Knoll, Florian Knoop, Franz Knuth, Simone Koecher, Gabrielle Koknat, Sebastian Kokott, Raul Laasner, Lucas Lang, Bjoern Lange, Marvin Lechner, Susi Lehtola, Maja-Olivia Lenz, Sergey Levchenko, Alan Lewis, Jiachen Li, Xinzheng Li, Kailai Lin, Xinyi Lin, Konstantin Lion, Yair Litman, Chi Liu, Andrew Logsdail, Andreas Marek, Thomas Markovich, Reinhard Maurer, Florian Merz, Joerg Meyer, Wenhui Mi, Evgeny Moerman, Christoph Muschielok, Mohammad Nakhaee, Lydia Nemec, Norbert Nemec, Kane O'Donnell, Harald Oberhofer, Berk Onat, Alberto Otero de la Rosa, Ramon L. Panades-Barrueta, Eszter Pos, Alastair Price, Thomas Purcell, Nathaniel Raimbault, Karsten Rasim, Xinguo Ren, Karsten Reuter, Norina Richter, Stefan Ringe, Patrick Rinke, Herzain Rivera, Matti Ropo, Mariana Rossi, Tuomas Rossi, Adrienn Ruzsinszky, Nikita Rybin, Georg Michelitsch, Andrea Sanfilippo, Matthias Scheffler, Markus Schneider, Christoph Schober, Franziska Schubert, Honghui Shang, Tonghao Shen, Markus Sinstein, Justin Clifford Smith, Ari-Pekka Soikkeli, Ruyi Song, Aloysius Soon, Pavel Stishenko, Muhammad Tahir, Alexandre Tkatchenko, Thomas Theis, Alvaro Vazquez Mayagoitia, Suzy Wallace, Tianlin Wang, Yanyong Wang, Jürgen Wieferink, Scott Woodley, Jianhang Xu,Yong Xu,Yi Yao, Mina Yoon,Ted Yu, Victor Yu, Zhenkun Yuan, Marios Zacharias, Guo-Xu Zhang, Igor Ying Zhang, Wenxing Zhang, Rundong Zhao, Ruiyi Zhou, Yuanyuan Zhou, Tong Zhu

MS1P
Molecular Simulations from First Principles e.V.

... and, with absolute certainty, more! Thank you!

ELSI: Connecting Electronic Structure Codes and Solvers

Yu et al., "ELSI – An Open Infrastructure for Electronic Structure Solvers"
Comp. Phys. Commun. 256, 107459 (2020).

Yu et al., Comput. Phys. Commun. 2018
http://elsi-interchange.org
http://git.elsi-interchange.org/elsi-devel/elsi-interface

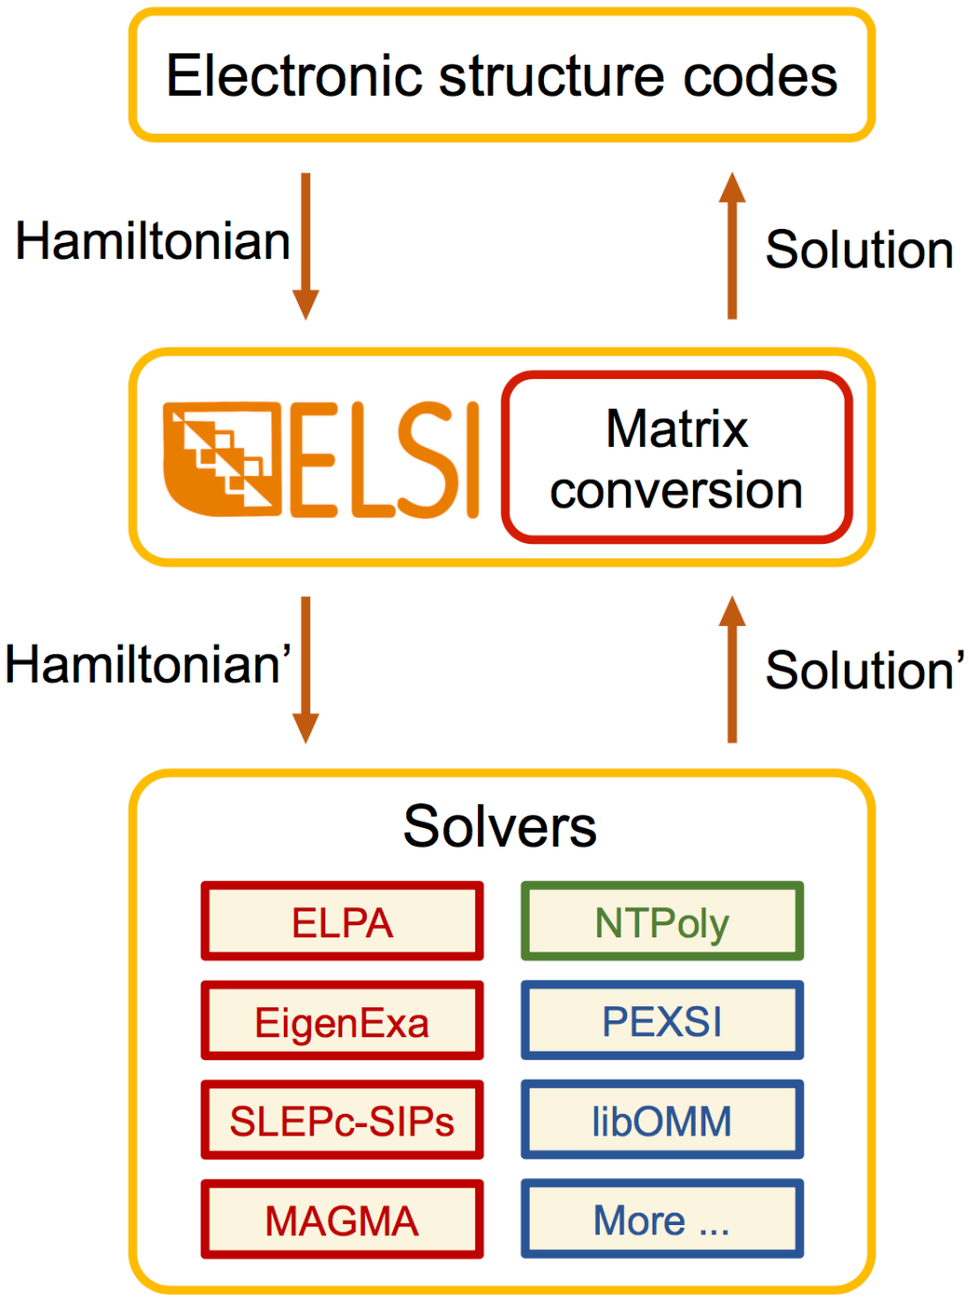

- Unified interface connecting KS-DFT codes and solvers
- Fast, automatic matrix format conversion and redistribution
- Recommendation of optimal solver based on benchmarks

NSF ACI-1450280

ELSI: Connecting Electronic Structure Codes and Solvers

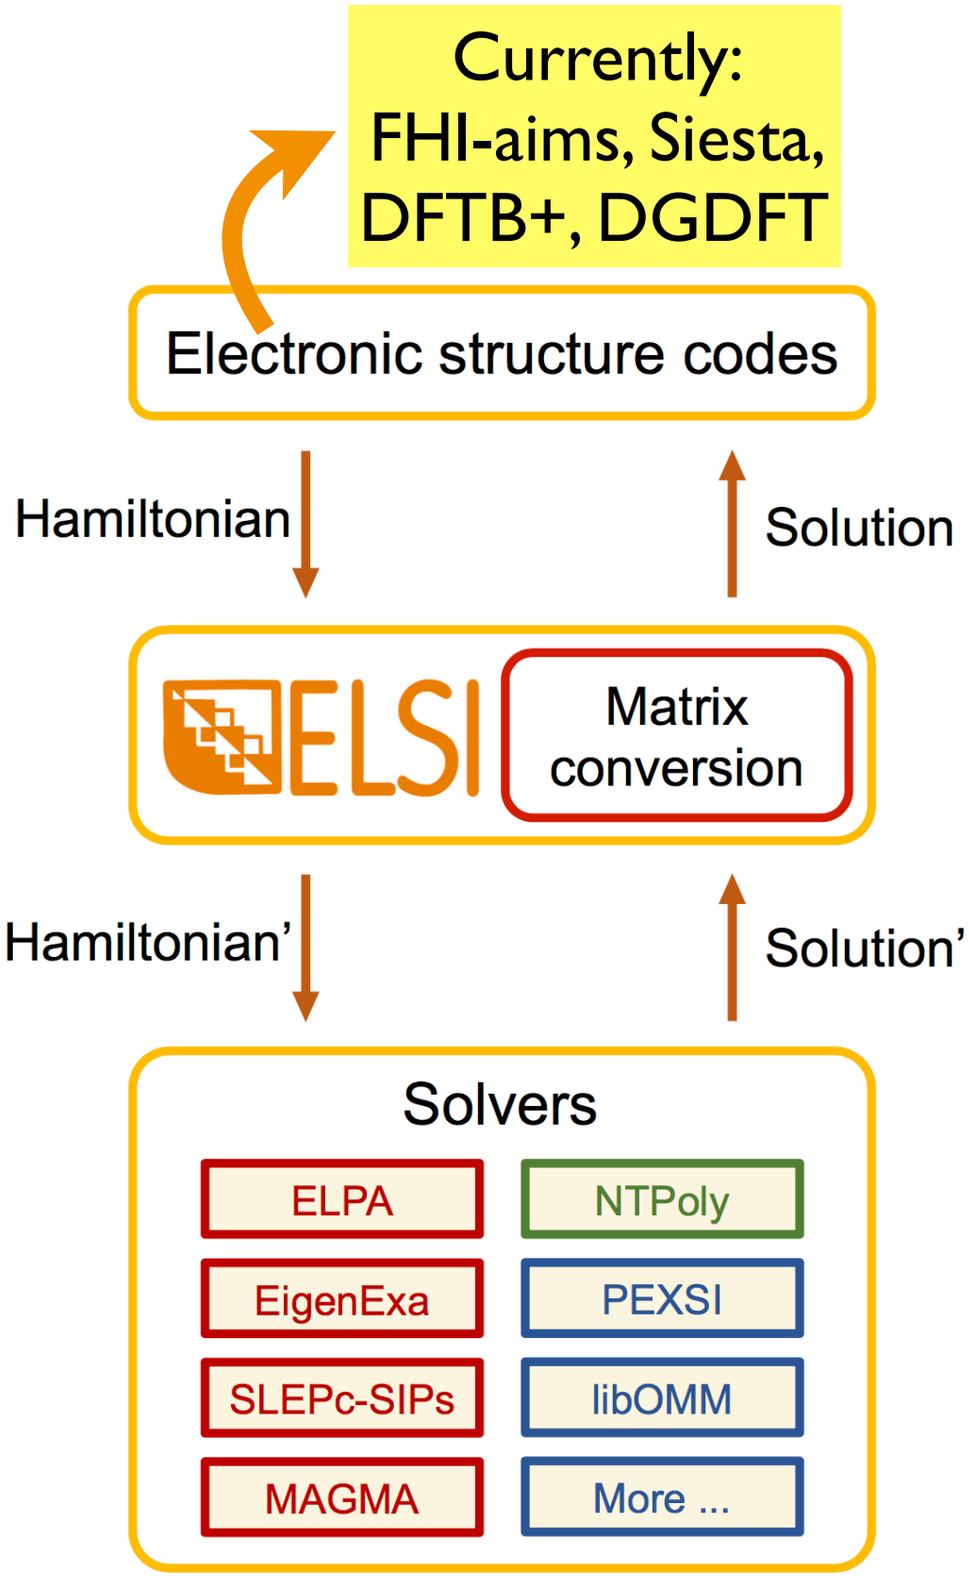

Yu et al., "ELSI – An Open Infrastructure for Electronic Structure Solvers"
Comp. Phys. Commun. 256, 107459 (2020).

Yu et al., Comput. Phys. Commun. 2018
http://elsi-interchange.org
http://git.elsi-interchange.org/elsi-devel/elsi-interface

- Unified interface connecting KS-DFT codes and solvers
- Fast, automatic matrix format conversion and redistribution
- Recommendation of optimal solver based on benchmarks

NSF ACI-1450280

Open-Source Graphical Interface for Materials Science

https://gims.ms1p.org

Browser-based graphical interface - no installation, open to more codes.

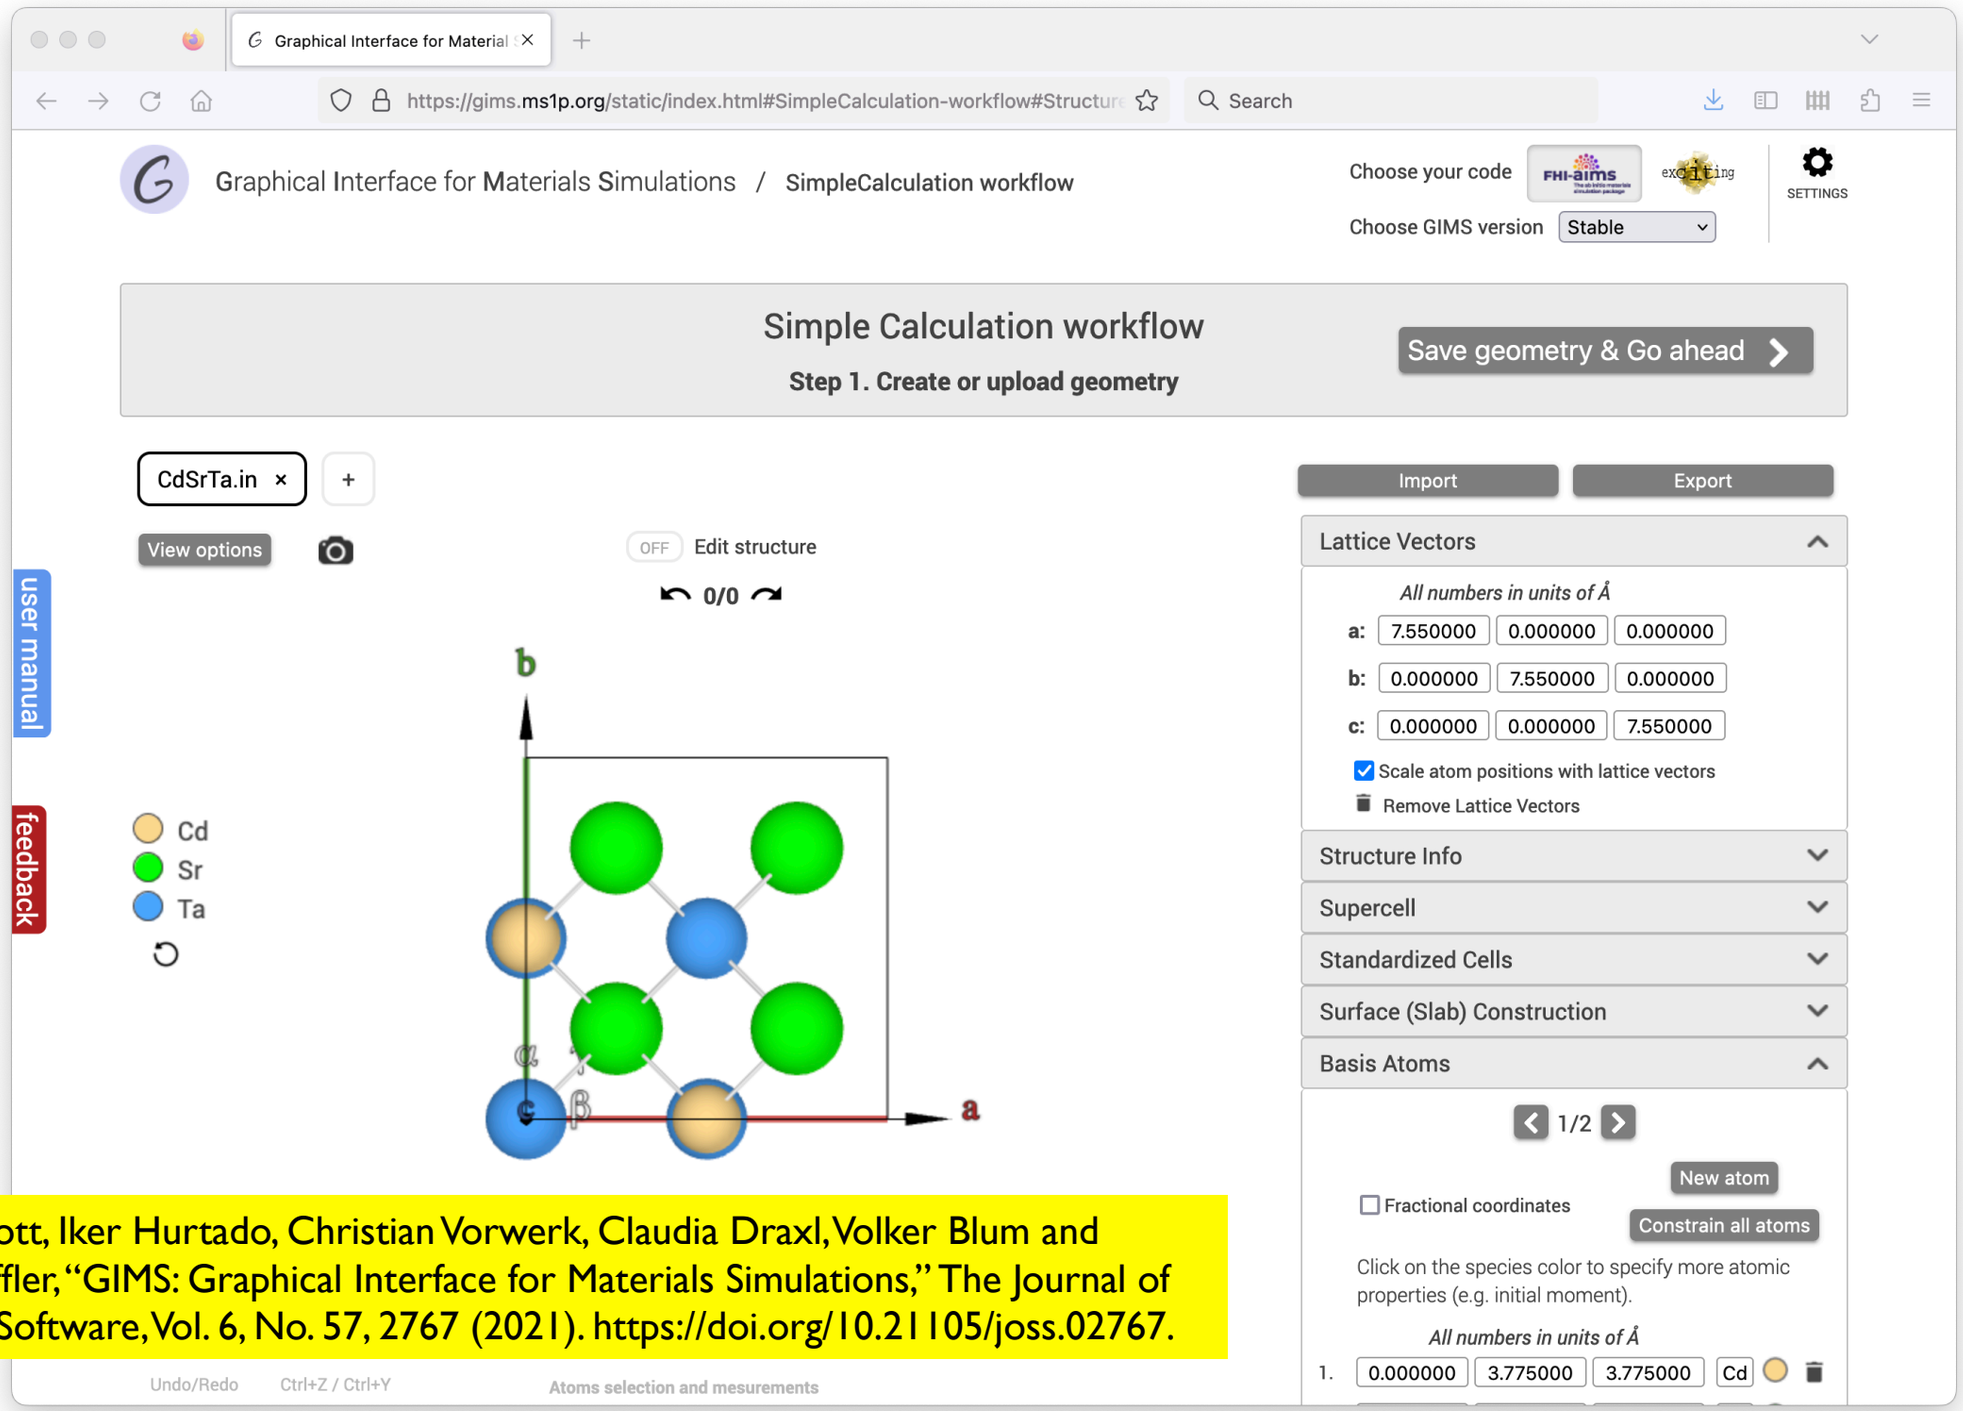

Sebastian Kokott, Iker Hurtado, Christian Vorwerk, Claudia Draxl, Volker Blum and Matthias Scheffler, "GIMS: Graphical Interface for Materials Simulations," The Journal of Open Source Software, Vol. 6, No. 57, 2767 (2021). https://doi.org/10.21105/joss.02767.

MatD³: Individualized Materials Database Software

Raul Laasner, Xiaochen Du, Aditya Tanikanti, Connor Clayton, Marco Govoni, Giulia Galli, Volker Blum

"MatD³: A Database and Online Presentation Package for Research Data Supporting Materials Discovery, Design, and Dissemination,"
Journal of Open-Source Software, Vol. 5, No. 45, 1945 (2020). DOI: 10.21105/joss.01945

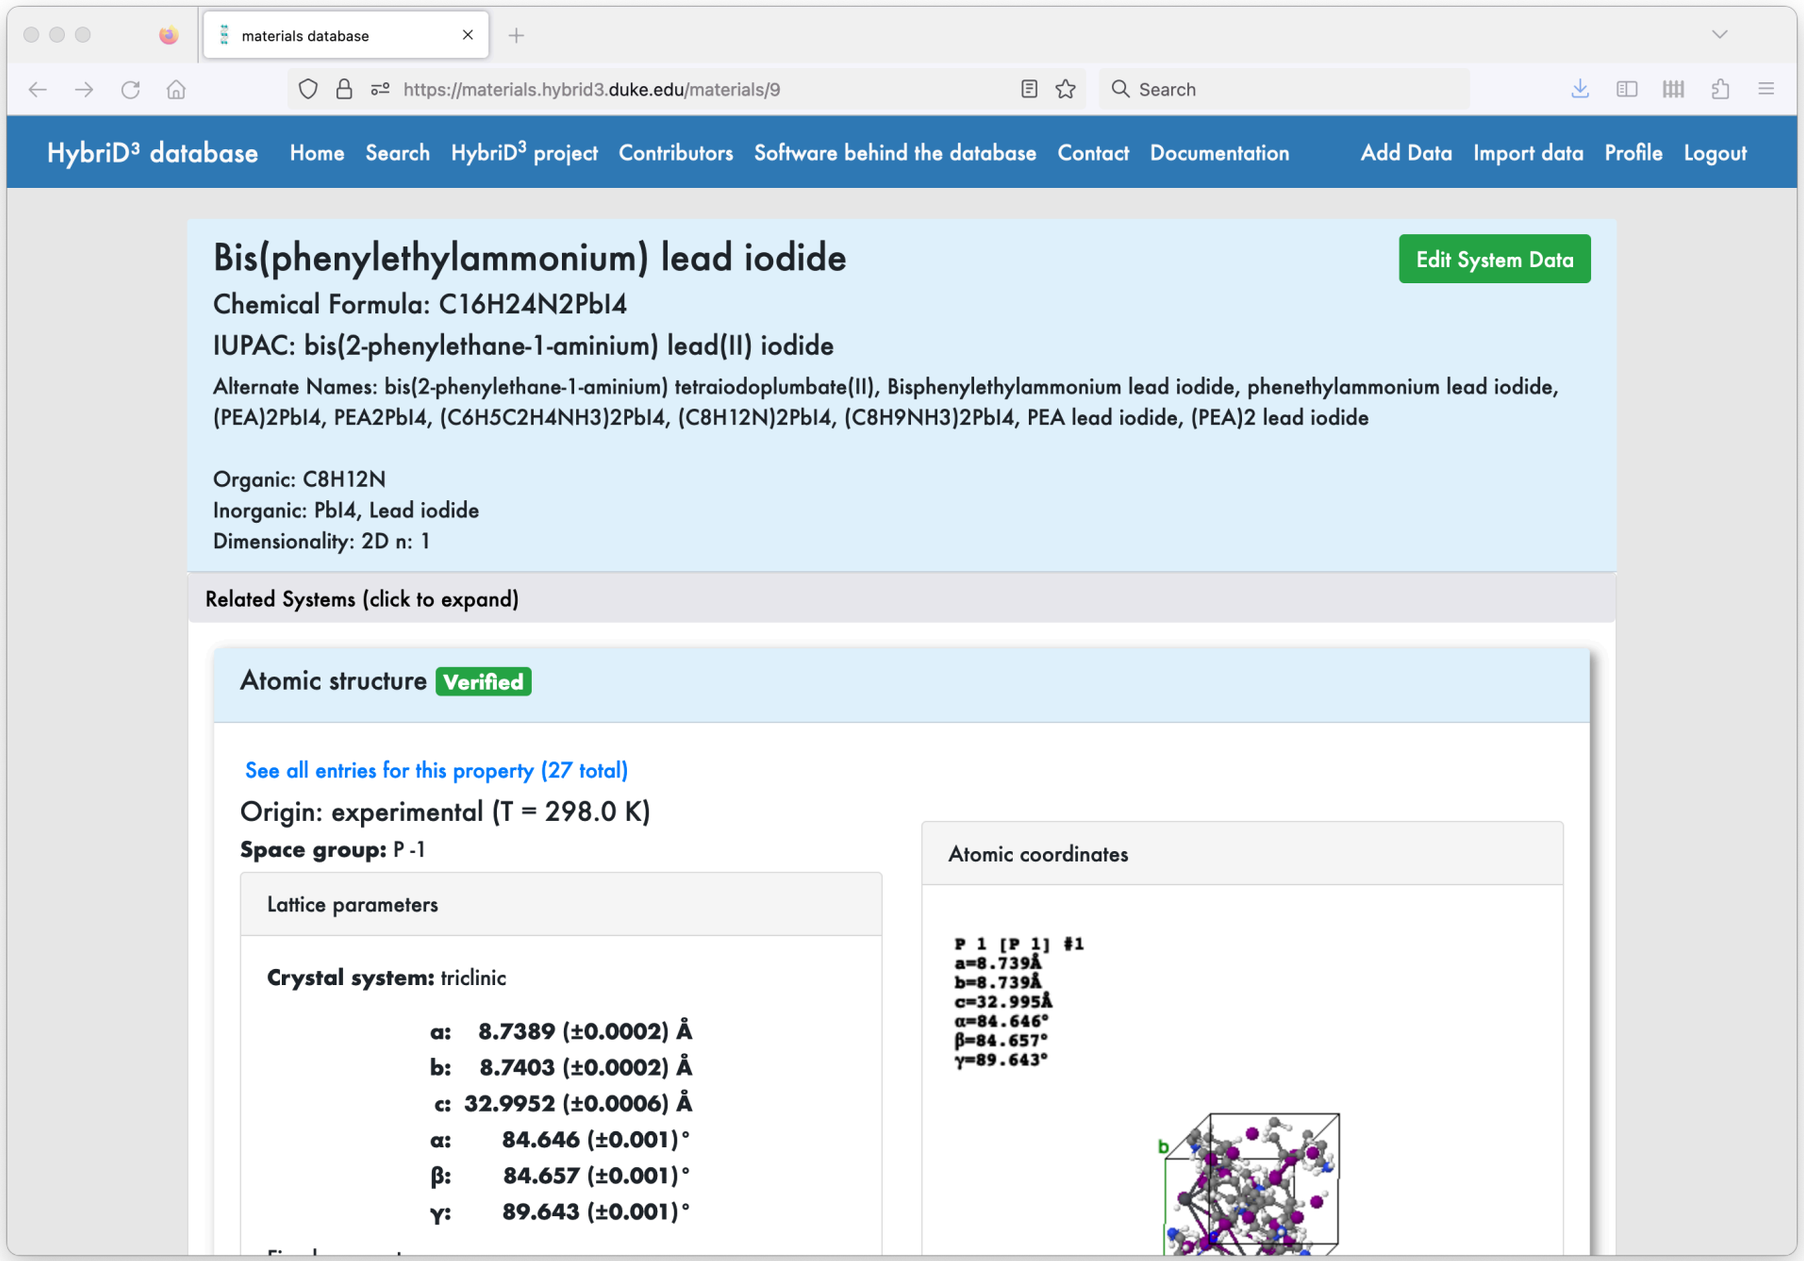

NSF DMR 1728921, 1729297, 1729383

CECAM Electronic Structure Library

Electronic Structure Theory underpins a broad swath of science,
with "engines" that live for decades: Much of our infrastructure is shared.

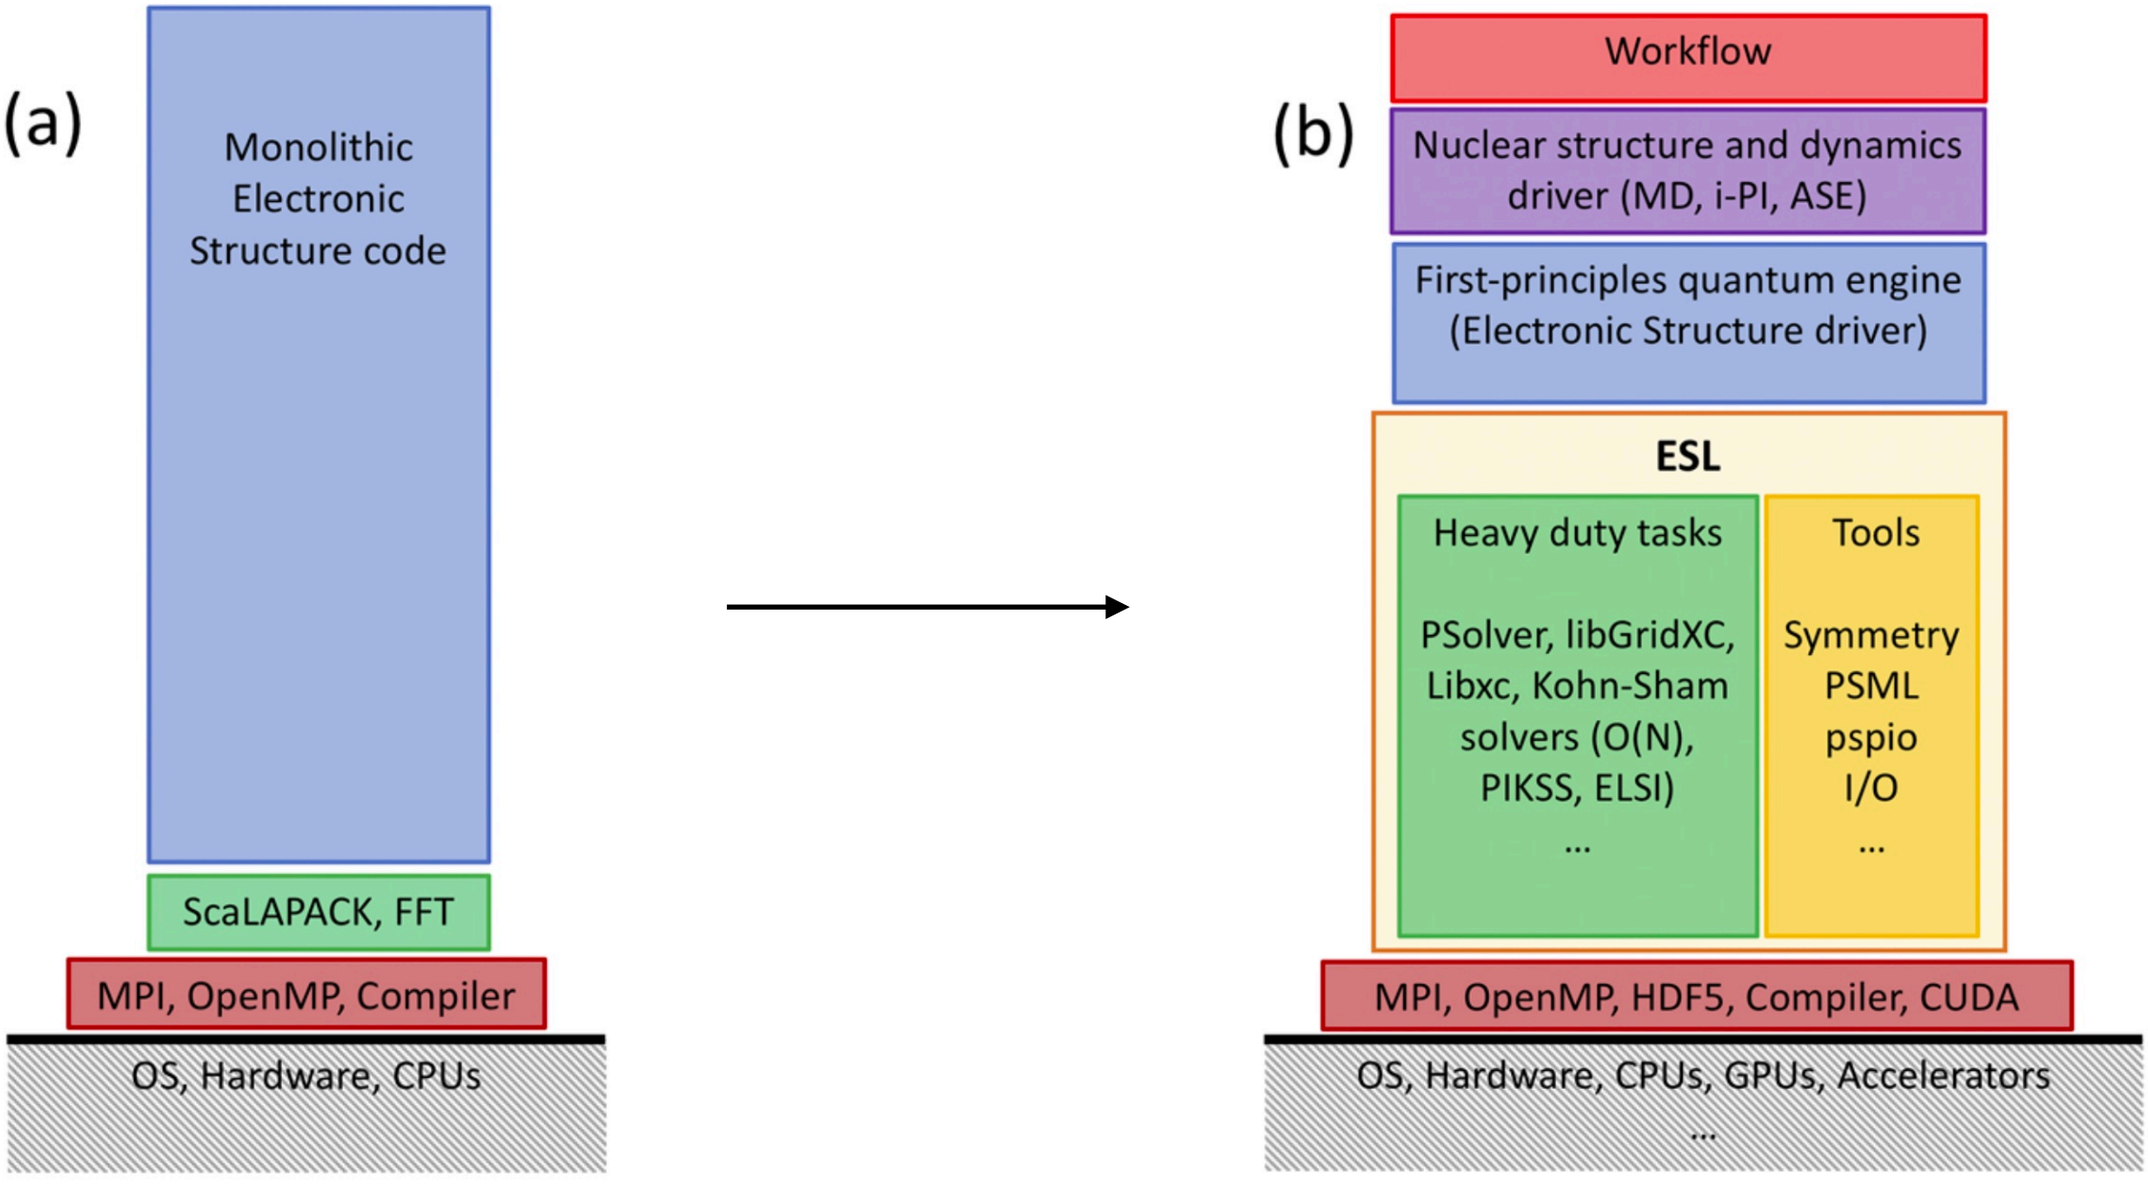

Oliveira et al., J. Chem. Phys. 153, 024117 (2020).

Next ESL Coding Workshop

Electronic Structure Software Development: Advancing the Modular Paradigm

https://www.cecam.org/workshop-details/1194

February 19 - March 1, 2024

CECAM-HQ, Lausanne, Switzerland

Centre Européen de Calcul Atomique et Moléculaire

Save the Date!

2 days of strategy / high-level discussions, 10 days of actual coding in Lausanne.

Standard Electronic Structure Theory

$$\left[-\frac{\nabla^2}{2} + v_{\text{ext}}(\boldsymbol{r}) + v_{\text{es}}(\boldsymbol{r}) + v_{\text{xc}}(\boldsymbol{r})\right]\psi_k(\boldsymbol{r}) = \epsilon_k\psi_k(\boldsymbol{r})$$

Kohn-Sham
1965

Standard Electronic Structure Theory

$$\left[-\frac{\nabla^2}{2} + v_{\text{ext}}(\boldsymbol{r}) + v_{\text{es}}(\boldsymbol{r}) + v_{\text{xc}}(\boldsymbol{r})\right]\psi_k(\boldsymbol{r}) = \epsilon_k \psi_k(\boldsymbol{r})$$

Kohn-Sham 1965

"As (almost) everyone does":

1. Pick *basis set* $\{|\varphi_i\rangle\}$:

$$\psi_k(\boldsymbol{r}) = \sum_i c_{ki}\varphi_i(\boldsymbol{r})$$

→generalized eigenvalue problem:

$$\underline{\underline{h}}\,\underline{c}_k = \epsilon_k\,\underline{\underline{s}}\,\underline{c}_k$$

$$h_{ij} = \langle\varphi_i|\hat{h}_{\text{KS}}|\varphi_j\rangle$$

$$s_{ij} = \langle\varphi_i|\varphi_j\rangle$$

Standard Electronic Structure Theory

$$\left[-\frac{\nabla^2}{2} + v_{\text{ext}}(\boldsymbol{r}) + v_{\text{es}}(\boldsymbol{r}) + v_{\text{xc}}(\boldsymbol{r})\right]\psi_k(\boldsymbol{r}) = \epsilon_k \psi_k(\boldsymbol{r})$$

Kohn-Sham 1965

"As (almost) everyone does":

1. Pick *basis set* $\{|\varphi_i\rangle\}$:

$$\psi_k(\boldsymbol{r}) = \sum_i c_{ki}\varphi_i(\boldsymbol{r})$$

→generalized eigenvalue problem:

$$\underline{\underline{h}}\,\underline{c}_k = \epsilon_k\,\underline{\underline{s}}\,\underline{c}_k$$

$$h_{ij} = \langle\varphi_i|\hat{h}_{\text{KS}}|\varphi_j\rangle$$

$$s_{ij} = \langle\varphi_i|\varphi_j\rangle$$

2. Self-consistency:

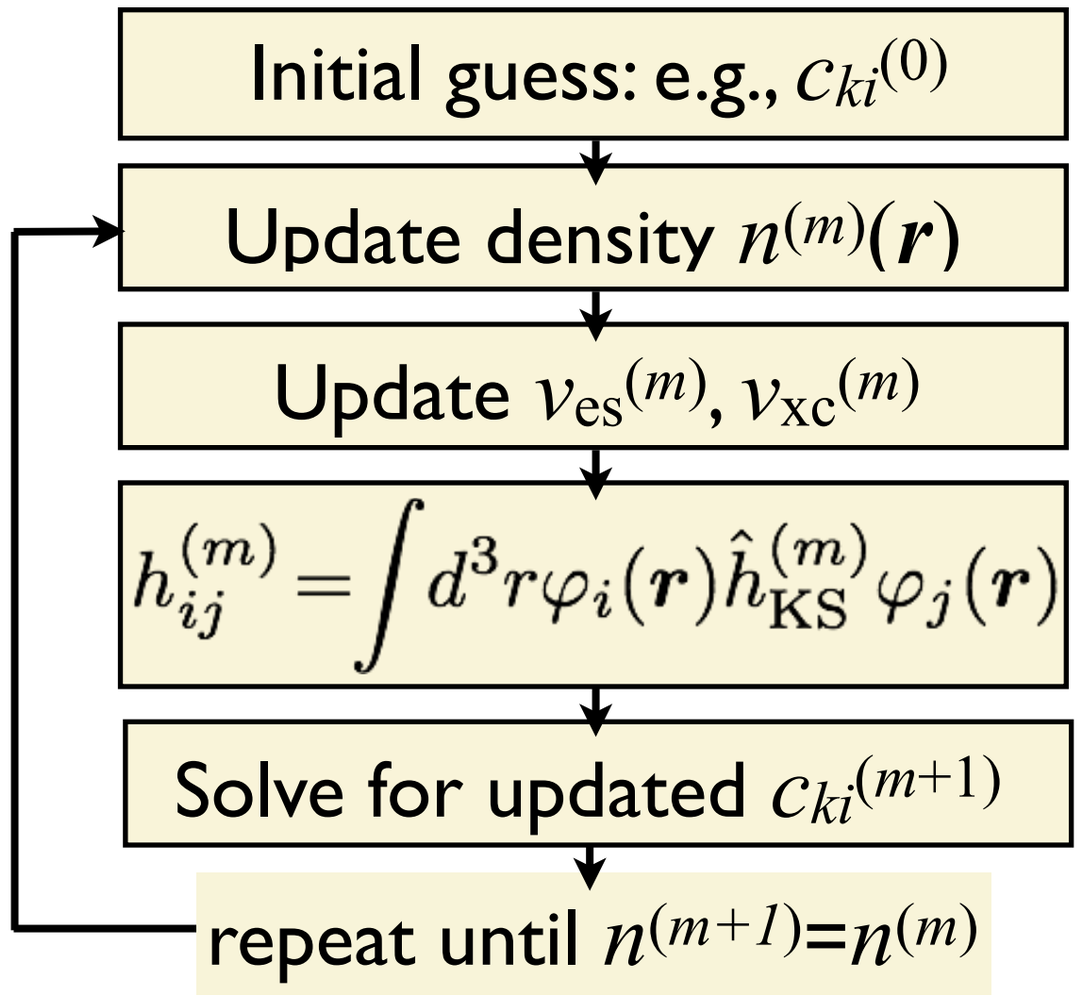

Numeric Atom-Centered Basis Sets for All-Electron DFT

$$\varphi_{i[lm]}(\boldsymbol{r}) = \frac{u_i(r)}{r} \cdot Y_{lm}(\Omega)$$

- $u_i(r)$: Flexible choice - "Anything you like."

Many popular implementations:
DMol³ (Delley), FPLO (Eschrig *et al.*), PLATO (Horsfield *et al.*), PAOs (Siesta, Conquest, OpenMX[2], Fireball, ABACUS, …)

Numeric Atom-Centered Basis Sets for All-Electron DFT

$$\varphi_{i[lm]}(\boldsymbol{r}) = \frac{u_i(r)}{r} \cdot Y_{lm}(\Omega)$$

Many popular implementations:
DMol³ (Delley), FPLO (Eschrig *et al.*), PLATO (Horsfield *et al.*), PAOs (Siesta, Conquest, OpenMX[2], Fireball, ABACUS, …)

- $u_i(r)$: Flexible choice - "Anything you like."

$$\left[-\frac{1}{2}\frac{d^2}{dr^2} + \frac{l(l+1)}{r^2} + v_i(r) + v_{\text{cut}}(r)\right] u_i(r) = \epsilon_i u_i(r)$$

Numeric Atom-Centered Basis Sets for All-Electron DFT

$$\varphi_{i[lm]}(\boldsymbol{r}) = \frac{u_i(r)}{r} \cdot Y_{lm}(\Omega)$$

Many popular implementations:
DMol³ (Delley), FPLO (Eschrig *et al.*), PLATO (Horsfield *et al.*), PAOs (Siesta, Conquest, OpenMX[2], Fireball, ABACUS, …)

- $u_i(r)$: Flexible choice - "Anything you like."

$$\left[-\frac{1}{2}\frac{d^2}{dr^2} + \frac{l(l+1)}{r^2} + v_i(r) + v_{\text{cut}}(r)\right] u_i(r) = \epsilon_i u_i(r)$$

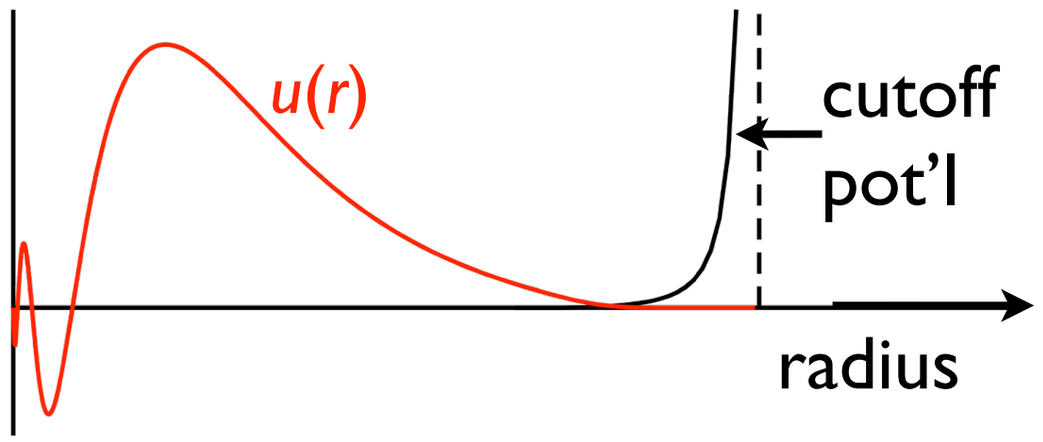

Numeric Atom-Centered Basis Sets for All-Electron DFT

$$\varphi_{i[lm]}(\boldsymbol{r}) = \frac{u_i(r)}{r} \cdot Y_{lm}(\Omega)$$

Many popular implementations:
DMol³ (Delley), FPLO (Eschrig *et al.*), PLATO (Horsfield *et al.*), PAOs (Siesta, Conquest, OpenMX[2], Fireball, ABACUS, …)

- $u_i(r)$: Flexible choice - "Anything you like."

$$\left[-\frac{1}{2}\frac{d^2}{dr^2} + \frac{l(l+1)}{r^2} + v_i(r) + v_{\text{cut}}(r)\right] u_i(r) = \epsilon_i u_i(r)$$

- free-atom like: $v_i(r) = v^{\text{DFT}}_{\text{free atom}}(r)$
- Hydrogen-like: $v_i(r) = z/r$
- free ions, harm. osc. (Gaussians), ...

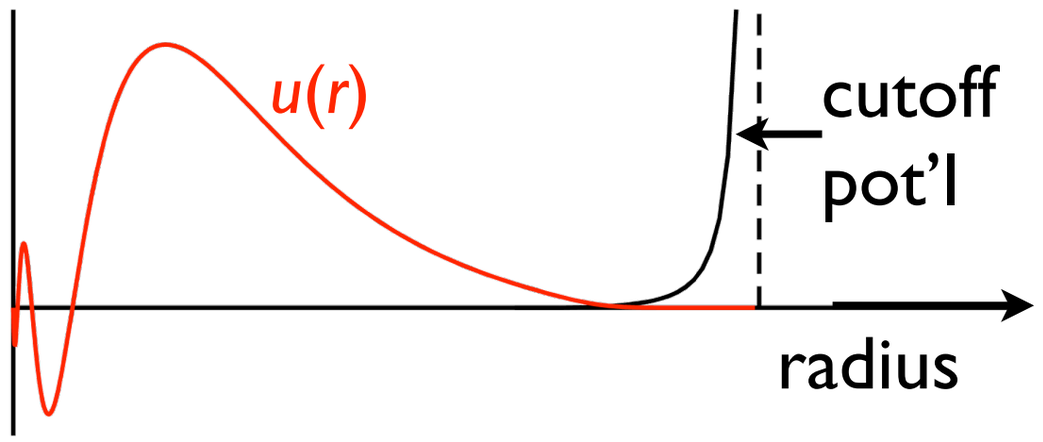

… and straightforward to substitute scalar- or fully-relativistic analogues

Numeric Atom-Centered Basis Sets for All-Electron DFT

$$\varphi_{i[lm]}(\boldsymbol{r}) = \frac{u_i(r)}{r} \cdot Y_{lm}(\Omega)$$

Many popular implementations:
DMol³ (Delley), FPLO (Eschrig *et al.*), PLATO (Horsfield *et al.*), PAOs (Siesta, Conquest, OpenMX[2], Fireball, ABACUS, …)

- $u_i(r)$: Flexible choice - "Anything you like."

→ Localized; "naturally" all-electron

→ The choice of efficient and of enough radial functions is obviously important

→ We have a "basis set library" - list of basis functions for all elements (1-102), from fast qualitative to meV-converged total energies (LDA/GGA/mGGA/hybrid DFA's)

V. Blum, R. Gehrke, F. Hanke, P. Havu, V. Havu, X. Ren, K. Reuter and M. Scheffler,
"Ab Initio Molecular Simulations with Numeric Atom-Centered Orbitals",
Computer Physics Communications ***180****, 2175-2196 (2009)*

Numeric Atom-Centered Basis Sets for All-Electron DFT

$$\varphi_{i[lm]}(\boldsymbol{r}) = \frac{u_i(r)}{r} \cdot Y_{lm}(\Omega)$$

Many popular implementations:
DMol³ (Delley), FPLO (Eschrig *et al.*), PLATO (Horsfield *et al.*),

- $u_i(r)$: Flexible choic

→ Localized; "naturall

→ The choice of effici
important

→ We have a "basis se
(1-102), from fast c
mGGA/hybrid DFA

V. Blum, R. Gehrk
"Ab Initio Molec
Comput

FHI-aims input example: Hydrogen, tight settings

```
################################################################################
#
#  Definition of "minimal" basis
#
################################################################################
#     valence basis states
    valence      1  s   1.
#     ion occupancy
    ion_occ      1  s   0.5
################################################################################
#
#  Suggested additional basis functions. For production calculations,
#  uncomment them one after another (the most important basis functions are
#  listed first).
#
#  Basis constructed for dimers: 0.5 A, 0.7 A, 1.0 A, 1.5 A, 2.5 A
#
################################################################################
#  "First tier" - improvements: -1014.90 meV to -62.69 meV
     hydro 2 s 2.1
     hydro 2 p 3.5
#  "Second tier" - improvements: -12.89 meV to -1.83 meV
     hydro 1 s 0.85
     hydro 2 p 3.7
     hydro 2 s 1.2
     hydro 3 d 7
#  "Third tier" - improvements: -0.25 meV to -0.12 meV
#     hydro 4 f 11.2
#     hydro 3 p 4.8
#     hydro 4 d 9
#     hydro 3 s 3.2
################################################################################
```

Precision in Community Wide Benchmark - "Delta Test"

Reproducibility in Density-Functional Calculations of Solids,
K. Lejaeghere, ... **68 coauthors!** ..., S. Cottenier,
Science **351**, aad3000 (2016).
https://molmod.ugent.be/deltacodesdft

E(V) for 71 elemental solids - 15 codes, all-electron & 40 pseudopot'l sets

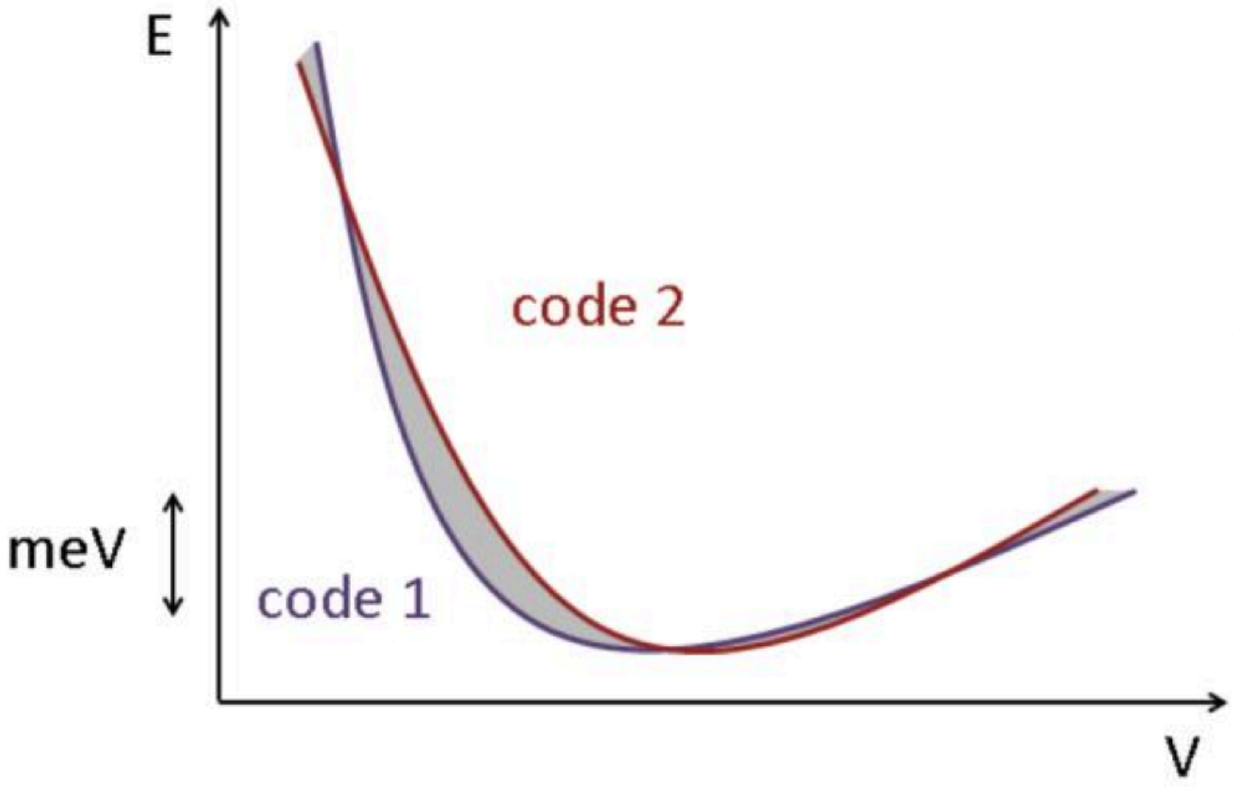

$$\rightarrow \quad \Delta_i(a,b) = \sqrt{\int_{0.94V_{0,i}}^{1.06V_{0,i}} \frac{\left(E_{b,i}(V) - E_{a,i}(V)\right)^2}{0.12V_{0,i}}\,\mathrm{d}V}$$

Precision in Community Wide Benchmark - "Delta Test"

Reproducibility in Density-Functional Calculations of Solids,
K. Lejaeghere, ... **68 coauthors!** ..., S. Cottenier,
Science **351**, aad3000 (2016).
https://molmod.ugent.be/deltacodesdft

E(V) for 71 elemental solids - 15 codes, all-electron & 40 pseudopot'l sets

FHI-aims: Test carried out independently by Dr. Marcin Dulak, DTU (Copenhagen)

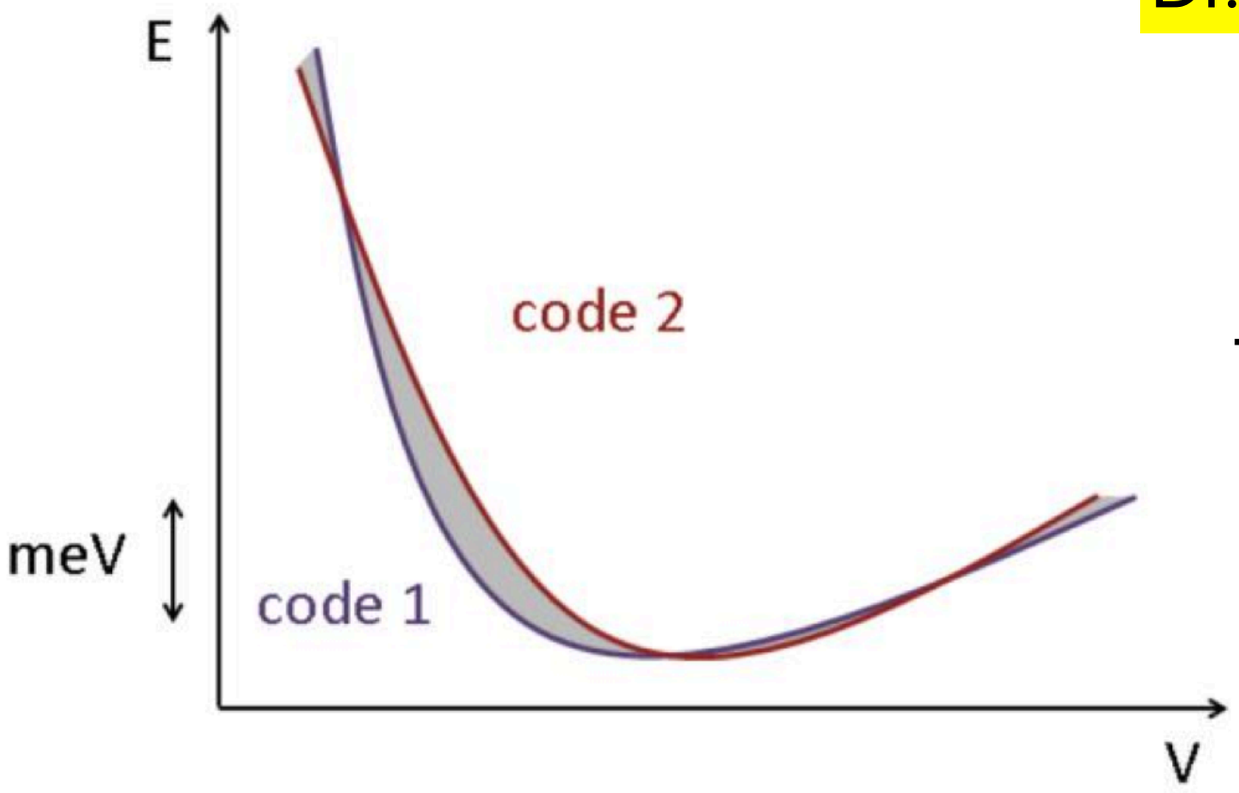

$$\rightarrow \quad \Delta_i(a,b) = \sqrt{\int_{0.94V_{0,i}}^{1.06V_{0,i}} \frac{\left(E_{b,i}(V) - E_{a,i}(V)\right)^2}{0.12V_{0,i}}\,\mathrm{d}V}$$

Precision in Community Wide Benchmark - "Delta Test"

Reproducibility in Density-Functional Calculations of Solids,
K. Lejaeghere, ... 68 coauthors! ..., S. Cottenier,
Science **351**, aad3000 (2016).
https://molmod.ugent.be/deltacodesdft

E(V) for 71 elemental solids - 15 codes, all-electron & 40 pseudopot'l sets

Code	Basis	Electron treatment	Delta (meV)
Wien2k 13.1	LAPW/APW+lo	All-electron	0
FHI-aims 081213*	NAO, tier2	All-electron (scalar rel. atomic ZORA)	0.2
Exciting (dev.)	LAPW+xlo	All-electron	0.2
Quantum Espresso 5.1	plane waves	SSSP accuracy (mixed NC/US/PAW library)	0.3
VASP 5.2.12	plane waves	PAW 2015	0.3
FHI-aims 081213*	NAO, tier2	All-electron (scalar rel., scaled ZORA)	0.3
ELK 3.1.5	APW+lo	All-electron	0.3
...			

**Results: Marcin Dulak, DTU (Copenhagen)*

Example - Bi-Doping of Layered HOIS, Hybrid DFT

Predicted Observable	High-precision atomic structure, energies and energy differences	Energy band structures for large, complex systems

U.S. DEPARTMENT OF ENERGY | Office of Science

CHOISE
Center for Hybrid Organic Inorganic Semiconductors for Energy

MAX-PLANCK-GESELLSCHAFT

NSF

XSEDE
Extreme Science and Engineering Discovery Environment

Example - Bi-Doping of Layered HOIS, Hybrid DFT

Predicted Observable	High-precision atomic structure, energies and energy differences	Energy band structures for large, complex systems
Level of Theory	Semilocal density functional theory with van der Waals corrections	Hybrid density functional theory including spin-orbit coupling

U.S. DEPARTMENT OF ENERGY | Office of Science
CHOISE Center for Hybrid Organic Inorganic Semiconductors for Energy
MAX-PLANCK-GESELLSCHAFT
NSF
XSEDE Extreme Science and Engineering Discovery Environment

Example - Bi-Doping of Layered HOIS, Hybrid DFT

Predicted Observable	High-precision atomic structure, energies and energy differences	Energy band structures for large, complex systems
Level of Theory	Semilocal density functional theory with van der Waals corrections	Hybrid density functional theory including spin-orbit coupling

Dr. Yi Yao

Gabrielle Koknat

Example:
Bi-doped $(PEA)_2PbI_4$
1,504 atoms
HSE06+SOC,
all-electron
High precision -
no tricks!

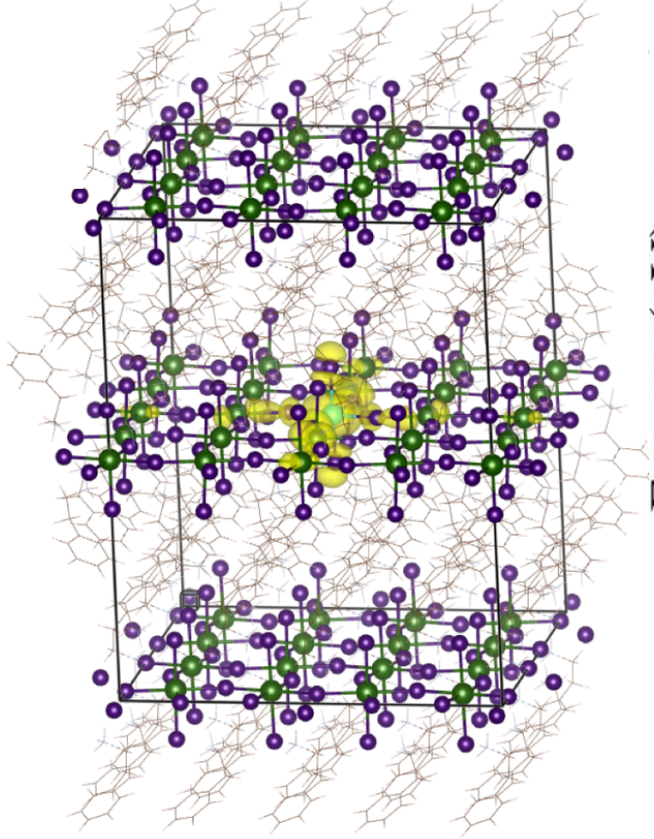

DFT-HSE06+SOC Band Structure

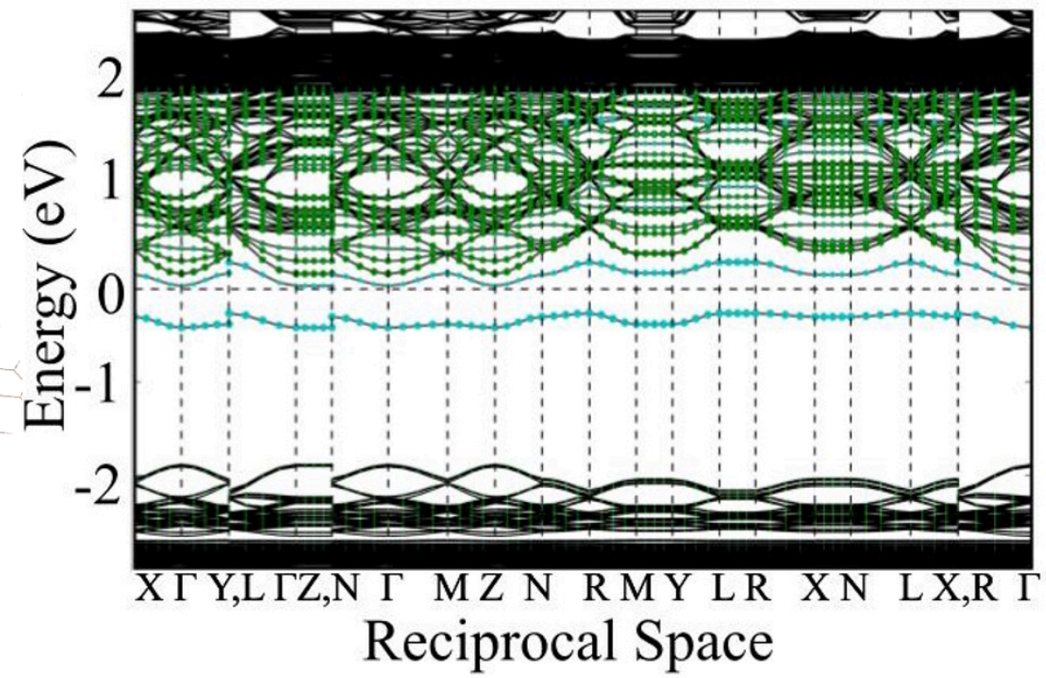

U.S. DEPARTMENT OF ENERGY | Office of Science

CHOISE
Center for Hybrid Organic Inorganic Semiconductors for Energy

MAX-PLANCK-GESELLSCHAFT

NSF

XSEDE
Extreme Science and Engineering Discovery Environment

This Calculation Was Not Cheap … But Valuable

Example:
Bi-doped $(PEA)_2PbI_4$
1,504 atoms
HSE06+SOC,
all-electron
High precision -
no tricks!

Dr. Yi Yao

Gabrielle Koknat

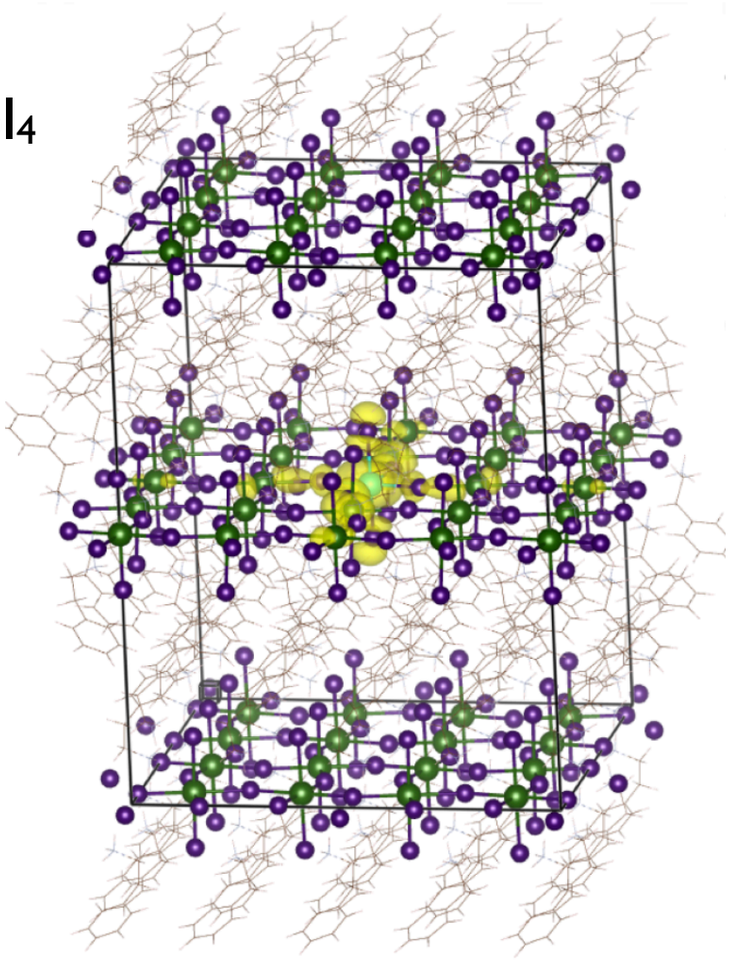

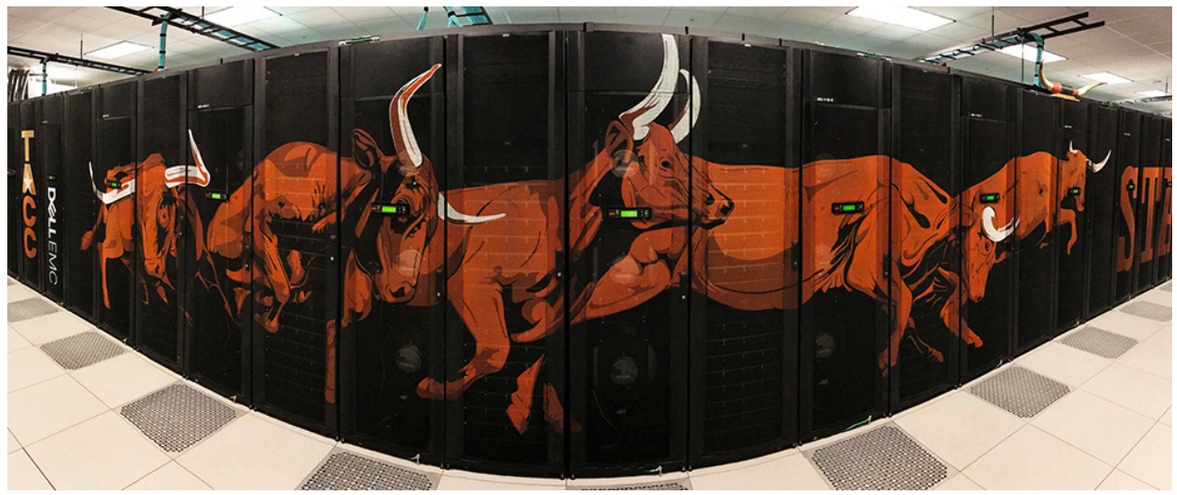

Stampede2 Supercomputer - 18 PFlop/s
NSF XSEDE / Texas Advanced Computing Center
No. 35, Top500 List of Supercomputers, June 2021

XSEDE
Extreme Science and Engineering Discovery Environment

This computation: 216 compute nodes
(out of ~6,000 total), ~4 hours

Drastically optimized hybrid DFT implementation (2020-22):
F. Merz (Lenovo), M. Rampp (Max Planck Computing and Data Facility, Munich)
S. Kokott, C. Carbogno, M. Scheffler (Fritz Haber Institute, Berlin)
Y. Yao, V. Blum (Duke University)

The Other Parallelism - Split Communicators

FHI-aims is a subroutine call, which can work cleanly with split MP communicators

```
program multi_aims
[...]
call initialize_mpi()
[...]
call MPI_Comm_split(mpi_comm_global,sub_class,sub_id,mpi_comm_subtask,ierr)
[...]
 mpi_comm_global = mpi_comm_subtask
 myid = sub_id
 n_tasks = sub_tasks
[...]
call change_directory(subfolder)
[...]
  if (myid == 0) then
     call aims(mpi_comm_global,myunit,use_mpi)
  else
     call aims(mpi_comm_global,6,use_mpi)
  end if
[...]
  call MPI_Comm_free(mpi_comm_subtask,ierr)
[...]
call finalize_mpi()
```

8,192 Full Geometry Relaxations, Split MPI Communicator

RESEARCH ARTICLE

ADVANCED OPTICAL MATERIALS
www.advopticalmat.de

Chiral Perovskite Nanoplatelets with Tunable Circularly Polarized Luminescence in the Strong Confinement Regime

Qinxuan Cao, Ruyi Song, Christopher C. S. Chan, Zhiyu Wang, Pui Ying Wong, Kam Sing Wong, Volker Blum, and Haipeng Lu**

Adv. Optical Mater. 2023, 2203125 - DOI: 10.1002/adom.202203125

Ruyi Song

4,096 monolayer flake conformations + 4,096 bilayer flake conformations.
106 atoms & 116 atoms each, periodic slabs
Full relaxation, DFT-PBE+TS, light settings, Perlmutter, batches of 1,024 structures each

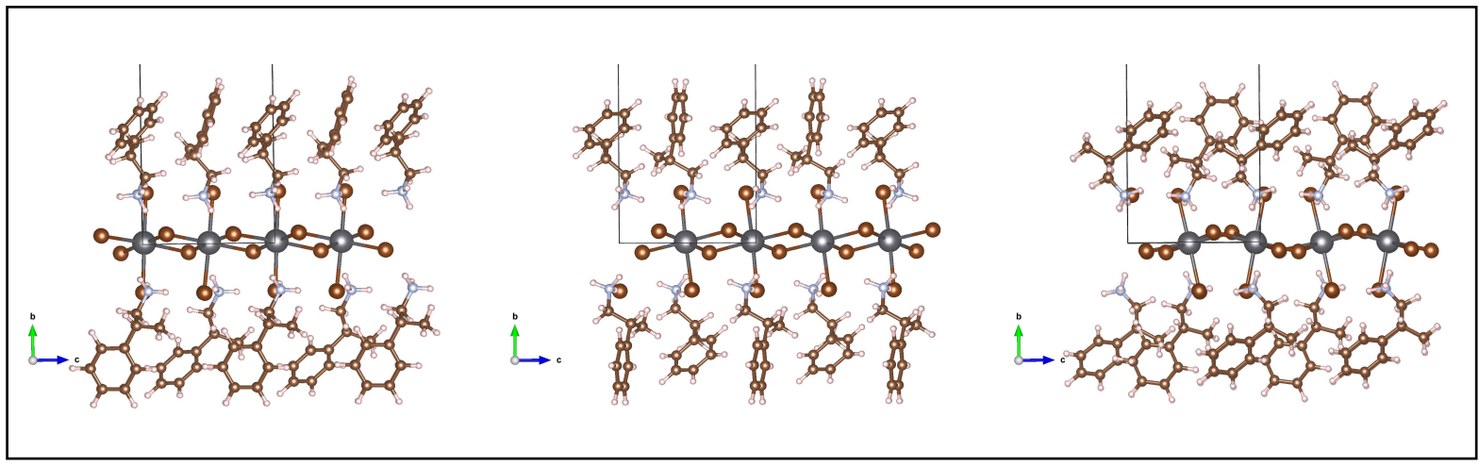

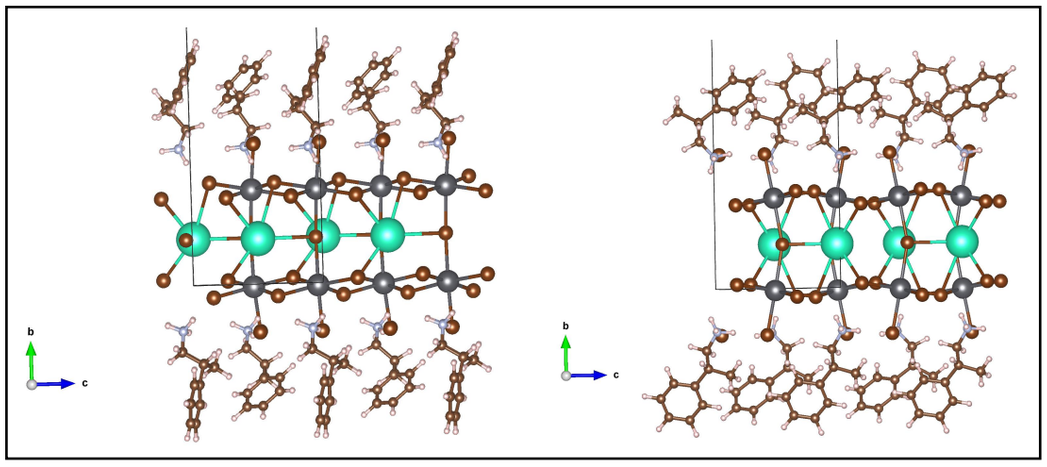

→ 5 low-energy models: spin-orbit coupled band structures, spin texture calculations.

Typical Scaling - $O(N^3)$ Wall

$$H\varphi = \epsilon S\varphi$$

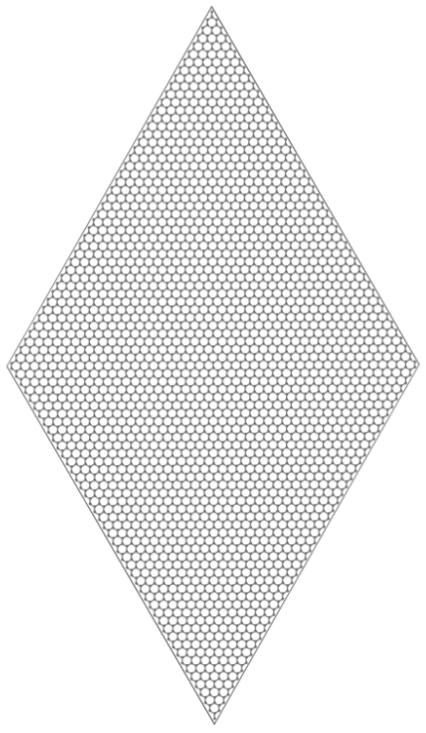

Graphene monolayer,
4050 atoms

FHI-aims, PBE, "light" settings

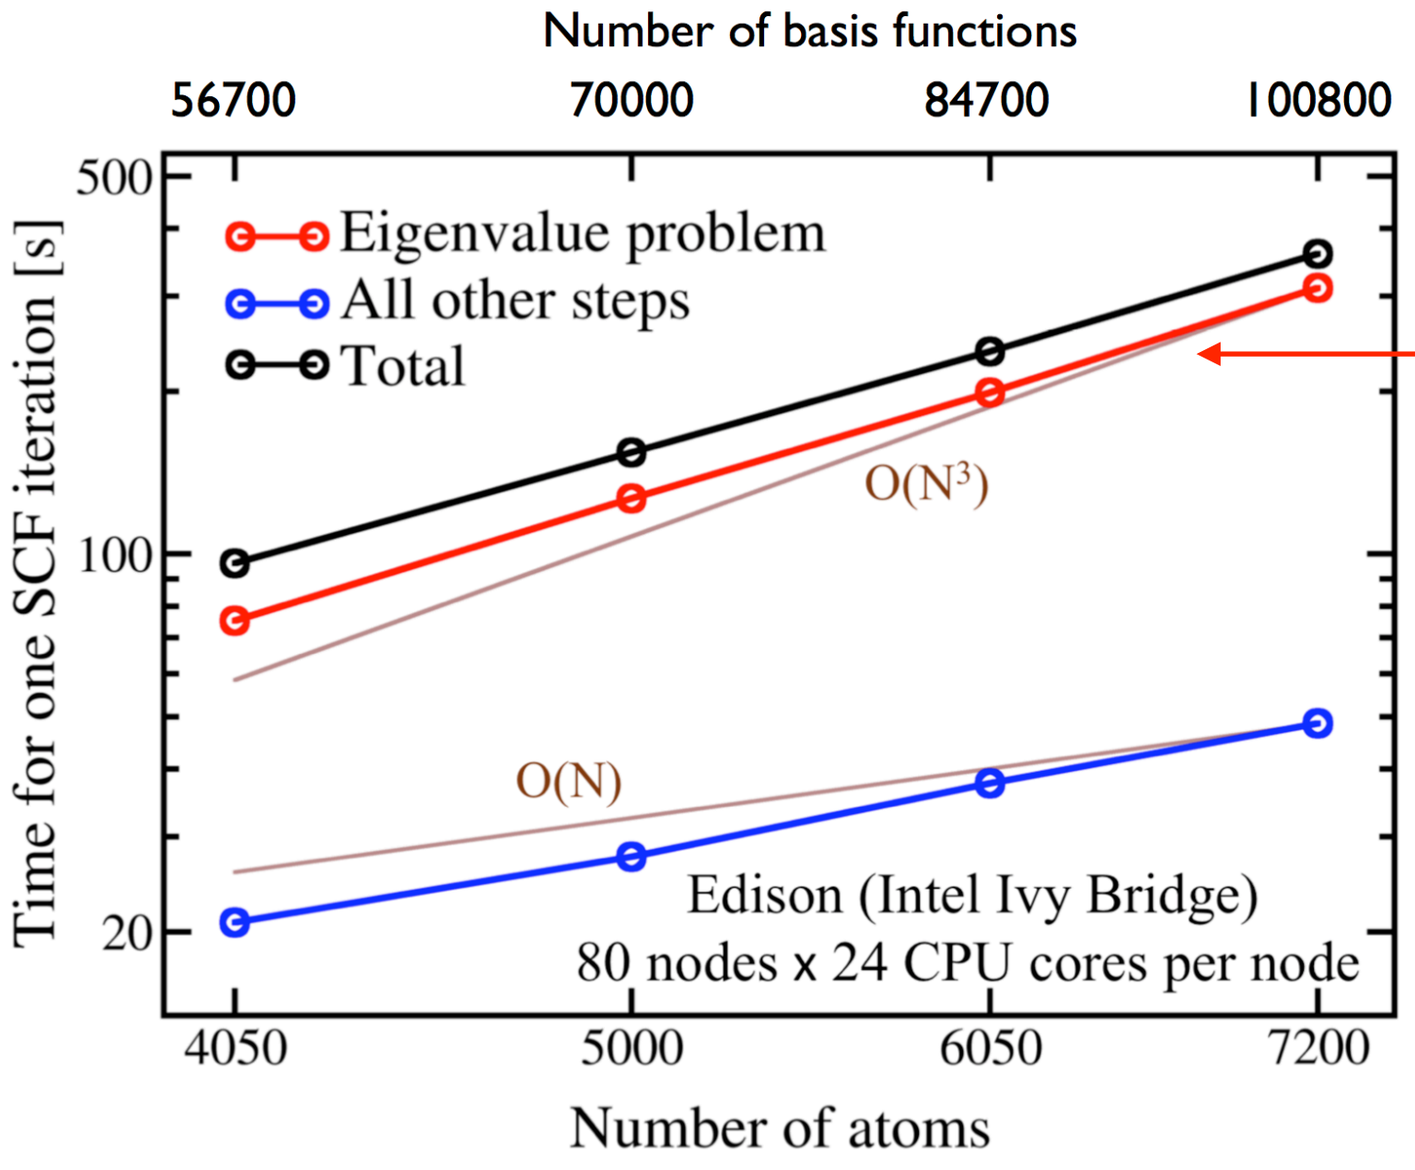

ELPA Library
http://elpa.rzg.mpg.de

Generic problem for any Kohn-Sham DFT code ... solution strategies?

ELSI - Infrastructure for High-Performance "Solvers"

Nucleus: Emilio Artacho, 2014: *"Dear all, There will be a workshop in CECAM at Lausanne … aiming to kick-start an electronic structure library. … I hope you are interested"*

ELSI - Infrastructure for High-Performance "Solvers"

Nucleus: Emilio Artacho, 2014: *"Dear all, There will be a workshop in CECAM at Lausanne … aiming to kick-start an electronic structure library. … I hope you are interested"*

Electronic Structure Library:

Micael Oliveira, Yann Pouillon, Fabiano Corsetti, Nick Papior, Alin Elena, Martin Lüders, Damien Caliste, many more.

https://esl.cecam.org

https://gitlab.com/ElectronicStructureLibrary

ELSI - Infrastructure for High-Performance "Solvers"

Nucleus: Emilio Artacho, 2014: *"Dear all, There will be a workshop in CECAM at Lausanne … aiming to kick-start an electronic structure library. … I hope you are interested"*

Electronic Structure Library:

Micael Oliveira, Yann Pouillon, Fabiano Corsetti, Nick Papior, Alin Elena, Martin Lüders, Damien Caliste, many more.

https://esl.cecam.org

https://gitlab.com/ElectronicStructureLibrary

Electronic Structure Infrastructure (ELSI): NSF-SI2 - ACI-1450280:

Volker Blum, Jianfeng Lu, Lin Lin, Chao Yang, Alvaro Vazquez-Mayagoitia, Fabiano Corsetti

Why ELSI Works:

Victor Yu (now ANL)

Yingzhou Li (Duke)

Will Huhn (now ANL)

Jonathan Moussa (MolSSI)

Yi Yao (Duke)

Ruyi Song (Duke)

MANY MORE!

Many Scalable Solution to Kohn-Sham Problem

$$\underline{\underline{h}}\,\underline{c}_k = \epsilon_k\,\underline{\underline{s}}\,\underline{c}_k$$

Different use cases (basis sets, physics), different "solvers".
Solve (eigenvectors, $O(N^3)$) or circumvent (density matrix)?

Many Scalable Solution to Kohn-Sham Problem

$$\underline{\underline{h}}\,\underline{c}_k = \epsilon_k\,\underline{\underline{s}}\,\underline{c}_k$$

Different use cases (basis sets, physics), different "solvers".
Solve (eigenvectors, $O(N^3)$) or circumvent (density matrix)?

Exact solvers	Iterative solvers	DM: O(N) solvers	Other DM-based approaches
Lapack Scalapack ELPA EigenExa Magma …	Davidson Projected Preconditioned Conjugate Gradient Chebychev Filtering Slepc-SIPS …	NTPoly Various code-internal and/or proprietary implementations	PEXSI Orbital Minimization Method FEAST …
Robust General	(Essentially) robust $N_{basis} >> N_{ev}$	Sparse H, S Nonmetallic systems	Sparse H, S can depend on XC

Many Different Solvers - How to Unify Access?

Electronic structure codes

?

Solvers

- ELPA
- NTPoly
- EigenExa
- PEXSI
- SLEPc-SIPs
- libOMM
- MAGMA
- More ...

? Replicated infrastructure to implement solvers efficiently

? Conversion between a variety of matrix storage formats

? Complexity in solver selection for different problems

ELSI: Connecting Electronic Structure Codes and Solvers

Yu et al., "ELSI – An Open Infrastructure for Electronic Structure Solvers"
Comp. Phys. Commun. 256, 107459 (2020).

Yu et al., Comput. Phys. Commun. 2018
http://elsi-interchange.org
http://git.elsi-interchange.org/elsi-devel/elsi-interface

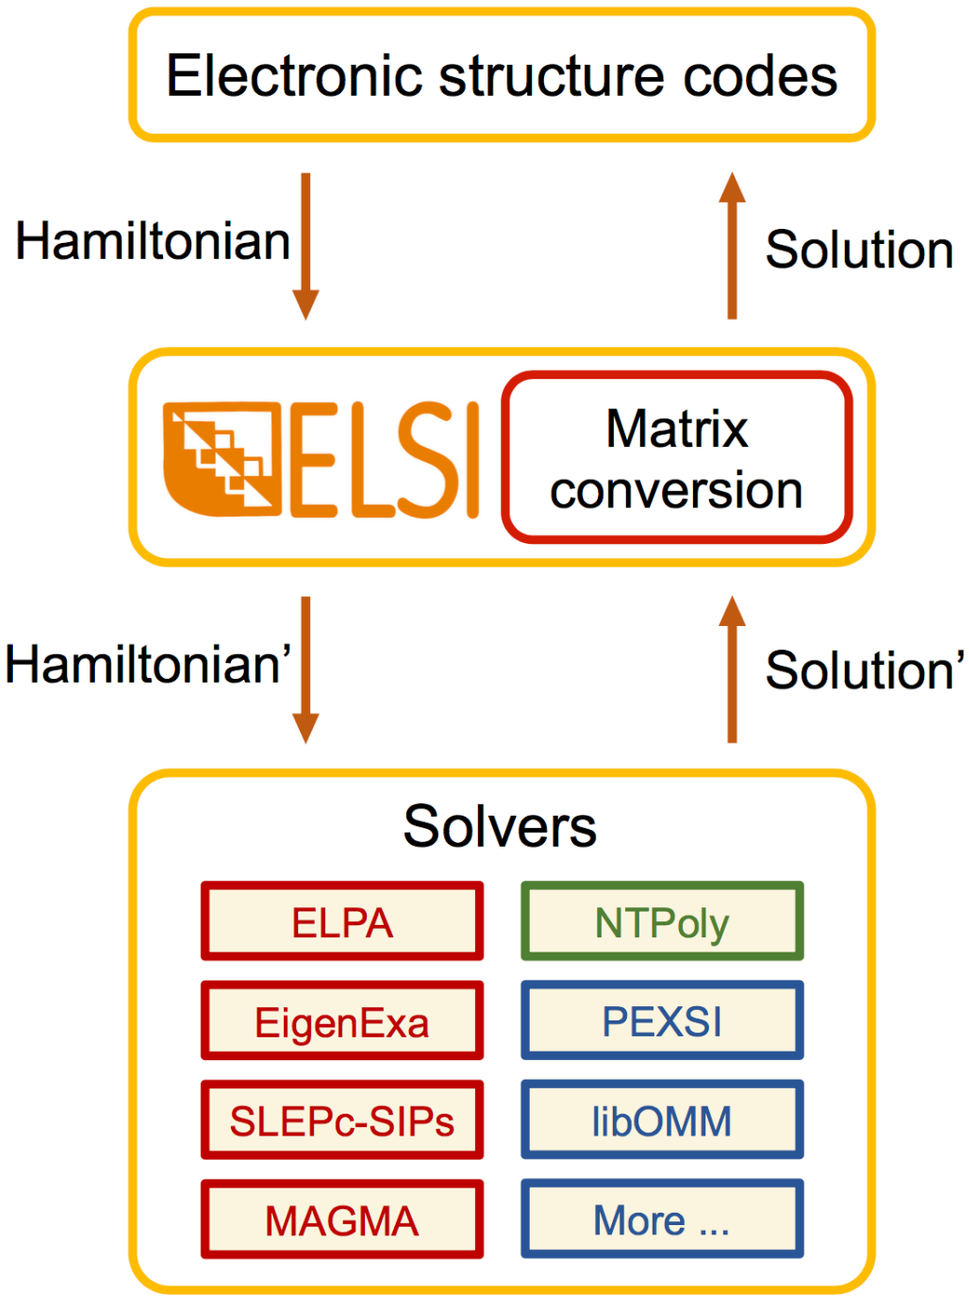

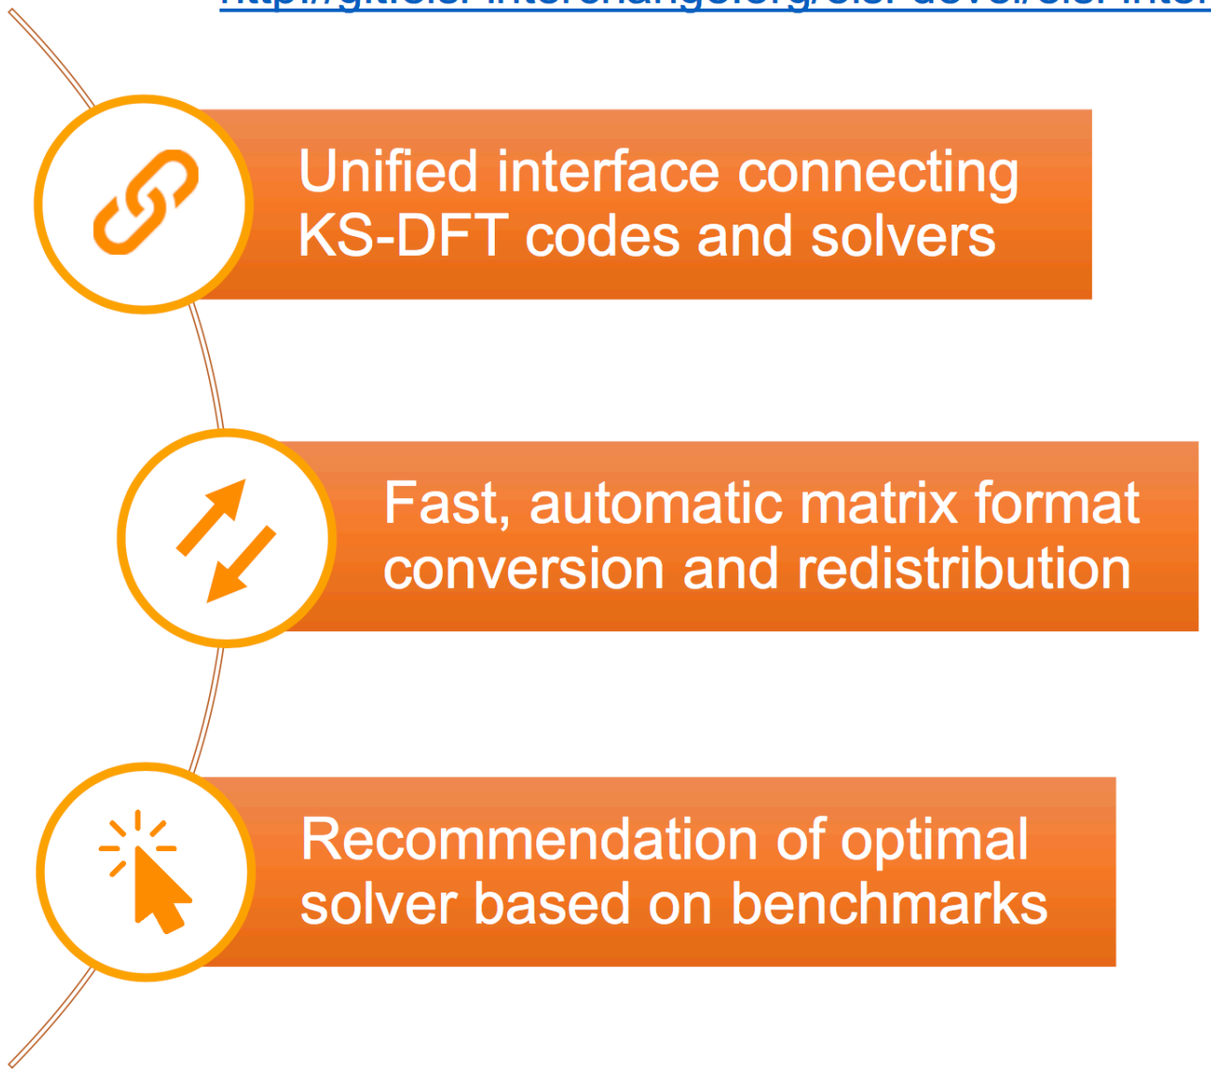

ELSI: Connecting Electronic Structure Codes and Solvers

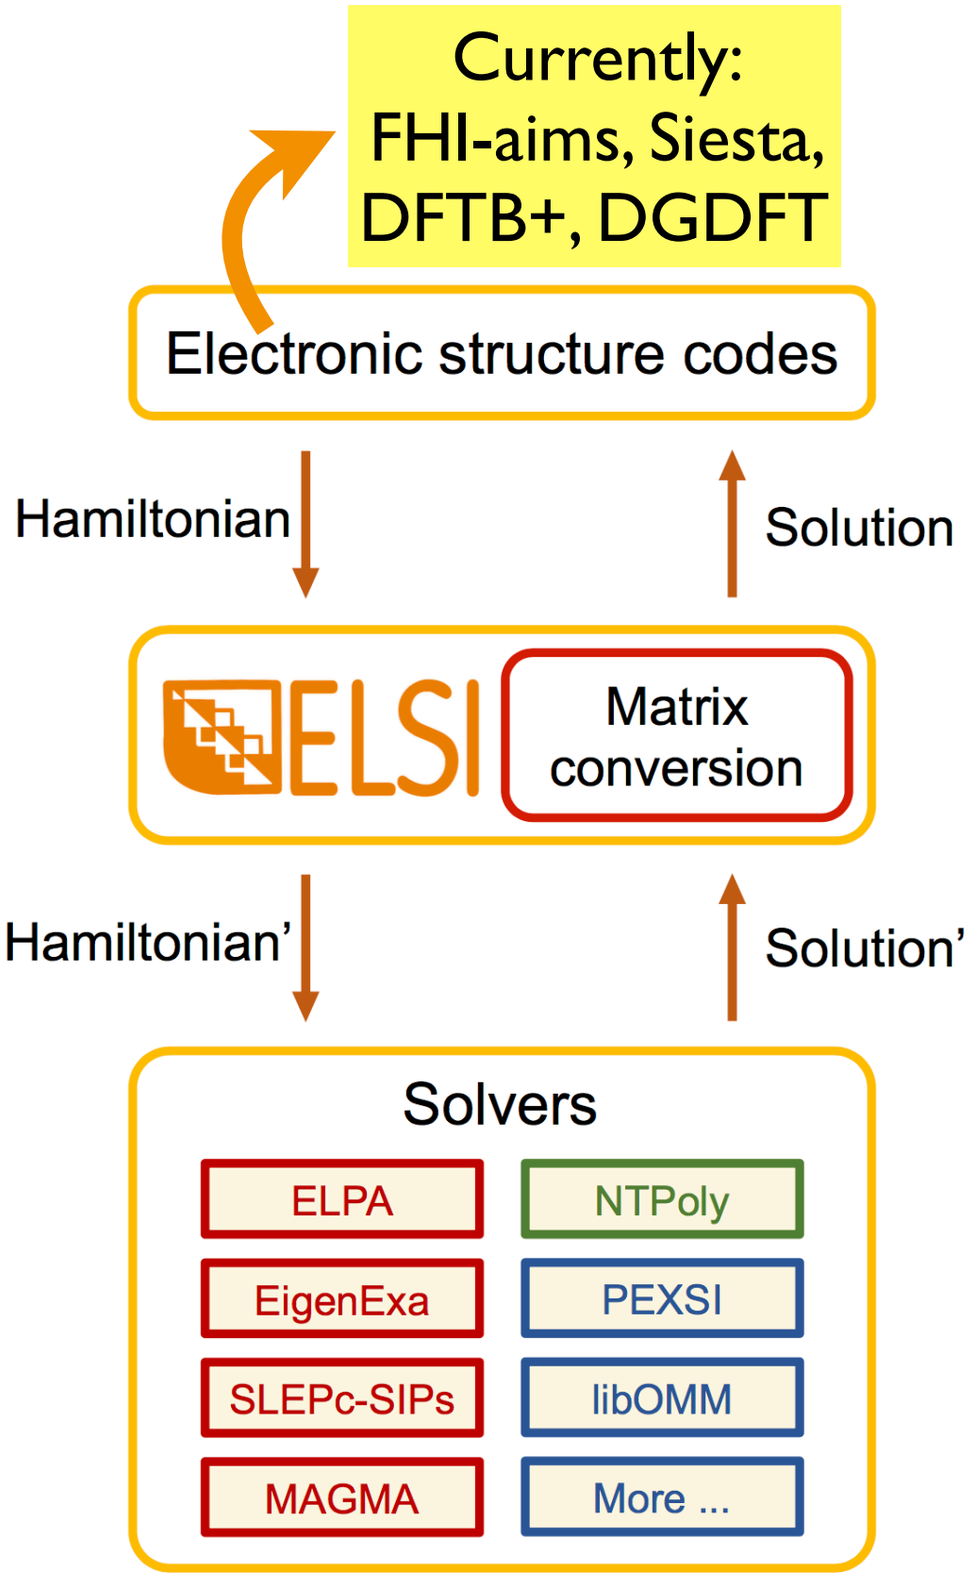

Yu et al., "ELSI – An Open Infrastructure for Electronic Structure Solvers"
Comp. Phys. Commun. 256, 107459 (2020).

Yu et al., Comput. Phys. Commun. 2018
http://elsi-interchange.org
http://git.elsi-interchange.org/elsi-devel/elsi-interface

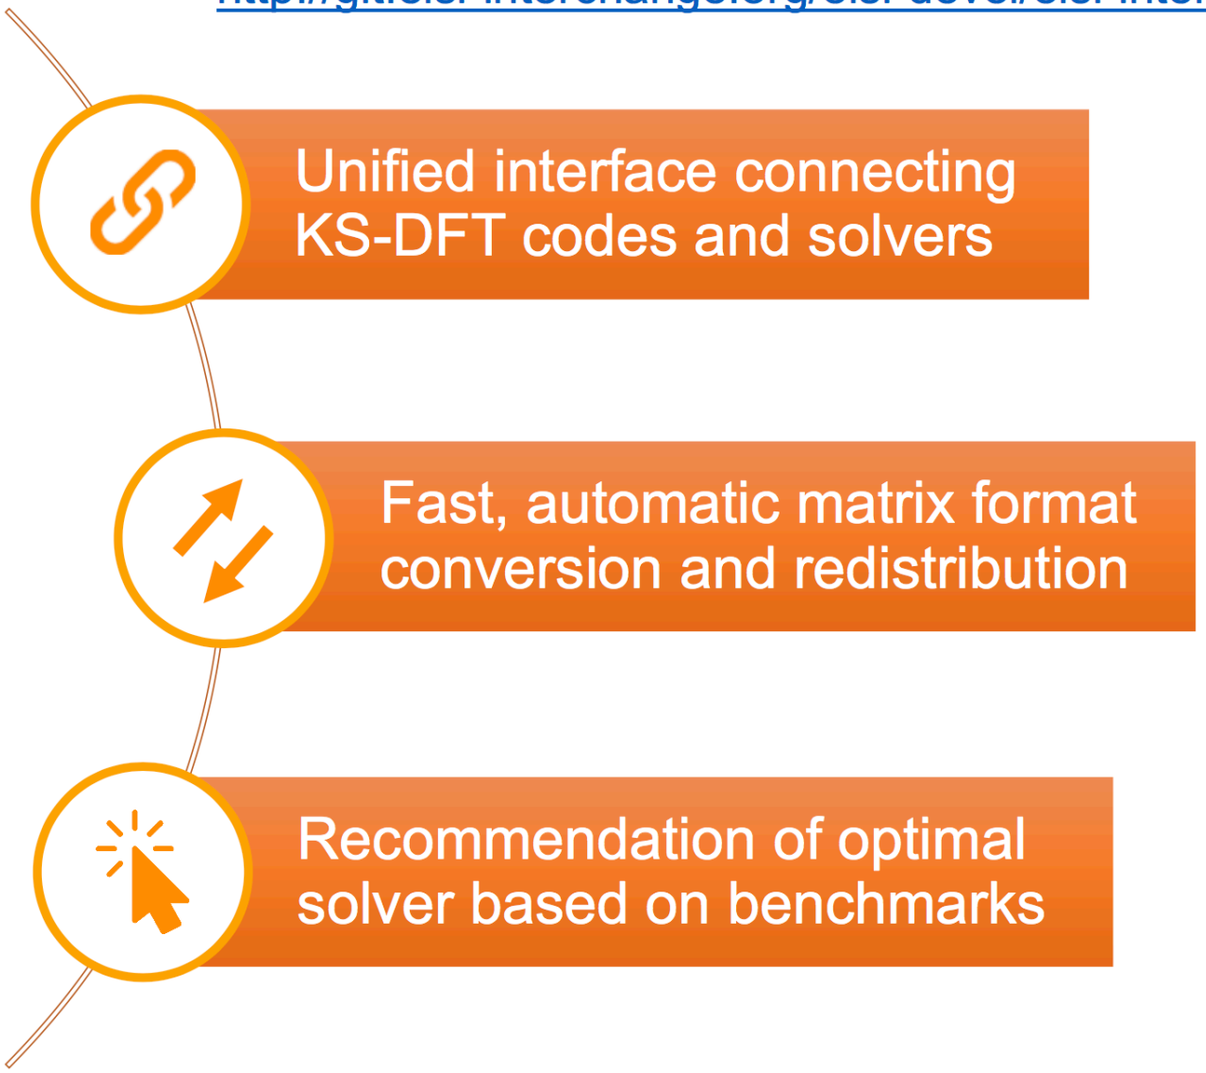

ELSI: Current Functionality

- Designed for rapid integration into a variety of electronic structure codes
- Compatible with common workflows
  - Single self-consistent field (SCF)
  - Multiple SCF cycles (geometry relaxation or molecular dynamics)
- Supports density matrix solvers and eigensolvers on equal footing
- Technical settings adjustable for experienced users
- Object-oriented: Concurrent instances

- **python interface (elsipy)**

Currently supported "solvers":

- Lapack
- ELPA
- EigenExa
- MAGMA
- PEXSI
- NTPoly
- libOMM
- SLEPc-SIPS
- BSEPACK
- **ChaSE - E. de Napoli et al.**

Distributed-Parallel GPU Solver: ELPA2-GPU

V.W.-z.Yu, J. Moussa, P. Kůs, A. Marek, P. Messmer, M.Yoon, H. Lederer, V. Blum, Comp. Phys. Commun. 262, 107808 (2021).

Victor Yu
(now ANL)

Jonathan
Moussa
(MolSSI)

Approach:

- cuBLAS / CUDA / MPI (all cores/node, no OpenMP)
- cuBLAS offloading wherever possible
- NVidia Multi-Process Service (MPS); data transfer, computation overlapped
- New (rewritten) CUDA kernel for tridiagonal-to-banded back transformation

$$(\boldsymbol{I} - \tau \boldsymbol{v}\boldsymbol{v}^*)\boldsymbol{x} = \boldsymbol{x} - \tau \boldsymbol{v}(\boldsymbol{v}^*\boldsymbol{x}).$$

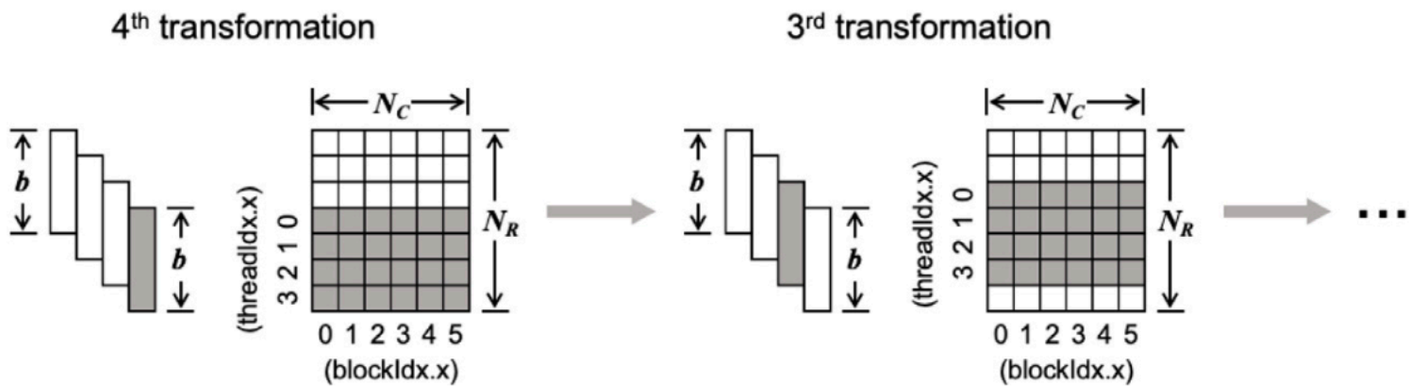

Fig. 4. Workflow of the Householder transformation CUDA kernel. The semi-bandwidth is $b = 4$. Four Householder transformations are applied to the local eigenvector matrix with $N_R = 7$ rows and $N_C = 6$ columns. As indicated by the block index (blockIdx.x) and thread index (threadIdx.x), a CUDA block works on a column of the eigenvector matrix, while a thread within a block works on a single element. From the n-th iteration to the $(n - 1)$-th iteration, the work set of the eigenvector matrix is shifted upward by one element.

Distributed-Parallel GPU Solver: ELPA2-GPU

V.W.-z.Yu, J. Moussa, P. Kůs, A. Marek, P. Messmer, M.Yoon, H. Lederer, V. Blum, Comp. Phys. Commun. 262, 107808 (2021).

Kus *et al.*, 2019: GPU port of ELPA1 - however, 2-stage solver approach not efficient?

<u>2-stage symmetric band reduction:</u>

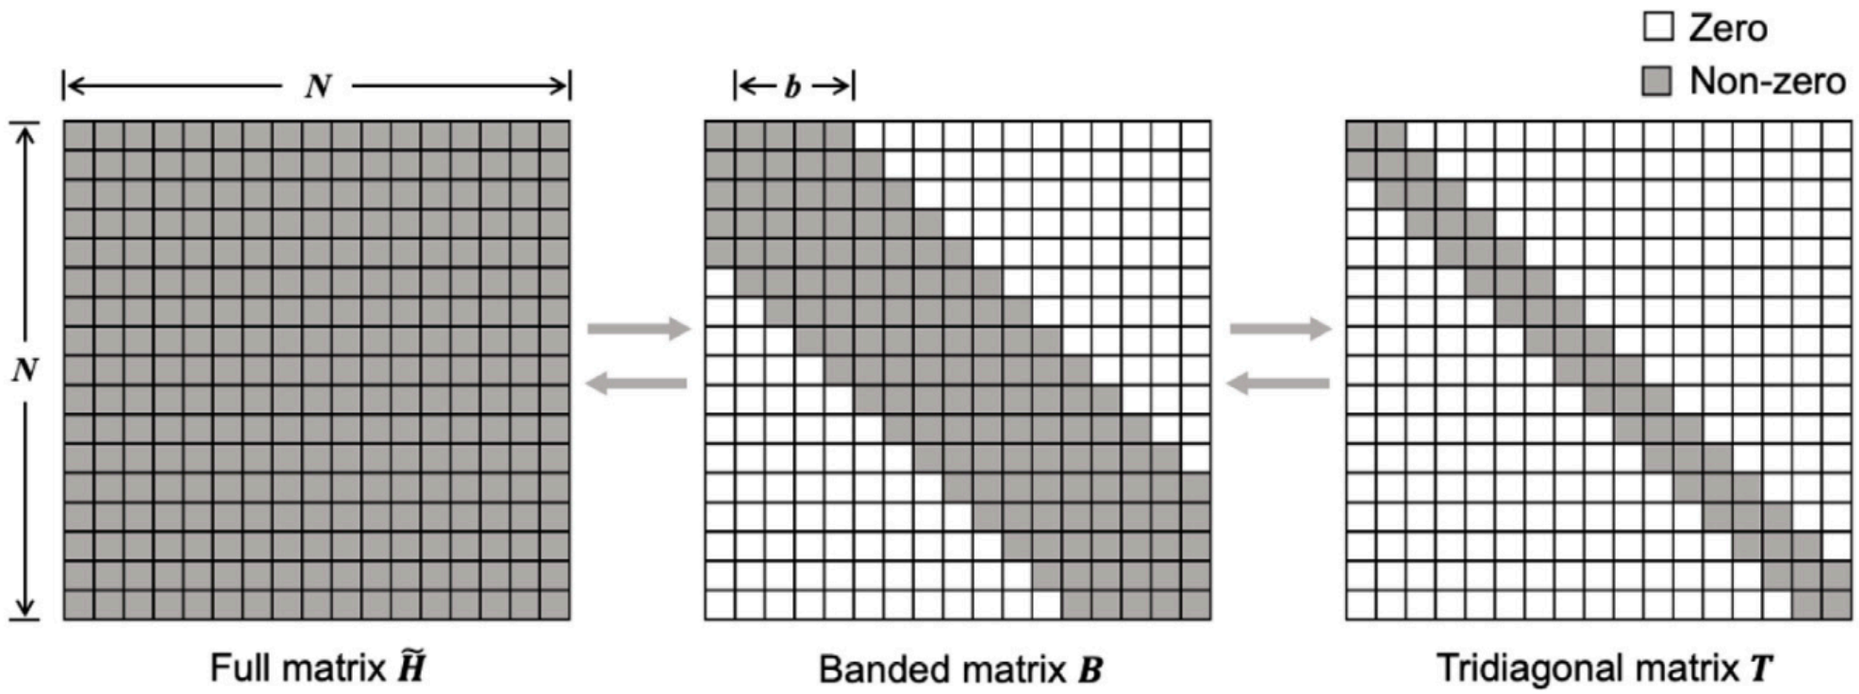

… in past, expected to scale to larger core / MPI rank counts than ELPA1.

… competitive advantage if M eigenpairs < N (dimension) sought.

Net Result: Speedup, Extended Scaling by ELPA2-GPU

ELPA2-GPU - all *M* eigenpairs (worst case)

Talos Cluster (MPCDF)
40 Intel cores + 2 V100 / node

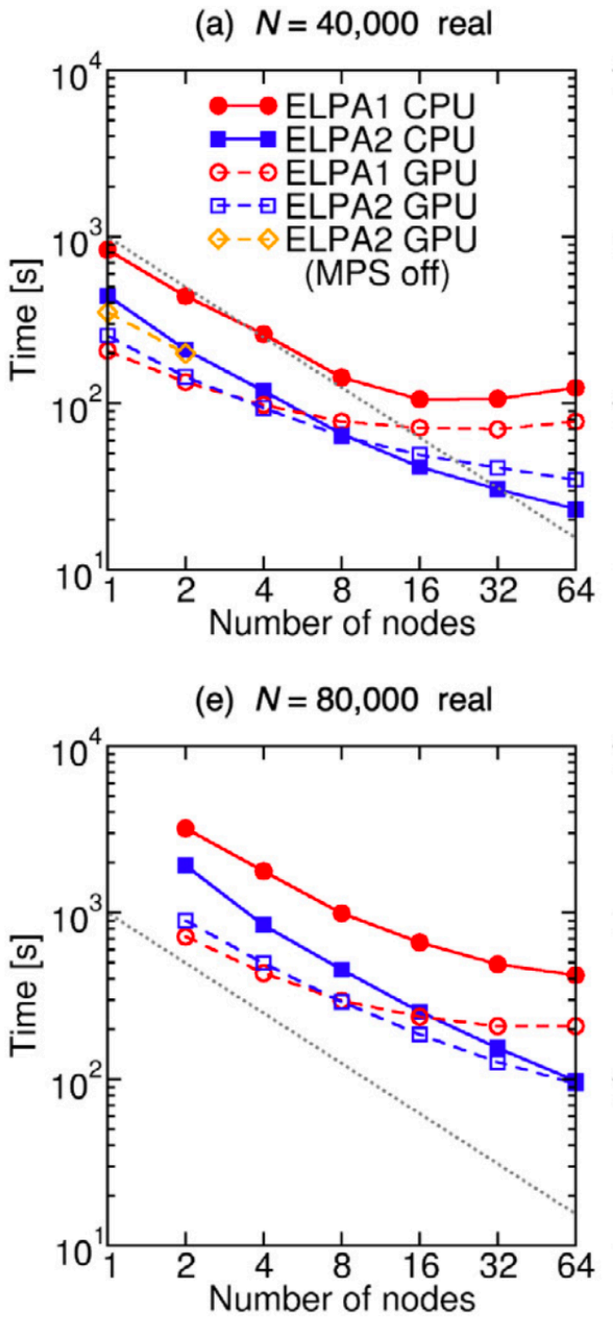

Summit (OLCF)
42 Power9 cores + 6 V100 / node

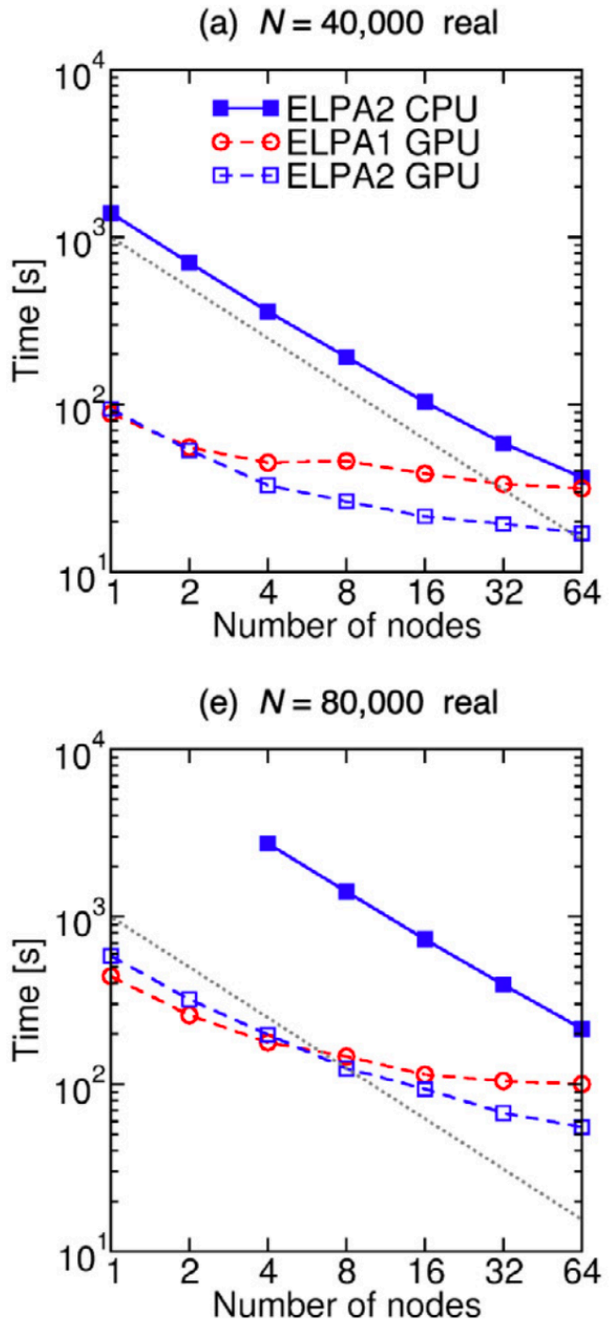

ELPA2-GPU Advantage for *M*<*N*

Summit (OLCF), 1 node

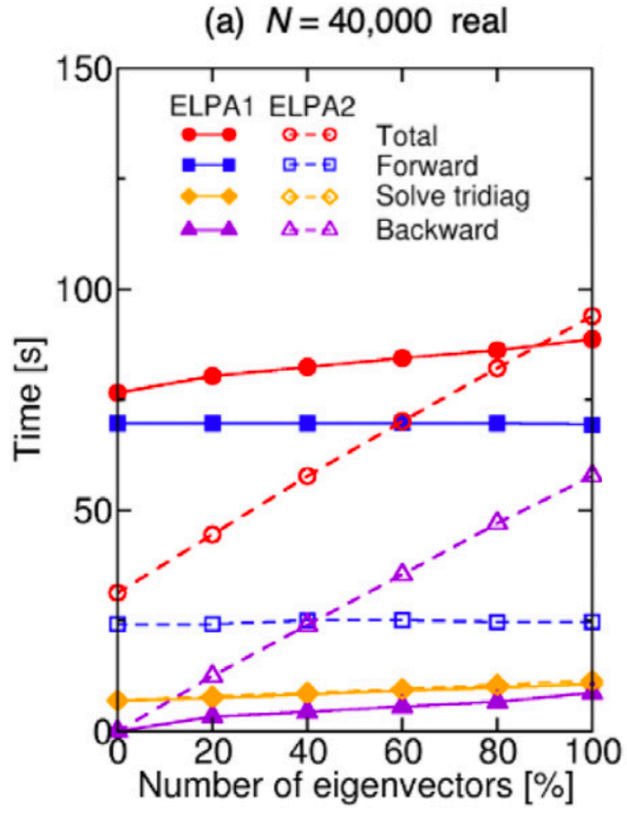

FHI-aims+ELSI Coupled with Google TPUs

JCTC Journal of Chemical Theory and Computation

J. Chem. Theory Comput. 2023, 19, 1, 25–32

pubs.acs.org/JCTC Letter

Large Scale Quantum Chemistry with Tensor Processing Units

Ryan Pederson,* John Kozlowski, Ruyi Song, Jackson Beall, Martin Ganahl, Markus Hauru, Adam G. M. Lewis, Yi Yao, Shrestha Basu Mallick, Volker Blum, and Guifre Vidal

Cite This: https://doi.org/10.1021/acs.jctc.2c00876 Read Online

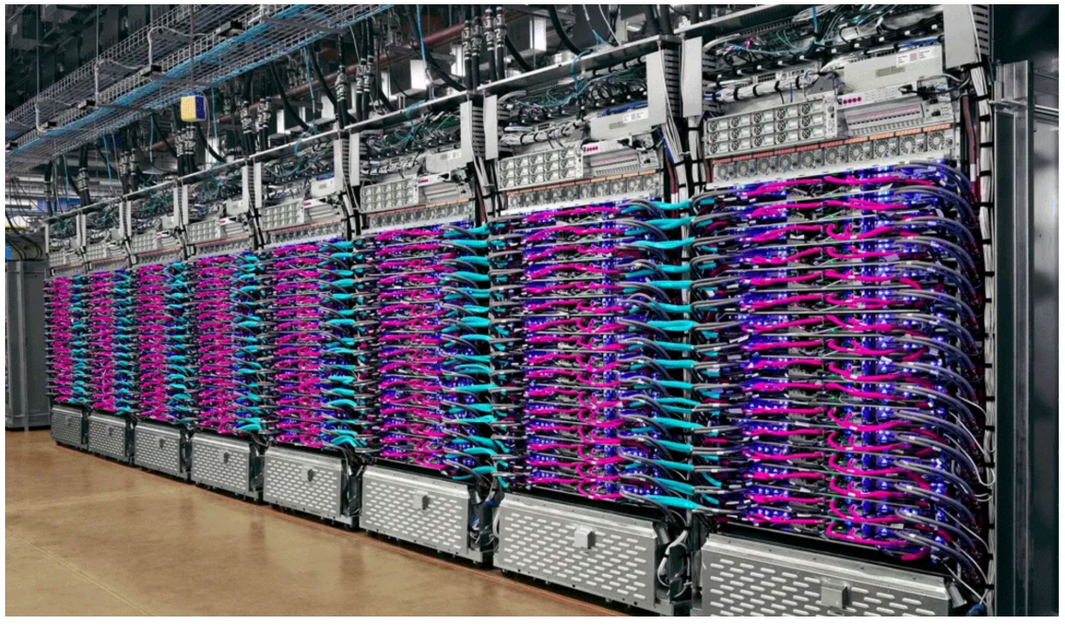

Google TPU pod v3, > 100 PFlops

Hamiltonian build

cost O(N)-O(N^2)

↔

Density matrix purification

cost O(N^3)

run on **CPU**s

FHI-aims & ELSI

run on **TPU**s

JAX & distributed linear algebra library

Prototype: Disk transfer of ELSI CSC stored matrices

FHI-aims+ELSI Coupled with Google TPUs

J. Chem. Theory Comput. 2023, 19, 1, 25–32

H_2O clusters up to 30,981 atoms

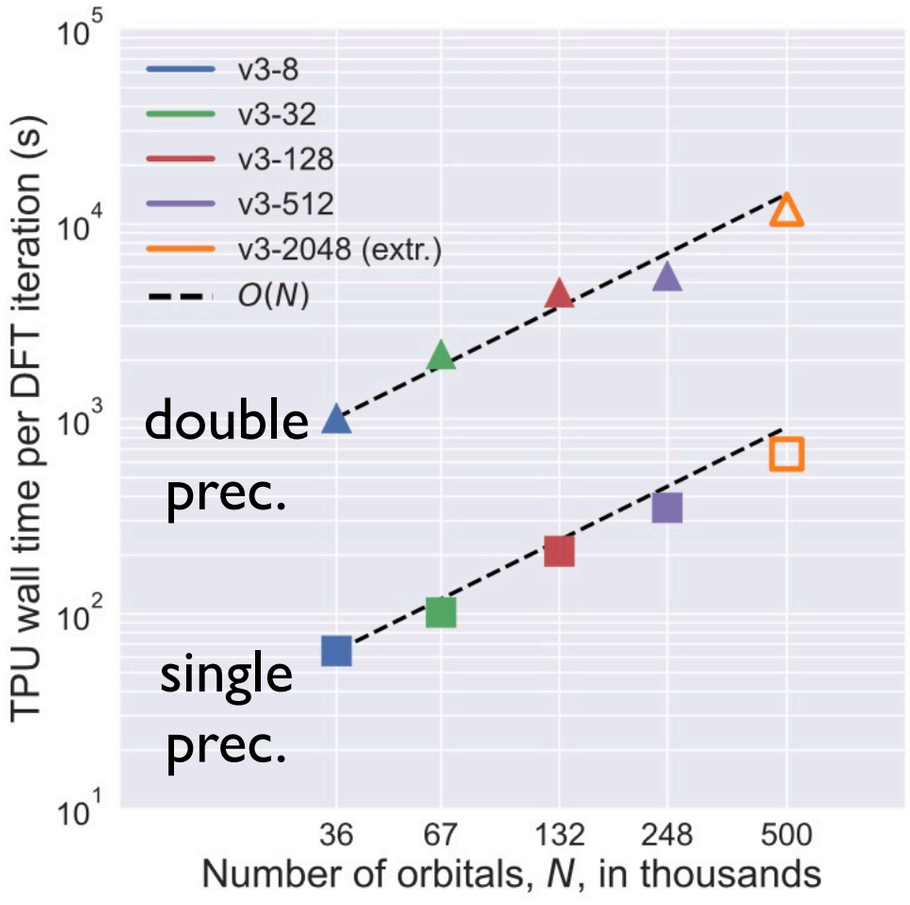

FIG. 1. TPU v3 wall time for $O(N^3)$ density matrix purification, Eqs. (5)-(7), as a function of the number N of orbitals, for clusters of water molecules, both in single (squares) and double (triangles) precision. A full TPU v3 pod with 2048 cores can handle $N \sim 500\,000$ orbitals (extrapolated results).

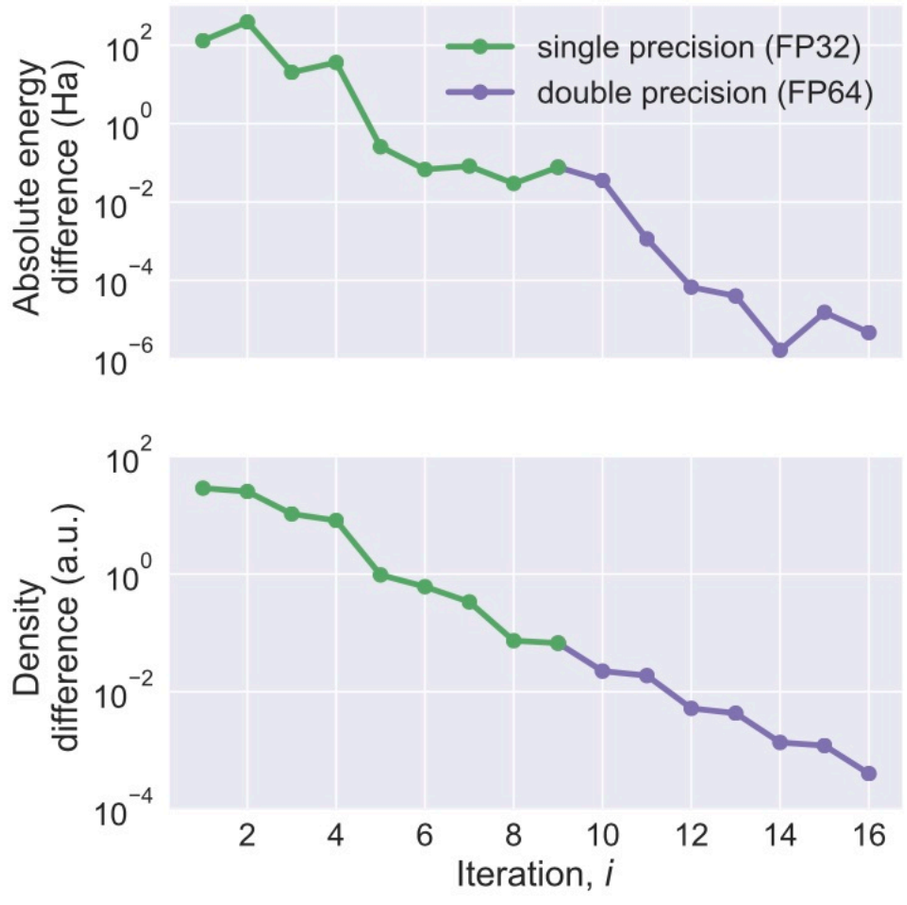

FIG. 4. Convergence trajectory of an end-to-end dynamic precision DFT calculation on a $(H_2O)_{10327}$ cluster. The absolute total energy differences between subsequent DFT iterations, i and $i-1$, are plotted (top). The corresponding difference in real-space densities within the L^1 norm is plotted (bottom).

Solvers Aren't Everything - Far From Exact Theory for Large Systems

DFT: So-called "band gap" problem.

Actually, quasiparticle states need many-body theory, not DFT - but not affordable beyond (maybe) few hundreds of atoms.

Many substitute hybrid DFT instead.

SCIENTIFIC DATA

Check for updates

OPEN

DATA DESCRIPTOR

A band-gap database for semiconducting inorganic materials calculated with hybrid functional

Sangtae Kim[1,3], Miso Lee[1,3], Changho Hong[1], Youngchae Yoon[1], Hyungmin An[1], Dongheon Lee[1], Wonseok Jeong[1], Dongsun Yoo[1], Youngho Kang[2], Yong Youn[1,4]✉ & Seungwu Han[1,4]✉

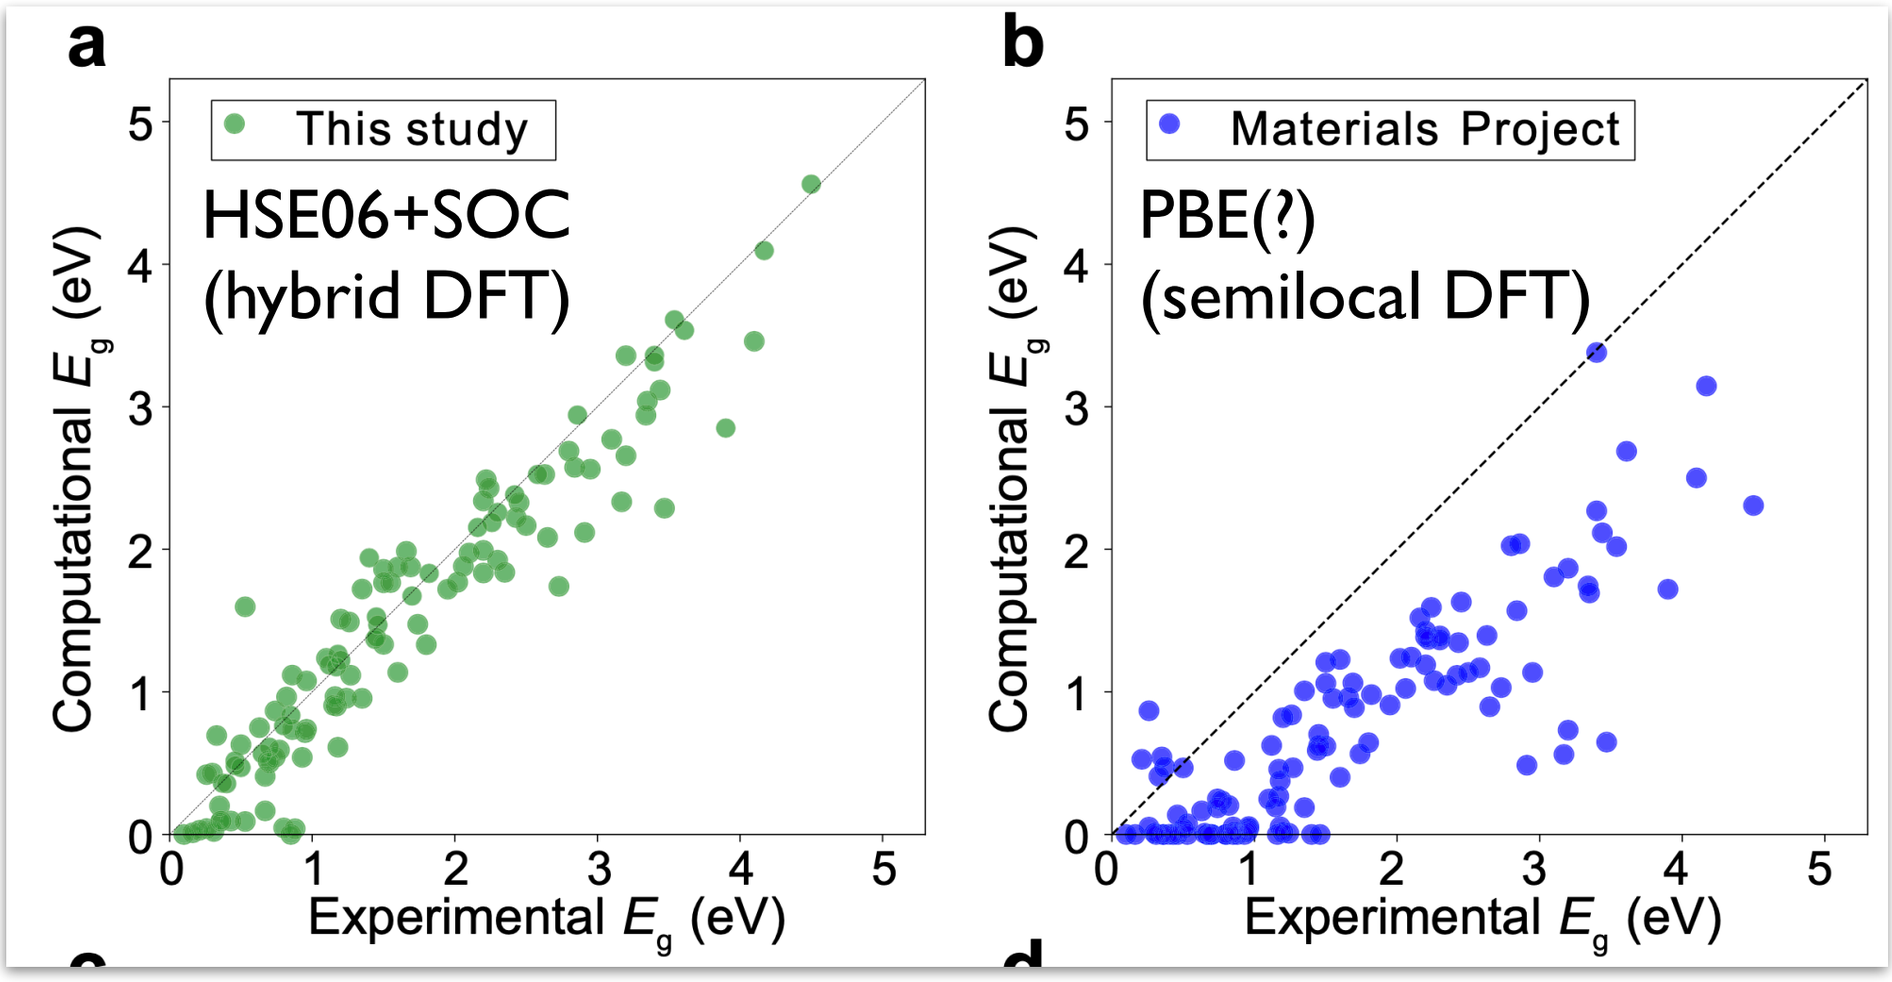

Localized "Resolution of Identity" (RI) for Two-Electron Terms

$$(ij|kl) = \int d^3r d^3r' \frac{\varphi_i(\boldsymbol{r})\varphi_j(\boldsymbol{r}')\varphi_k(\boldsymbol{r})\varphi_l(\boldsymbol{r}')}{|\boldsymbol{r}-\boldsymbol{r}'|}$$

$$\varphi_i(\boldsymbol{r})\varphi_j(\boldsymbol{r}) = \sum_\mu C_{ij}^\mu P_\mu(\boldsymbol{r})$$

$$C_{ij}^\mu = \sum_\nu (ij|\nu) V_{\nu\mu}^{-1}$$

Problem: Full RI-V delocalizes *C* across entire system

In FHI-aims: *Ihrig, Wieferink, Zhang, Ropo, Ren, Rinke, Scheffler, Blum, New J. Phys.* ***17****, 093020 (2015)*

Localized "Resolution of Identity" (RI) for Two-Electron Terms

$$(ij|kl) = \int d^3r d^3r' \frac{\varphi_i(\boldsymbol{r})\varphi_j(\boldsymbol{r}')\varphi_k(\boldsymbol{r})\varphi_l(\boldsymbol{r}')}{|\boldsymbol{r}-\boldsymbol{r}'|}$$

$$\varphi_i(\boldsymbol{r})\varphi_j(\boldsymbol{r}) = \sum_\mu C_{ij}^\mu P_\mu(\boldsymbol{r})$$

$$C_{ij}^\mu = \sum_\nu (ij|\nu) V_{\nu\mu}^{-1}$$

Problem: Full RI-V delocalizes *C* across entire system

Solution: For each $C_{ij}{}^{\mu}$, restrict μ only to atoms *I* and *J* at which *i* and *j* are centered.

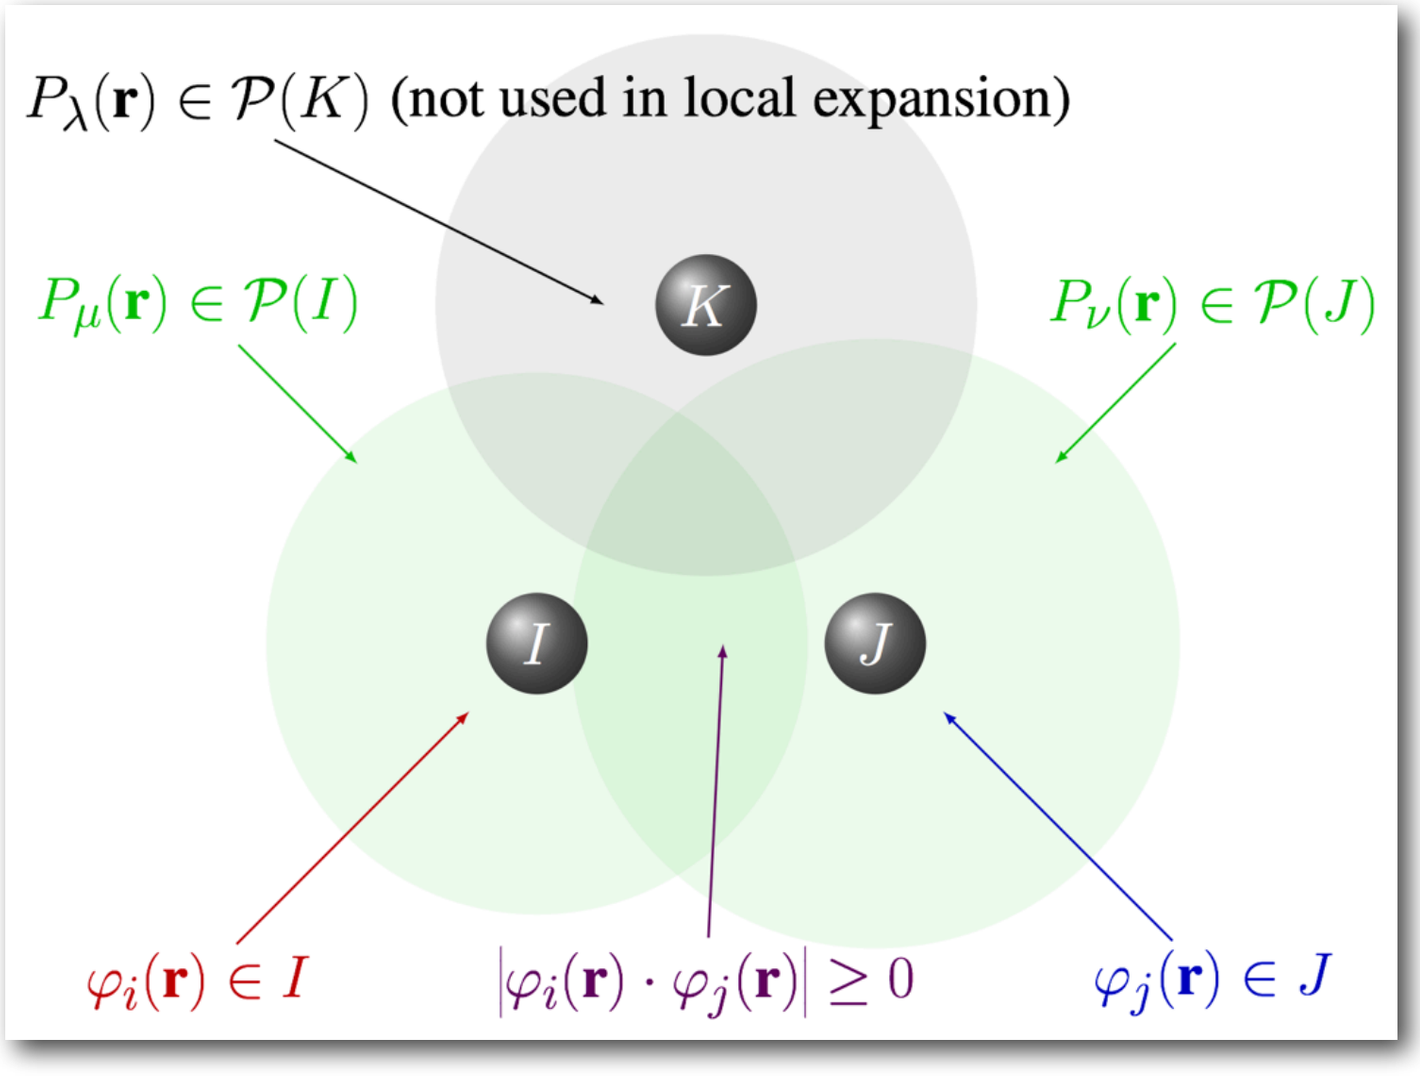

In FHI-aims: *Ihrig, Wieferink, Zhang, Ropo, Ren, Rinke, Scheffler, Blum, New J. Phys.* ***17****, 093020 (2015)*

Hybrid DFT in FHI-aims 2015

*Ihrig, Wieferink, Zhang, Ropo, Ren, Rinke, Scheffler, Blum, New J. Phys. **17**, 093020 (2015)*
*Levchenko, Ren, Wieferink, Rinke, Johanni, Blum, Scheffler, Comp. Phys. Commun. **192**, 60-69 (2015).*

O(N) scaling implementation, localized resolution of identity.

Note: Localized RI also works for MP2, RPA, GW, etc. (but not O(N))

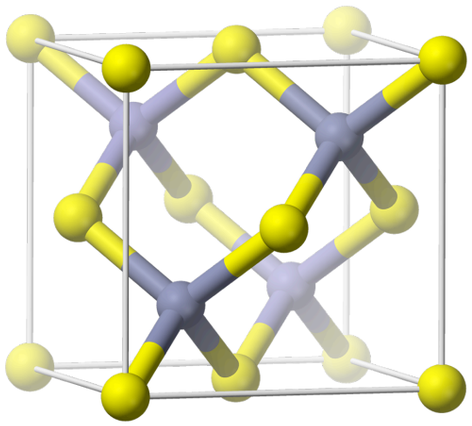

Zincblende GaAs

Computational Scaling of Periodic GaAs, "tight" HSE06 Hybrid Density-Functional Theory

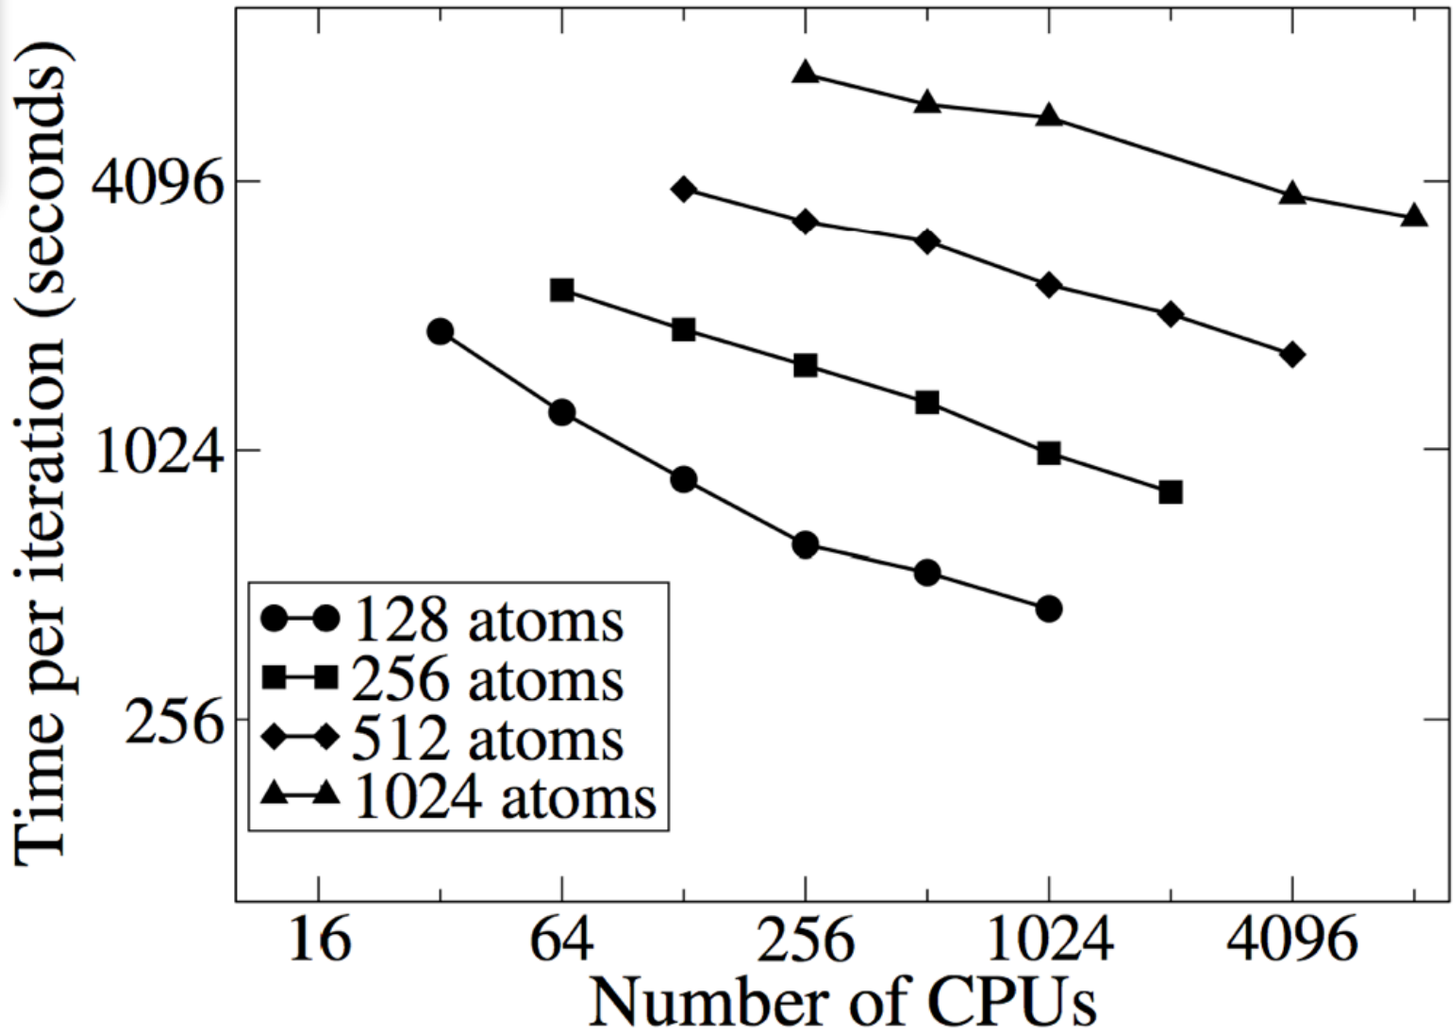

… proved incredibly useful, but not enough. Can we be faster, larger?

Hybrid DFT in FHI-aims 2023

Florian Merz[1], Andreas Marek[2], **Sebastian Kokott**[3], Christian Carbogno[3], Yi Yao[3,5], Mariana Rossi[4], Markus Rampp[2], Matthias Scheffler[3], and Volker Blum[5]

Charge Localization

Saline Liquide = $Na^+ + Cl^- + (H_2O)_{360}$
PBE0(α=0.5), T=300K
In green: Highest occ. state localized at Cl

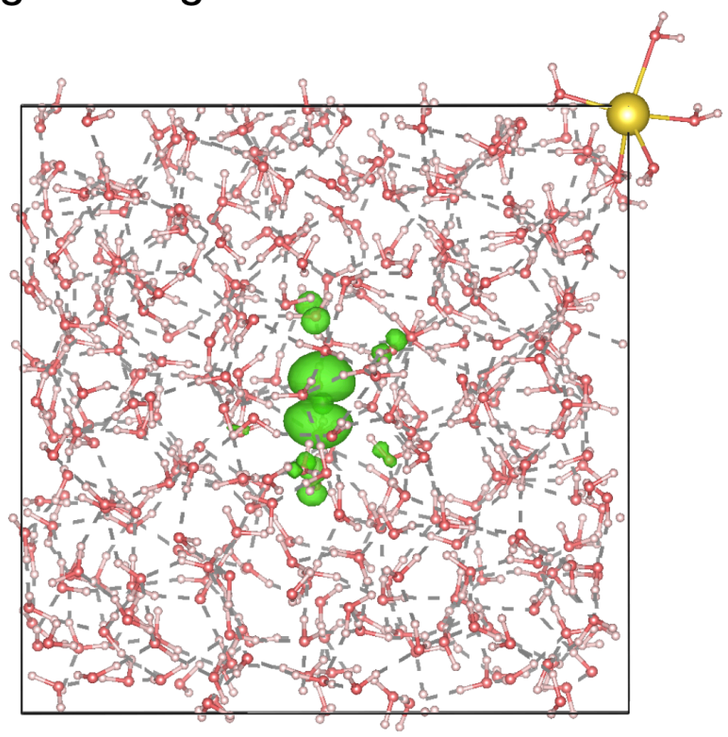

Band Gaps/Defect Levels

phenylethylammonium lead iodide

$PEA_{64}(Pb_{29}Bi_2\square)I_{128}$

=1503 atoms

HSE06 + pSOC

Haipeng Lu, Gabrielle Koknat et al., submitted (2022)

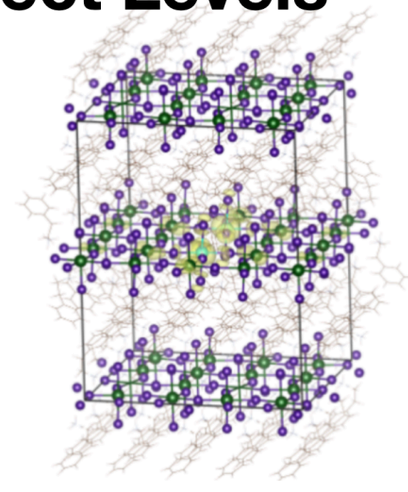

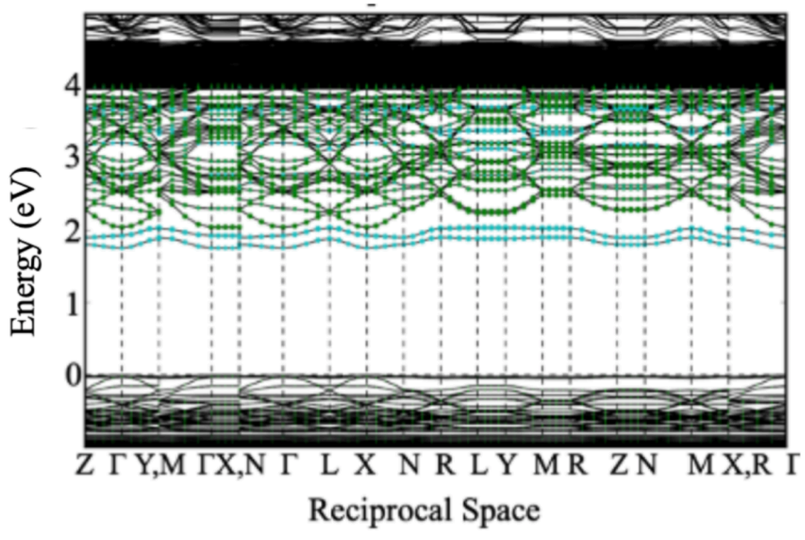

[1]*Lenovo HPC Innovation Center, Stuttgart*
[2]*Max Planck Computing and Data Facility, Garching*
[3]*The NOMAD Laboratory at the FHI-MPG and IRIS-HU, Berlin*
[4]*MPI for the Structure and Dynamics of Matter, Hamburg*
[5]*Duke University, Durham, North Carolina, USA*

Molecular Simulations from First Principles e.V.

Code Improvements (Hard Work But Effective)

$$(ij|kl) = \int d^3r d^3r' \frac{\varphi_i(\boldsymbol{r})\varphi_j(\boldsymbol{r}')\varphi_k(\boldsymbol{r})\varphi_l(\boldsymbol{r}')}{|\boldsymbol{r}-\boldsymbol{r}'|}$$

$$\varphi_i(\boldsymbol{r})\varphi_j(\boldsymbol{r}) = \sum_\mu C_{ij}^\mu P_\mu(\boldsymbol{r})$$

$$C_{ij}^\mu = \sum_\nu (ij|\nu) V_{\nu\mu}^{-1}$$

CPUs only!

FHI-aims
The ab initio materials simulation package

- MPI-3 intra-node shared memory arrays
  - Facilitate data redistribution via one-sided MPI routines; e.g. Coulomb matrix
- Compression of the RI-coefficients; exploiting sparsity in real-space
  - Same screening mechanism as for Coulomb matrix (2015)
- Auto-tuning mechanism for blocking the data
  - Minimize communication, maximize memory/node usage
- Additional parallelization layers
  - Finer granulation on the level of basis functions (previously: on the level of atoms)
- Auto-tuning for an optimal workload balance

Memory Optimizations

GaAs supercell, 512 atoms: HSE06, tight settings

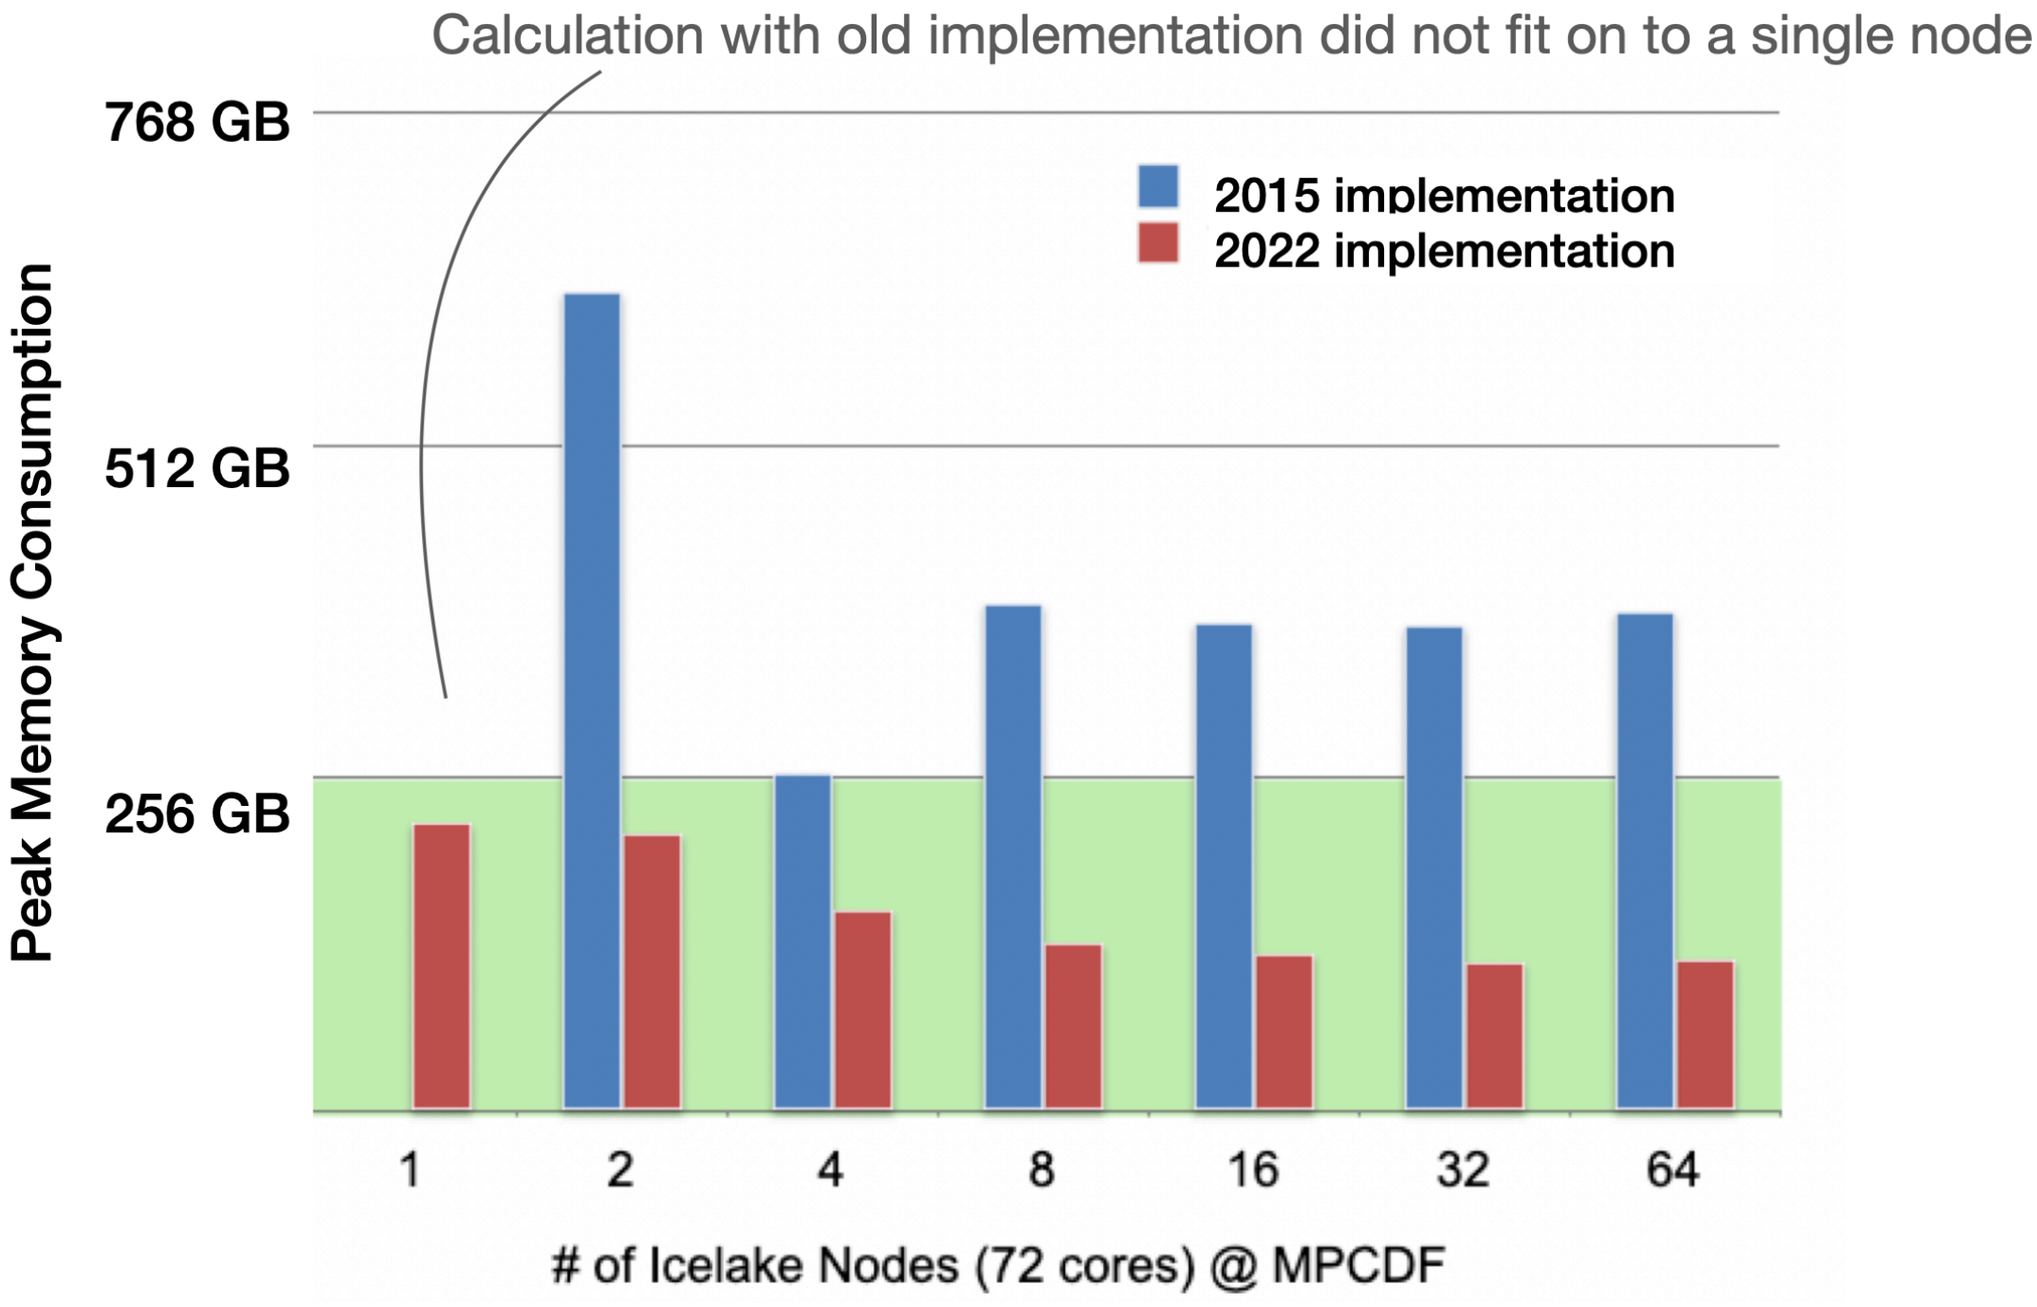

Drastic Scalability Improvements*

GaAs supercells: HSE06, tight settings

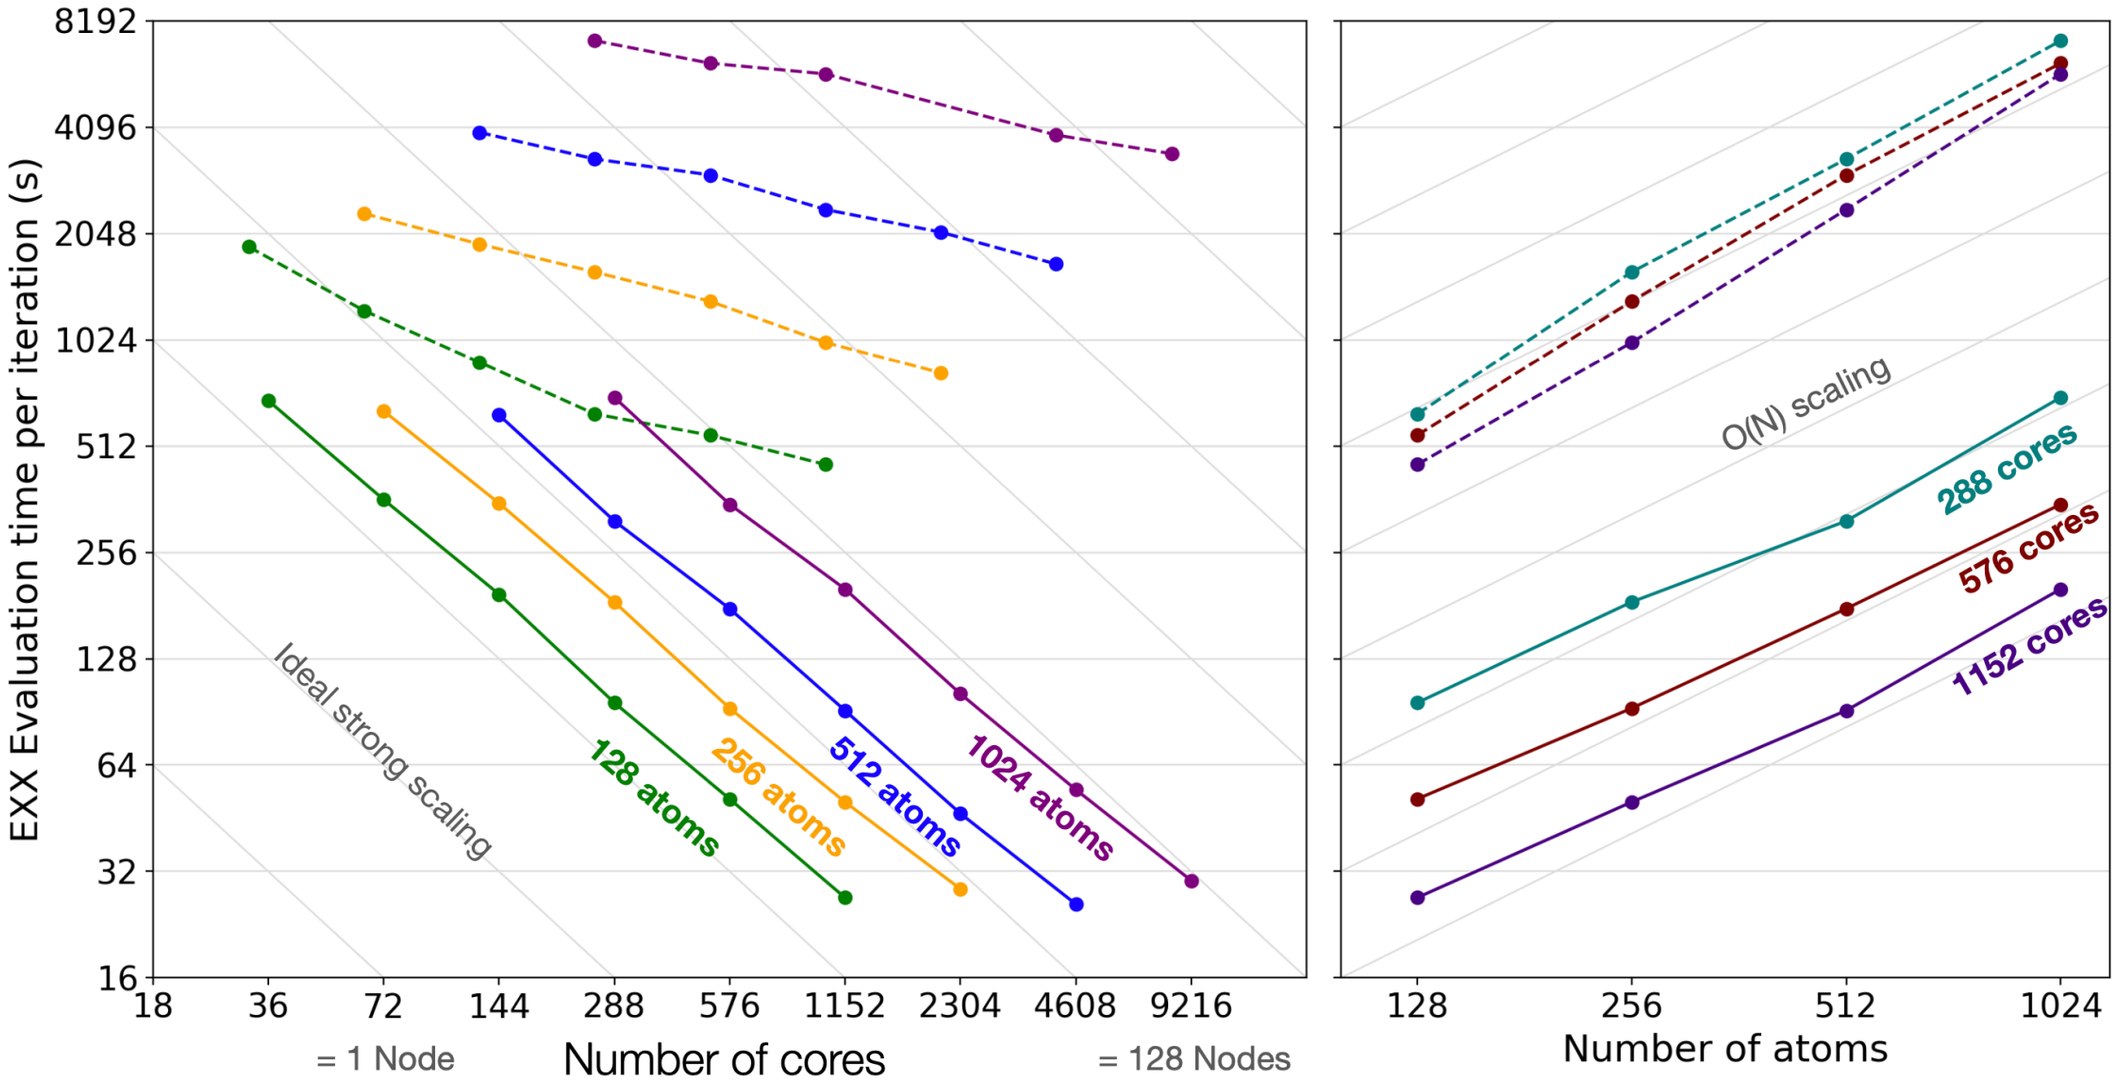

Hydra @ MPCDF
Intel Sandy Bridge
32 cores / node

Raven @ MPCDF
Intel Xeon IceLake
(Platinum 8360Y)
72 cores / node
256 GB / node

MAX PLANCK
COMPUTING & DATA FACILITY

*Not so simple. Our 2015 implementation had already seen some years of work by several very smart and focused individuals.

Scalability: Hybrid DFT for Very Large Systems

Raven @ MPCDF
Intel Xeon IceLake
(Platinum 8360Y)
72 cores / node
256 GB / node

MAX PLANCK
COMPUTING & DATA FACILITY

Saline Liquide (PBE0)
Intermediate species defaults
1,082 atoms
576 cores @ 8 nodes
18,780 basis functions

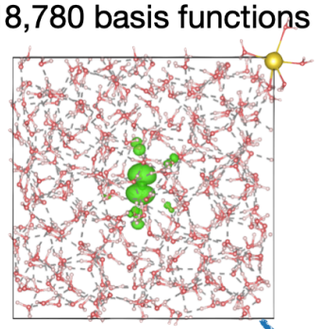

(6x6) PEPI $Bi_2\square$
Intermediate species defaults
3,383 atoms
2,304 cores @ 32 nodes
75,100 basis functions

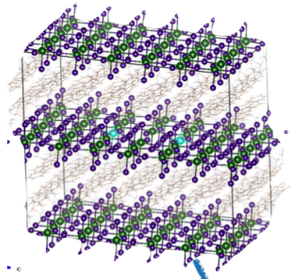

Scalability: Hybrid DFT for Very Large Systems

Raven @ MPCDF
Intel Xeon IceLake
(Platinum 8360Y)
72 cores / node
256 GB / node

MAX PLANCK
COMPUTING & DATA FACILITY

Saline Liquide (PBE0)
Intermediate species defaults
1,082 atoms
576 cores @ 8 nodes
18,780 basis functions

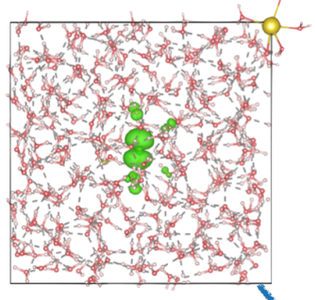

(6x6) PEPI $Bi_2\square$
Intermediate species defaults
3,383 atoms
2,304 cores @ 32 nodes
75,100 basis functions

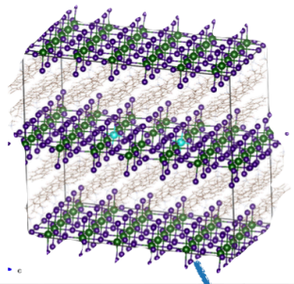

Paracetamol Supercell
Light species defaults
10,240 atoms
2,304 cores @ 32 nodes
101,888 bas. functions

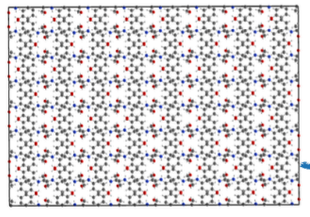

Scalability: Hybrid DFT for Very Large Systems

Raven @ MPCDF
Intel Xeon IceLake
(Platinum 8360Y)
72 cores / node
256 GB / node

MAX PLANCK
COMPUTING & DATA FACILITY

Saline Liquide (PBE0)
Intermediate species defaults
1,082 atoms
576 cores @ 8 nodes
18,780 basis functions

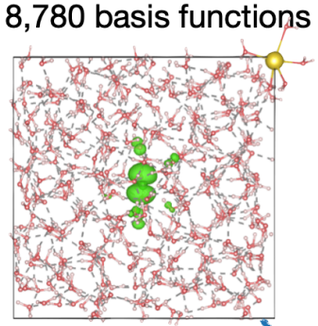

(6x6) PEPI $Bi_2\square$
Intermediate species defaults
3,383 atoms
2,304 cores @ 32 nodes
75,100 basis functions

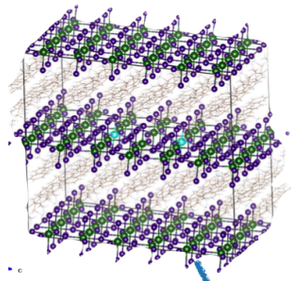

Ice XI $(H_20)_{5096}$
Light species defaults
15,288 atoms
2,304 cores @ 32 nodes
122,304 basis functions

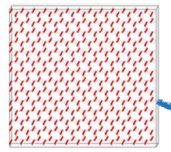

Paracetamol Supercell
Light species defaults
10,240 atoms
2,304 cores @ 32 nodes
101,888 bas. functions

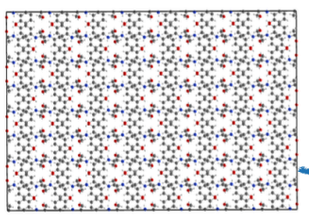

Scalability: Hybrid DFT for Very Large Systems

Raven @ MPCDF
Intel Xeon IceLake
(Platinum 8360Y)
72 cores / node
256 GB / node

MAX PLANCK
COMPUTING & DATA FACILITY

Saline Liquide (PBE0)
Intermediate species defaults
1,082 atoms
576 cores @ 8 nodes
18,780 basis functions

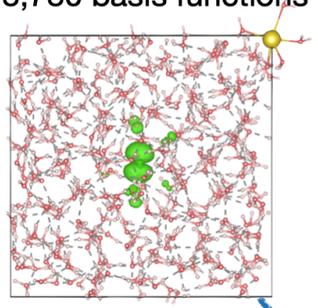

(6x6) PEPI $Bi_2\square$
Intermediate species defaults
3,383 atoms
2,304 cores @ 32 nodes
75,100 basis functions

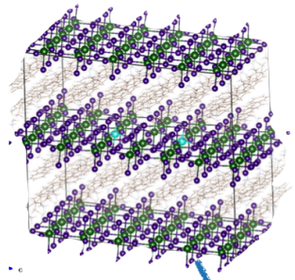

Ice XI $(H_20)_{5096}$
Light species defaults
15,288 atoms
4,608 cores @ 64 nodes
122,304 basis functions

HSE-X Including Force Evaluation

Ice XI $(H_20)_{5096}$
Light species defaults
15,288 atoms
2,304 cores @ 32 nodes
122,304 basis functions

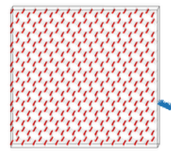

Paracetamol Supercell
Light species defaults
10,240 atoms
2,304 cores @ 32 nodes
101,888 bas. functions

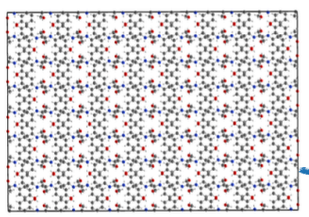

Scalability: Hybrid DFT for Very Large Systems

Raven @ MPCDF
Intel Xeon IceLake
(Platinum 8360Y)
72 cores / node
256 GB / node

MAX PLANCK
COMPUTING & DATA FACILITY

Saline Liquide (PBE0)
Intermediate species defaults
1,082 atoms
576 cores @ 8 nodes
18,780 basis functions

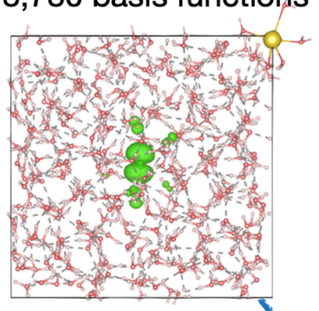

(6x6) PEPI $Bi_2\square$
Intermediate species defaults
3,383 atoms
2,304 cores @ 32 nodes
75,100 basis functions

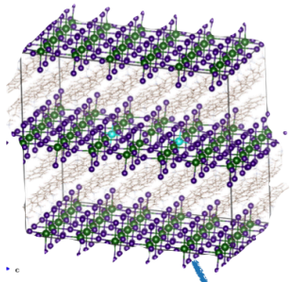

Ice XI $(H_20)_{5096}$
Light species defaults
15,288 atoms
4,608 cores @ 64 nodes
122,304 basis functions

HSE-X Including Force Evaluation

Ice XI $(H_20)_{5096}$
Light species defaults
15,288 atoms
2,304 cores @ 32 nodes
122,304 basis functions

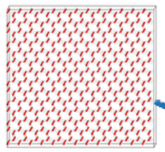

Ice XI $(H_20)_{10,192}$
Light species defaults
30,576 atoms
4,320 cores @ 60 nodes
244,608 basis functions

Paracetamol Supercell
Light species defaults
10,240 atoms
2,304 cores @ 32 nodes
101,888 bas. functions

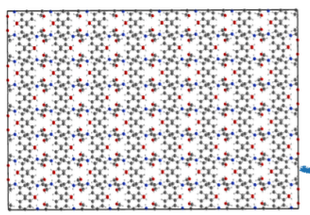

So Are We There Yet?

HSE06+SOC. No tricks.

Lu, Koknat et al., submitted (2022).

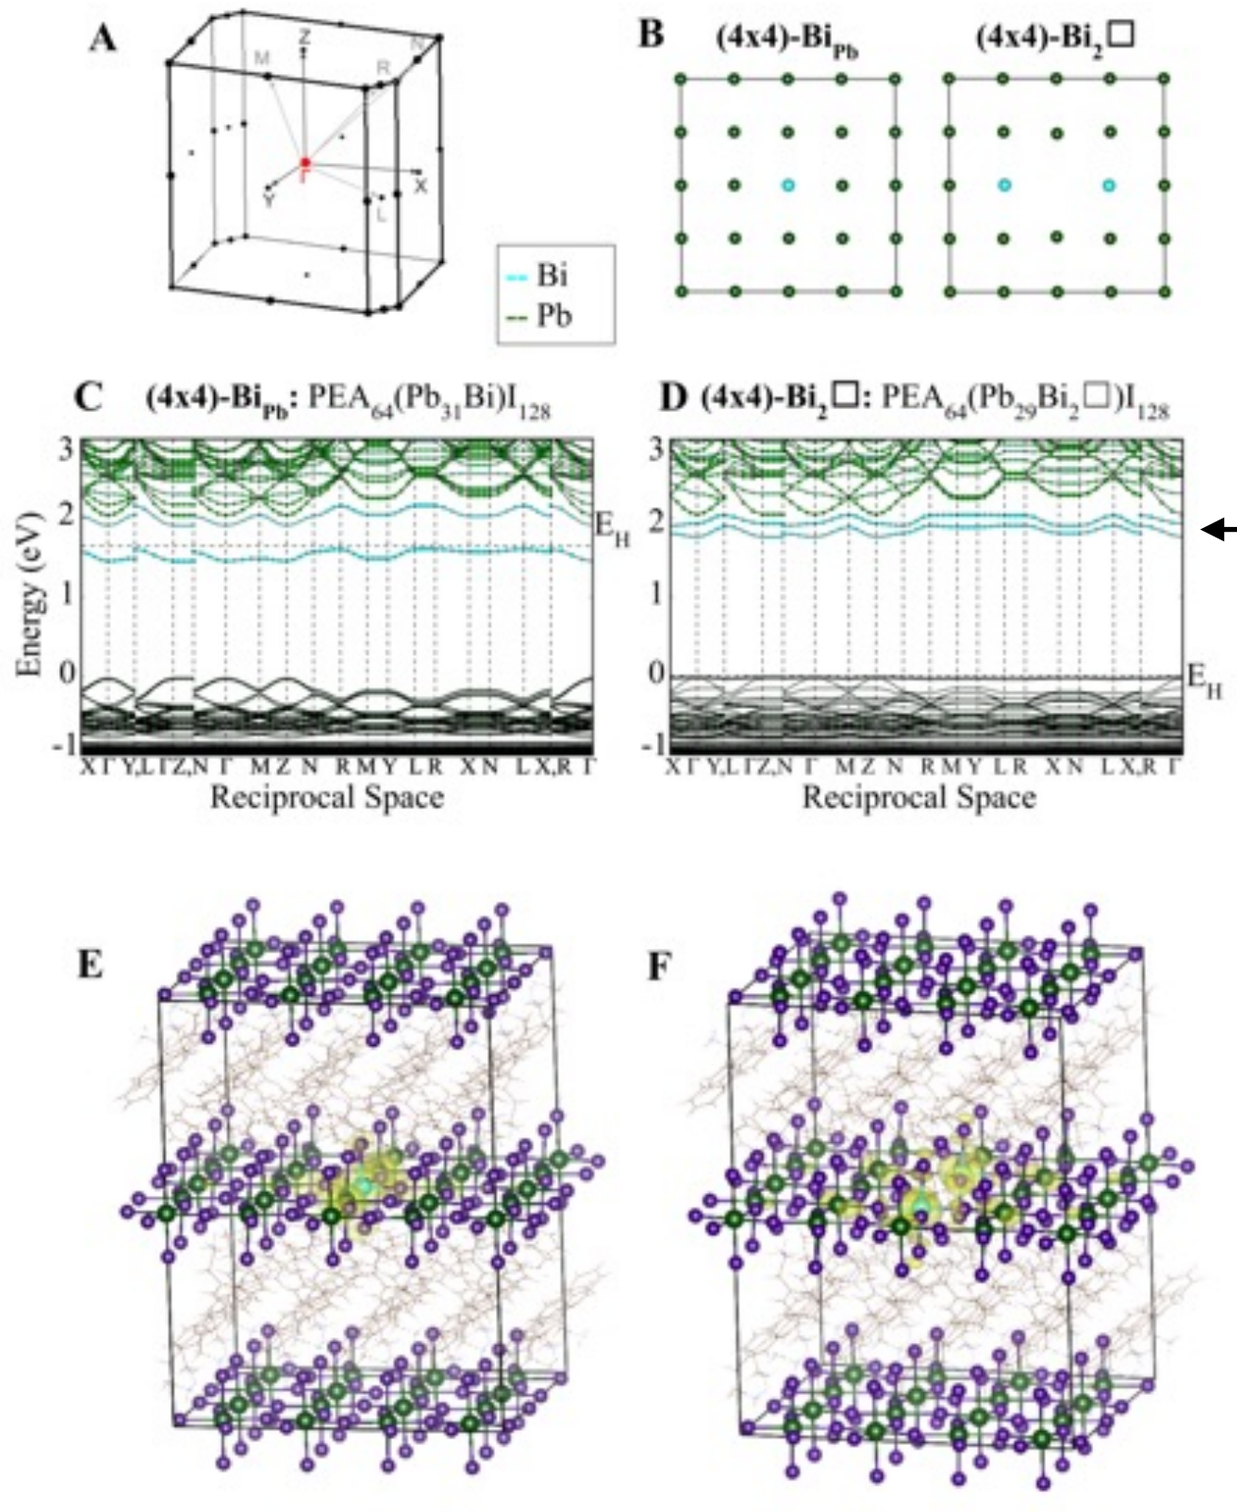

Anonymous reviewer:

"Finally, I actually worry that in some cases the supercell is not large enough despite its large number of atoms. In particular, for the Bi2 case shown in Fig. 3B, within the utilized supercell the Bi atoms are actually second-nearest neighbors with each other, likely indicating that the supercell is too small. A hint of this can be seen in Fig. 3D, which shows that the defect orbitals are not flat as expected but show some dispersion."

I'm not sure what to say.

So How About 3,383 Atoms.

Lu, Koknat et al., submitted (2023).

HSE06+SOC. Still no tricks.

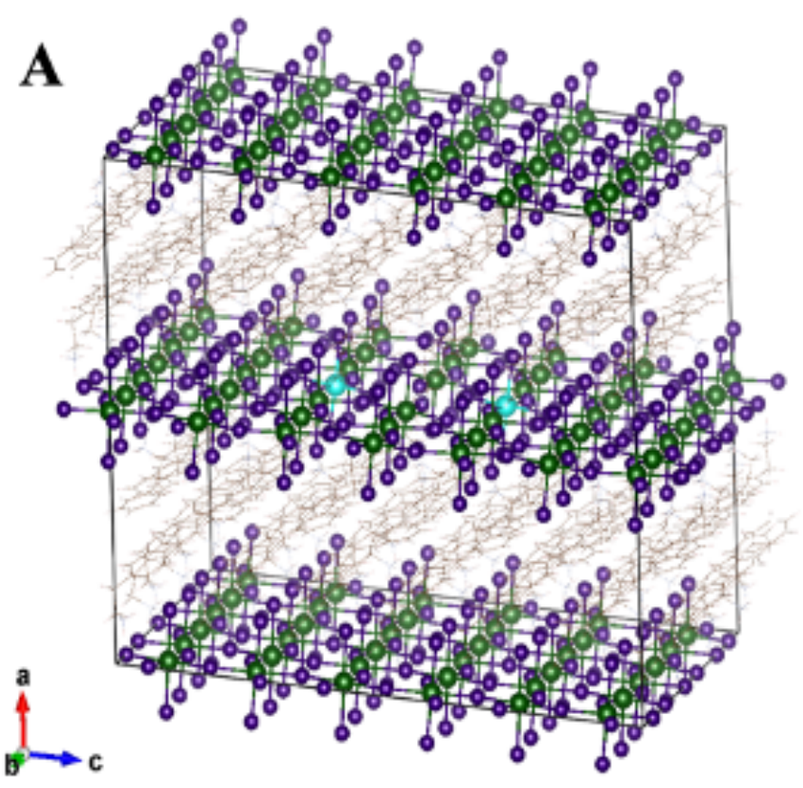

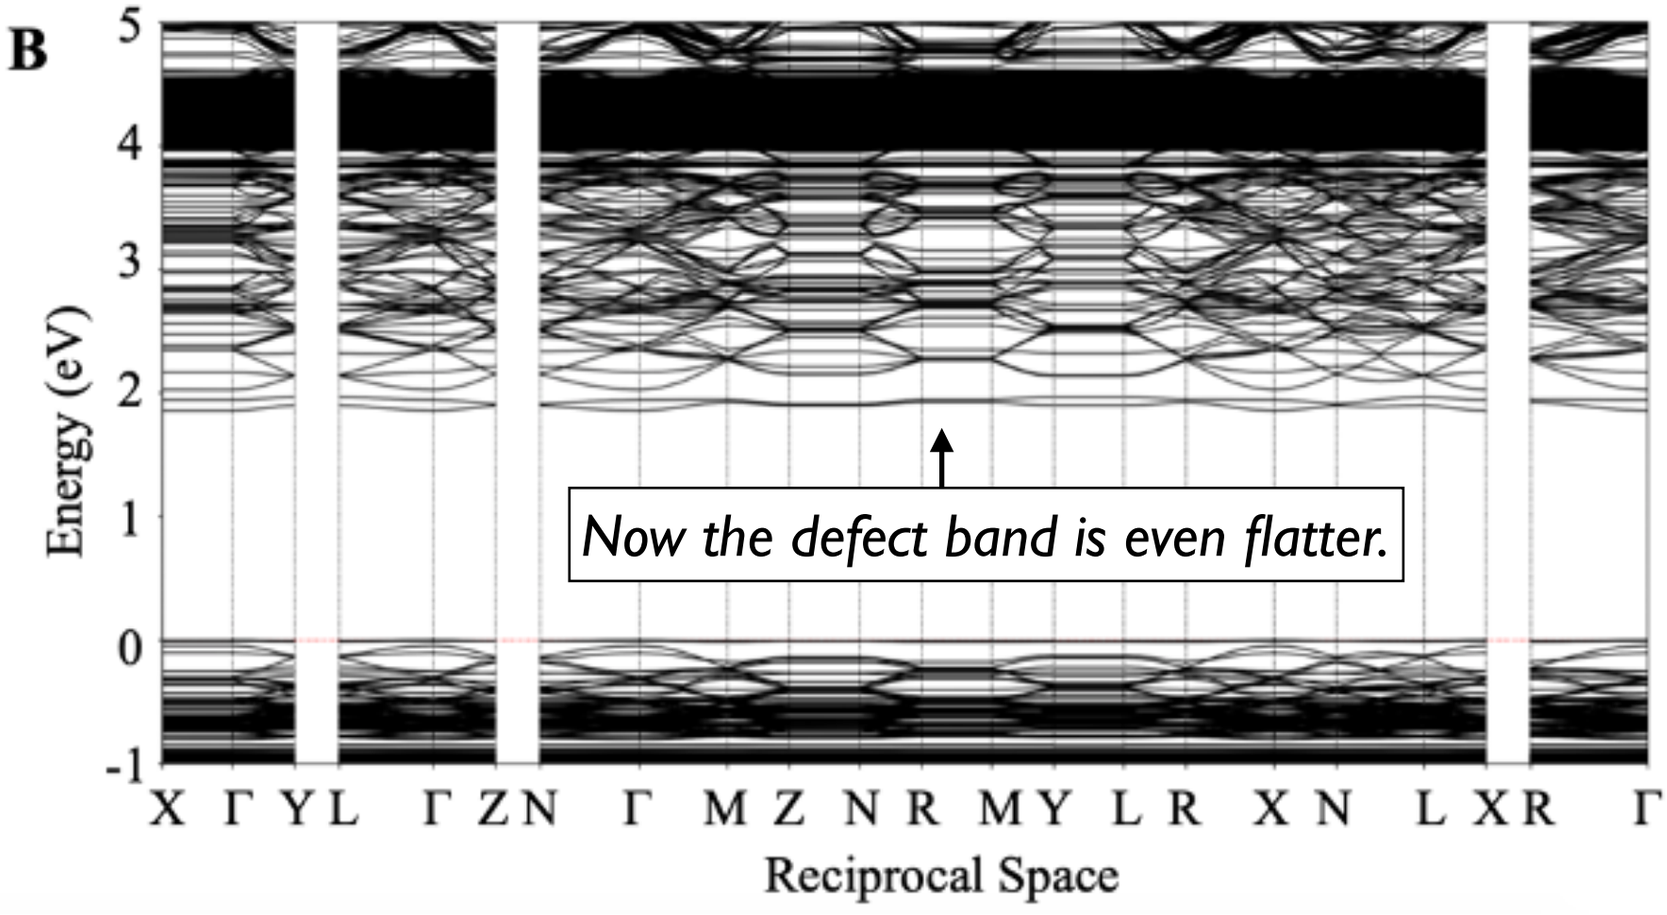

Electronic structure of the (6x6)-Bi_2□ structure.

Summary

FHI-aims
The ab initio materials simulation package

Efficient, scalable first-principles approach to structure, electron band structure, relativity in complex materials.

ELSI

ELPA

Significant investment in large scale eigenvalue / density matrix solver strategies.
Library infrastructure for electronic structure theory.

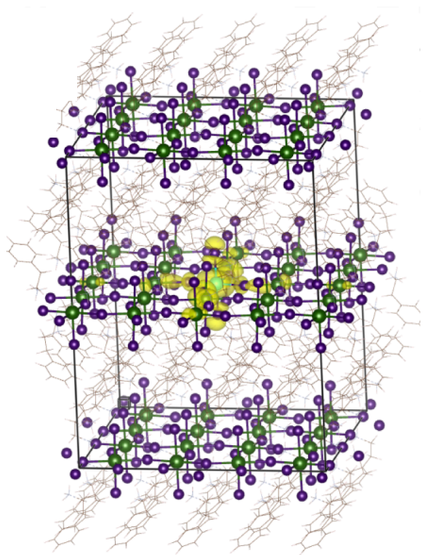

Very large-scale hybrid DFT, now feasible.

NSF

CHOISE
Center for Hybrid Organic Inorganic Semiconductors for Energy

U.S. DEPARTMENT OF ENERGY
Office of Science

U.S. DEPARTMENT OF ENERGY
INCITE
LEADERSHIP COMPUTING

Argonne NATIONAL LABORATORY
ALCF

XSEDE
Extreme Science and Engineering Discovery Environment

NERSC

NREL
Transforming ENERGY

Performance: Solver Benchmarks on Equal Footing

Victor Yu

http://www.nersc.gov/edison

http://www.nersc.gov/cori

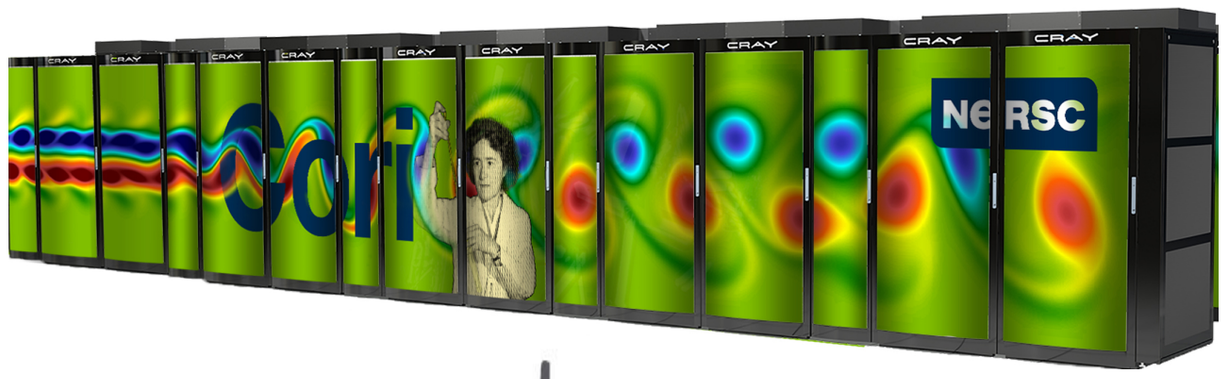

[Edison Cray XC30]

Processor: Intel Ivy Bridge
Interconnect: Cray Aries

5,586 compute nodes
134,064 processing cores
2.57 Petaflops

[Cori-Haswell Cray XC40]

Processor: Intel Haswell
Interconnect: Cray Aries

2,388 compute nodes
76,416 processing cores
2.81 Petaflops

[Cori-KNL Cray XC40]

Processor: Intel Knights Landing
Interconnect: Cray Aries

9,688 compute nodes
658,784 processing cores
29.5 Petaflops

Performance: Solver Benchmarks on Equal Footing

1D	2D	3D
(a) Ge	(b) MoS_2	(c) Cu_2BaSnS_4

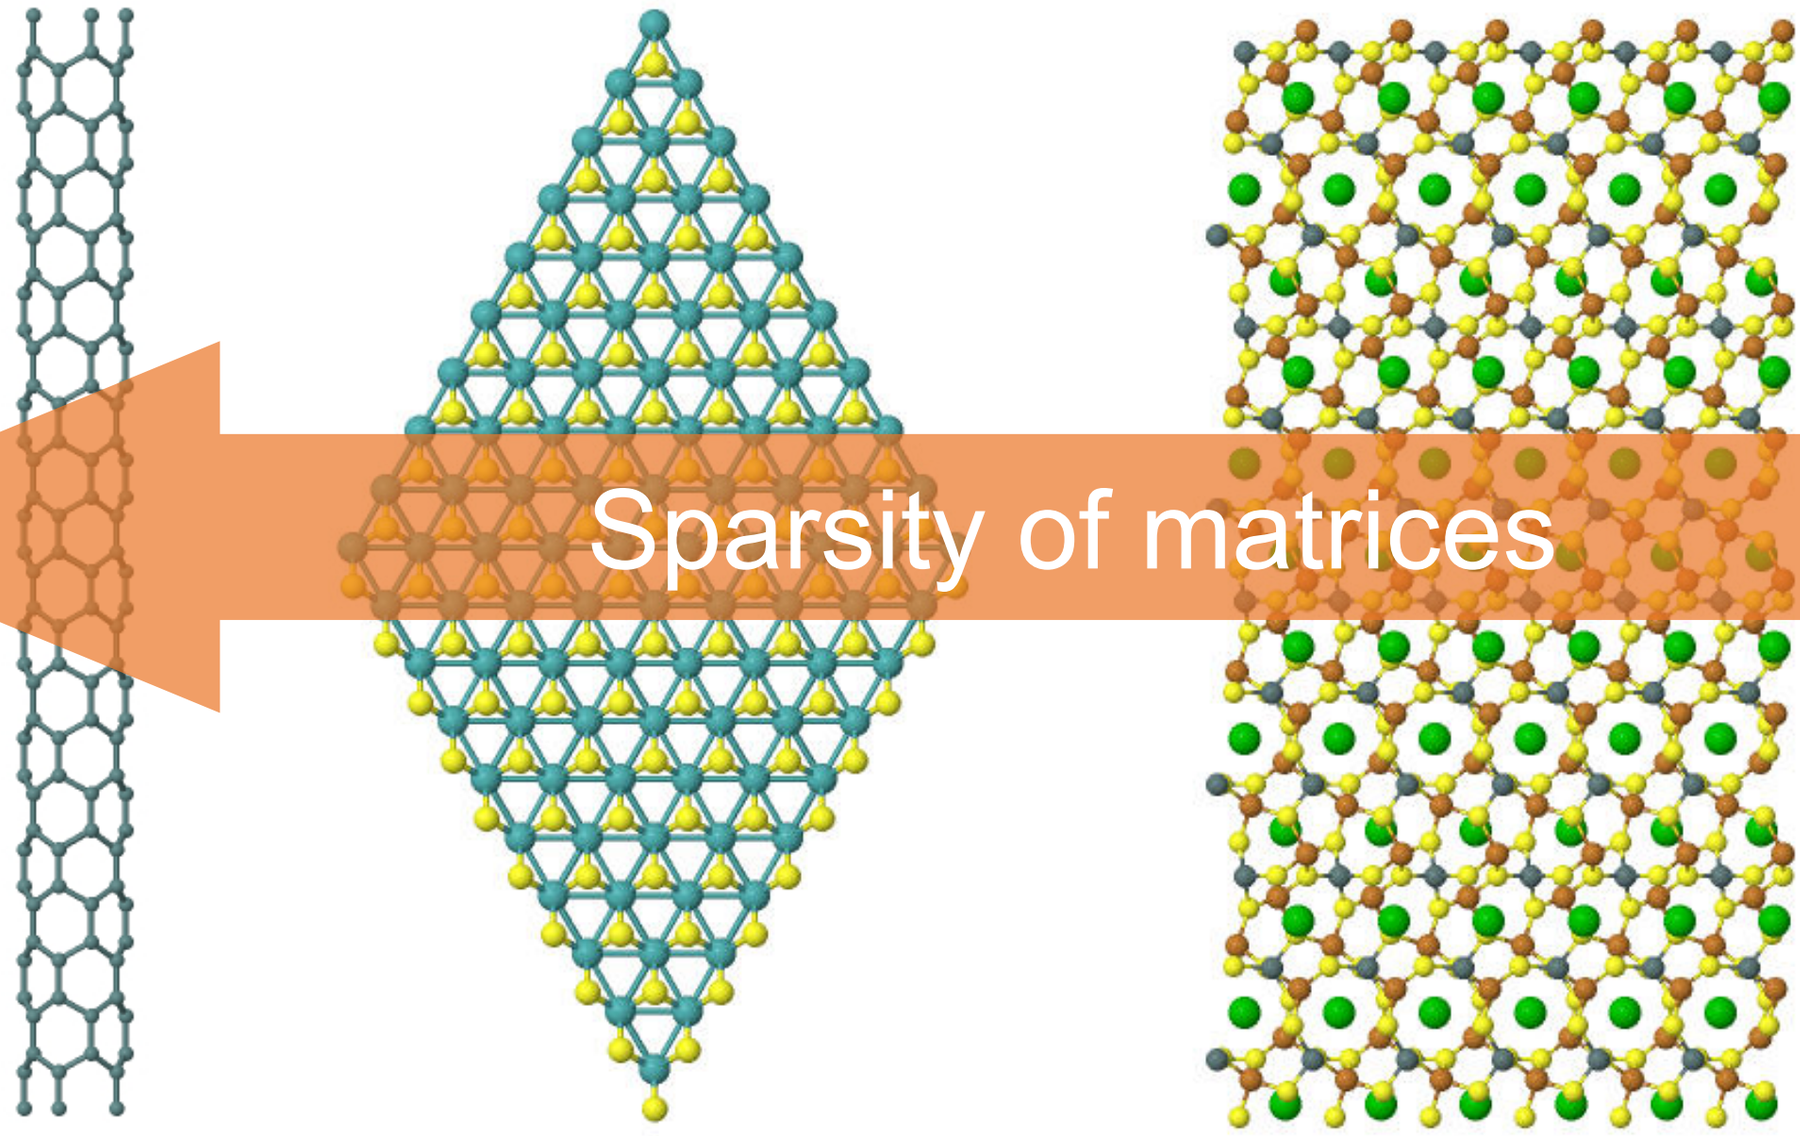

Sparsity of matrices

Example: FHI-aims Basis Sets - ELPA vs. PEXSI

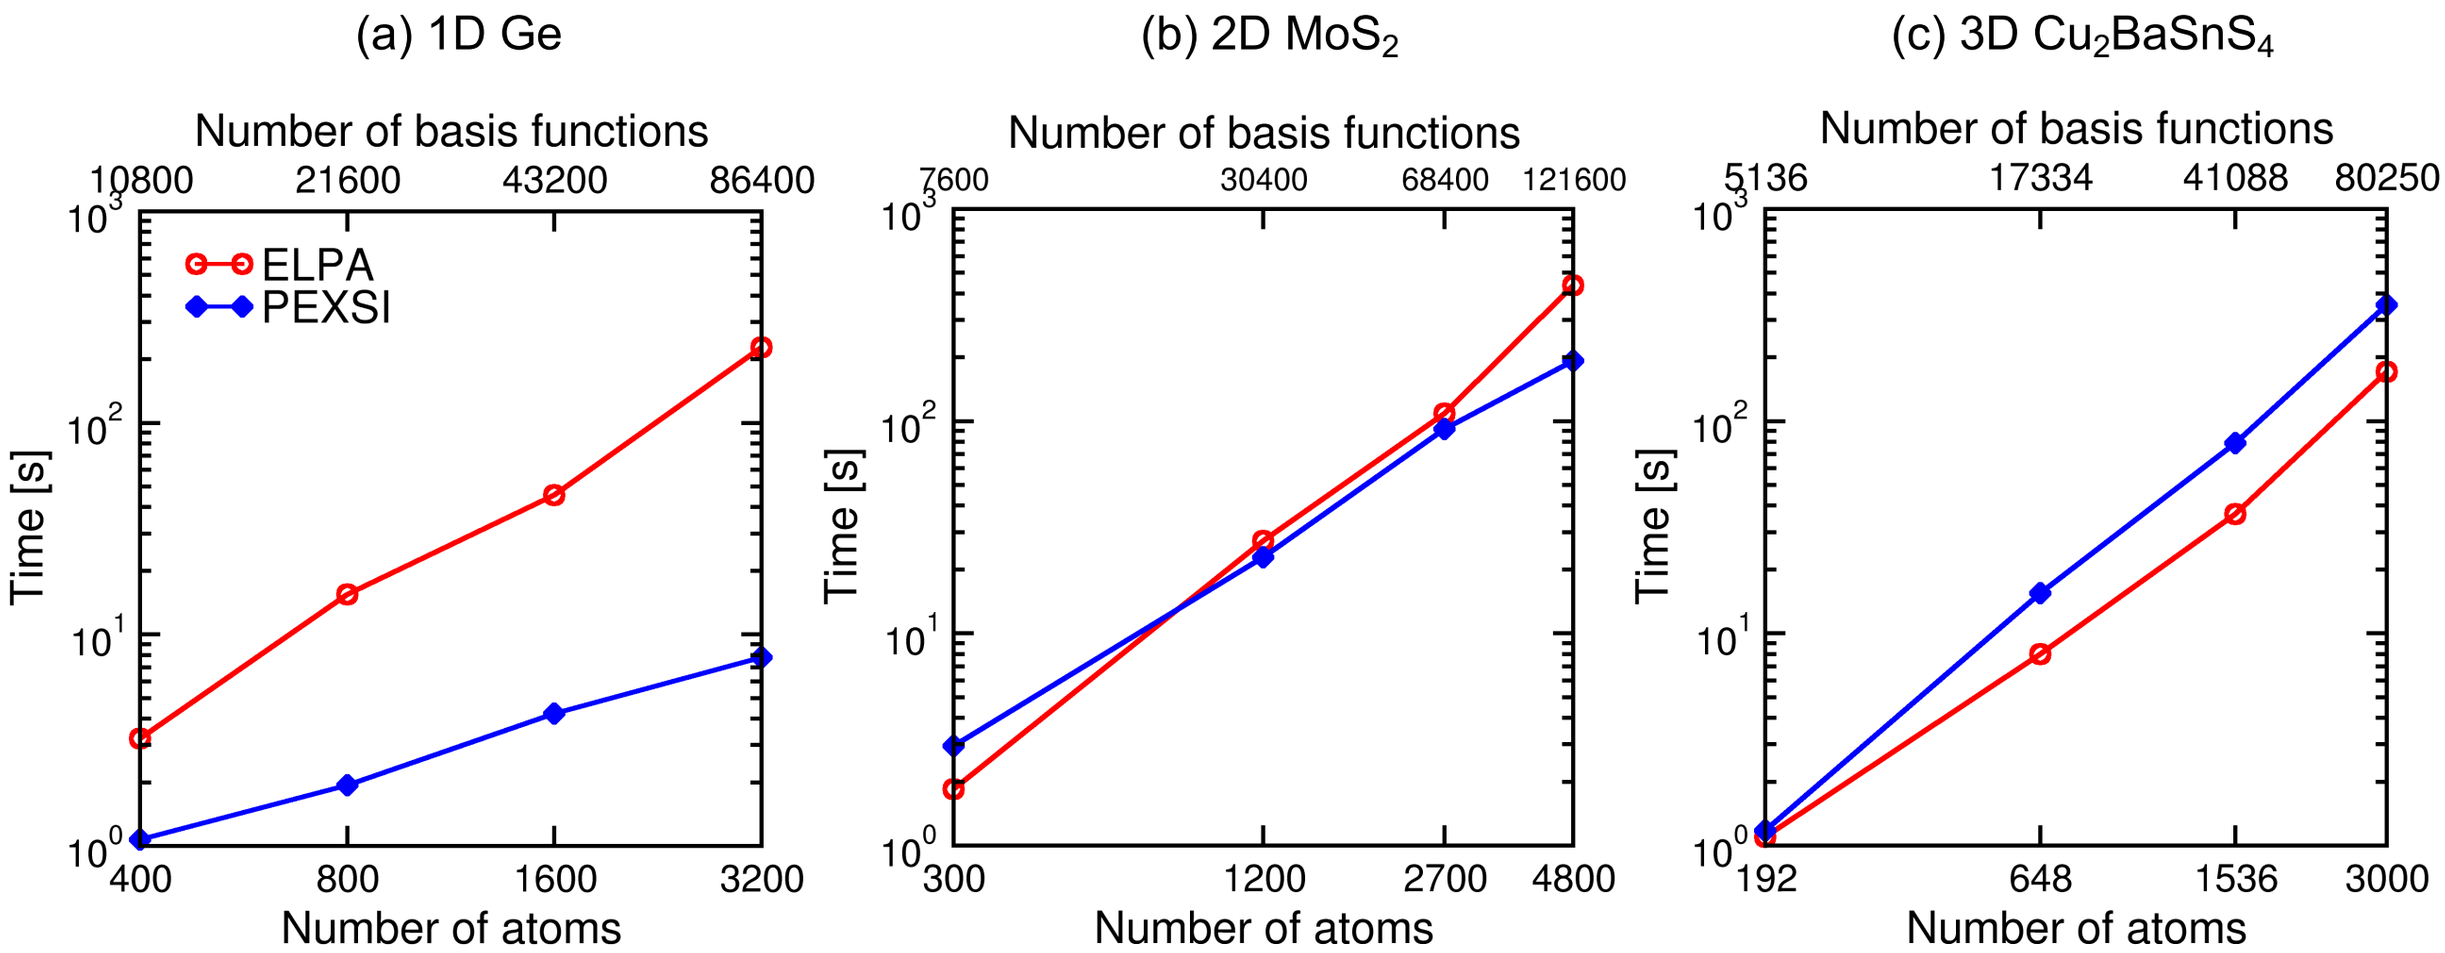

- DFT-PBE
- FHI-aims (all-electron)
- 2,560 CPU cores on Cori-Haswell

PEXSI faster for large low-dimensional (sparse) systems

PEXSI: Semilocal DFT, $O(N)$ - $O(N^2)$ for large systems

Lin et al., Commun. Math. Sci. 7, 755 (2009); Lin et al., J. Phys.: Condens. Matter 25, 295501 (2013); Lin et al., J. Phys: Condens. Matter 26, 305503 (2014)

Performance: Solver Benchmarks on Equal Footing

1D
(a) CNT

2D
(b) Graphene

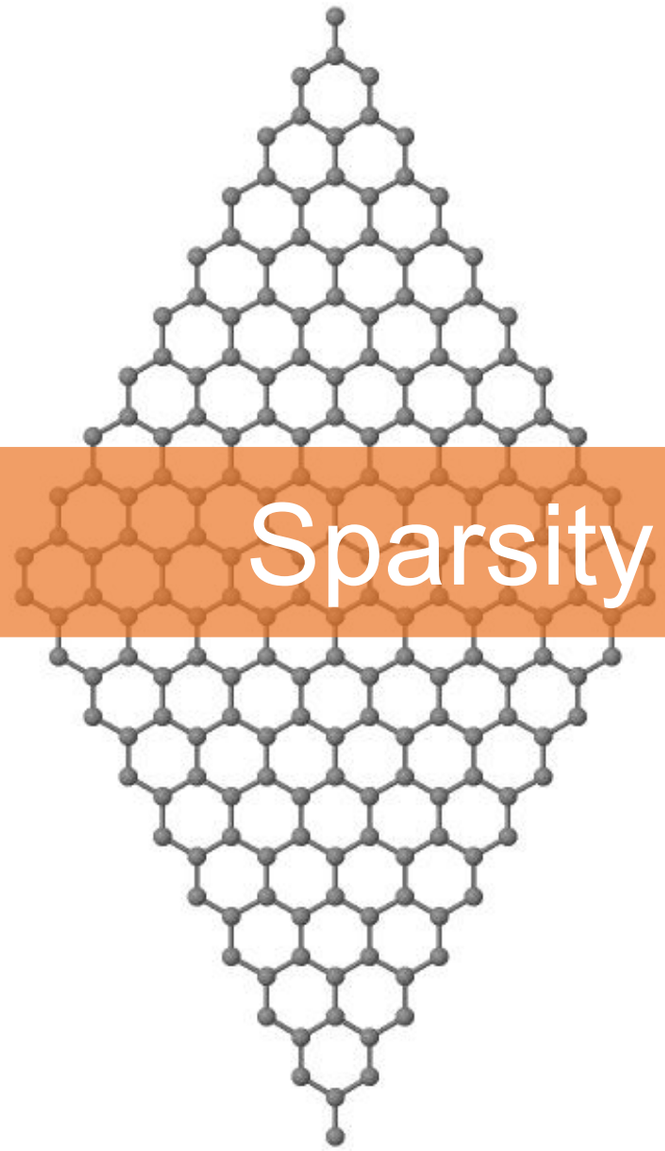

3D
(c) Graphite

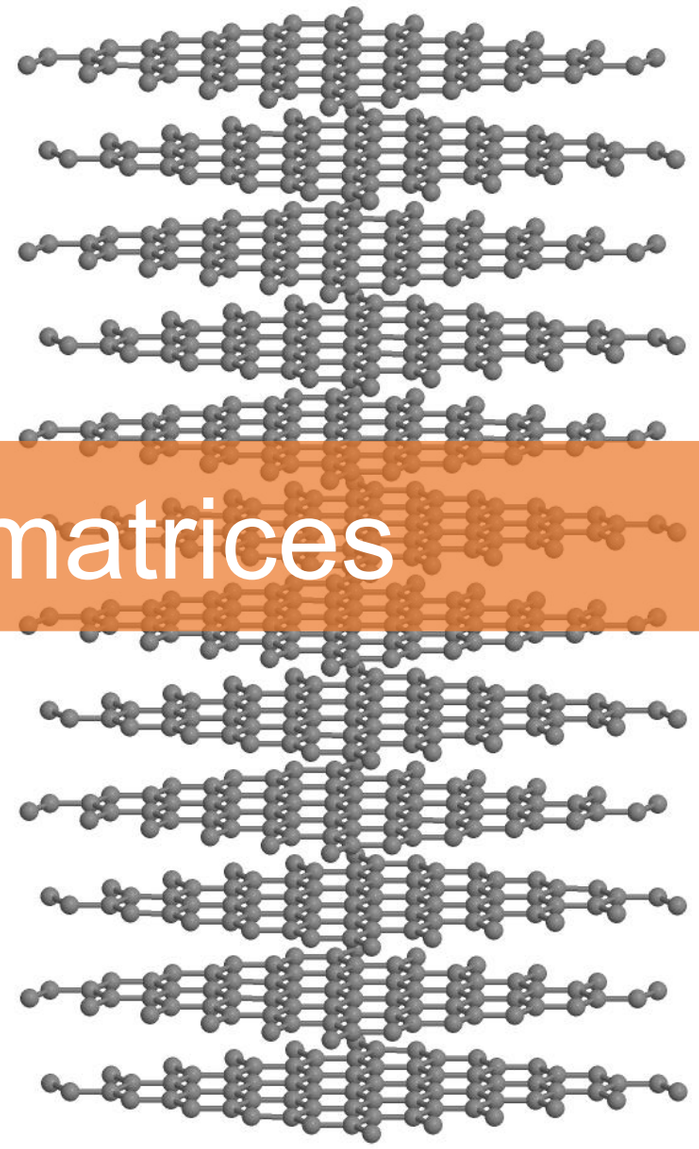

Sparsity of matrices

Example: Siesta Basis Sets - ELPA vs. PEXSI

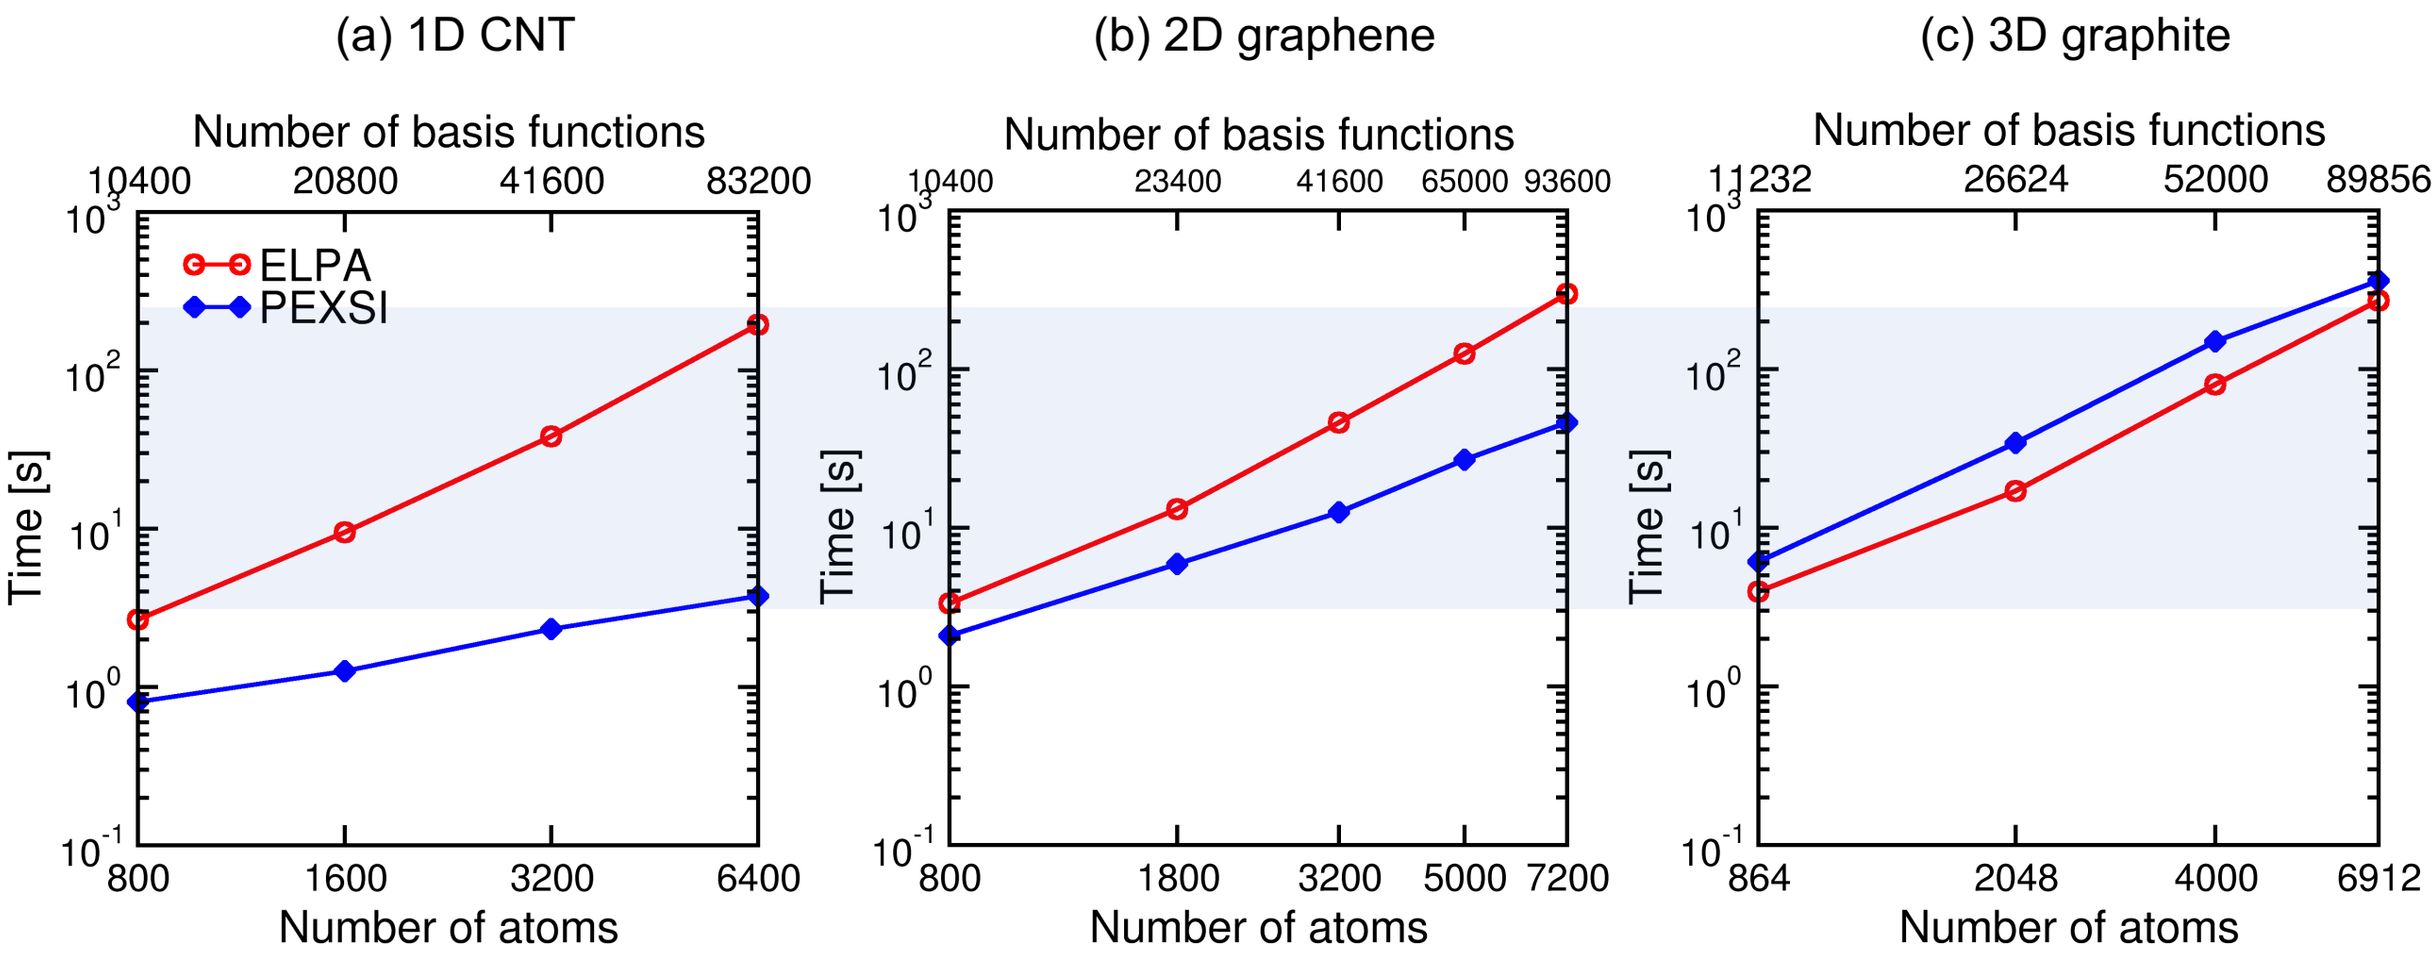

- DFT-PBE
- SIESTA (pseudopotential)
- 1,920 CPU cores on Edison

PEXSI faster for large low-dimensional (sparse) systems

PEXSI: Semilocal DFT, $O(N)$ - $O(N^2)$ for large systems

Lin et al., Commun. Math. Sci. 7, 755 (2009); Lin et al., J. Phys.: Condens. Matter 25, 295501 (2013); Lin et al., J. Phys: Condens. Matter 26, 305503 (2014)

Frozen Core Constrained Eigensolver in ELSI

diagonal

$$H = \begin{bmatrix} H_{cc} & H_{cv} \\ H_{vc} & H_{vv} \end{bmatrix}$$

$$S = \begin{bmatrix} S_{cc} & S_{cv} \\ S_{vc} & S_{vv} \end{bmatrix}$$

Frozen core approach - known for ~8 decades.

Fairly recent summary: Koepernik/Eschrig (1999)

Assume:

1. H_{cc} diagonal (good assumption)
2. S_{cc} = identity
3. $H_{vc} = S_{vc} H_{cc}$,
   $H_{cv} = H_{cc} S_{cv}$.

$$\tilde{H}_{vv} = L_{vv}^{-1}(H_{vv} - S_{vc} H_{cc} S_{cv})(L_{vv}^{*})^{-1}$$

$$H_{cc} C_{cc} = S_{cc} C_{cc} \Sigma_{cc},$$
$$\tilde{H}_{vv} \tilde{C}_{vv} = \tilde{S}_{vv} \tilde{C}_{vv} \Sigma_{vv}$$

… separate core, valence problems; core diagonal.

Frozen Core Constrained Eigensolver in ELSI

V.W.-z.Yu, J. Moussa, V. Blum, J. Chem. Phys. 154, 224107 (2021).

Victor Yu
(now ANL)

Jonathan
Moussa
(MolSSI)

Problem: ELPA hard to beat by lower scaling solvers below many 1,000s of atoms. Way out?

FHI-aims code (also others) - deliberately "all electron" (no shape approximation to potential; high precision)

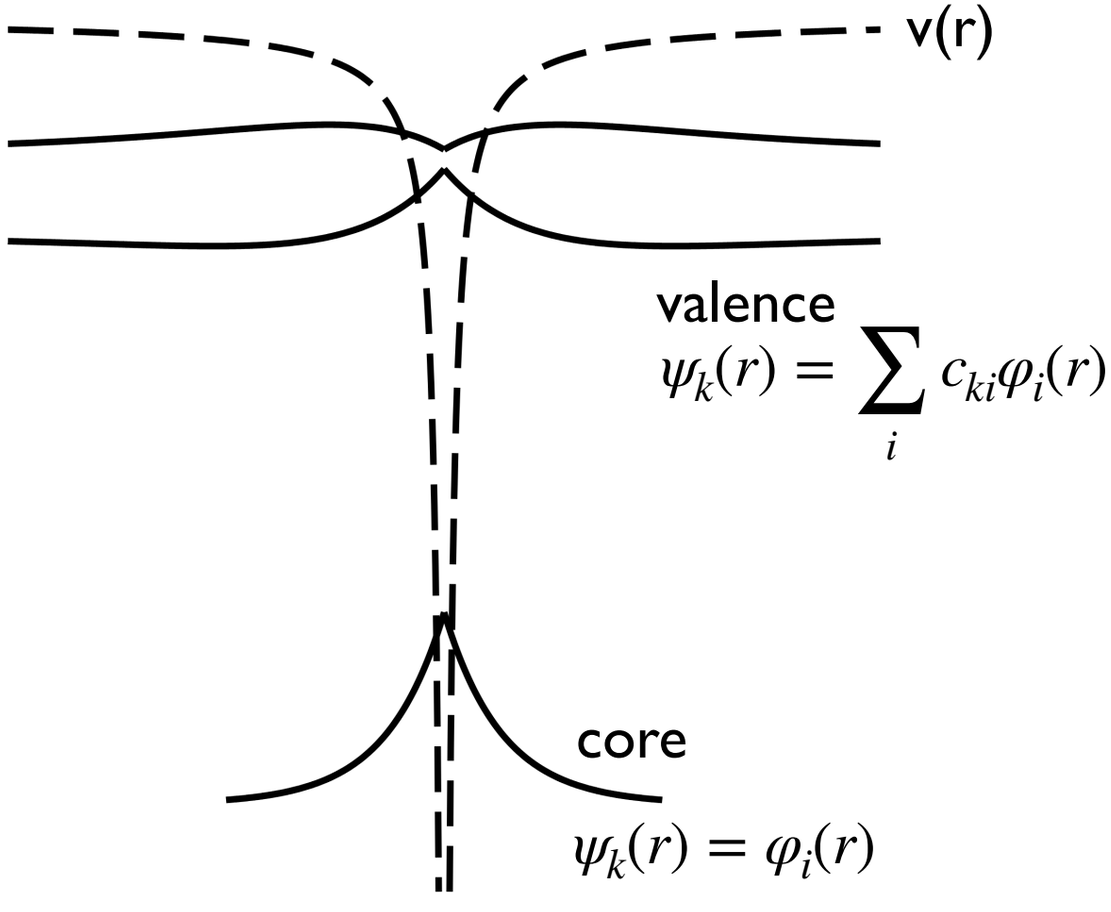

Extended valence states, localized deep core states on equal footing

Advantage - no irrevocable approximations. However: Core states practically known.

diagonal

$$\boldsymbol{H} = \begin{bmatrix} \boldsymbol{H}_{\mathrm{cc}} & \boldsymbol{H}_{\mathrm{cv}} \\ \boldsymbol{H}_{\mathrm{vc}} & \boldsymbol{H}_{\mathrm{vv}} \end{bmatrix}$$

Frozen Core Eigensolver: Speedup

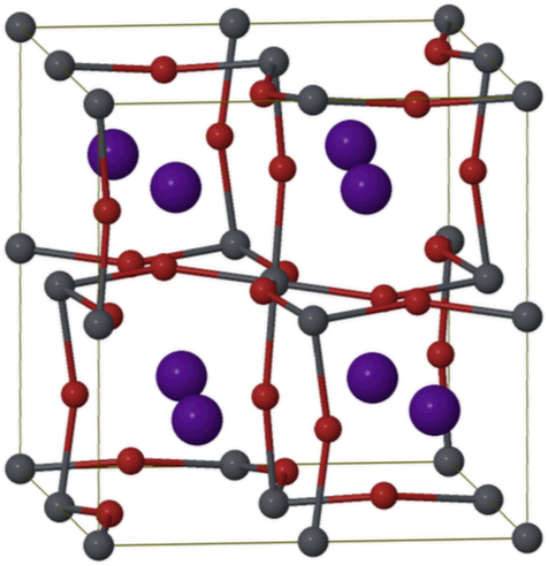

$CsPbBr_3$

Overall speedup, $CsPbBr_3$ supercell (2,560 atoms)

Cori Haswell
N=94,208
N_c=28,672
Implemented in ELSI.

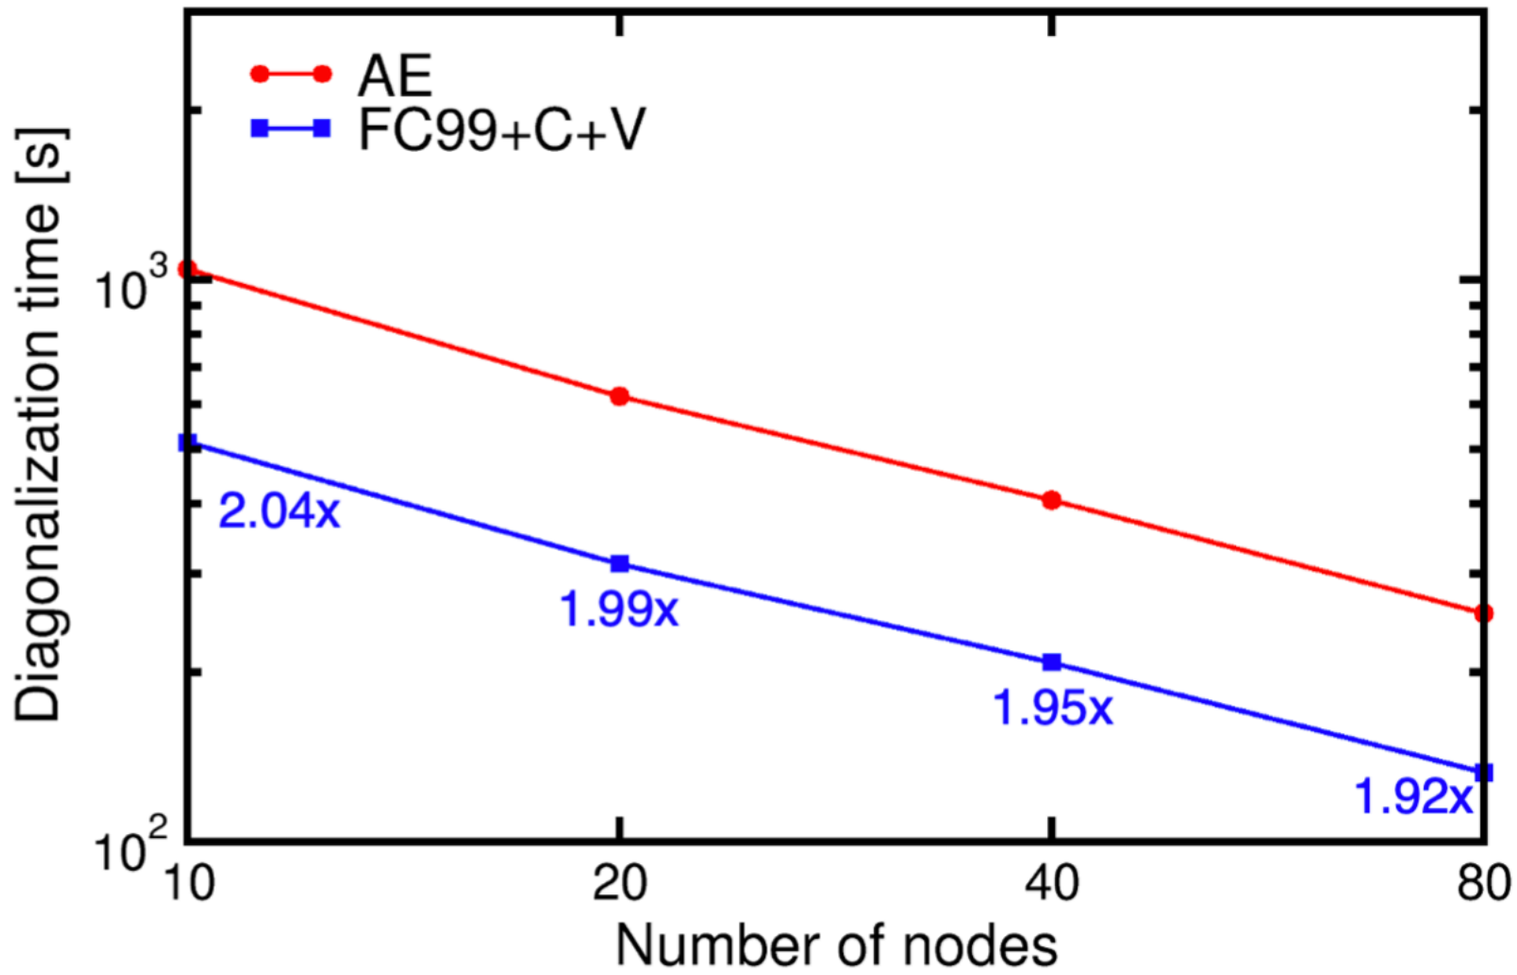

V.W.-z.Yu, J. Moussa, V. Blum, J. Chem. Phys. 154, 224107 (2021).

Next Steps: Mixed Precision (Single-precision ELPA)

TACC Stampede2: 16 Skylake nodes (768 cores)

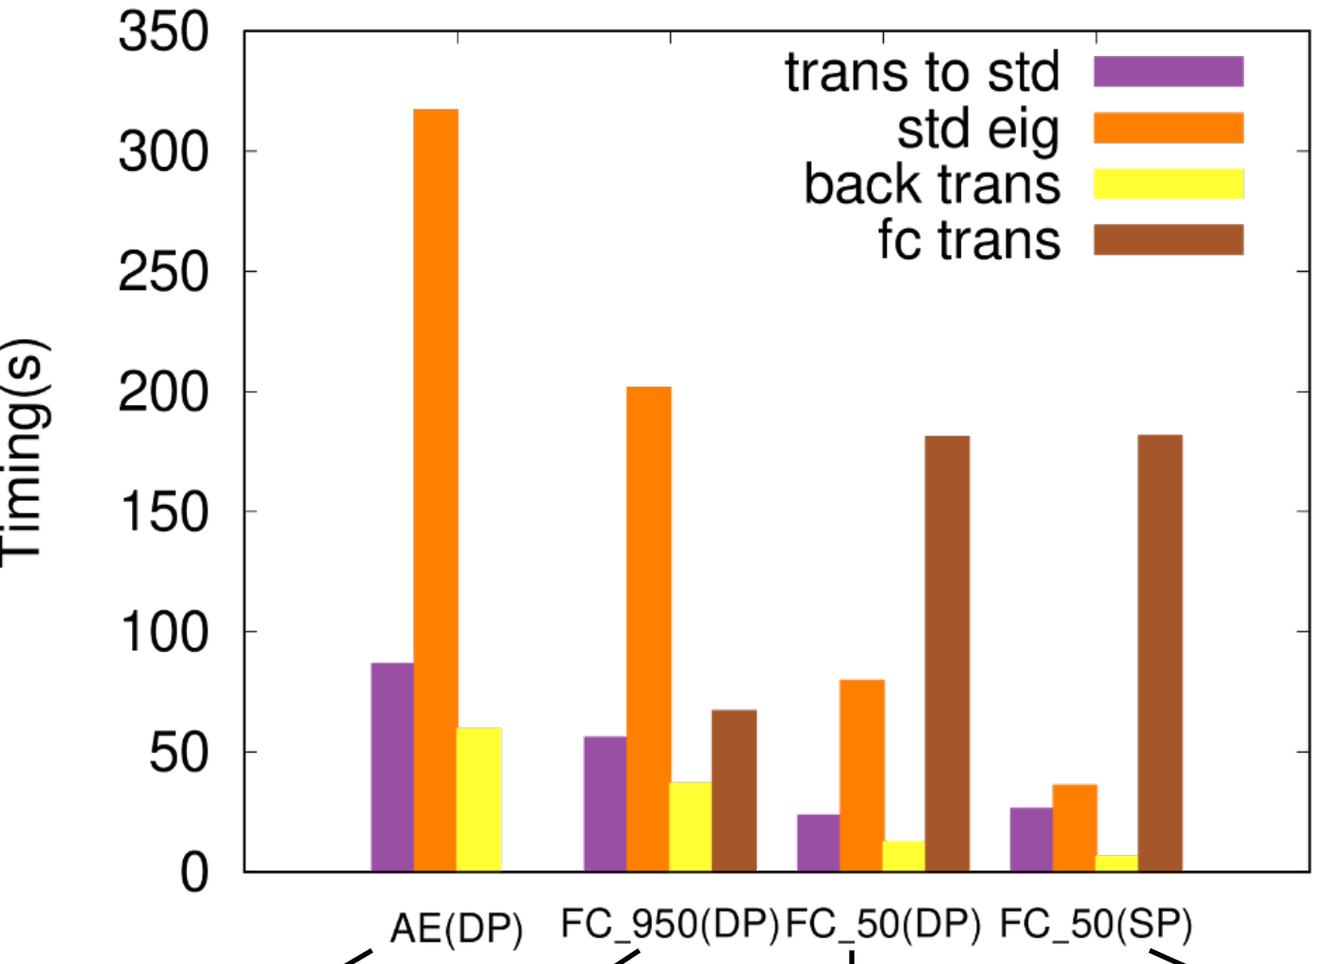

All-electron — 462.5 s

Deep core frozen (below -950 eV) — 362.2 s

Shallow core frozen (below -50 eV) — 295.2 s

Shallow core frozen single precision — 247.4 s

Yi Yao (Duke)

$(CH_3NH_3)PbBr_3$
8x8x8 supercell

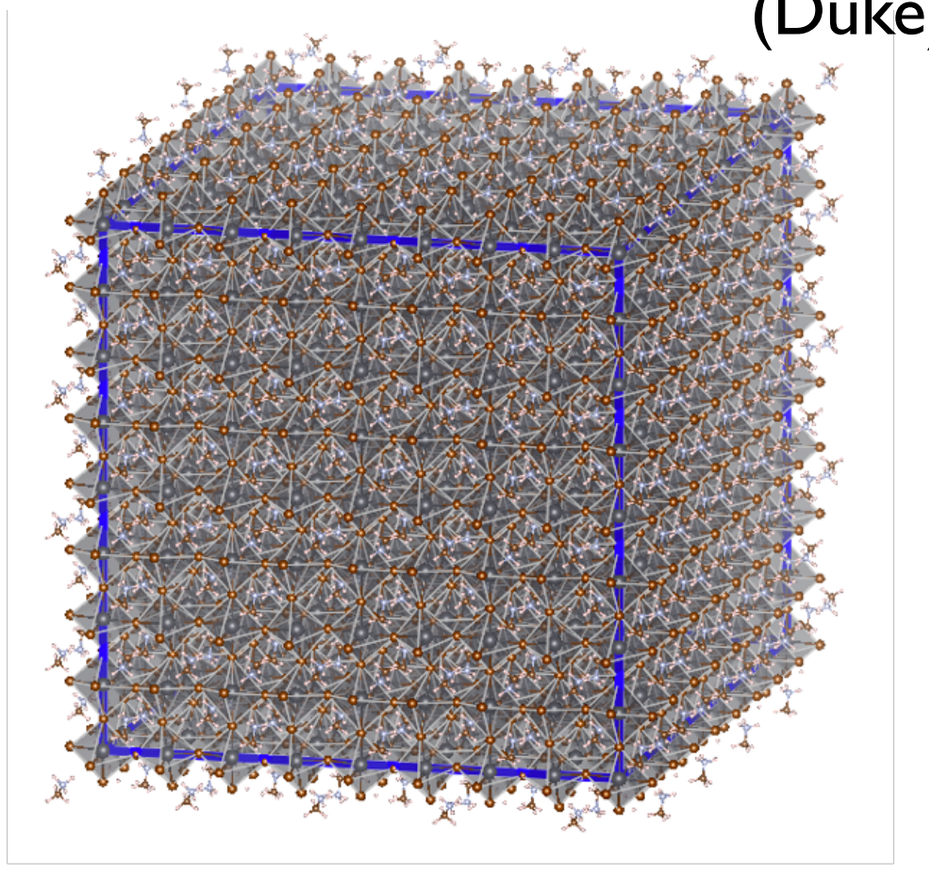

6144 atoms
105472 electrons
101376 basis functions
kpoints 1 1 1

Self-Consistency: Frozen Core + Mixed Precision

Mixed-precision SCF:
Orbitals frozen below −50 eV

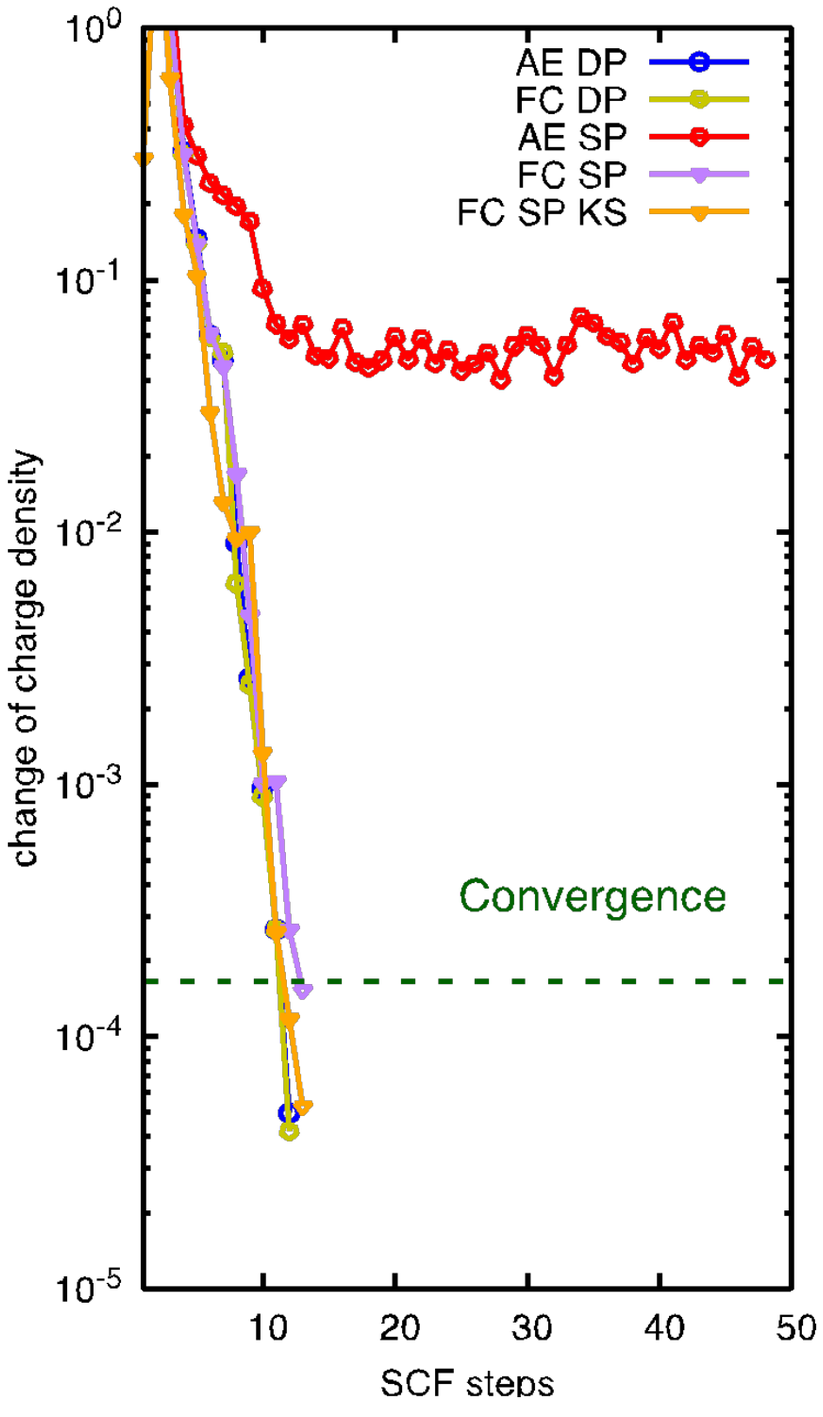

$MAPbBr_3$ supercell -
2,592 atoms, 44,496 electrons

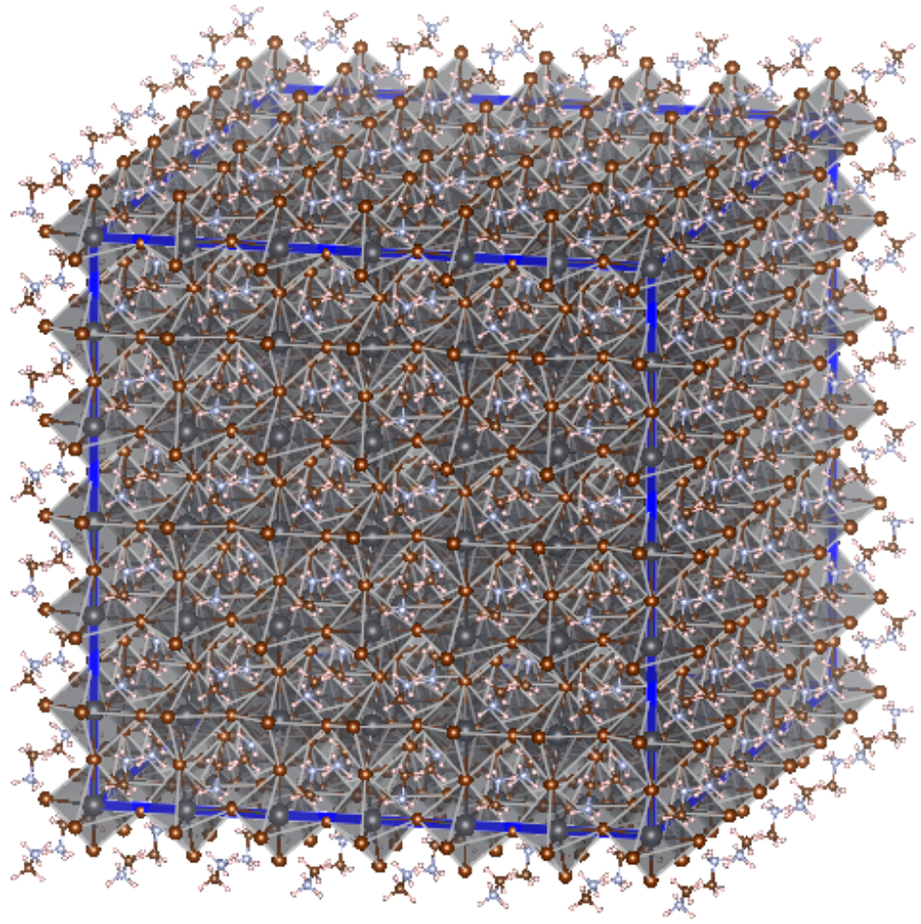

<u>Solver time per s.c.f. iteration:</u>

AE	129 s
FC	81 s
AE SP	61 s
FC SP	71 s - note: AE <-> FC transformation matters

Eigenvalue Problem(s) in Electronic Structure Theory

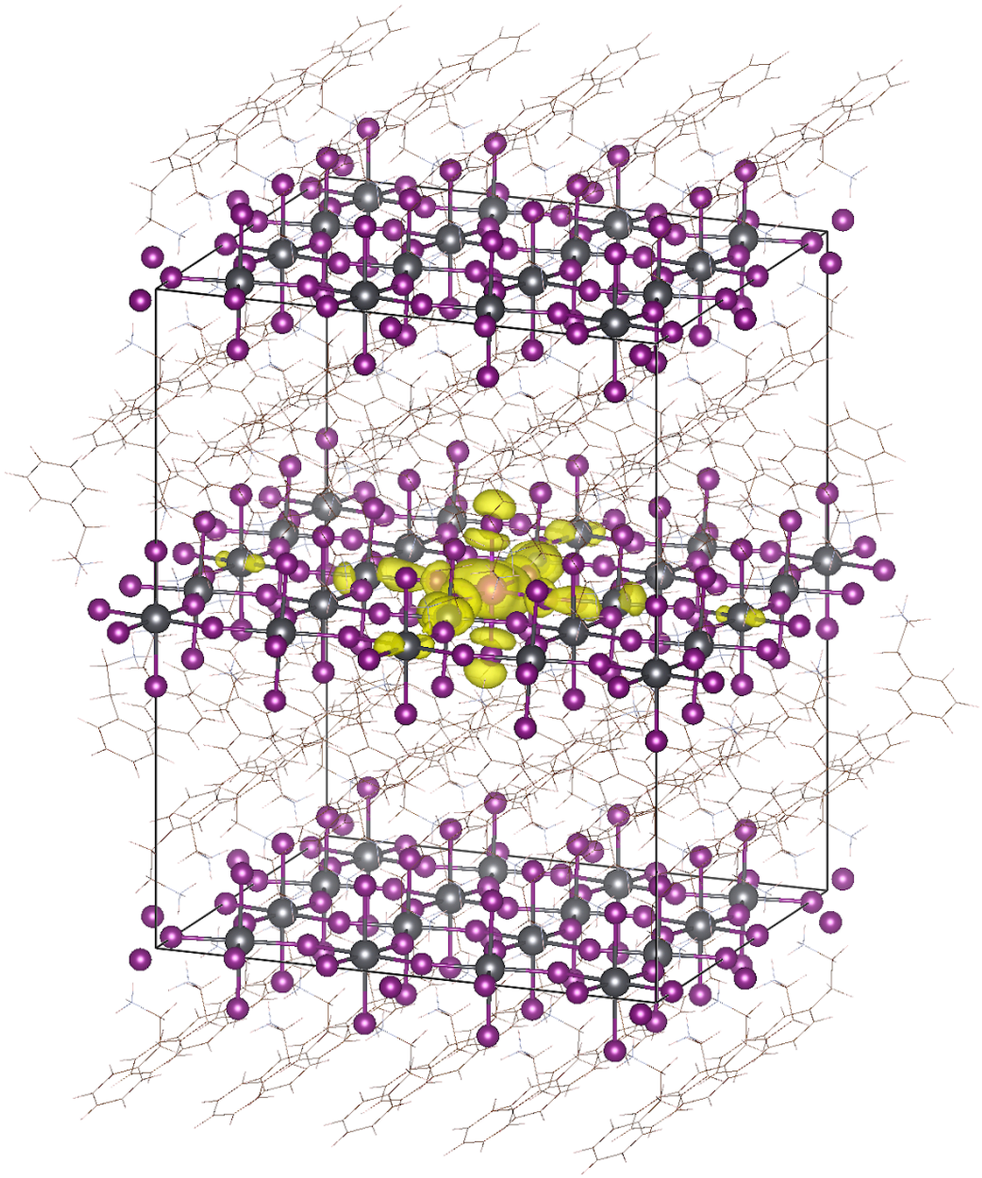

Dopant level, Bi/$(PEA)_2PbI_4$
1,504 atom supercell
Hybrid DFT, FHI-aims code
Stampede2 (NSF XSEDE)

Generalized Kohn-Sham Equations

$$\left(\hat{t} + \hat{v}_{KS}[n(r)]\right)|\psi_k\rangle = \epsilon_k|\psi_k\rangle$$

$$n(r) = \sum_k f_k |\psi_k(r)|^2$$

$$\psi_k(r) \approx \sum_i c_{ki}\varphi_i(r)$$

$$\underline{\underline{H}}\,\underline{c}_k = \epsilon_k \underline{\underline{S}}\,\underline{c}_k$$

N=33,000-40,000 basis functions
M=13,698 electrons / eigenstates

Limited by Eigenvalue Problem