Linear scaling Self-Consistent Field Calculations with Millions of Atoms in the Condensed Phase

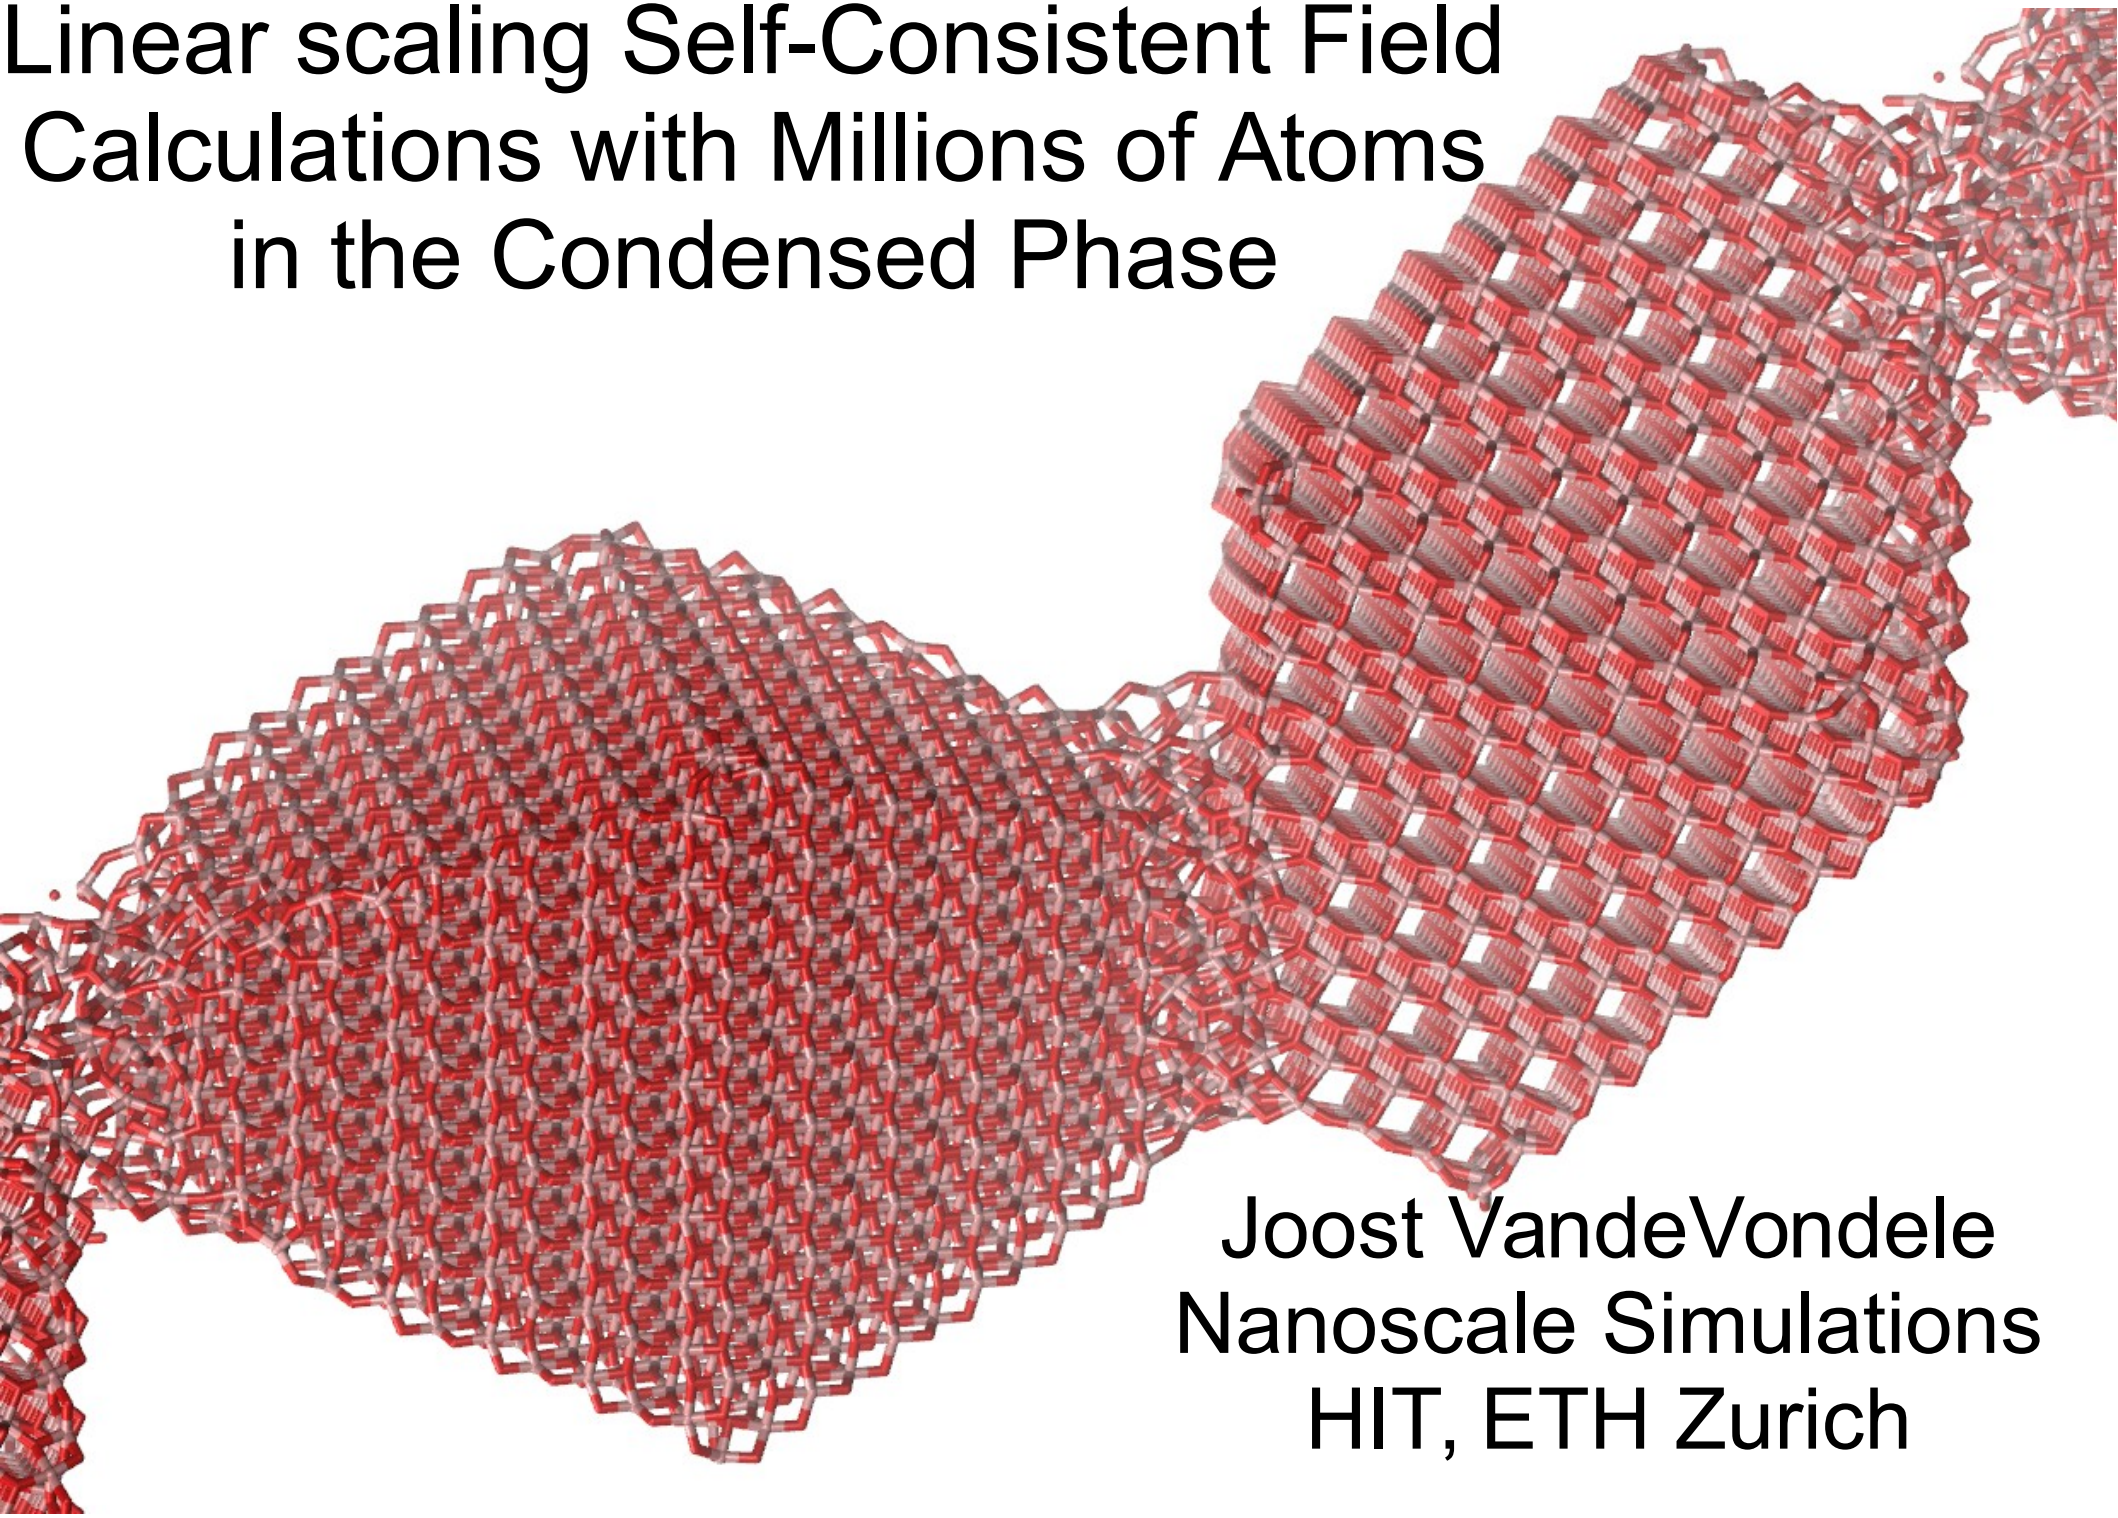

Joost VandeVondele
Nanoscale Simulations
HIT, ETH Zurich

CP2K: the swiss army knife of atomistic simulation

- A wide variety of models Hamiltonians
  - Empirical (classical)
  - semi-empirical
  - local and non-local DFT
  - MP2 & RPA
  - Combinations (e.g. QM/MM)
- Various sampling/dynamics algorithms
  - Molecular dynamics & Monte Carlo
    - NVE, NVT, NPT
  - Free energy and PES tools
  - Ehrenfest MD
- Properties
  - Vibrational
  - NMR, EPR, XAS, TDDFT
- Open source & rapid development
  - 1.000.000 lines of code

Made available as open source software to the community at www.cp2k.org

CP2K: algorithms & implementation

cscs

Research & co-design: Hardware vendors & scientists look together for the best solution (both soft- and hardware)

How can we program

for 10'000 – 1'000'000 cores ?

for 'emerging' architectures ?

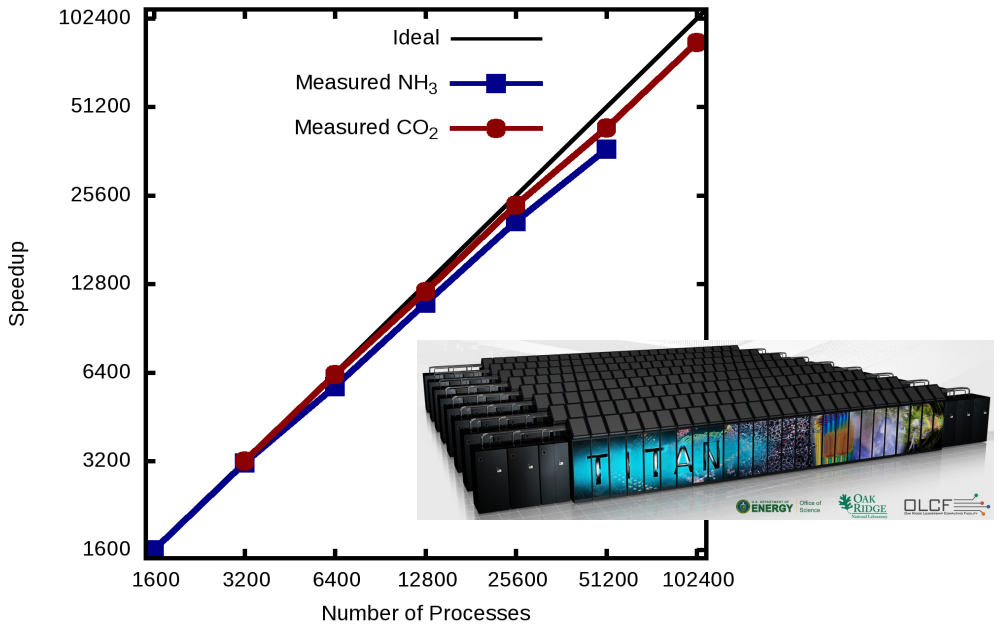

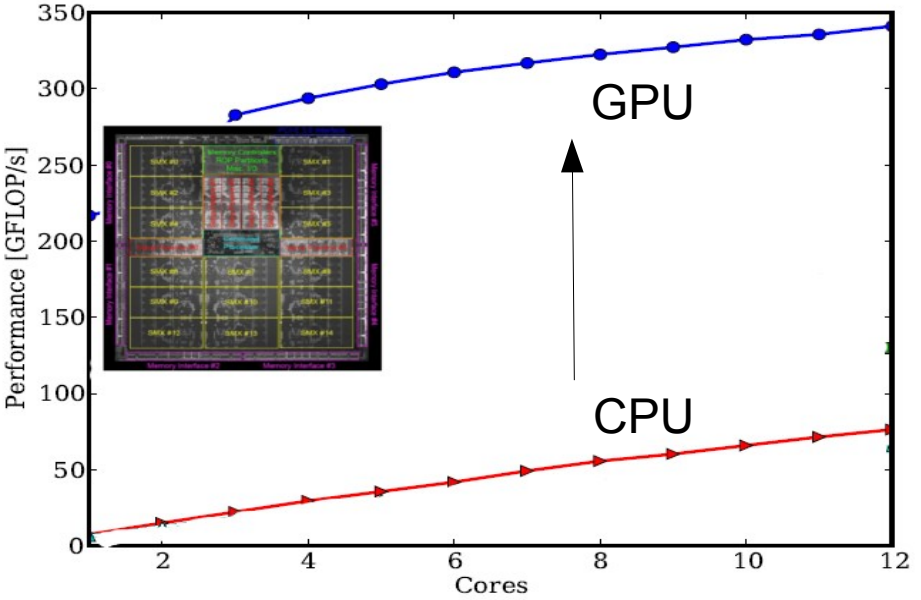

Could save 300 MWh/yr for our group.

M. Del Ben, J. Hutter, J. VandeVondele, J. Chem. Theory Comput 8, 4177-4188 (2012)

Example from ~10 years ago

The two algorithms that enabled CP2K to do new science

- GPW in QS:
  - Combine the computational approaches (basis sets) from chemistry and physics, gas and condensed phases.
- OT
  - New approach to robustly and efficiently obtain electronic structure

G. Lippert, J. Hutter, M. Parrinello, Theor. Chem. Acc. 103 (2), 124 (1999)
J. VandeVondele, M. Krack, F. Mohamed, M. Parrinello, T. Chassaing and J. Hutter, Comp. Phys. Comm. 167, 103 (2005).
J. VandeVondele, J. Hutter, *J. Chem. Phys.,* 118 (10), 4365-4369 (2003)

http://www.cp2k.org/science/

LETTER

doi:10.1038/nature14015

Proton transport through one-atom-thick crystals

S. Hu[1,2], M. Lozada-Hidalgo[1], F. C. Wang[3], A. Mishchenko[1], F. Schedin[2], R. R. Nair[1], E. W. Hill[2], D. W. Boukhvalov[4], M. I. Katsnelson[4], R. A. W. Dryfe[5], I. V. Grigorieva[1], H. A. Wu[3] & A. K. Geim[1,2]

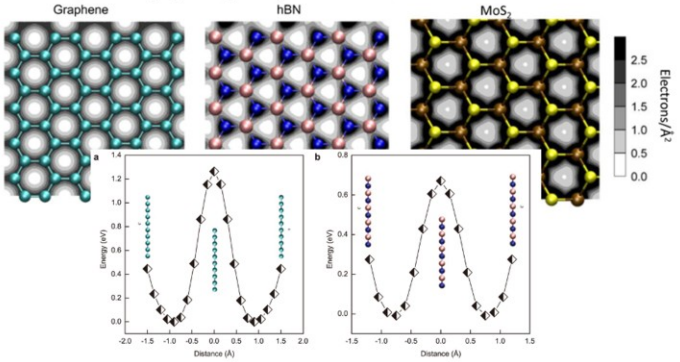

NANO LETTERS

Letter

pubs.acs.org/NanoLett

Friction of Water on Graphene and Hexagonal Boron Nitride from *Ab Initio* Methods: Very Different Slippage Despite Very Similar Interface Structures

Gabriele Tocci,[†,‡,§] Laurent Joly,[∥] and Angelos Michaelides*[,†,‡,§]

[†]Thomas Young Centre, [‡]London Centre for Nanotechnology, University College London, London WC1E 6BT, United Kingdom
[§]Department of Chemistry, University College London, London WC1H 0AJ, United Kingdom
[∥]Institut Lumière Matière, UMR5306 Université Lyon 1-CNRS, Université de Lyon 69622 Villeurbanne, France

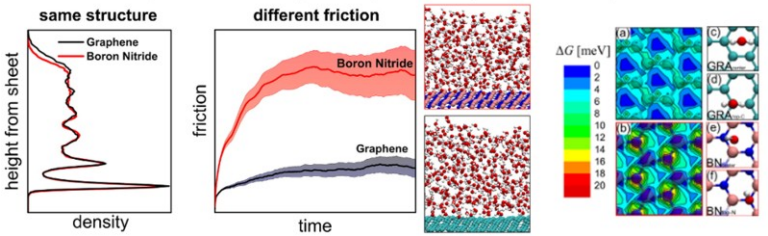

Thermodynamics

DOI: 10.1002/anie.201405648

Aligning Electronic and Protonic Energy Levels of Proton-Coupled Electron Transfer in Water Oxidation on Aqueous TiO_2**

Jun Cheng, Xiandong Liu, John A. Kattirtzi, Joost VandeVondele, and Michiel Sprik*

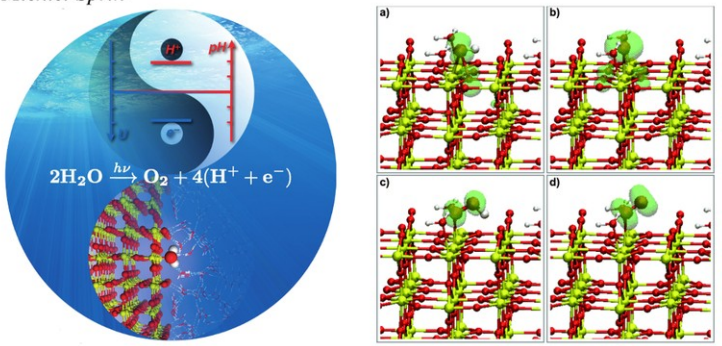

Amorphous Materials Very Important Paper

DOI: 10.1002/anie.201404223

Bonding Nature of Local Structural Motifs in Amorphous GeTe**

Volker L. Deringer, Wei Zhang, Marck Lumeij, Stefan Maintz, Matthias Wuttig, Riccardo Mazzarello, and Richard Dronskowski**

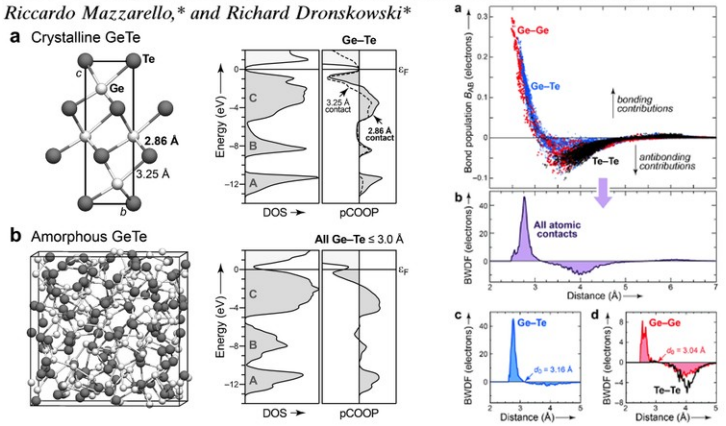

http://www.cp2k.org/science/

nature COMMUNICATIONS

ARTICLE

Received 22 Jan 2014 | Accepted 2 Jul 2014 | Published 1 Aug 2014

DOI: 10.1038/ncomms5578

Liquid-metal electrode to enable ultra-low temperature sodium–beta alumina batteries for renewable energy storage

Xiaochuan Lu[1], Guosheng Li[1], Jin Y. Kim[1], Donghai Mei[2], John P. Lemmon[1], Vincent L. Sprenkle[1] & Jun Liu[1]

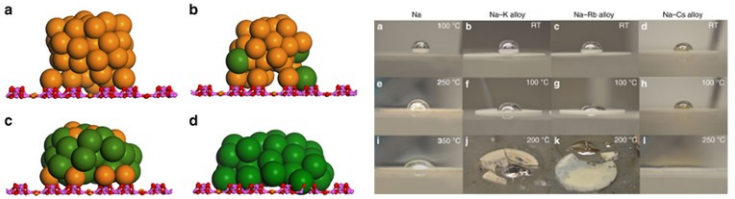

JACS

JOURNAL OF THE AMERICAN CHEMICAL SOCIETY

Article

pubs.acs.org/JACS

First-Principles Study of Phenol Hydrogenation on Pt and Ni Catalysts in Aqueous Phase

Yeohoon Yoon,[†] Roger Rousseau,[*,†] Robert S. Weber,[†] Donghai Mei,[*,†] and Johannes A. Lercher[*,†,‡]

[†]Institute for Integrated Catalysis, Pacific Northwest National Laboratory, Richland, Washington 99352, United States

[‡]Department of Chemistry and Catalysis Research Institute, Technische Universität München, Garching 85747, Germany

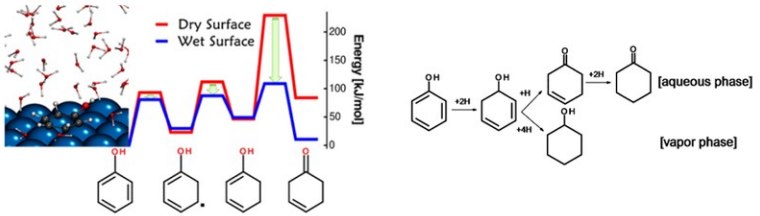

ADVANCED MATERIALS

www.advmat.de

Materials Views

www.MaterialsViews.com

Synthesis of a Covalent Monolayer Sheet by Photochemical Anthracene Dimerization at the Air/Water Interface and its Mechanical Characterization by AFM Indentation

*Payam Payamyar, Khaled Kaja, Carlos Ruiz-Vargas, Andreas Stemmer, Daniel J. Murray, Carey J. Johnson, Benjamin T. King, Florian Schiffmann, Joost VandeVondele, Alois Renn, Stephan Götzinger, Paola Ceroni, Andri Schütz, Lay-Theng Lee, Zhikun Zheng, Junji Sakamoto, and A. Dieter Schlüter**

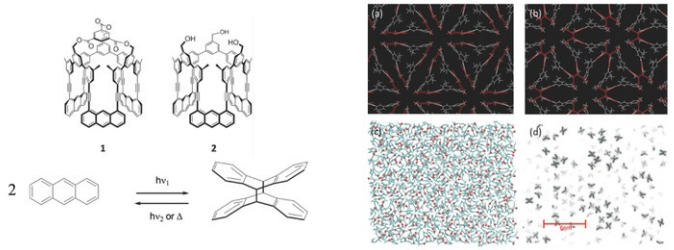

ARTICLES

PUBLISHED ONLINE: 8 JUNE 2014 | DOI: 10.1038/NMAT4007

nature materials

Air-stable n-type colloidal quantum dot solids

Zhijun Ning[1], Oleksandr Voznyy[1], Jun Pan[2], Sjoerd Hoogland[1], Valerio Adinolfi[1], Jixian Xu[1], Min Li[3], Ahmad R. Kirmani[2], Jon-Paul Sun[4], James Minor[1], Kyle W. Kemp[1], Haopeng Dong[1], Lisa Rollny[1], André Labelle[1], Graham Carey[1], Brandon Sutherland[1], Ian Hill[4], Aram Amassian[2], Huan Liu[3], Jiang Tang[5], Osman M. Bakr[2] and Edward H. Sargent[1]*

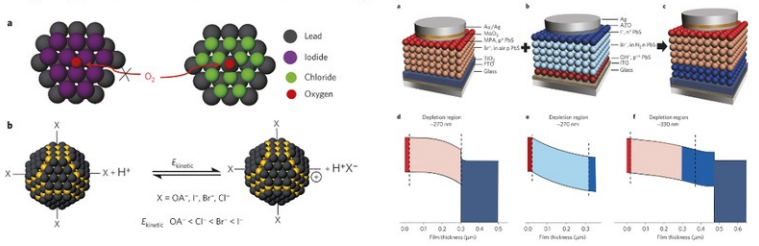

http://www.cp2k.org/science/

nature COMMUNICATIONS

ARTICLE

Received 19 Feb 2013 | Accepted 21 Aug 2013 | Published 25 Sep 2013

DOI: 10.1038/ncomms3481

Stable platinum nanoparticles on specific $MgAl_2O_4$ spinel facets at high temperatures in oxidizing atmospheres

Wei-Zhen Li[1], Libor Kovarik[1], Donghai Mei[1], Jun Liu[1], Yong Wang[1,2] & Charles H.F. Peden[1]

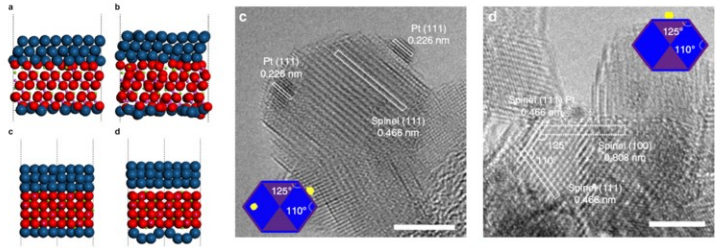

ARTICLE

Received 2 Apr 2012 | Accepted 11 Jun 2012 | Published 17 Jul 2012

DOI: 10.1038/ncomms1949

Self-assembling subnanometer pores with unusual mass-transport properties

Xibin Zhou[1,*], Guande Liu[2,*], Kazuhiro Yamato[3], Yi Shen[4], Ruixian Cheng[1], Xiaoxi Wei[3], Wanli Bai[1], Yi Gao[4,5], Hui Li[5], Yi Liu[1], Futao Liu[1], Daniel M. Czajkowsky[4], Jingfang Wang[2], Michael J. Dabney[3], Zhonghou Cai[6], Jun Hu[4], Frank V. Bright[3], Lan He[1], Xiao Cheng Zeng[5], Zhifeng Shao[2] & Bing Gong[1,3]

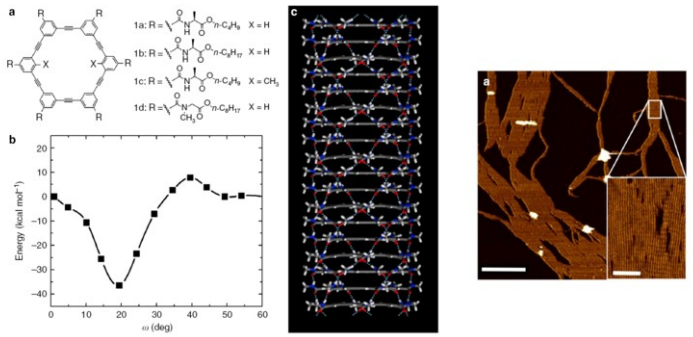

ARTICLES

PUBLISHED ONLINE: 14 OCTOBER 2012 | DOI: 10.1038/NMAT3456

nature materials

Role of vacancies in metal–insulator transitions of crystalline phase-change materials

W. Zhang[1], A. Thiess[2,3], P. Zalden[4], R. Zeller[2], P. H. Dederichs[2], J-Y. Raty[5], M. Wuttig[4,6]*, S. Blügel[2,6] and R. Mazzarello[1,6]*

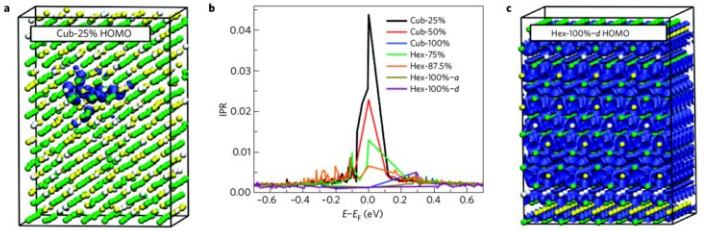

Modular and predictable assembly of porous organic molecular crystals

James T. A. Jones[1], Tom Hasell[1], Xiaofeng Wu[1], John Bacsa[1], Kim E. Jelfs[1], Marc Schmidtmann[1], Samantha Y. Chong[1], Dave J. Adams[1], Abbie Trewin[1], Florian Schiffman[2], Furio Cora[2], Ben Slater[2], Alexander Steiner[1], Graeme M. Day[1] & Andrew I. Cooper[1]

16 JUNE 2011 | VOL 474 | NATURE | 367

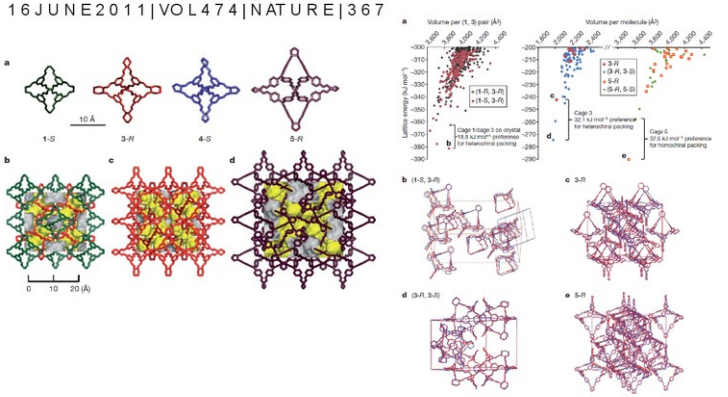

http://www.cp2k.org/science/

nature chemistry
ARTICLES
PUBLISHED ONLINE: 12 SEPTEMBER 2010 | DOI: 10.1038/NCHEM.827

Synthesis of glycine-containing complexes in impacts of comets on early Earth

Nir Goldman*, Evan J. Reed†, Laurence E. Fried, I.-F. William Kuo and Amitesh Maiti

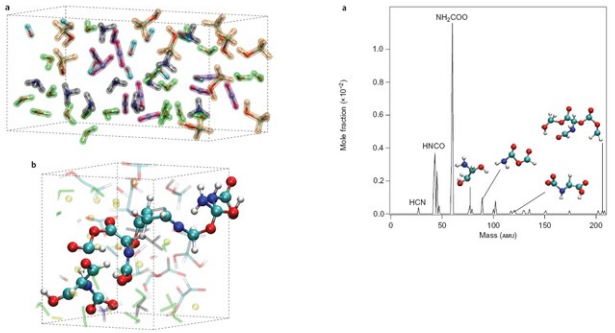

Large variation of vacancy formation energies in the surface of crystalline ice

M. Watkins[1,2,3], D. Pan[4], E. G. Wang[5], A. Michaelides[1,2,3], J. VandeVondele[6] and B. Slater[1,3]*

[1]Department of Chemistry, Christopher Ingold Building, 20 Gordon Street, University College London, London WC1H 0AJ, UK, [2]London Centre for Nanotechnology, University College London, London WC1H 0AJ, UK, [3]TYC@UCL, University College London, London WC1H 0AJ, UK, [4]Institute of Physics, Chinese Academy of Sciences, PO Box 603, Beijing 100190, China, [5]School of Physics, Peking University, Beijing 100871, China, [6]Institute of Physical Chemistry, University of Zurich, Winterthurerstrasse 190, CH-8057 Zurich, Switzerland. *e-mail: b.slater@ucl.ac.uk.

NATURE MATERIALS | VOL 10 | OCTOBER 2011

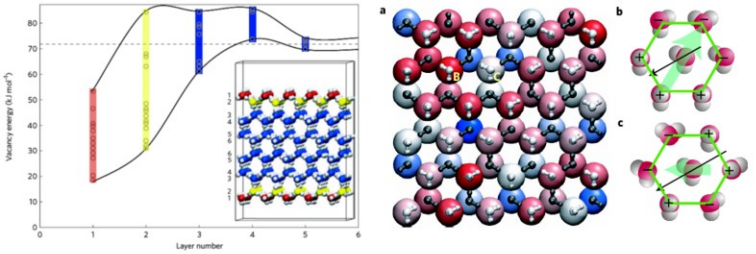

Far-IR Spectroscopy

DOI: 10.1002/anie.201311189

Gas-Phase Peptide Structures Unraveled by Far-IR Spectroscopy: Combining IR-UV Ion-Dip Experiments with Born–Oppenheimer Molecular Dynamics Simulations**

Sander Jaeqx, Jos Oomens, Alvaro Cimas, Marie-Pierre Gaigeot, and Anouk M. Rijs**

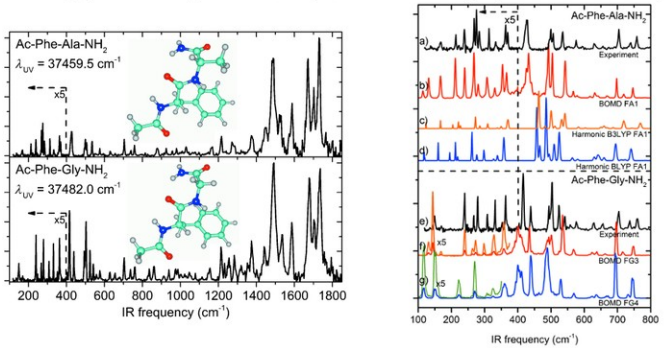

nature chemistry
ARTICLES
PUBLISHED ONLINE: 7 NOVEMBER 2010 | DOI: 10.1038/NCHEM.891

Surface-assisted cyclodehydrogenation provides a synthetic route towards easily processable and chemically tailored nanographenes

Matthias Treier[1†], Carlo Antonio Pignedoli[1], Teodoro Laino[2†], Ralph Rieger[3], Klaus Müllen[3], Daniele Passerone[1] and Roman Fasel[1,4]*

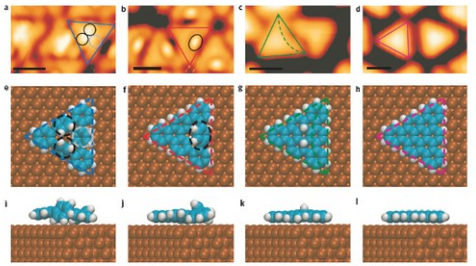

Figure 3 | High-resolution STM images of reactant, intermediates and final product on Cu(111). a-d, Triangles and circles are used to highlight

Petascale supercomputing

1 petaflops = solve 100'000 coupled equations for 100'000 unknowns in 1 sec.
= 1'000'000'000'000'000 multiplications/additions per sec.

#1 = 34 petaflops (June 2014), Switzerland: 6 petaflops (rank 6, 1st in europe)

The 37 fastest computers in the world have peak petaflop performance

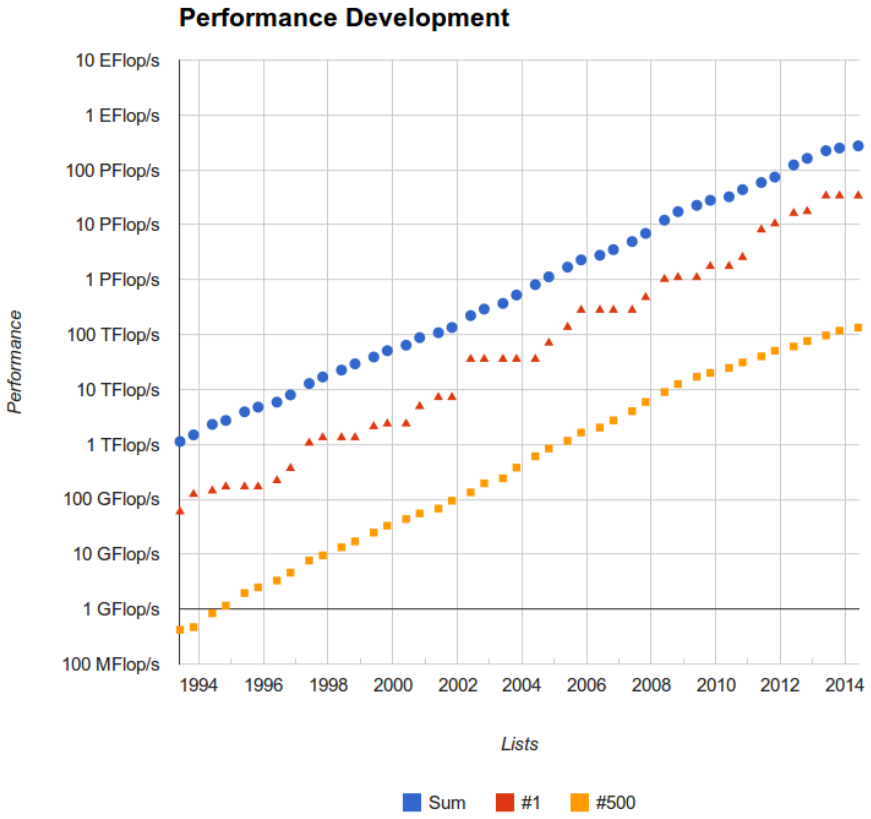

Parallel computers have followed a path of sustained *exponential growth* for 20 years

- Serial computers.... do not exist anymore
  Serial programs become irrelevant
- 1 month now = 1 day in 5 years
- Few experimental techniques show exponential increases in throughput, performance, or resolution

Improving the predictive nature of atomistic simulations

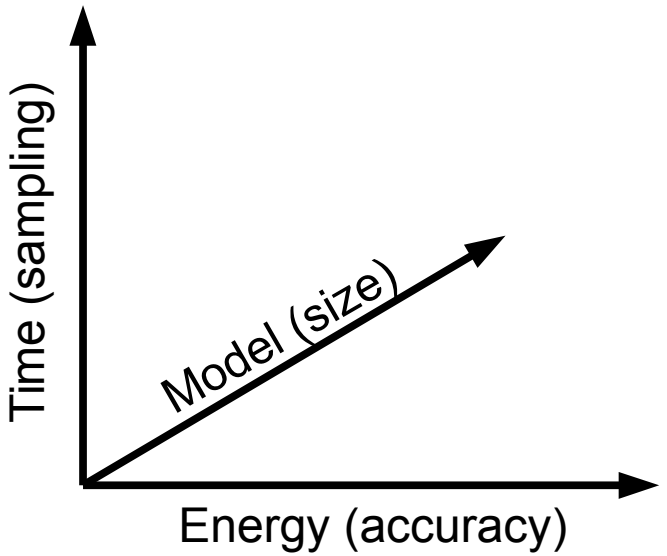

Time:
- Longer simulation
- Sampling (Entropy)
- Parameter scans
- Uncertainty quantification

Energy:
- 'eliminate' technicalities (basis)
- beyond GGA

Model:
- reduce size effects (small unit cells?)
- include explicit solvents
- nanoparticles vs. slabs

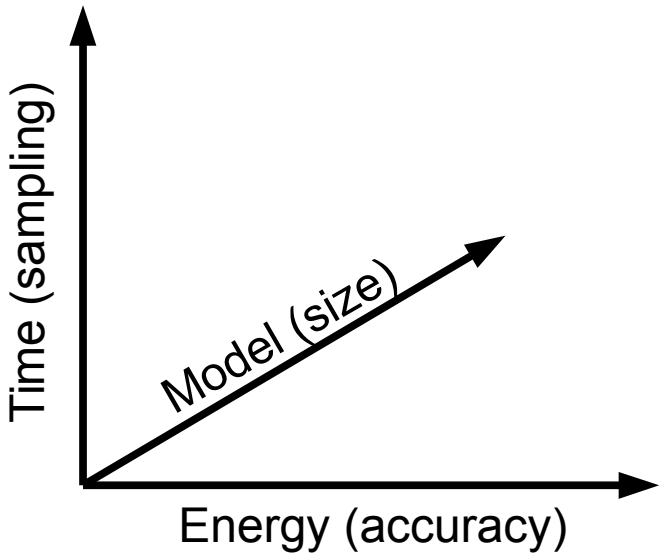

MODEL

A million atoms and nanometers

Small electronic devices, heterostructures, interfaces, nano-particles, a small virus.

Solvated STMV: 1M

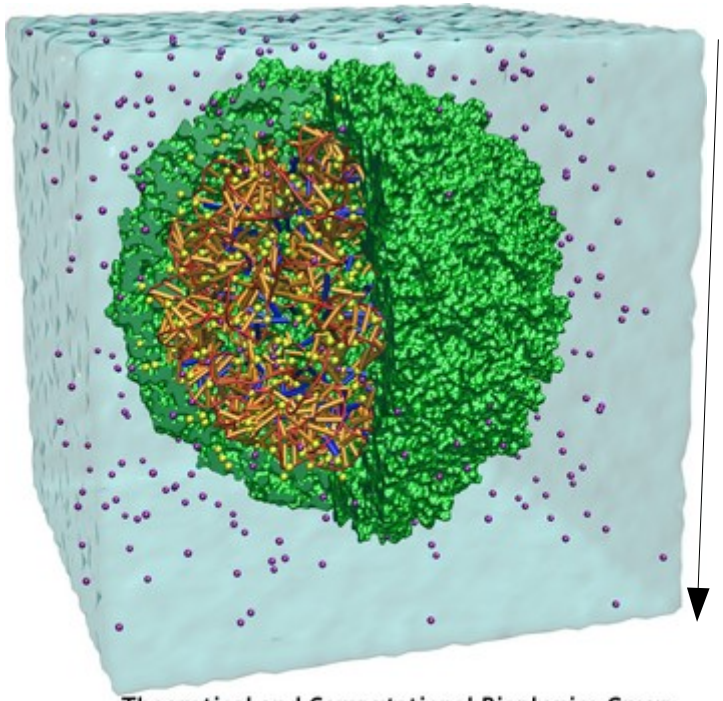

Gate-all-around FET (<22nm)

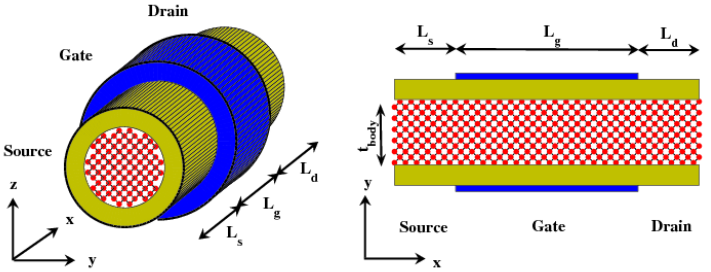

With Mathieu Luisier

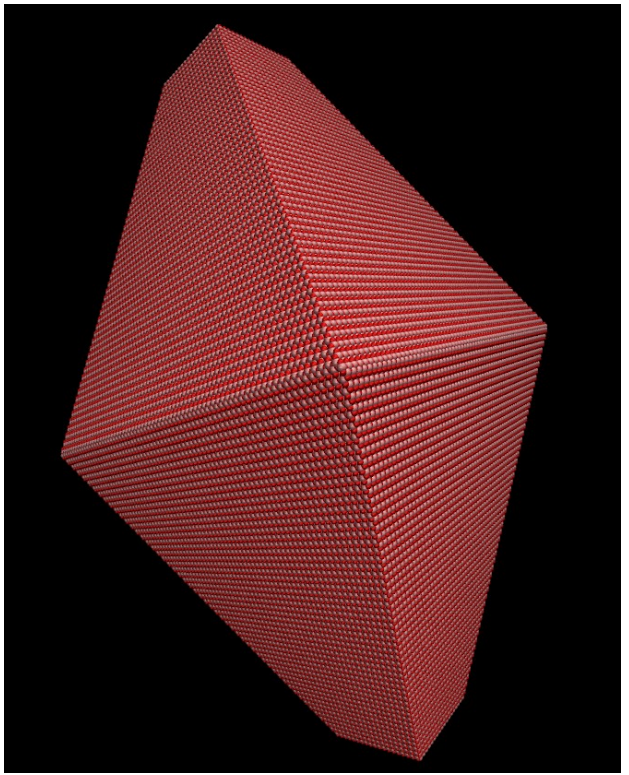

1.5M atoms
Anatase nanocrystal

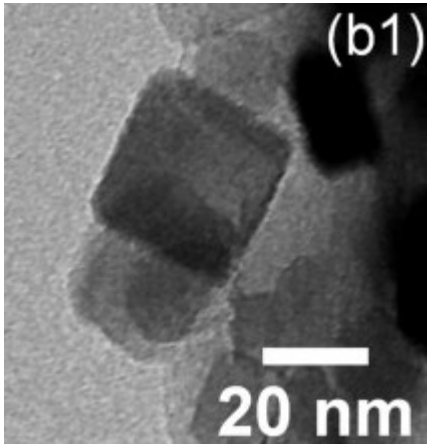

Caplovicova et al.
App. Cat. B, 224, 117

Linear Scaling SCF

Traditional approaches to solve the self-consistent field (SCF) equations are $O(N^3)$ limiting system size significantly.

New algorithms are O(N), allowing for far larger systems to be studied.

Avoid finding 20% lowest eigenvectors of a 10'000'000 x 10'000'000 matrix:

Sparse linear algebra needed

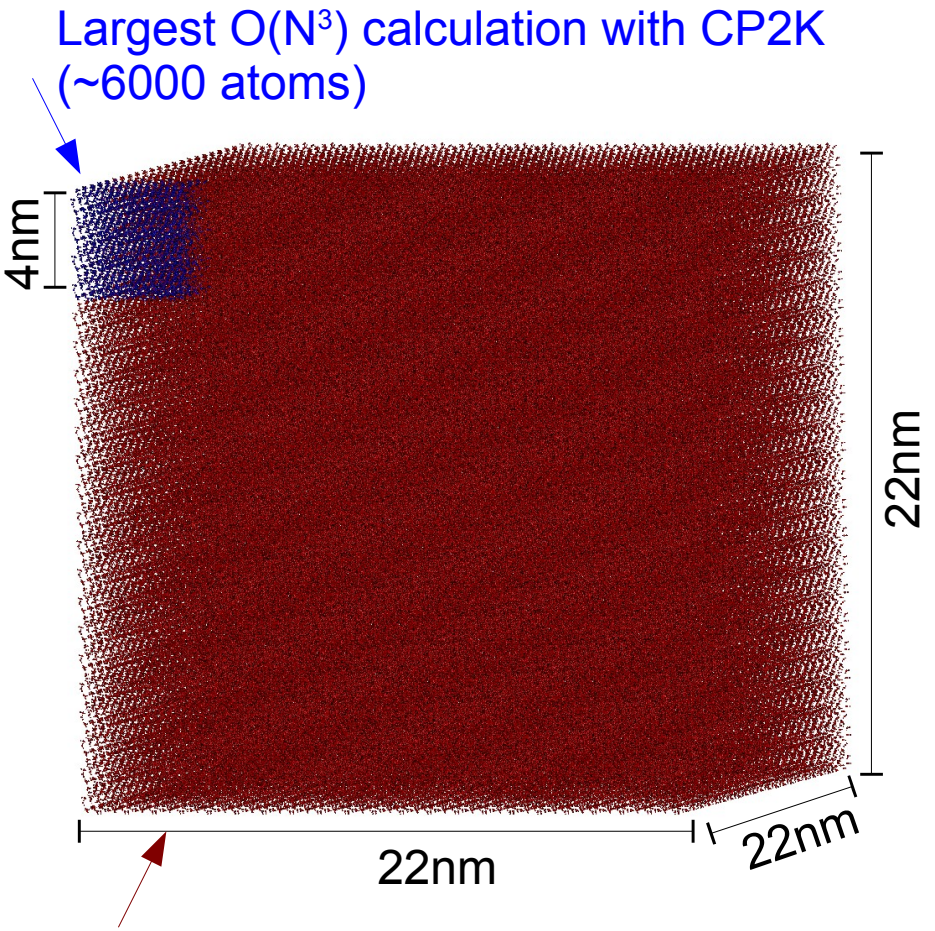

Linear scaling techniques do not speedup calculations on small systems.
They enable calculations on large systems

Diagonalization vs. Linear Scaling

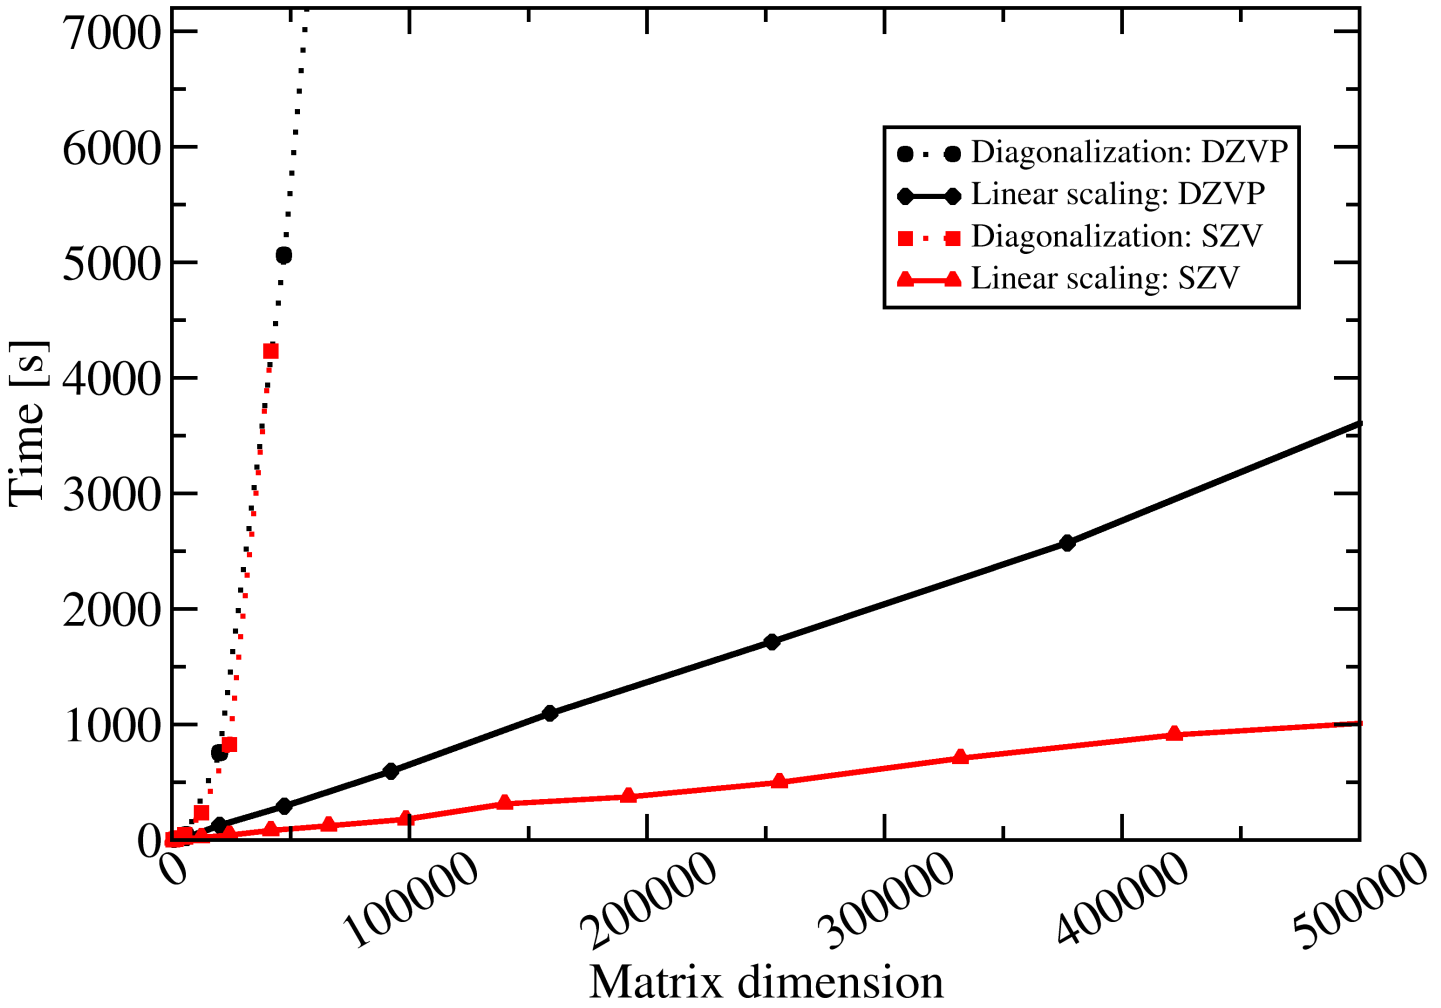

Bulk liquid water, traditional diagonalization vs. Linear scaling algorithms, $\varepsilon=10^{-5}$
Typical crossover point (vs. OT) still a few thousand atoms.

DFT: computational aspects

A two step iterative procedure (simplified) is needed:

SCF

Compute the matrix elements of H_{KS}

- Computational procedure depends on the choice of basis, code, etc.
- Dominant term for small systems (< 100s atoms)
- >10 years of development in the current CP2K code, efficient

Compute P from H_{KS}

- Dominant term for large systems (>100s atoms)
- Diagonalization-like procedures standard
- 'Interesting' methods for large systems.

Some numbers
for traditional simulations DFT

- Matrix dimension: the number of basis functions
  - Typically ~ 10-20 functions per atom (LCAO)
  - 2'000 – 40'000
- Rank P: the number of electrons
  - Typically ~ 4 per atom
  - 200 – 5'000
  - 10% - 50% of the eigenvectors of H needed
- # of 'diagonalizations' needed for 'science'
  - 100s for static calculations
  - 100'000s for dynamic calculations (ab initio MD)

Gaussian basis: The sparsity of H and S

$$S_{\alpha\beta} = \int \varphi_\alpha(r)\varphi_\beta(r)\,dr$$

$$H_{\alpha\beta} = \int \varphi_\alpha(r)v(r)\varphi_\beta(r)\,dr$$

The sparsity pattern of S and H depends on the basis and the spatial location of the atoms, but not on the chemical properties of the system in GGA DFT.

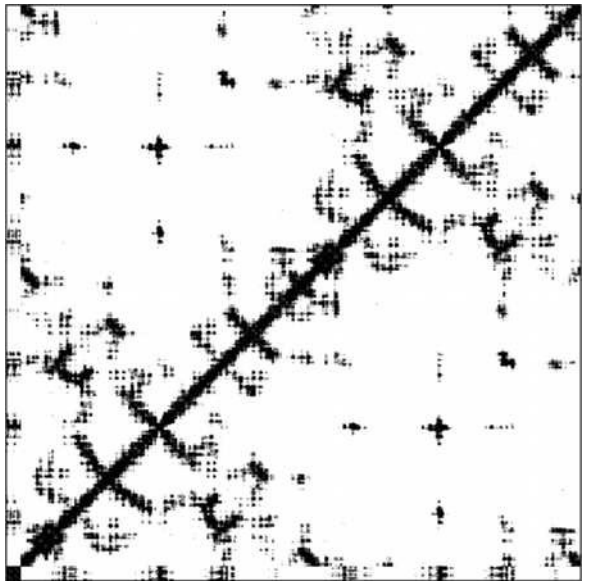

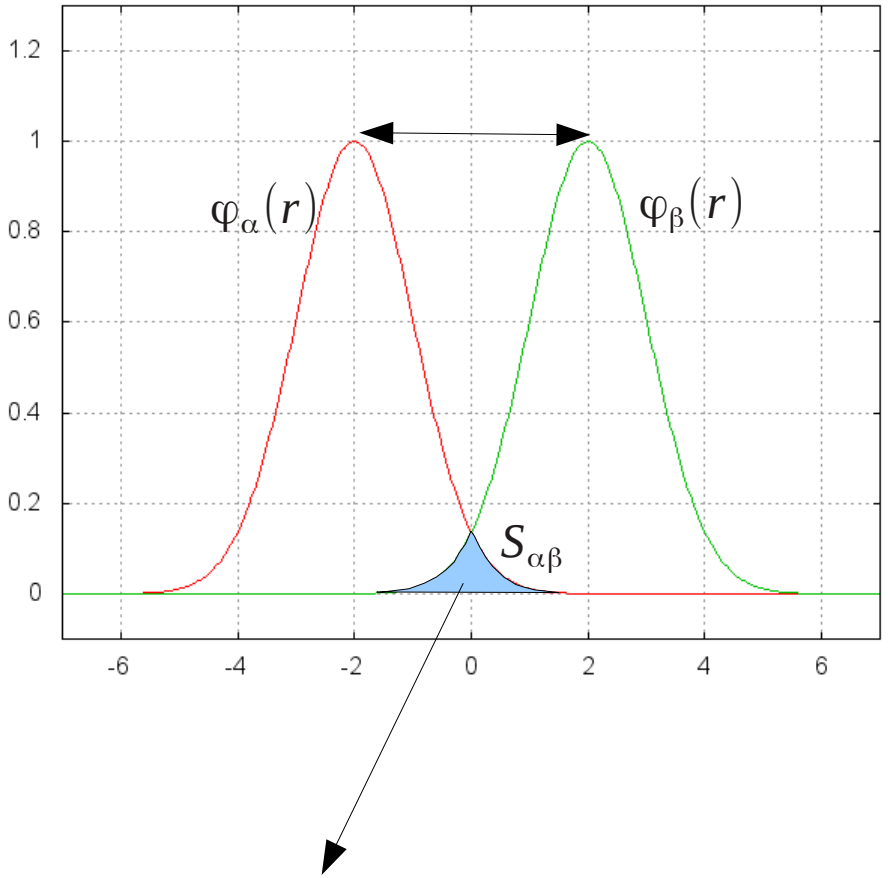

The overlap (integral of the product) rapidly decays with the spatial separation of the basis functions.

Gaussian and plane waves: GPW in CP2K

- Primary basis: Gaussians → Chemistry
  - compact
  - sparse H^{ks} (and P)
  - Many terms analytic
- Auxiliary basis: Plane waves → Physics
  - regular grid for e^- density
  - FFT for Poisson equation
  - No four center integrals needed (GGA)

The GPW algorithm : compute the GGA Kohn-Sham matrix in O(N) time, PBC are natural.

J. VandeVondele, M. Krack, F. Mohamed, M. Parrinello, T. Chassaing and J. Hutter, Comp. Phys. Comm. 167, 103 (2005).
Lippert, G; Hutter, J; Parrinello, M. Mol. Phys., 92 (3), 477-487 (1997).

GPW: O(N) Coulomb solver

1) Transform the density: Gaussians → PW
2) Use (few) FFTs to get the Hartree energy and potential

$$P_{\alpha\beta} \xrightarrow[\text{collocate}]{} \tilde{\rho}(\mathbf{R}) \xrightarrow[\text{FFT}]{} \tilde{\rho}(\mathbf{G}) \rightarrow V_H(\mathbf{G}) \xrightarrow[\text{FFT}^{-1}]{} V_H(\mathbf{R}) \xrightarrow[\text{integrate}]{} K^H_{\alpha\beta}$$

$$\tilde{\rho}(\mathbf{R}) = \sum_{\alpha\beta} P_{\alpha\beta}\phi_\alpha(\mathbf{R})\phi_\beta(\mathbf{R})$$

All steps are O(N) and efficient

Some properties of P

Orthonormality of the C's is equivalent to idempotency of P

$$C_i^T S C_j = \delta_{ij} \Leftrightarrow PS\,PS = PS$$

P is a projector on the subspace spanned by the C_i

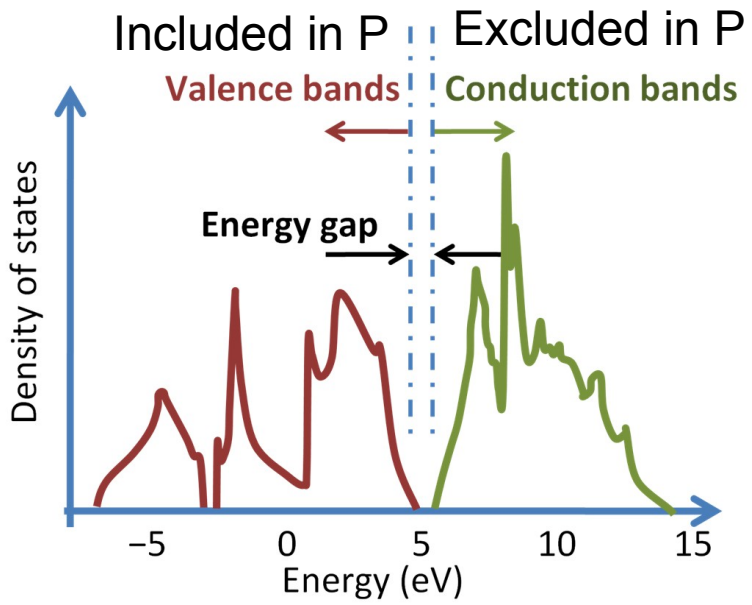

The sparsity of P depends on the chemistry of the system or the spectral properties of (H,S)

Systems with an energy gap (semiconductors) exhibit exponential decay for the density matrix.

Systems without an energy gap (metals), at zero temperature, show a polynomial decay.

Fig. John R. Brews

1D example: Hückel Theory

Metal (c=1)

$$H = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & c \\ 1 & 0 & c & 0 & 0 & 0 \\ 0 & c & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & c & 0 \\ 0 & 0 & 0 & c & 0 & 1 \\ c & 0 & 0 & 0 & 1 & 0 \end{pmatrix}$$

Insulator (c=0.5)

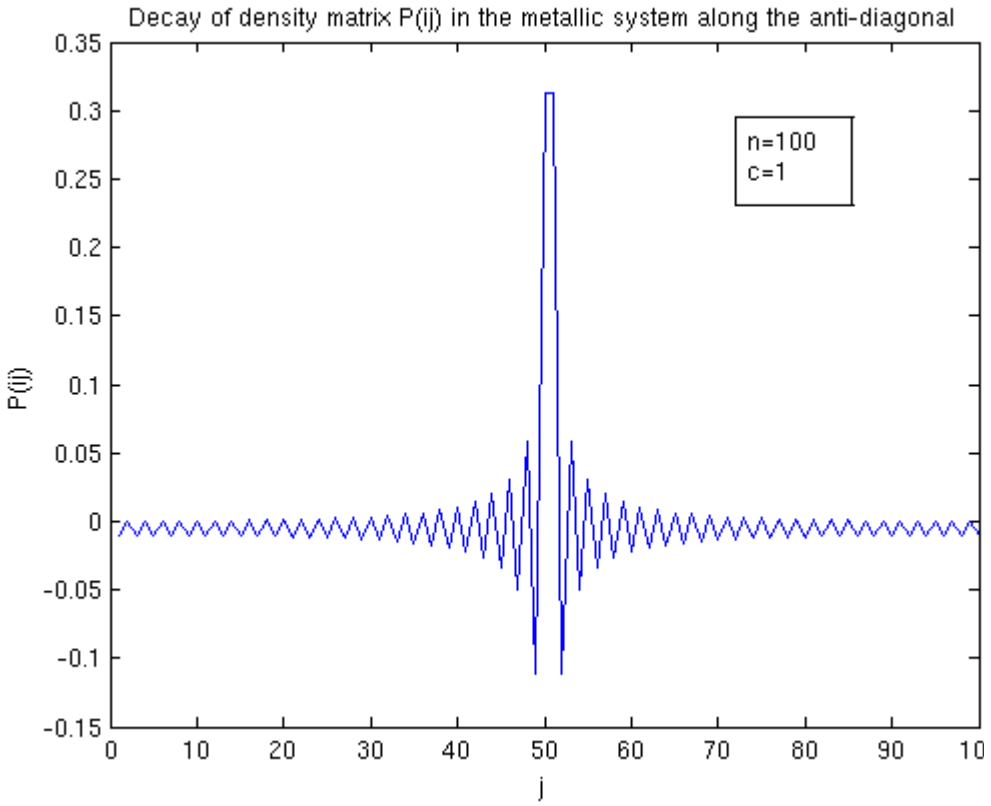

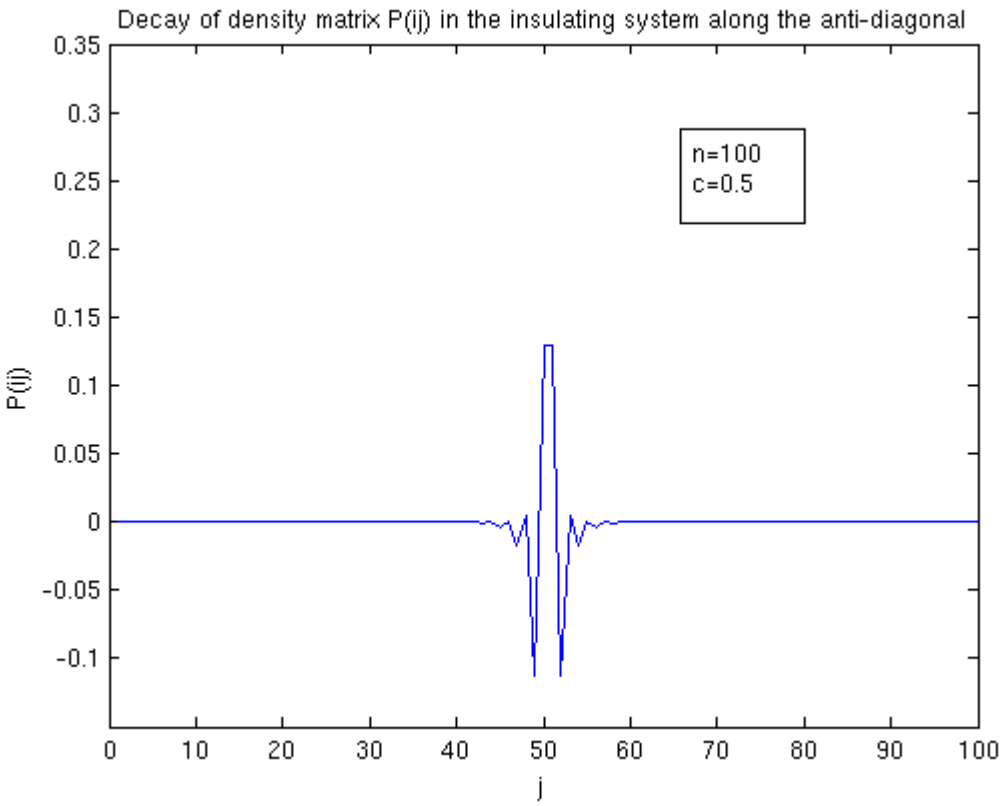

Decay as sin(|i-j|)/|i-j|

Decay as exp(-a|i-j|)

Some numbers for real examples

Atoms: 13'846
Basis functions: 133'214
Basis quality: DZVP

At a threshold 10^{-5}
the percentage non-zero elements is:

H, S : 2%
P : 15%
Inv(S) : 20%

Typical: 20'000 non-zeros per row

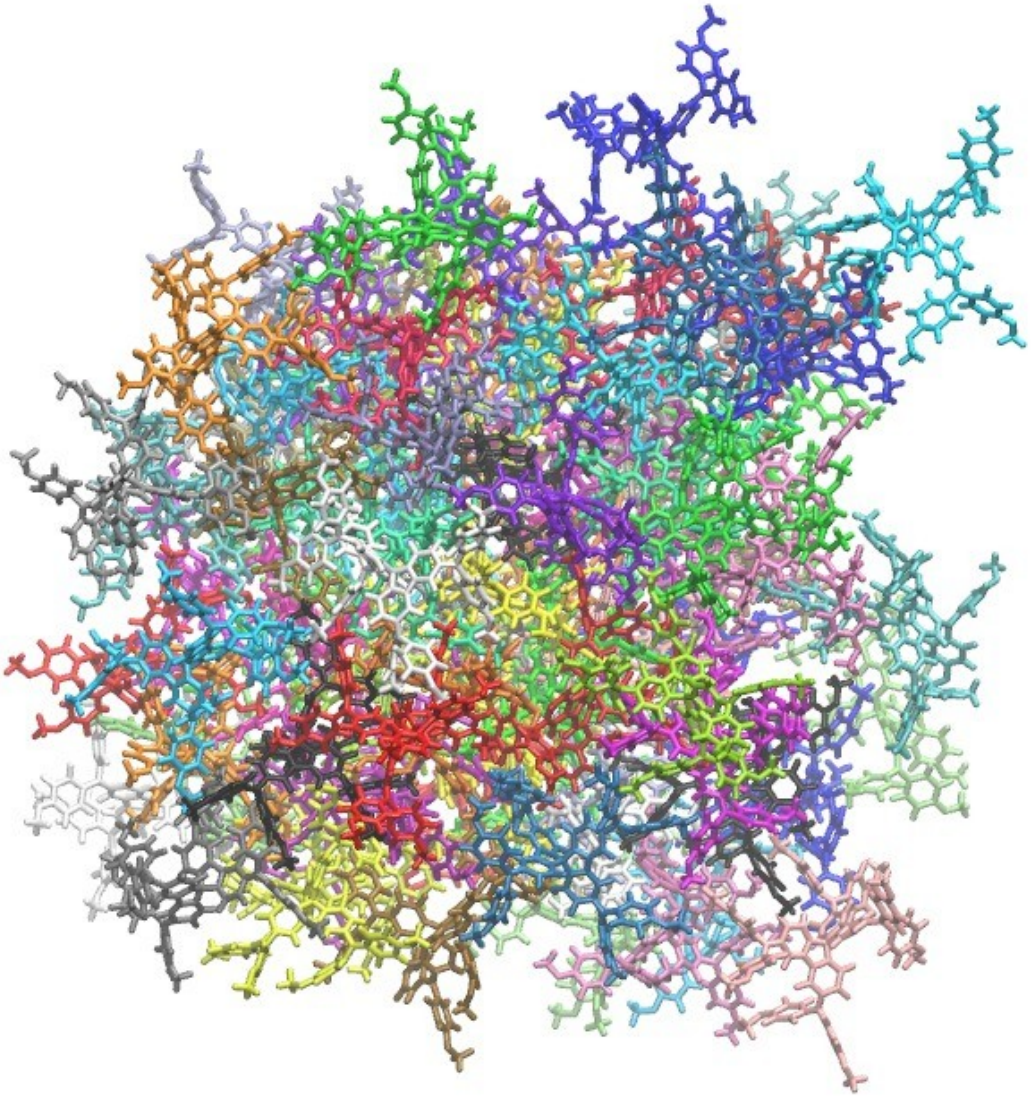

An amorphous hole conducting material used in solar cells

An complementary view on P as f(H)

The density matrix can also be seen as a (matrix) function of H

$$PS = \frac{1}{1+\exp\left(\frac{S^{-1}H-\mu I}{kT}\right)}$$

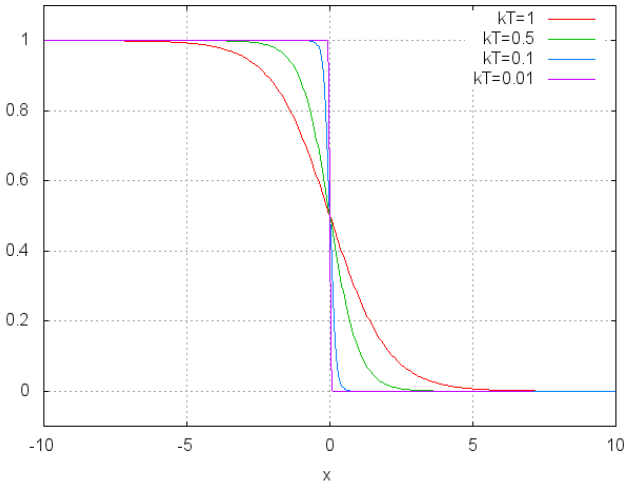

Fermi function

In the limit of small kT a step function is obtained, conveniently written as

$$PS = \frac{1}{2}\left(1-sign\left(S^{-1}H-\mu I\right)\right)$$

If we can exploit the sparsities of H, S and P computing the matrix functions, we can be more efficient than diagonalization based approaches

Almost any way to compute a matrix function can be (has been) tried...
Chebyshev expansion, contour integrals, recursions, minimizations

The matrix sign function

$$\operatorname{sign}(A) = A(A^2)^{-\frac{1}{2}}.$$

For diagonalizable A, eigenvectors of A are eigenvectors of sign(A), with eigenvalues of -1 and 1 respectively

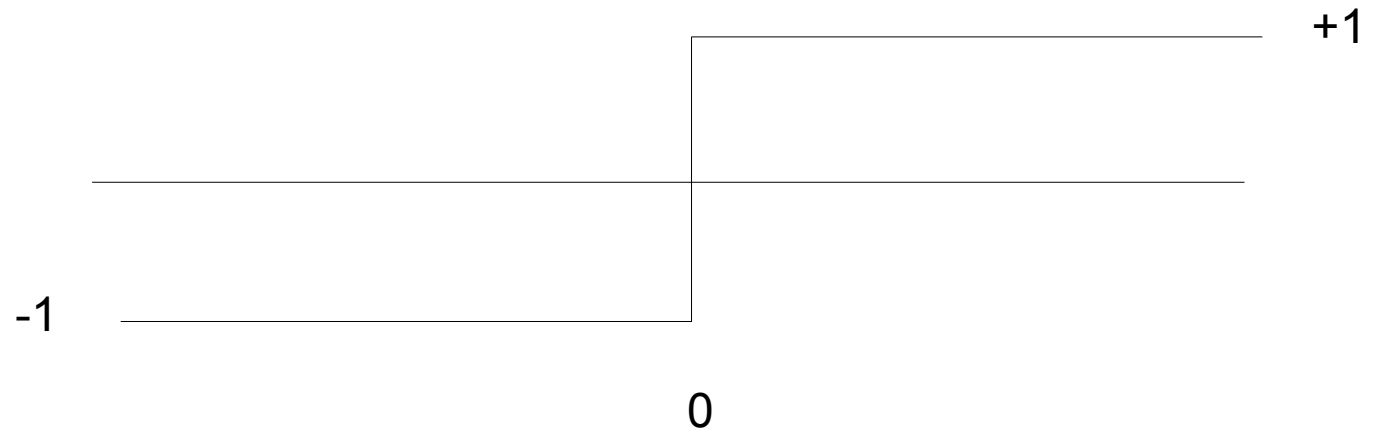

Sign matrix iterations

Various iterative schemes exist to compute sign(A), the simplest is

$$X_{n+1} = \frac{1}{2} X_n (3I - X_n^2).$$

Newton Schulz iteration, requires only matrix multiplications

$$X_0 = cA \quad \xrightarrow[\cdot\, c = ||A||^{-1}.]{} \quad X_\infty = \mathrm{sign}(A)$$

In exact arithmetic convergence is quadratic:
The number of correct digits of X_n is doubled for each iteration

Matrix inversion from sign iterations

$$\operatorname{sign}\left(\begin{bmatrix} 0 & A \\ I & 0 \end{bmatrix}\right) = \begin{bmatrix} 0 & A^{\frac{1}{2}} \\ A^{-\frac{1}{2}} & 0 \end{bmatrix}$$

So, can be used to compute inv(S), sqrt(S), inv(sqrt(S)) ....

Millions of atoms in the condensed phase

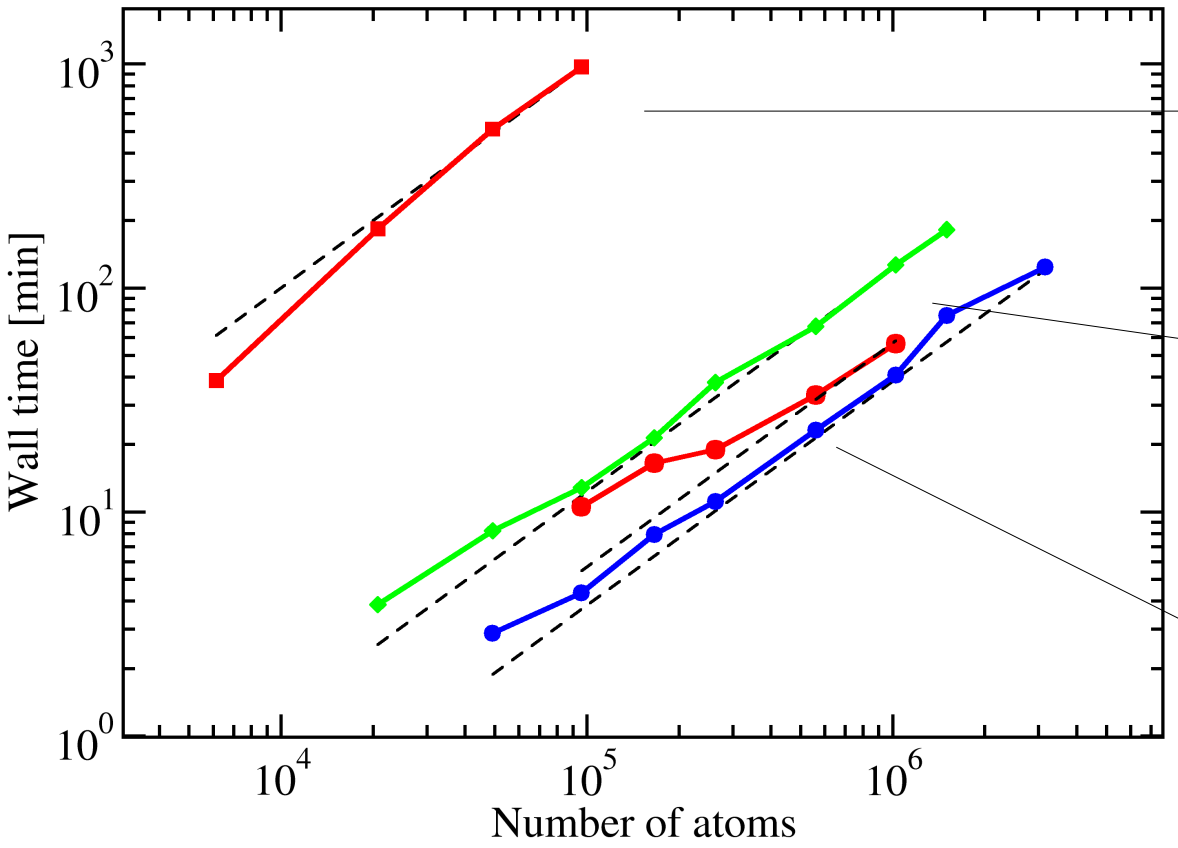

Accurate basis sets, DFT
46656 cores

The electronic structure
$O(10^6)$ atoms in < 2 hours

Minimal basis sets:
DFT, NDDO, DFTB
9216 cores

Bulk liquid water. Dashed lines represent ideal linear scaling.

VandeVondele, Borstnik, Hutter, JCTC, 8, 3565 (2012)

Filtering threshold

The cost of the calculation depends on the criterion (threshold) used to decide if matrix elements are zero.
Roughly: 10x more accurate means 2x the cost.

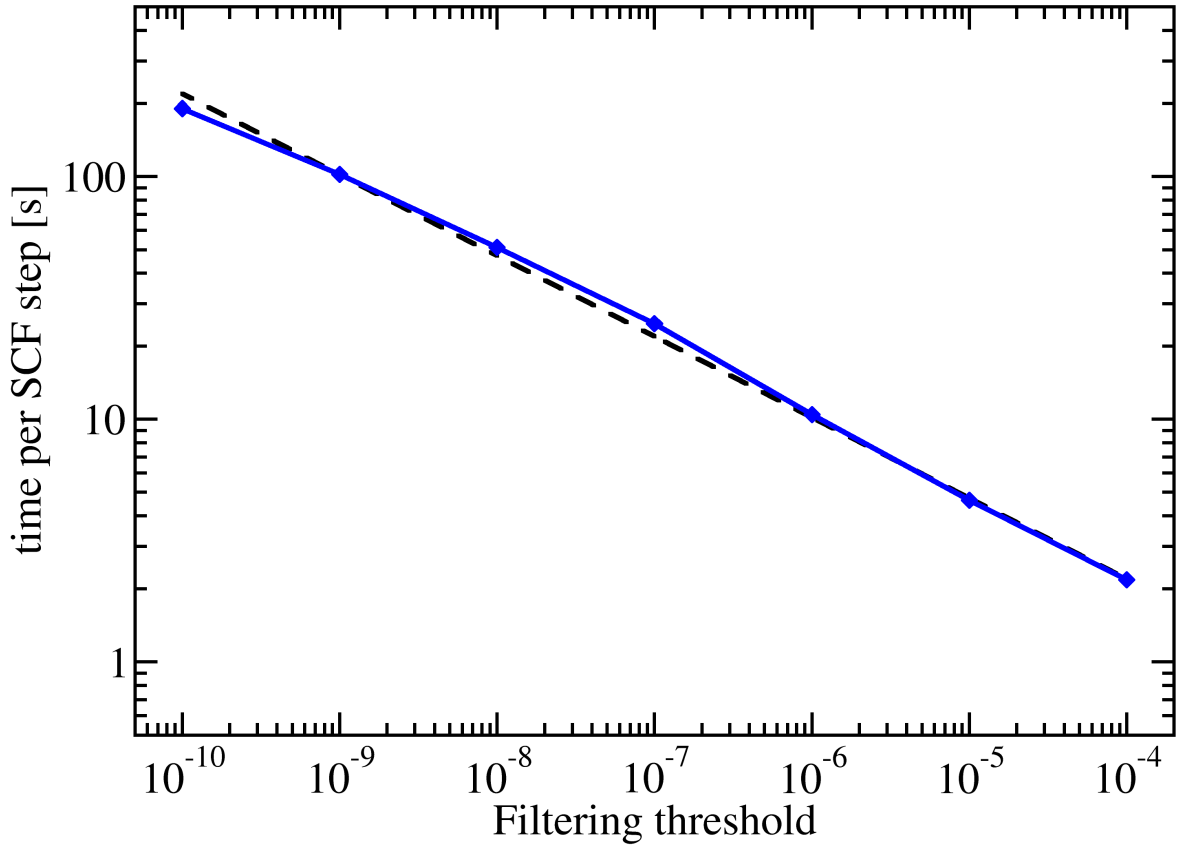

Cost per SCF step for a dftb calculation on 6912 water molecules as a function of the filtering threshold for the matrix multiplication. The dashed line represent as fit using the functional form $a\varepsilon^{-1/3}$. The error in the trace and the total energy at full SCF convergence is linear in the threshold.

More advanced SCF methods

Available in CP2K (among others)

- TRS4 :
  trace resetting purification (Niklasson et al. JCP (2003))
  Automatically determines the chemical potential

- Curvy Steps (Shao et al. JCP (2003)):
  Robust minimization approach based on

$$\mathbf{P}_{k+1} = e^{-\mathbf{\Delta S}} \mathbf{P}_k e^{\mathbf{S\Delta}} \qquad \frac{\partial E}{\partial \mathbf{\Delta}} = [\mathbf{F}, \mathbf{P}_k]$$

- PEXSI (Lin, JP Cond. Matter 2013):
  O(N)-O(N**2) depending on dimensionality, suitable for metals

Further refinements are still possible (we have a backlog of only 10 years) and are being worked on:
The prefactor is all that remains in O(N) methods.

Input section

TRS4

```
! linear scaling SCF
&LS_SCF
  ! TRS4, does not need an estimate for the chemical potential
  PURIFICATION_METHOD TRS4
  ! threshold used to determine sparsity and thus speed and accuracy
  EPS_FILTER 1E-7
  ! convergence for the SCF
  EPS_SCF    1E-5
  S_PRECONDITIONER ATOMIC
&END
```

Curvy steps

```
   ! linear scaling SCF
   &LS_SCF
     ! TRS4, does not need an estimate for the chemical potential
     PURIFICATION_METHOD TRS4
     ! threshold used to determine sparsity and thus speed and accuracy
     EPS_FILTER 1E-7
     ! convergence for the SCF
     EPS_SCF    1E-5
     S_PRECONDITIONER NONE
     &CURVY_STEPS
     &END CURVY_STEPS
   &END
```

Full example on the wiki:

http://www.cp2k.org/exercises:2015_pitt:ls

A sparse matrix matrix multiplication library is crucial for linear scaling SCF

- fully O(N)
- distributed parallel
- suitable for large number of non-zeros per row (20'000)
- 'optimal' for high fill-in (e.g. 10-20%)
- exploiting the natural structure (atomic blocks)
- ...

DBCSR: a sparse matrix library

Distributed Blocked Compressed Sparse Row
Distributed Blocked Cannon Sparse Recursive

Optimized for the science case: 10000s of non-zeros per row.
The dense limit as important as the sparse limit.

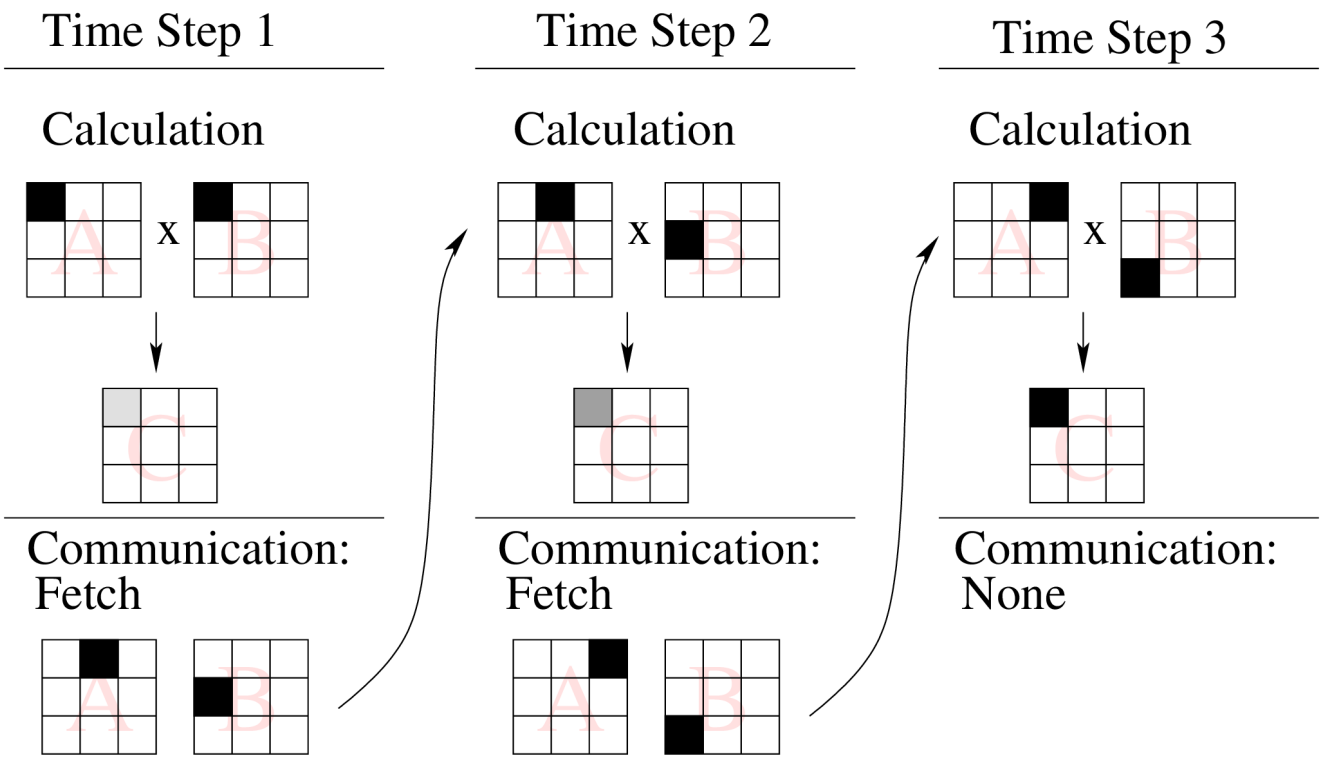

Cannon style communication on a homogenized matrix for strong scaling

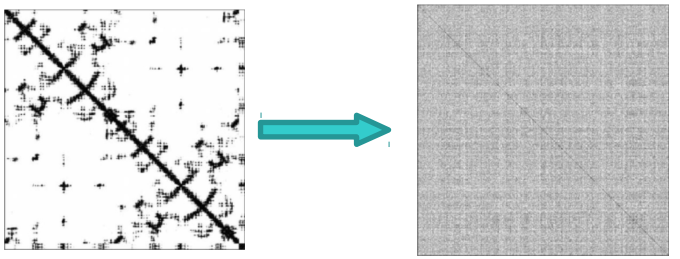

Borstnik et al. : parallel computing (2014)

DBCSR software layout

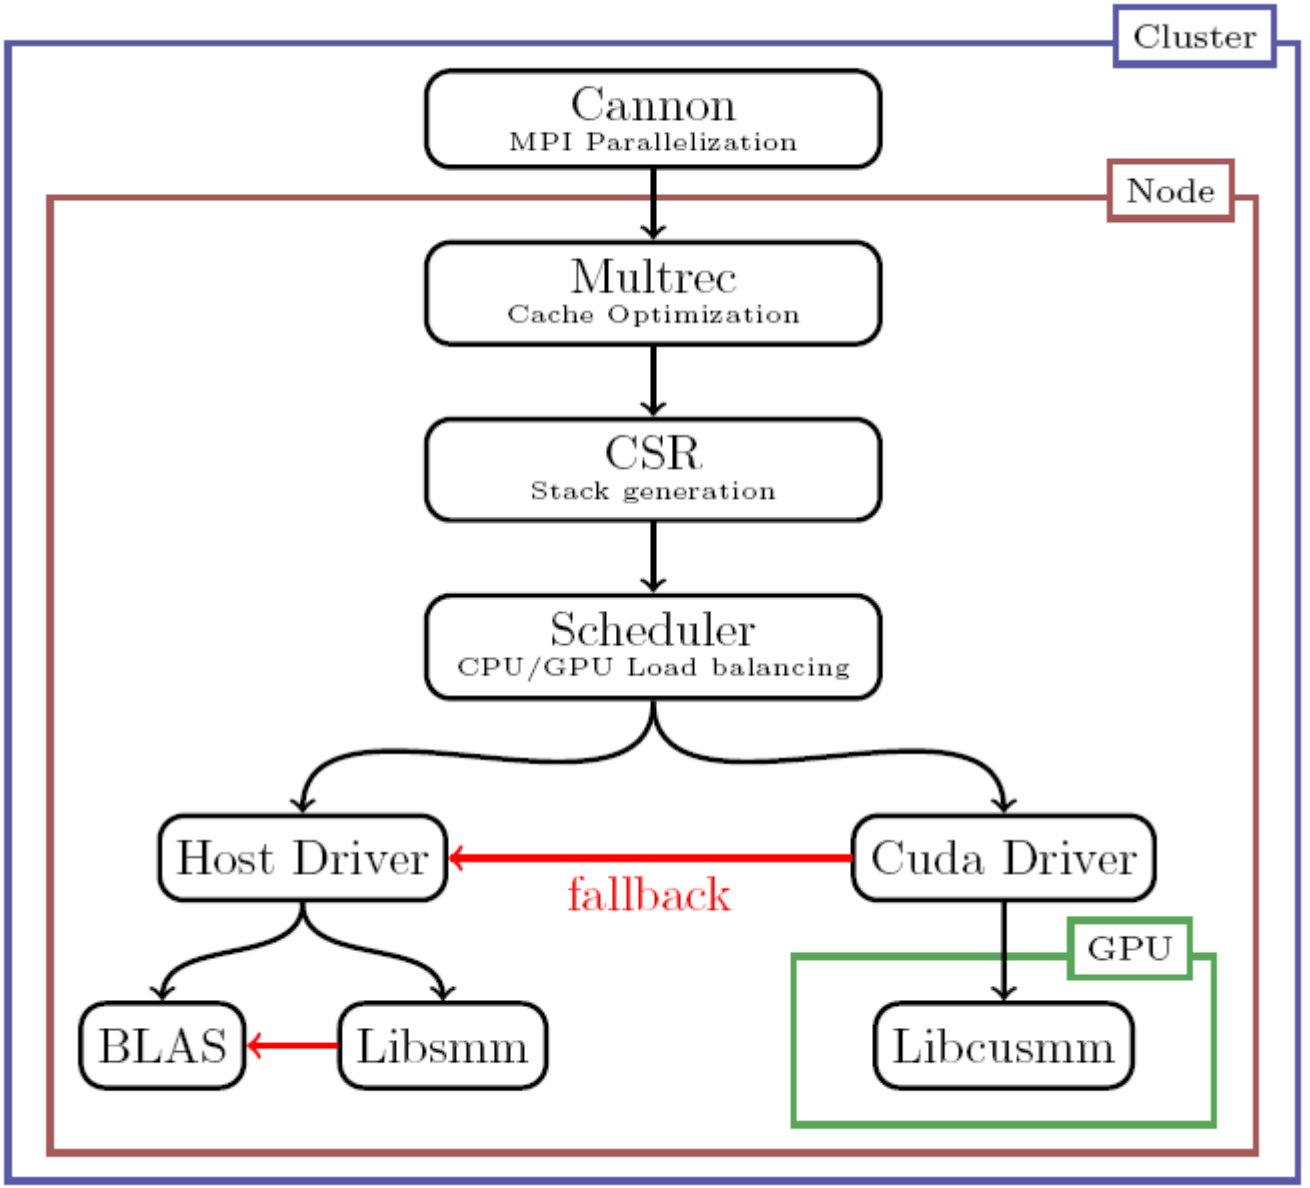

Cannon's Algorithm

2-D grid => √P communication steps

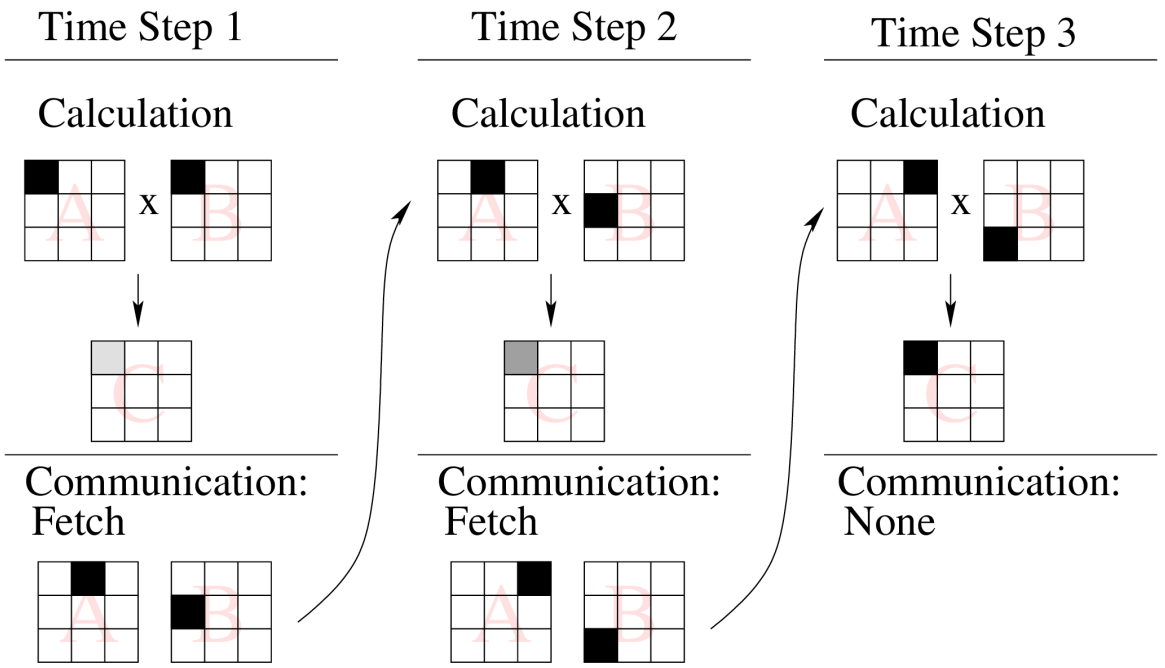

- Reduces to the known, efficient, algorithms in the dense case.
- Avoids worst-case all-to-all communication.
- Communication volume scales as 1/sqrt(P).
- Best performance when the number of ranks is a 'square' number

Performance: strong scaling

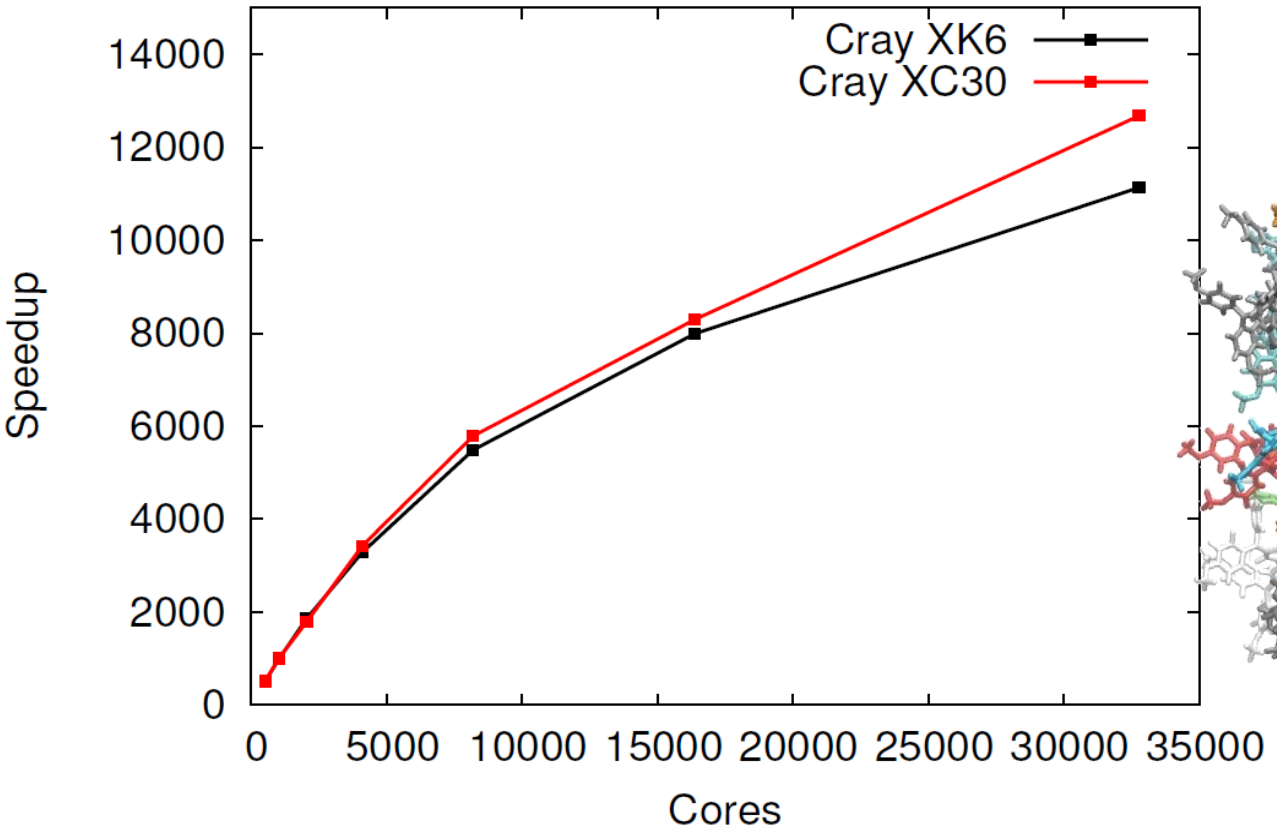

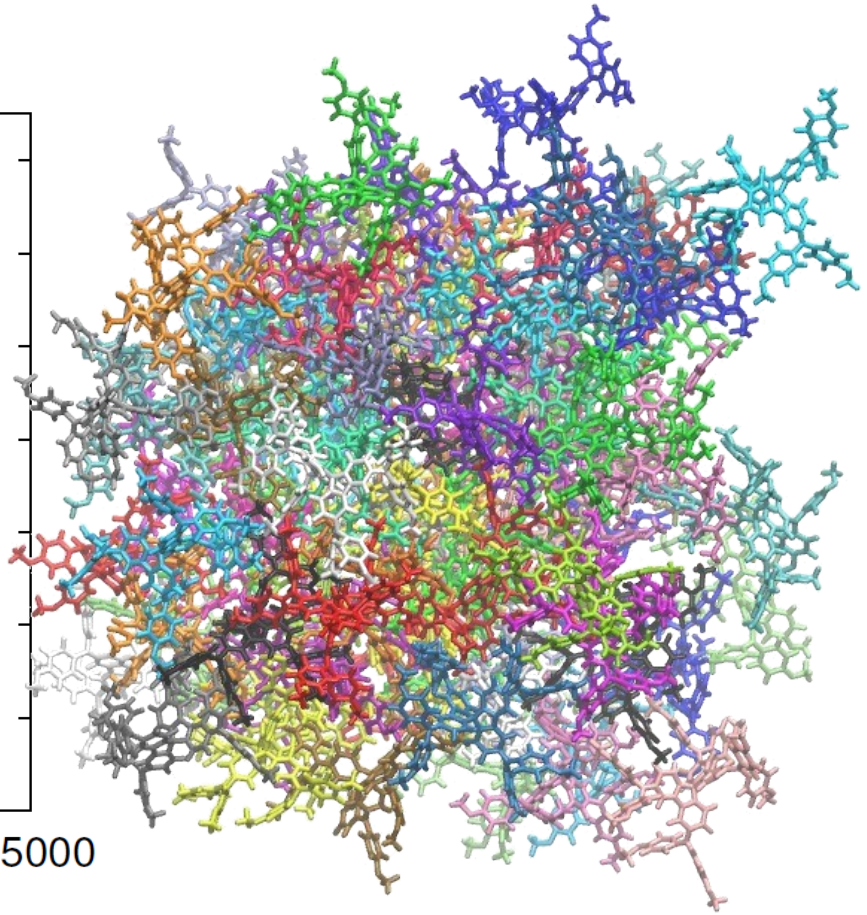

13846 atoms and 39560 electrons (cell 53.84 A), 133214 basis functions.

At full scale-out on the XC30 one multiplication takes less than 0.5s on average, one SCF step 24s.

Towards O(1) :
constant walltime with proportional resources

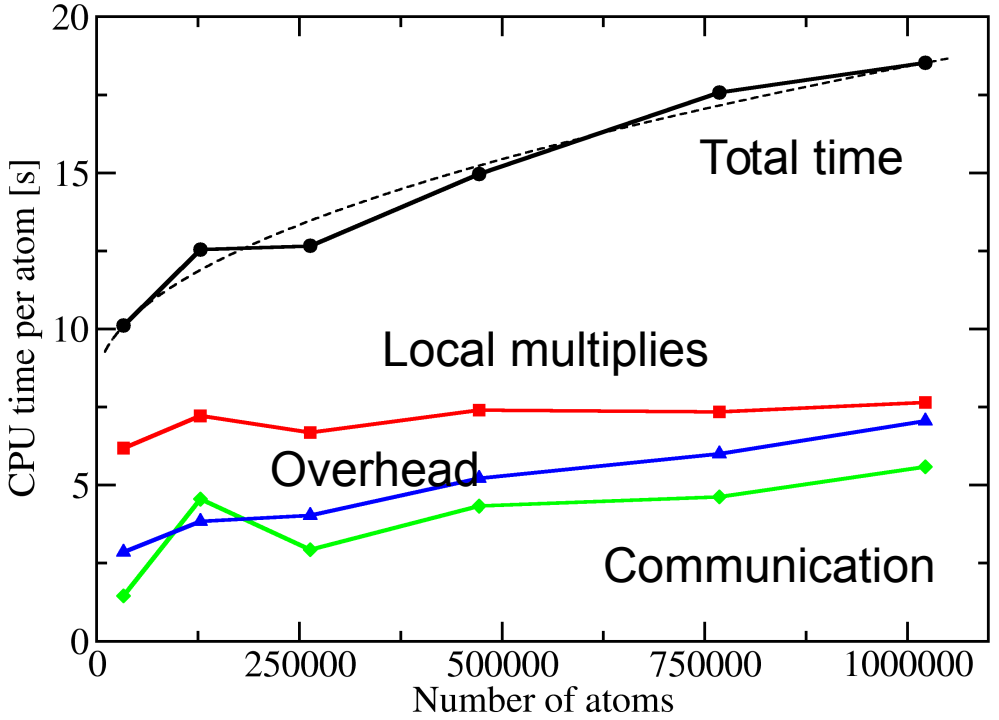

Stringent test:
- Small blocks, large overhead
- Very sparse matrices
- Running with 200 atoms / MPI task

Local multiplies constant (OK!).

Overhead & Communication
- Grows with sqrt(N)
- Needs a replacement for Cannon

Work is underway to replace the Cannon algorithm with something new!
Retain the sqrt(N) max comm, yield constant comm in the limit.

Host: LibSMM

Autotuning framework for small matrix matrix multiplications

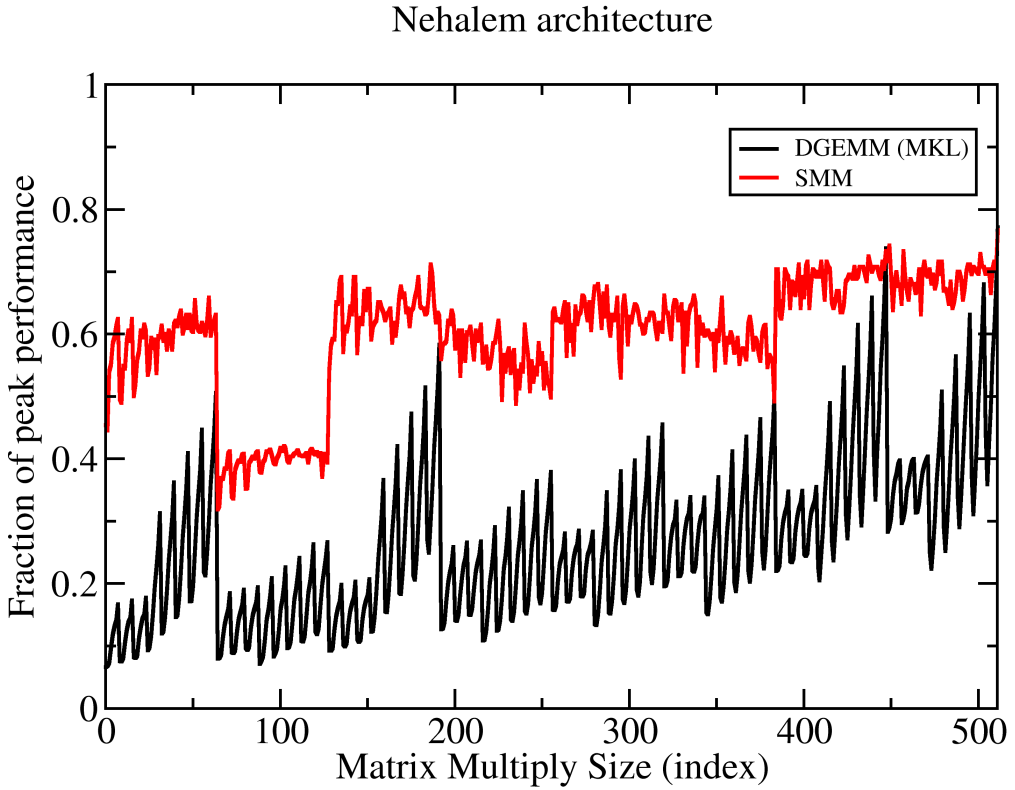

Chemical block sizes: 1,4,5,6,13,23,26
Generate autotuned kernels
Similar to ATLAS, FFTW,...
Unrolling, loop ordering, ….
Compiler based (no asm).

LibCUSMM performance

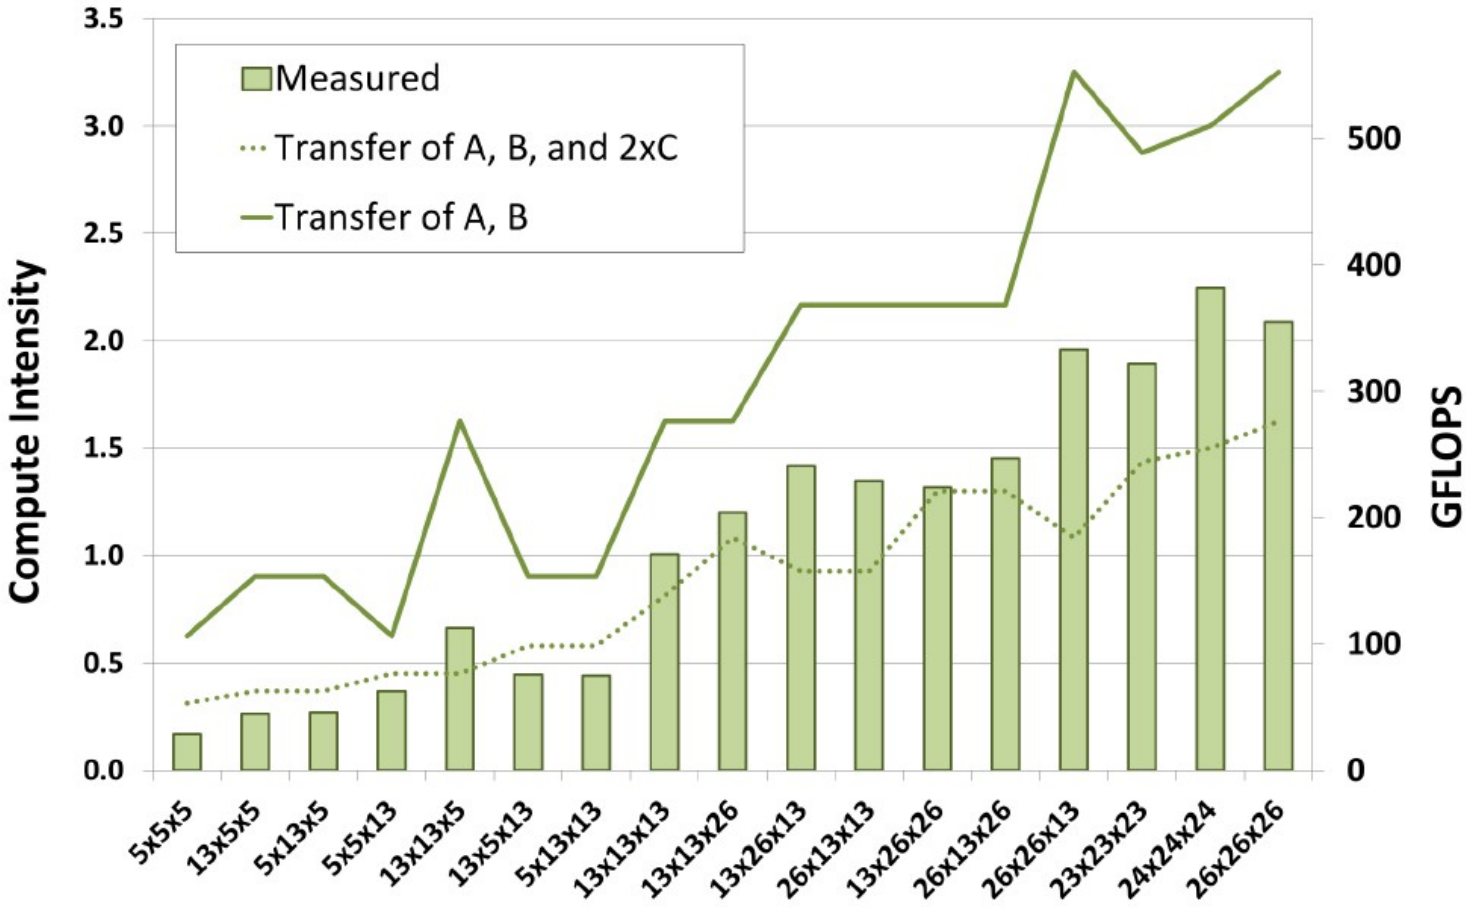

Measured performance against a roofline model based on memory transfer

1st CPU-GPU comparison

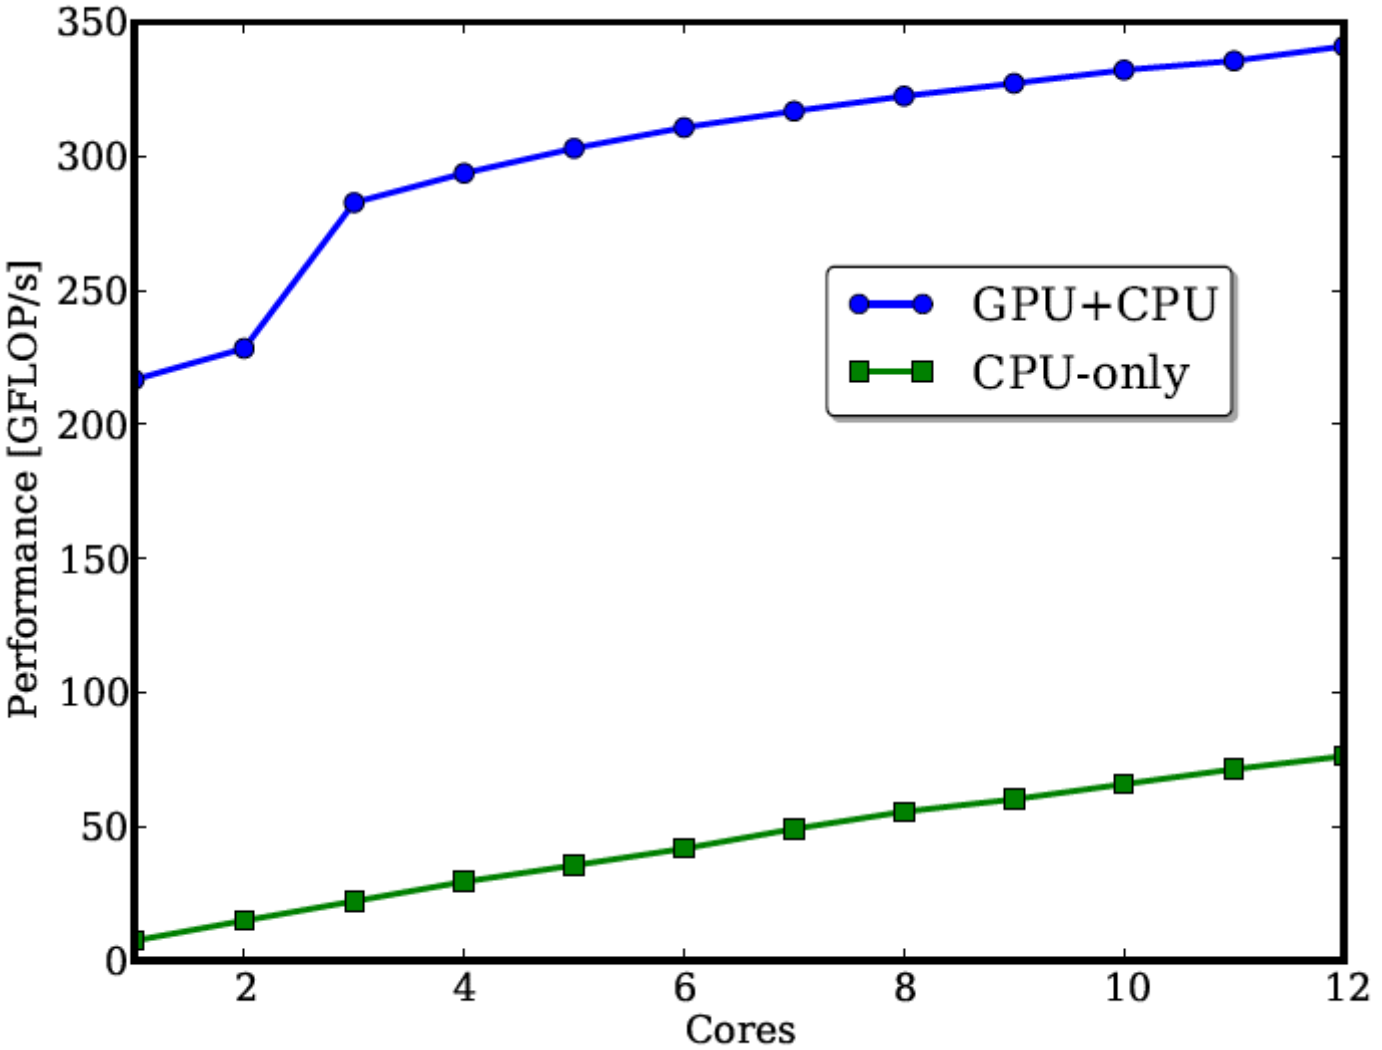

Performance comparison of the multi-threaded DBCSR library based on 23x23 matrix blocks, and was not using the MPI capabilities. The benchmark was run on a dual Sandy Bridge (E5-2620, 2.0GHz, 6 cores) machine, equipped with one NVIDIA Tesla K20 card.

Hybrid Daint vs dual SB Daint

Daint: XC30, 8 cores + 1 GPU / per node, ~5200 nodes
Fastest computer in Europe.

Three science benchmarks:
various block sizes, CPU loads, and communication.

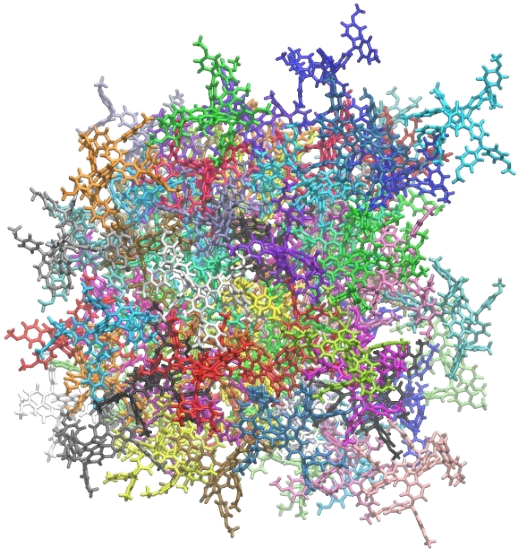

Amorph... a hole conducting solar cell material

	Amorph	H2O	TiO2
2 SB	372	275	446
1 SB + 1K20X	272	187	263
Ratio	1.37	1.47	1.70
GPU flop %	92	99	88

Balanced

Comm. limited

Small blocks

Canonical benchmarks on 169 nodes
(slightly old code, in particular 2 SB)

CPU-GPU on hybrid Daint

# nodes	1 CPU-only	1 CPU + 1 GPU	CPU+GPU Blocked	CPU / GPU ratio
3844	617s	459s	406s	1.5
1024	2208s	1351s	1054s	2.1
512	4046s	2566s	1341s	3.0
256	7124s	4686s	OOM-GPU	N/A
128	14268s	OOM-GPU	OOM-GPU	N/A

As expected, GPU benefit decreases as communication becomes important.

GPU memory limit (~6 Gb) is triggered in these tests,
no further memory vs. speed trading possible (i.e. 2.5D/3D multiplication)

Blocking groups 'atoms into molecules', improves data-locality but increases total data and flops: 85 PFLOP vs 132 PFLOP, 256 vs 512 nodes needed

Testcase 'H2O-dft-ls-orig' : 20'000 atoms

Historical comparison

1) Run on Jaguarpf (XT5, 2011-01-01), 3888 nodes (12 cores)
2) Run on Daint (XC30, 2013-11-17), 3844 nodes (8 cores + 1 GPU)

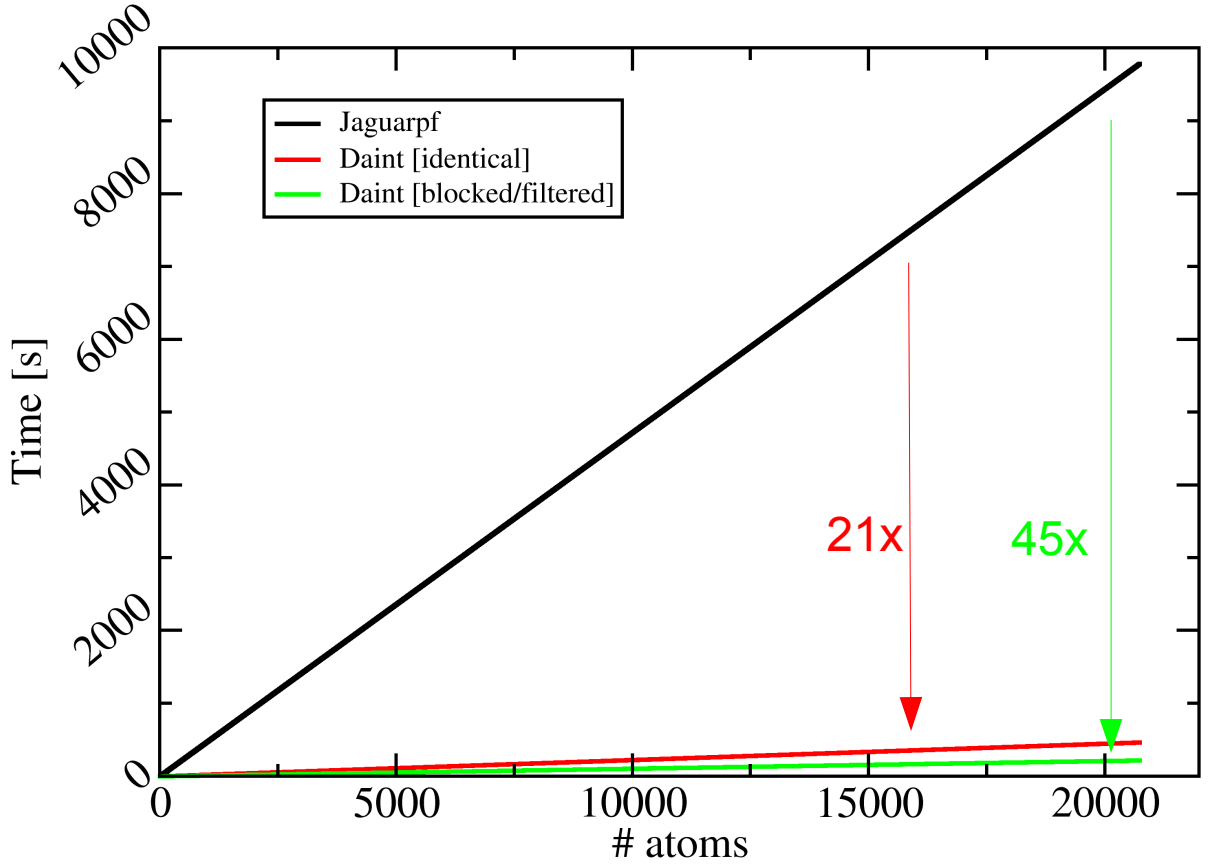

Testcase 'H2O-dft-ls-orig' : 20'000 atoms

Original run published in : VandeVondele, Borstnik, Hutter, JCTC, 2012

Bridging from linear scaling SCF to materials properties

2D polymers: synthetically tailored 2D materials beyond graphene

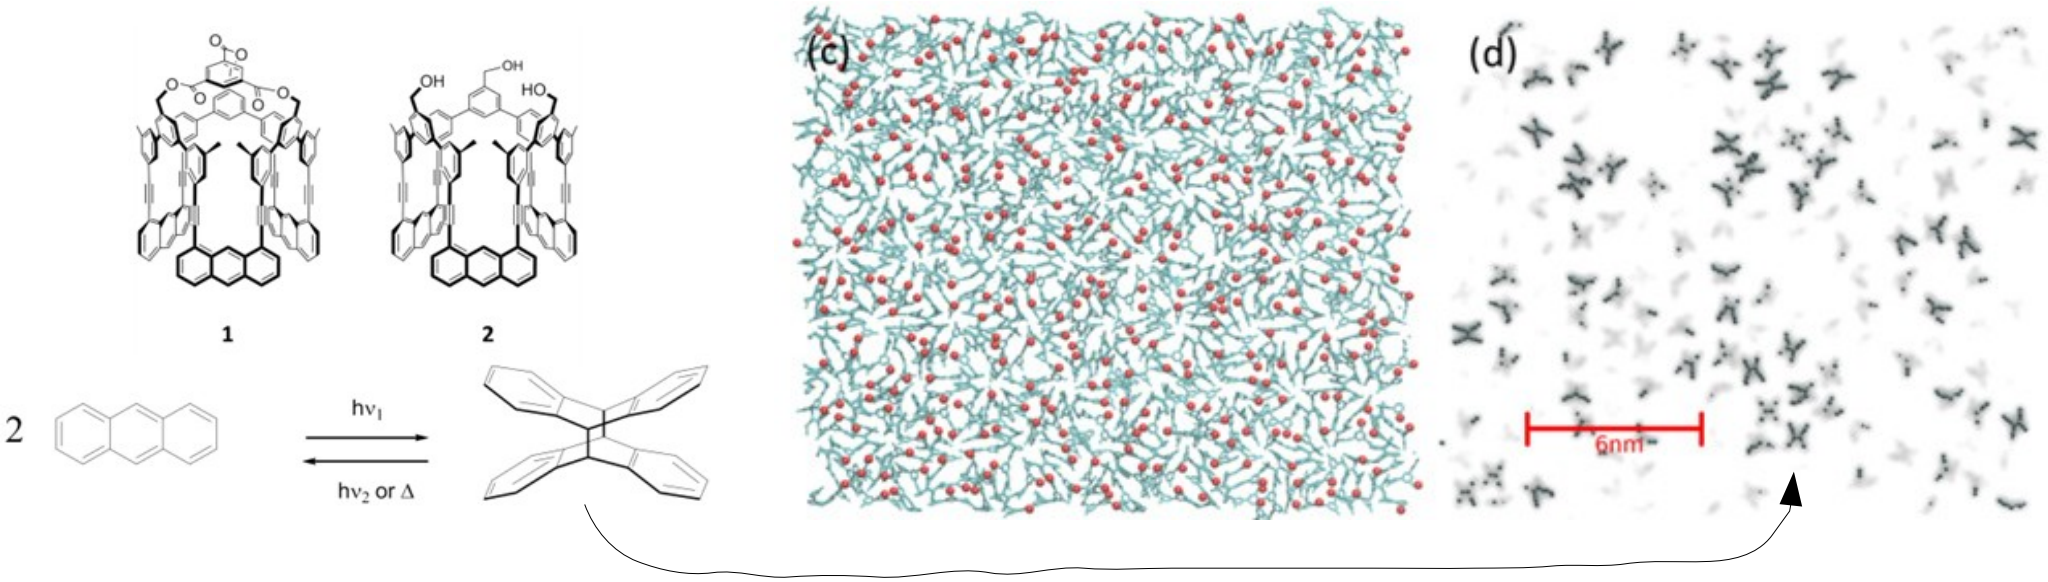

Based on <u>linear scaling MD simulations for 10'000s of atoms</u>, the morphology and properties of the proposed 2D polymer sheets has been investigated

Payam Payamyar, Khaled Kaja, Carlos Ruiz Vargas, Andreas Stemmer, Daniel J. Murray, Carey Johnson, Benjamin T. King, Florian Schiffmann, Joost VandeVondele, Alois Renn, Paola Ceroni, Andri Schütz, Lay-Theng Lee, Zhikun Zheng, Junji Sakamoto, A. Dieter Schlüter, Accepted in ADVANCED MATERIALS (2013).

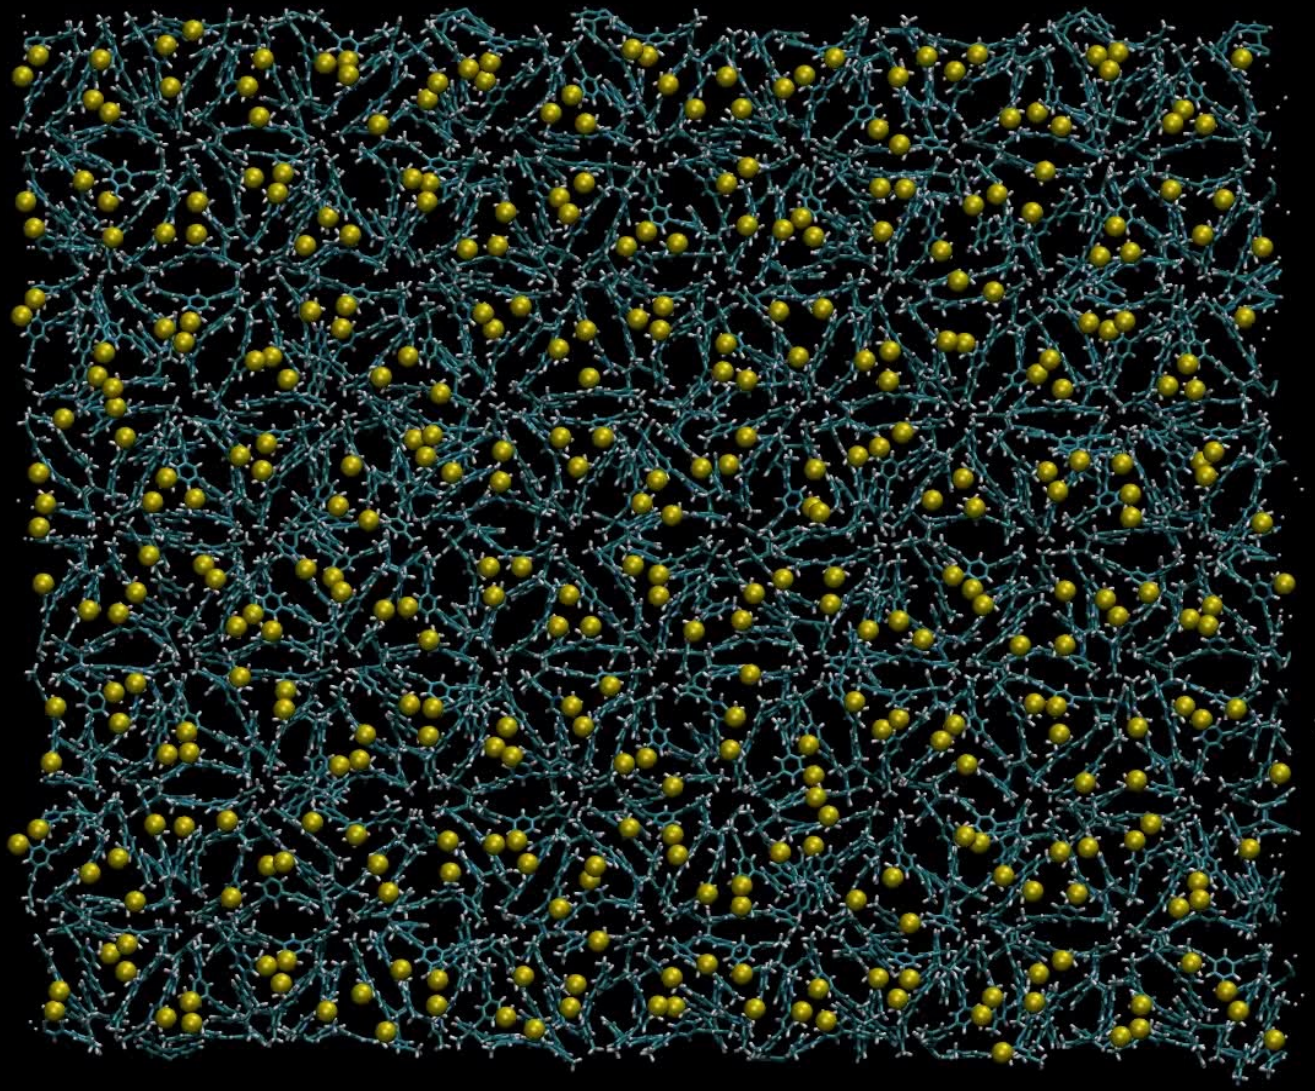

Electronic properties of TiO_2 nanocrystals

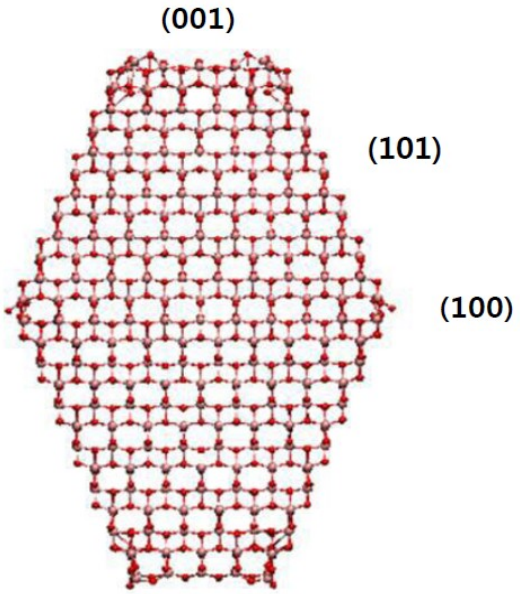

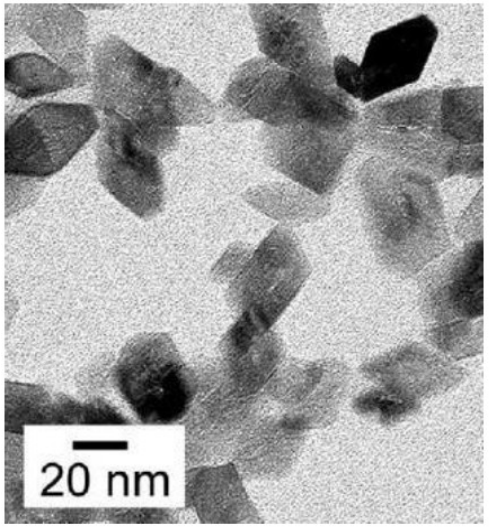

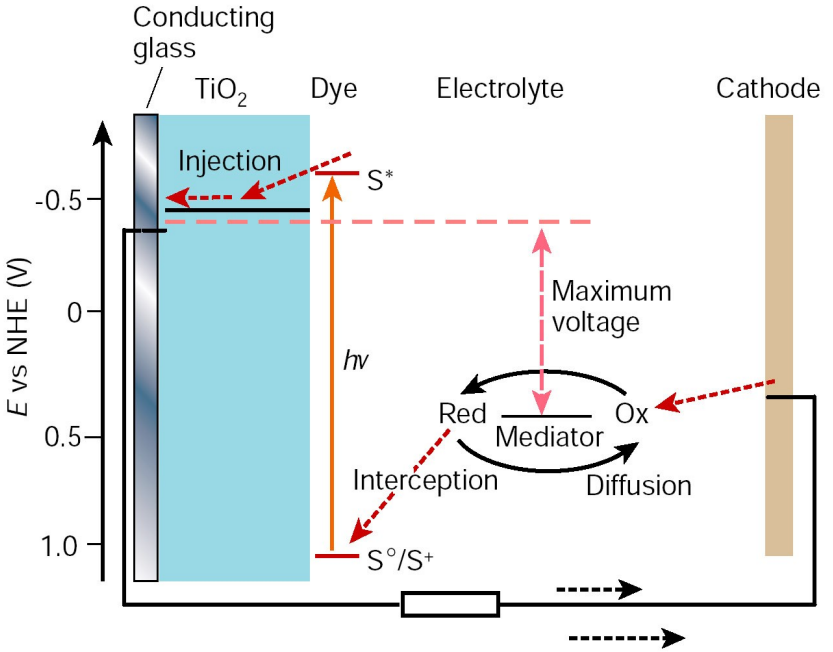

Grätzel, Nature (1991,2001)

TiO2 nanoparticles are a key ingredient in various systems, including Dye Sensitized Solar Cells.

Models in explicit solvent

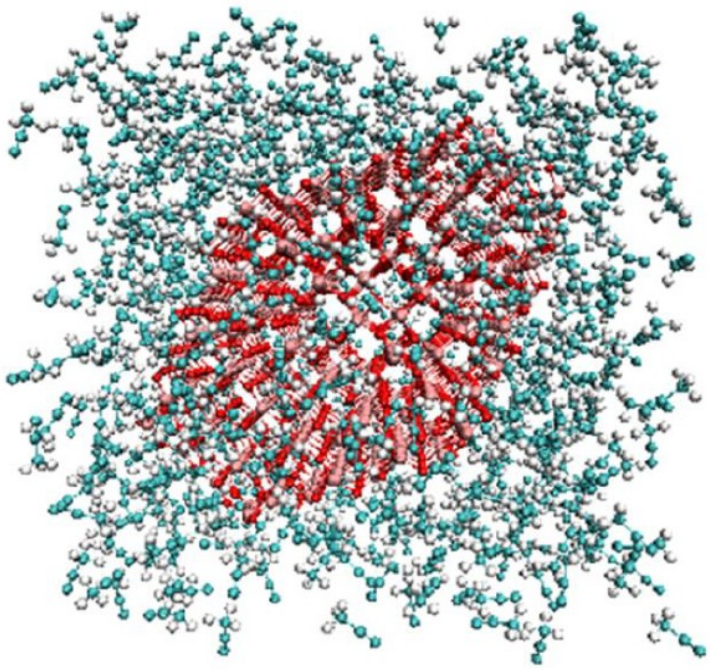

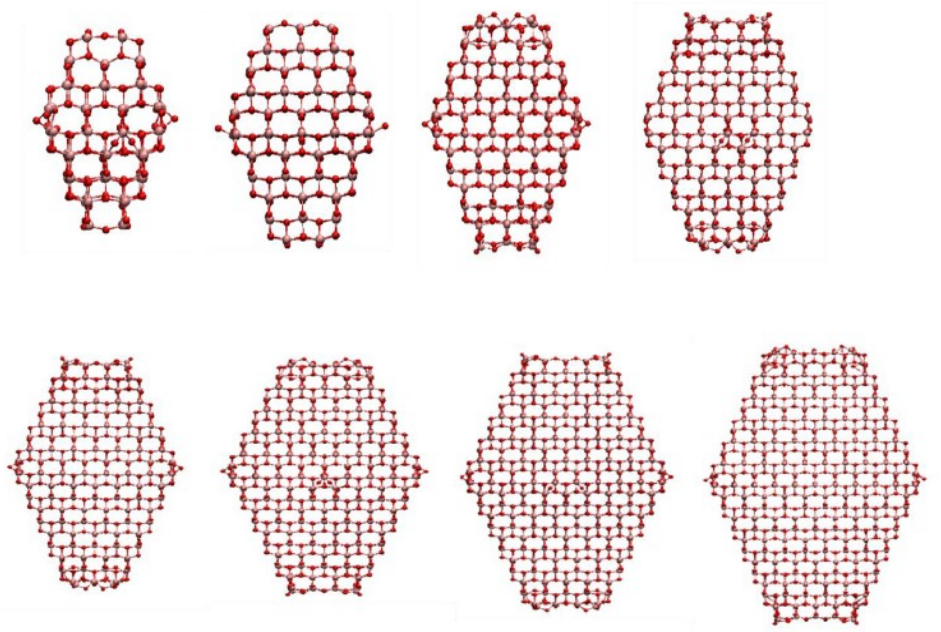

H = **2nm, 2.6nm, 3.2nm, 3.7nm, 4.4nm, 5nm**, 5.6nm, **6nm**

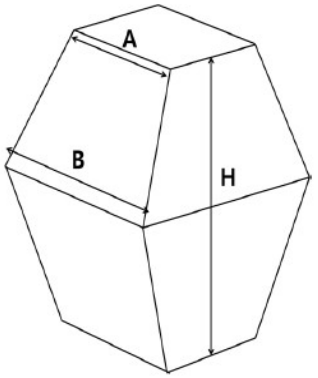

A/B = 0.33~0.5

Sizes ranging from 44k to 118k basis functions

ACN solvent treated with the Kim-Gordon DFT model: naturally suited for linear scaling

Full system science case

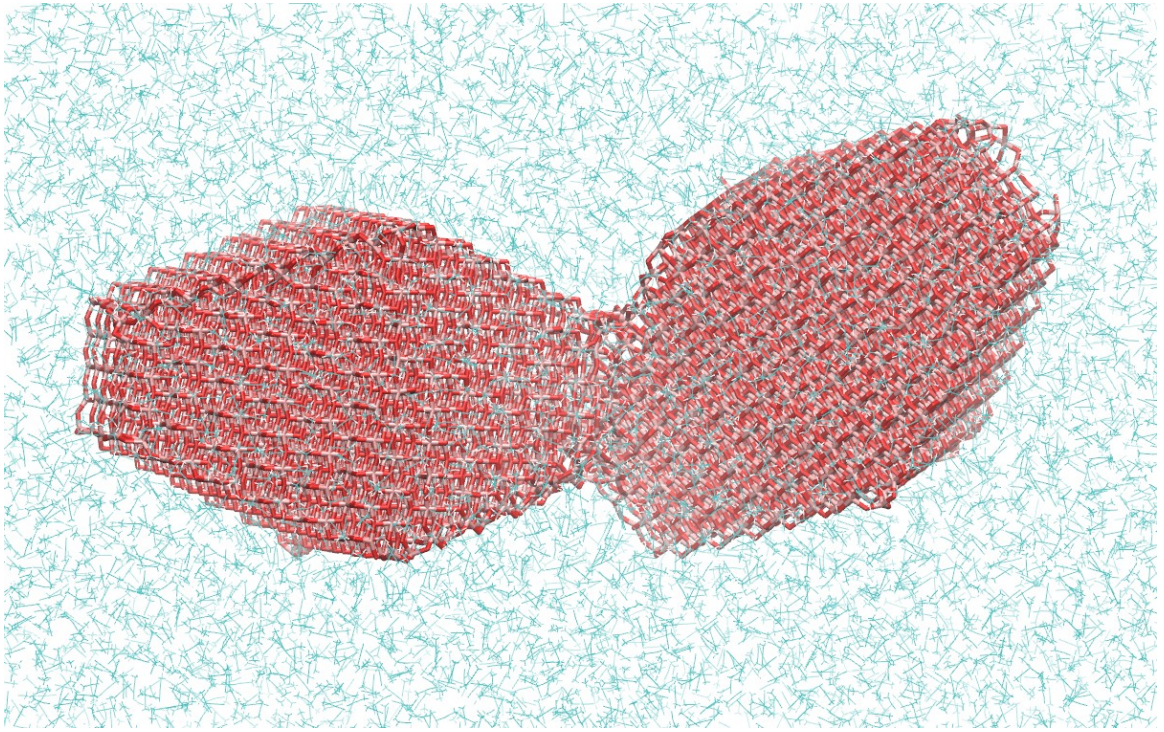

80'000 atoms DFT, high accuracy settings
Aggregated nanoparticles in explicit solution
Relevant for 3rd generation solar cells

Matrix dims ~ 772868 x 772868
Threshold ~1E-6
% non-zero ~ 4%
SCF steps ~ 50
multiplies needed ~ 2000

Dense flops needed =
184661334367982412800O

Actual flops needed =
84992840373629580 2

Sparsity boost = 2172x

GPU flop % = 99.4

Time on 5184 nodes = 6264s

Sustained actual flops = 0.13 PF

Sustained dense flops = 294.7 PF

With explicit MD based equilibration

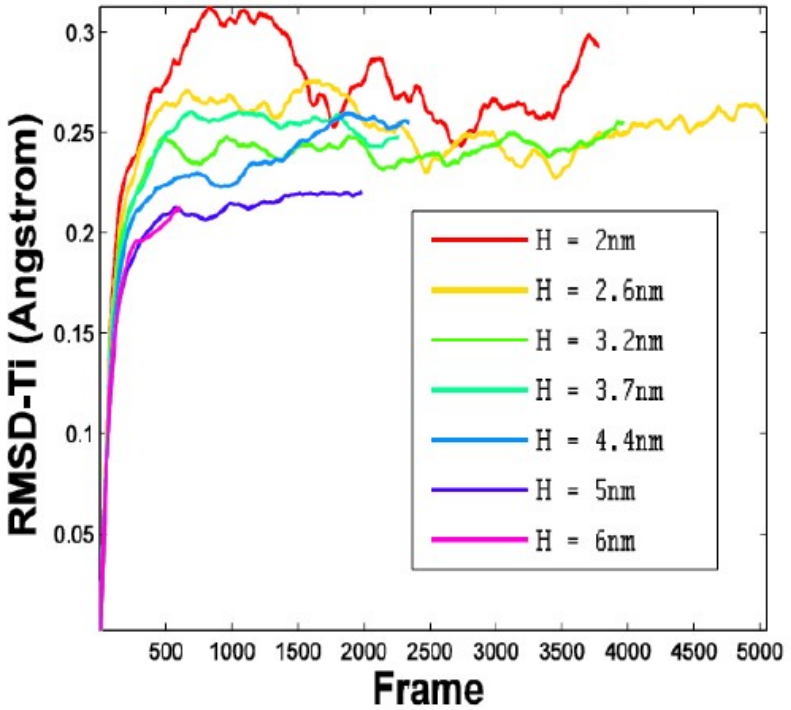

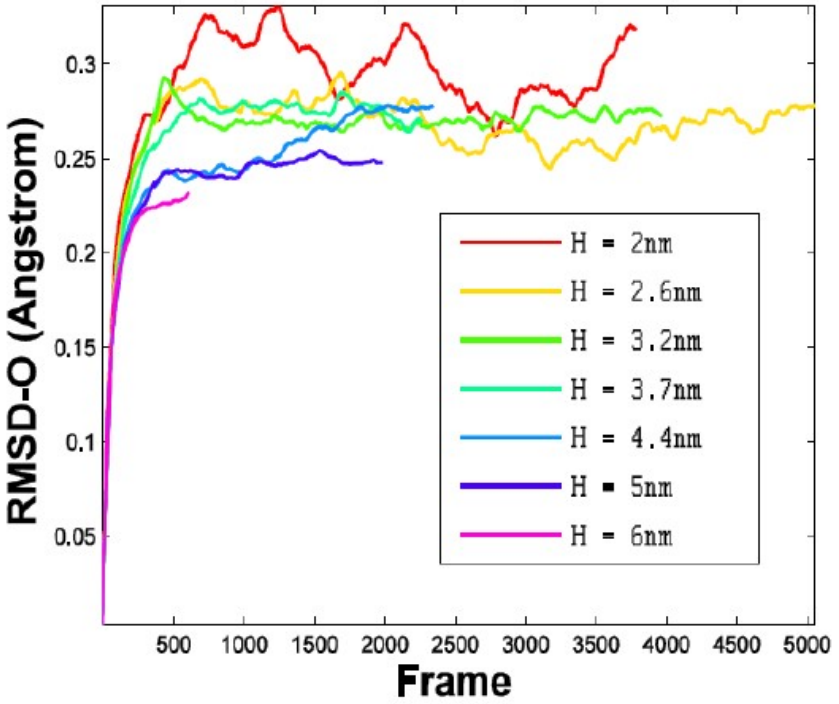

A crucial aspect for linear scaling calculations... how to get reliable structures. Typical empirical models are often not good enough.

Yielding detailed geometric information

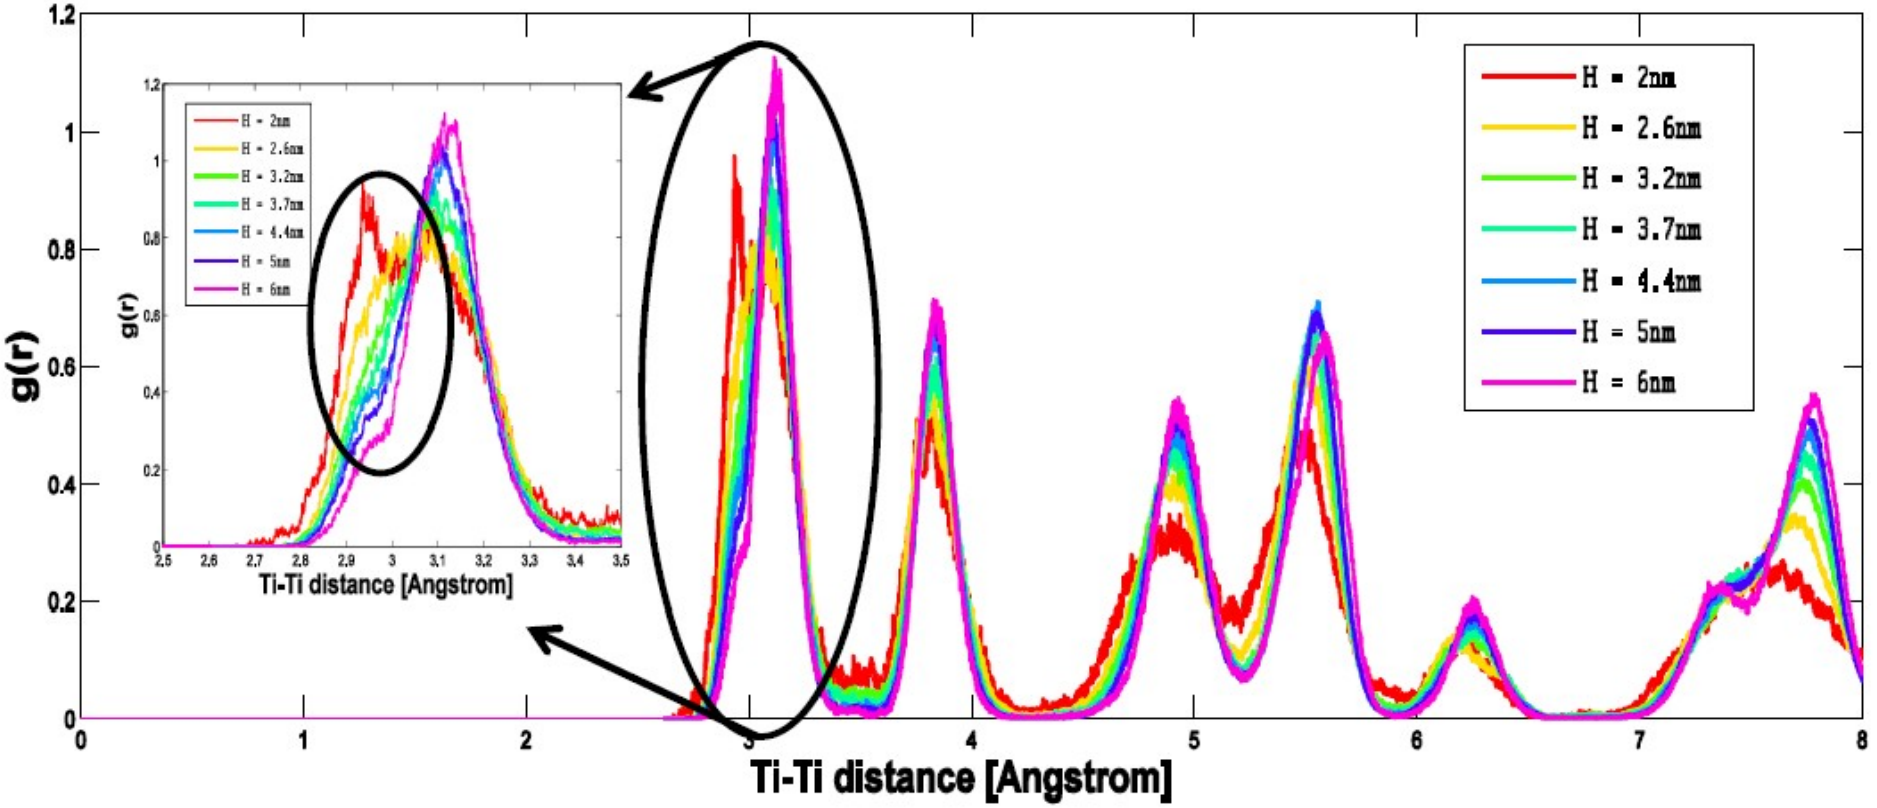

Large fluctuations in smaller crystals, and compressed surfaces.

Electronic properties

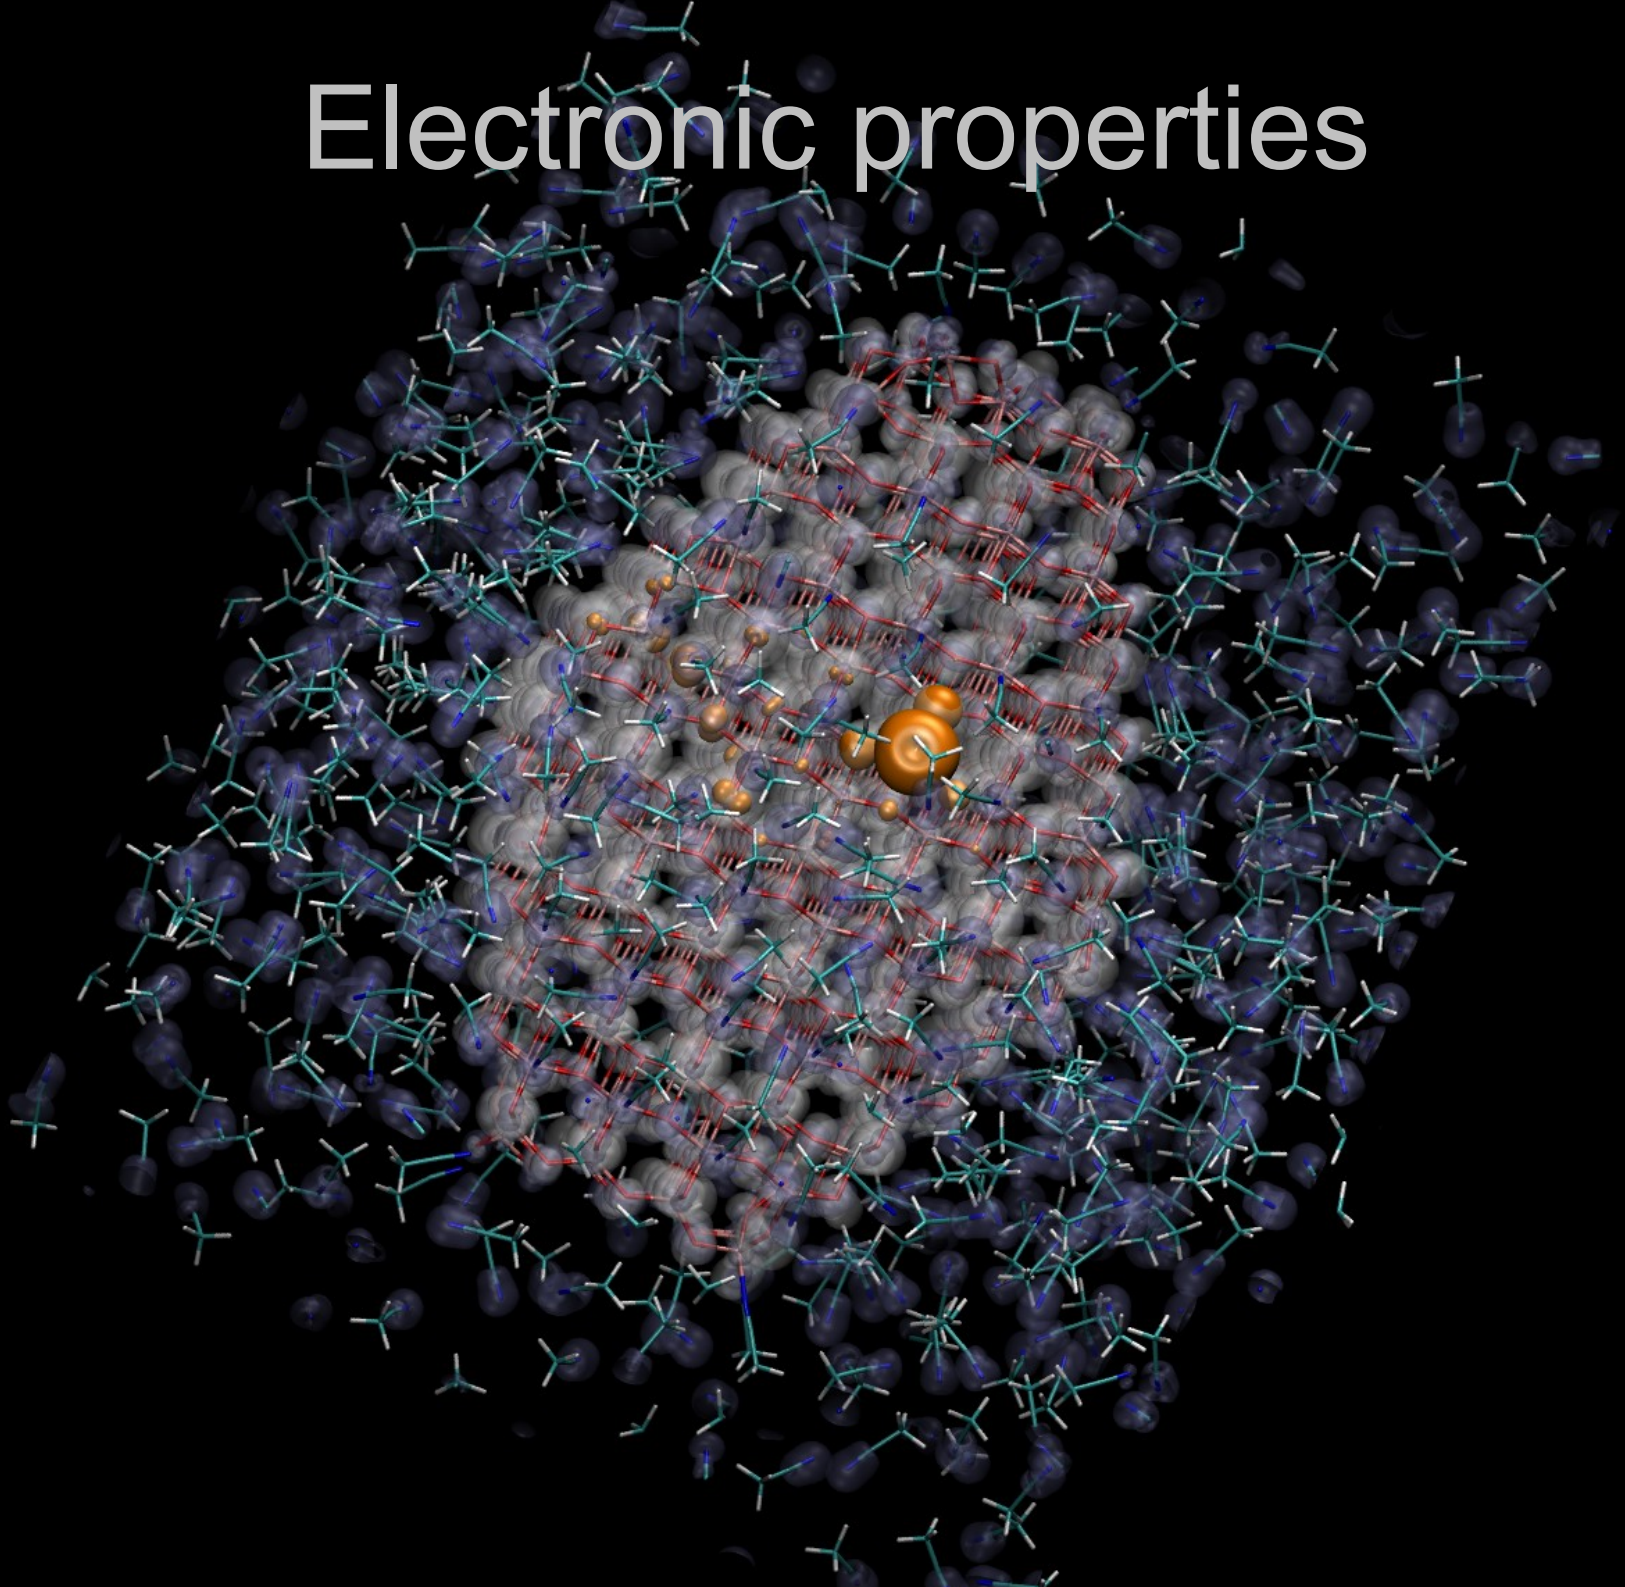

Towards ab initio device simulations

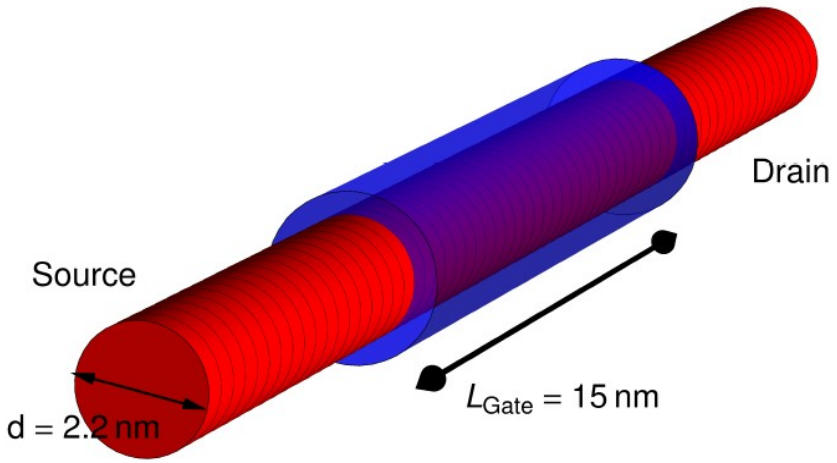

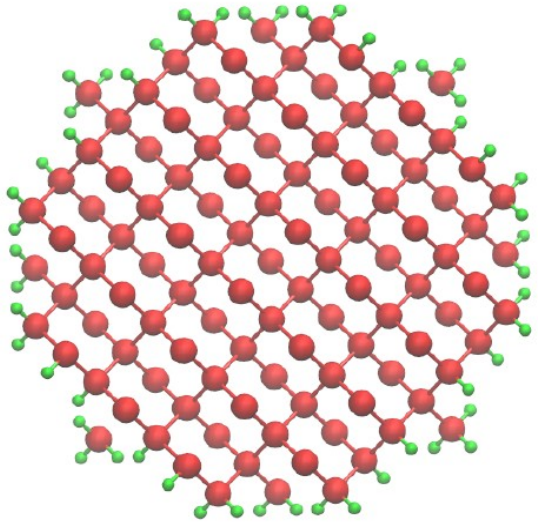

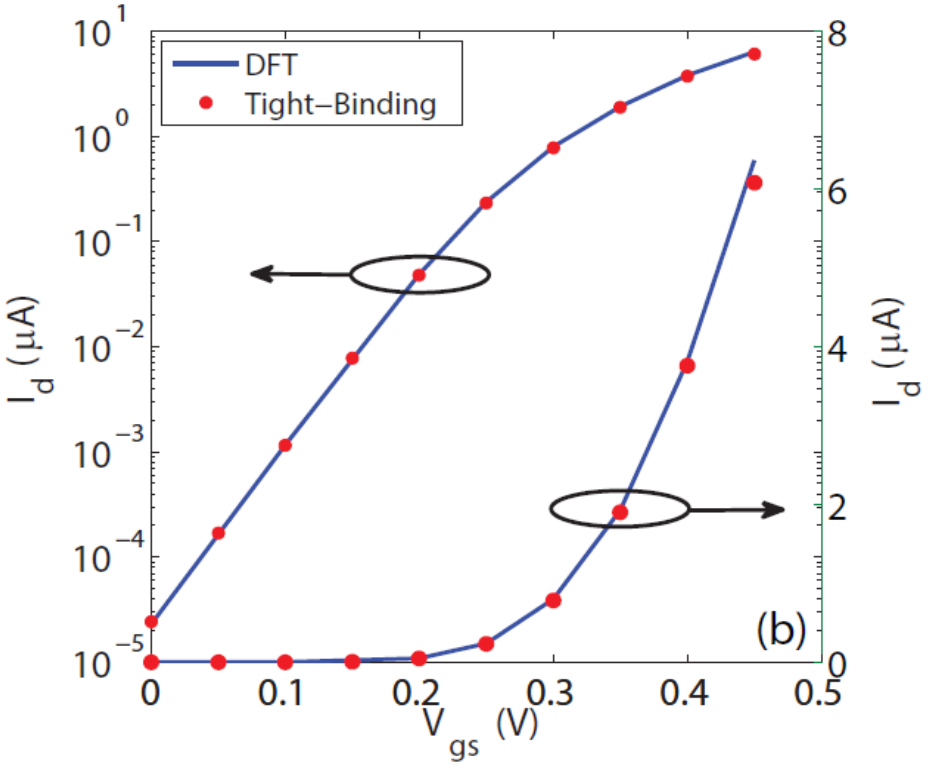

10000 atoms NEGF calculations on Si NWFET, a coupling between OMEN and CP2K

Brück S. Calderara M Bani-Hashemanian MH, VandeVondele J, Luisier M Proc. Int. W. Comput. Electronics (IWCE). (2014).

Electronic dynamics

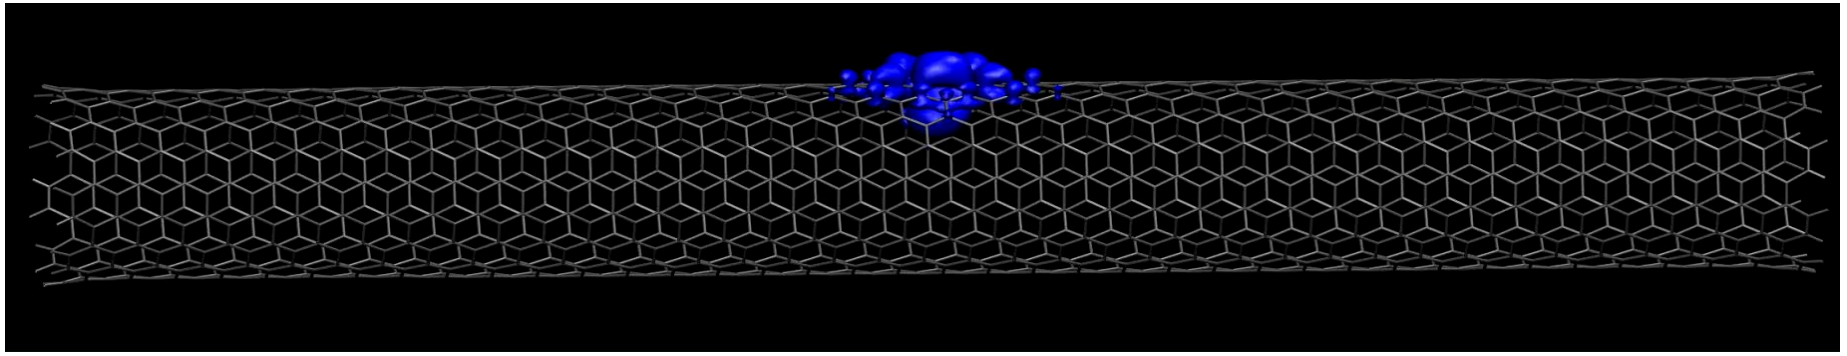

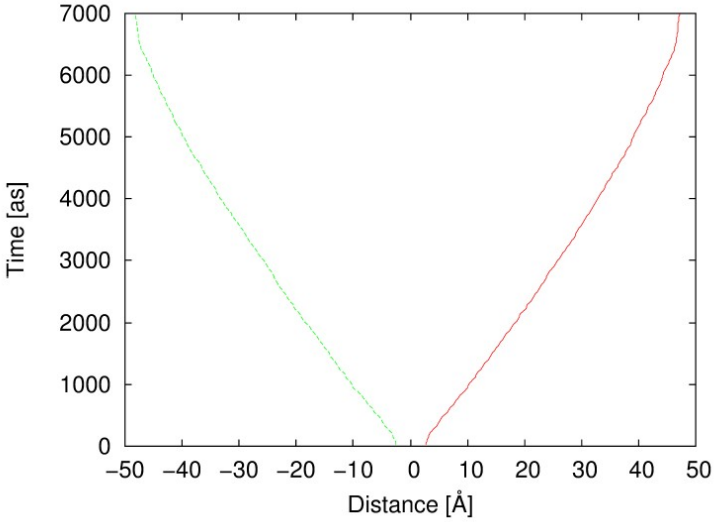

Ehrenfest dynamics:
- nuclear and electronic motion.
- Timescale : 9370as (1874 steps)
- 10nm tube (1440 Carbon atoms)
- Computed in 3 days (144 nodes XC30)

Quantify the electronic spreading by the motion of the front.

Conclusions

- Full linear scaling is possible, the prefactor is key
- With current parallel computer we can start probing a regime where linear scaling is the only option
- A dedicated, specialized matrix multiplication library has been developed.
- Nanoparticles of interesting sizes have become within reach

Acknowledgements

ETH
Florian Schiffmann
Ole Schütt
Samuel Andermatt
Martin Häufel
Hossein Bani Hashemian
Clelia Spreafico
Florian Thöle
Oliver Gindele
Cha Jinwoong
Luisier Mathieu
Sascha Brück

UZH
Juerg Hutter
Urban Borstnik
Christiane Pousa
Valery Weber

Chris Mundy (PNL)
Iain Bethune (EPCC)
Thomas Schulthess (CSCS)
Neil Stringfellow (CSCS)
Peter Messmer (NVIDIA)
Roberto Ansaloni (CRAY)

EU FP7

HP2C
INCITE
PRACE

CSCS
ORNL

You for your attention!