SAND2018-0843PE

Electronic structure software development: something old & something new

Jonathan Moussa

Sandia National Laboratories

Jan. 19th, 2018

Overview

1. HSE functional implementation in SeqQuest
2. Some thoughts on scientific software development
3. Localized & randomized linear-scaling electronic structure

History of SeqQuest

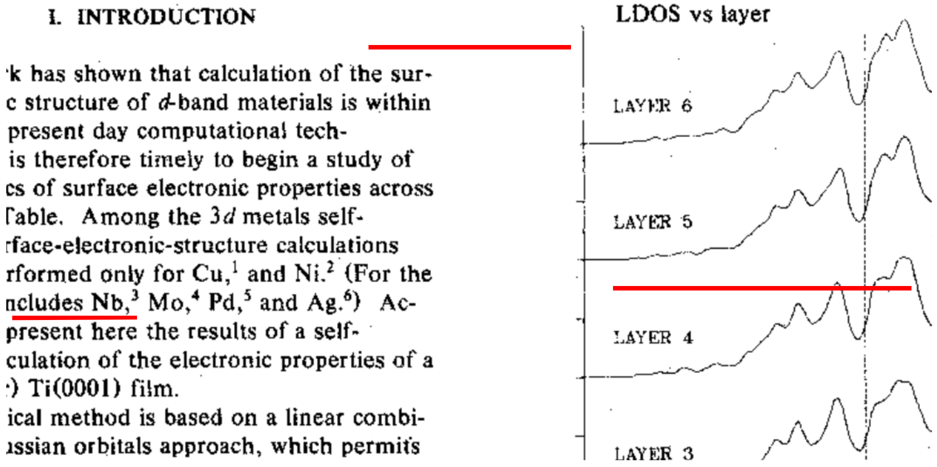

I. INTRODUCTION

ʼk has shown that calculation of the sur-
c structure of *d*-band materials is within
present day computational tech-
is therefore timely to begin a study of
cs of surface electronic properties across
Table. Among the 3*d* metals self-
rface-electronic-structure calculations
rformed only for Cu,[1] and Ni.[2] (For the
ncludes Nb,[3] Mo,[4] Pd,[5] and Ag.[6]) Ac-
present here the results of a self-
culation of the electronic properties of a
) Ti(0001) film.
ical method is based on a linear combi-
ussian orbitals approach, which permits

[Phys. Rev. B **20**, 1433 (1979)]

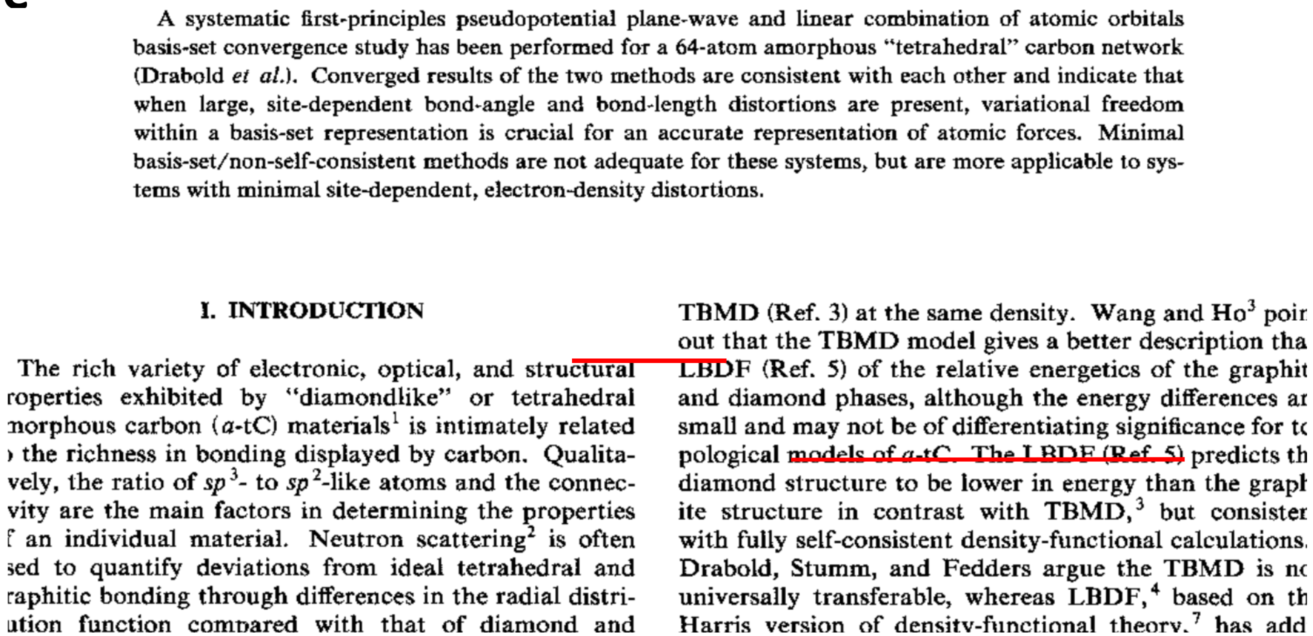

A systematic first-principles pseudopotential plane-wave and linear combination of atomic orbitals basis-set convergence study has been performed for a 64-atom amorphous "tetrahedral" carbon network (Drabold *et al.*). Converged results of the two methods are consistent with each other and indicate that when large, site-dependent bond-angle and bond-length distortions are present, variational freedom within a basis-set representation is crucial for an accurate representation of atomic forces. Minimal basis-set/non-self-consistent methods are not adequate for these systems, but are more applicable to systems with minimal site-dependent, electron-density distortions.

I. INTRODUCTION

The rich variety of electronic, optical, and structural
roperties exhibited by "diamondlike" or tetrahedral
norphous carbon (*a*-tC) materials[1] is intimately related
› the richness in bonding displayed by carbon. Qualita-
vely, the ratio of sp^3- to sp^2-like atoms and the connec-
vity are the main factors in determining the properties
f an individual material. Neutron scattering[2] is often
sed to quantify deviations from ideal tetrahedral and
raphitic bonding through differences in the radial distri-
ution function compared with that of diamond and

TBMD (Ref. 3) at the same density. Wang and Ho[3] poir
out that the TBMD model gives a better description tha
LBDF (Ref. 5) of the relative energetics of the graphit
and diamond phases, although the energy differences ar
small and may not be of differentiating significance for tc
pological ~~models of *a*-tC. The LBDF (Ref. 5)~~ predicts th
diamond structure to be lower in energy than the grapł
ite structure in contrast with TBMD,[3] but consister
with fully self-consistent density-functional calculations
Drabold, Stumm, and Fedders argue the TBMD is nc
universally transferable, whereas LBDF,[4] based on th
Harris version of density-functional theory,[7] has add

[Phys. Rev. B **52**, 9254 (1995)]

Mix of physics & chemistry methods:

- *spd* Gaussian atomic orbitals (DZP)
- ECP pseudopotentials (large core)
- local Δ-potentials on a uniform grid
- semilocal DFT (**no** Fock exchange)

- http://dft.sandia.gov/Quest & http://nanoHUB.org
- **primary uses:** surface science & point-defect chemistry
- 100k lines of Fortran 77, 4k lines for main control flow
- MPI parallelization w/ few-node scalability

Heyd-Scuseria-Ernzerhof density functional

Perdew’s “Jacob’s ladder” of universal density functionals [Perdew & Schmidt, AIP Conf. Proc. **577**, 1 (2001)]:

electron density (LDA) → density gradient (GGA) → kinetic energy (meta-GGA) → occupied orbitals (Fock exchange) → unoccupied orbitals (RPA)

↓

exchange screening
(≈1/ε-fraction hybrids)

↙ *d*-shells in solid state
(short-ranged hybrids)

↘ molecules in vacuum
(long-ranged hybrids)

PBE/HSE exchange-hole model:
[Heyd, Scuseria, & Ernzerhof, J. Chem. Phys. **118**, 8207 (2003)]

$$
\begin{aligned}
E_x &= \int d\mathbf{r}\,\rho(\mathbf{r},\mathbf{r})\epsilon_x(\mathbf{r},\cdots) \\
&= \int d\mathbf{r}d\mathbf{r}'\rho(\mathbf{r},\mathbf{r})^2 J(\mathbf{r},|\mathbf{r}-\mathbf{r}'|,\cdots)V(|\mathbf{r}-\mathbf{r}'|) \\
&= -\frac{1}{2}\int d\mathbf{r}d\mathbf{r}'\rho(\mathbf{r},\mathbf{r}')^2 V(|\mathbf{r}-\mathbf{r}'|)
\end{aligned}
$$

← *d*-shells in solid state (short-ranged hybrids)

Fock exchange: physics vs. chemistry

implicit ← $h_x(\mathbf{r},\mathbf{r}') = -\rho(\mathbf{r},\mathbf{r}')V(|\mathbf{r}-\mathbf{r}'|)$ → explicit

$$\psi_{\text{out}}(\mathbf{r}) := \int d\mathbf{r} h_x(\mathbf{r},\mathbf{r}')\psi_{\text{in}}(\mathbf{r}')$$

$$\rho(\mathbf{r},\mathbf{r}') = \sum_{i\in\text{occ}} \psi_i(\mathbf{r})\psi_j(\mathbf{r})$$

$$v_i(\mathbf{r}) = \int d\mathbf{r}' V(|\mathbf{r}-\mathbf{r}'|)\psi_i(\mathbf{r}')\psi_{\text{in}}(\mathbf{r}')$$

$$\phi_{\text{out}}(\mathbf{r}) = -\sum_{i\in\text{occ}} \phi_i(\mathbf{r})v_i(\mathbf{r})$$

cost (per vector) ∝ $N_{\text{occ}}\, N_{\text{grid}} \log(N_{\text{grid}})$

VASP, Quantum ESPRESSO, CASTEP, …

$$H_{ij} := \int d\mathbf{r} d\mathbf{r}' \phi_i(\mathbf{r})h_x(\mathbf{r},\mathbf{r}')\phi_j(\mathbf{r}')$$

$$\rho(\mathbf{r},\mathbf{r}') = \sum_{i,j\in\text{basis}} \phi_i(\mathbf{r})P_{ij}\phi_j(\mathbf{r}')$$

$$V_{ijkl} = \int d\mathbf{r} d\mathbf{r}' \phi_i(\mathbf{r})\phi_j(\mathbf{r})V(|\mathbf{r}-\mathbf{r}'|)\phi_k(\mathbf{r}')\phi_l(\mathbf{r}')$$

$$H_{ij} = -\sum_{k,l\in\text{basis}} V_{ikjl}P_{kl}$$

cost ∝ $N_{\text{basis}}\, (n_{\text{overlap}})^2\, n_{\text{kernel}}$

GAUSSIAN, CRYSTAL, PySCF (solid state)

Recalibrating HSE

Simultaneously optimize HSE fraction (0.25) & range (9.1 bohr) to:

understand physical significance, reduce range/cost, & improve accuracy

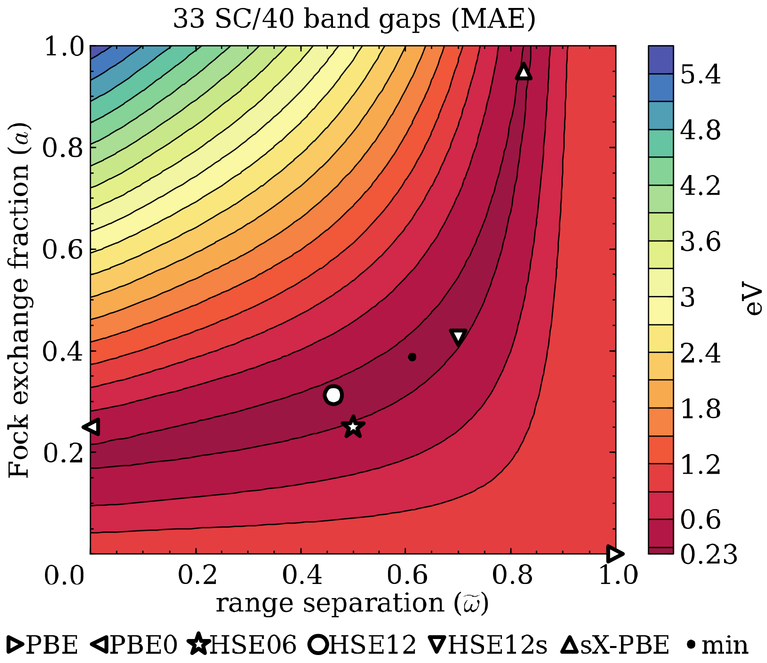

33 SC/40 band gaps (MAE)

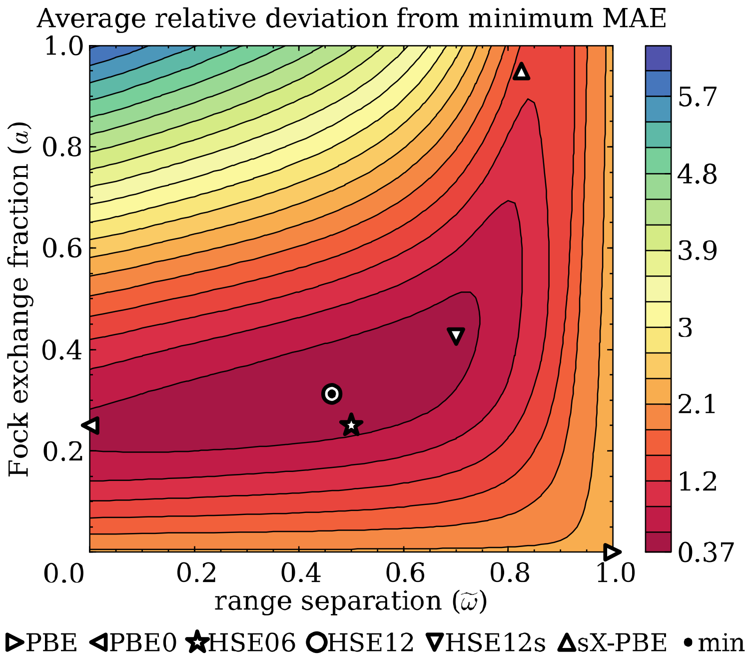

Average relative deviation from minimum MAE

tests: SC/40, G3/99, BH42/04, & T-96R

"HSE12": (0.313, 10.2 bohr)

"HSE12s": (0.425, 4.6 bohr)

min-range: (1.0, 2.5 bohr)

(narrower than atomic orbitals)

[**Moussa**, Shultz, & Chelikowsky, J. Chem. Phys. **136**, 204117 (2012)]

*Truhlar's "Minnesota" density functionals constructed from a similar methodology

Refactoring HSE

$$\epsilon_x(\mathbf{r},\cdots) = \rho(\mathbf{r},\mathbf{r}) \int d\mathbf{r}' J(\mathbf{r}, |\mathbf{r}-\mathbf{r}'|, \cdots) V(|\mathbf{r}-\mathbf{r}'|)$$

- numerical & regularization problems in standard implementation
- replaced integration quadrature with semi-analytic evaluation
- refactored twice (we could afford to make this into a reliable black box)

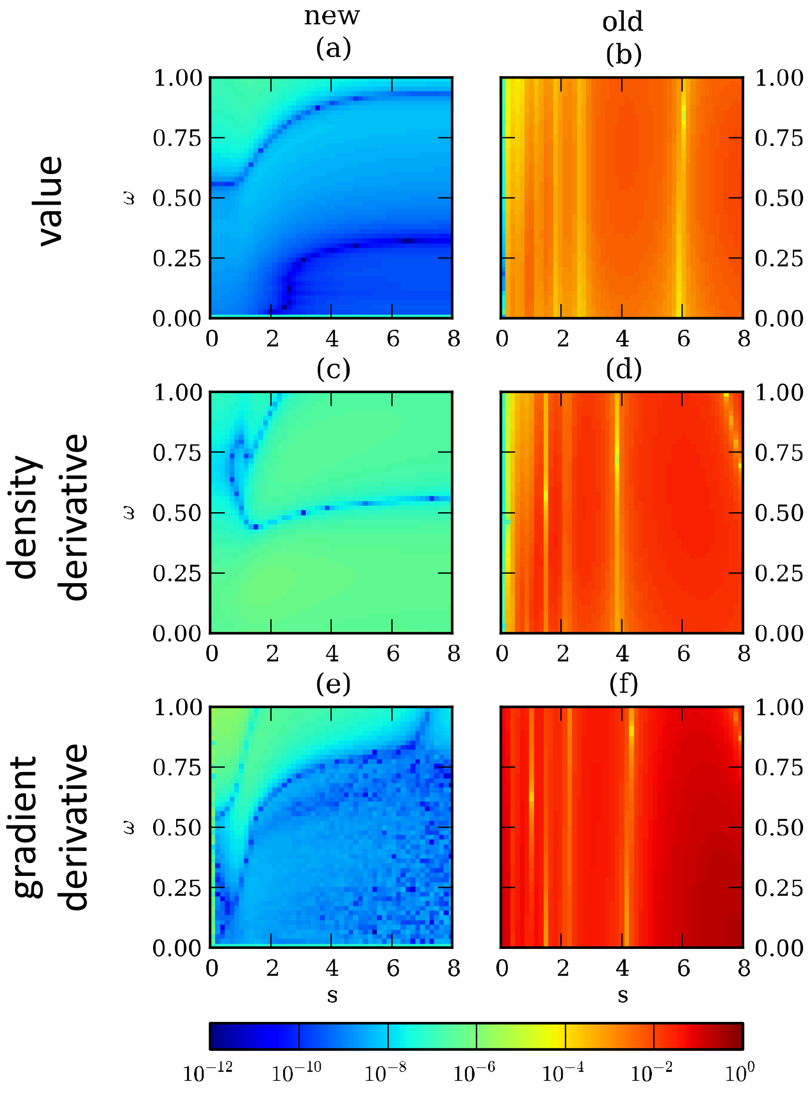

[**Moussa** & Schultz, arXiv:1210.8233 (2012)]

Verifying HSE

- Coulomb integrals & derivatives: Libint 2.3.1 w/ added C wrapper
- no pre-computed integrals ("direct") or screening (envelop tests)
- SeqQuest basis sets & ECPs interoperable* w/ chemistry codes
- G2/97 test set for numerical validation:

	PBE	**HSE**
experiment – all-electron:	0.20	0.05
all-electron – DZP:	0.07	0.07
DZP – DZP+ECP:	0.23	0.18
GAUSSIAN – SeqQuest (DZP+ECP):	0.02	0.02

RMS errors (eV/atom)

Screening HSE

```
LOOP atomic shell $I$ (central cell):
  LOOP atomic shell $J$:
    screen $(I,J)$ pairs
    LOOP atomic shell $K$:
      screen $(I,J,K)$ triplets
      LOOP atomic shell $L$:
        screen $(I,J,K,L)$ quartets
        read 1RDM blocks
        screen w/ 1RDM blocks
        Fock exchange on $(I,J,K,L)$
        write Hamiltonian blocks
```

bulk Al, DZP+ECP, 2.5 bohr exchange length

$\varepsilon = 10^{-7}$ (E_x = 122.371)	**$\varepsilon = 10^{-8}$** (E_x = -0.309)
$1\text{x}10^4$ pairs	$1\text{x}10^4$ pairs
$3\text{x}10^3$ unscreened pairs	$4\text{x}10^3$ unscreened pairs
$7\text{x}10^6$ triplets	$8\text{x}10^6$ triplets
$3\text{x}10^5$ unscreened triplets	$6\text{x}10^5$ unscreened triplets
$6\text{x}10^8$ quartets	$1\text{x}10^9$ quartets
$1\text{x}10^7$ unscreened quartets	$3\text{x}10^7$ unscreened quartets
$2\text{x}10^7$ pre-screened quartets	$4\text{x}10^7$ pre-screened quartets

ill-conditioned summation →

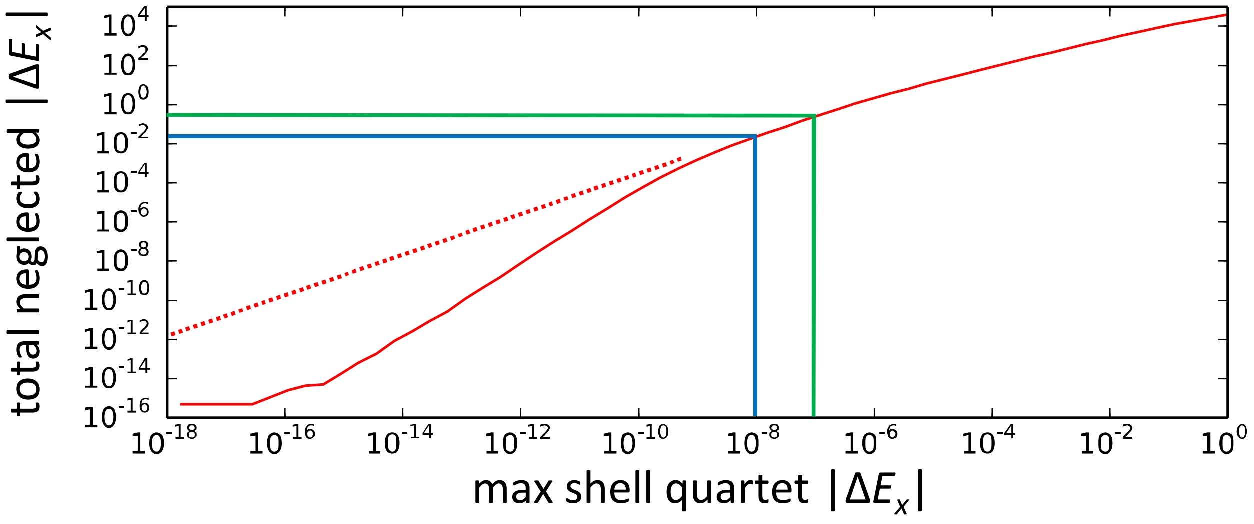

Work in progress

- Precompute/accelerate screening to mitigate cost
- Fully verify periodic HSE implementation
- Compatible MPI parallelization
- Coulomb tensor decompositions?

I am a “software scientist”

- Some custom software for almost every paper
- Firm believer in open/reproducible computing
- Attempts at pedagogical software w/ methods papers
- Monolithic methods/codes should exist in a primordial sea of software fragments that explore/communicate new ideas

Pros & cons of emphasizing software

+ accelerate development cycle
+ maximize hardware performance
+ maximize algorithm performance
+ enable quantitative comparisons

– attachment to mature software
– difficult to follow HPC evolution
– methodological ossification
– can’t fix unhealthy competition

Will this *help* or *hinder* the adoption of new, asymptotically better ideas?

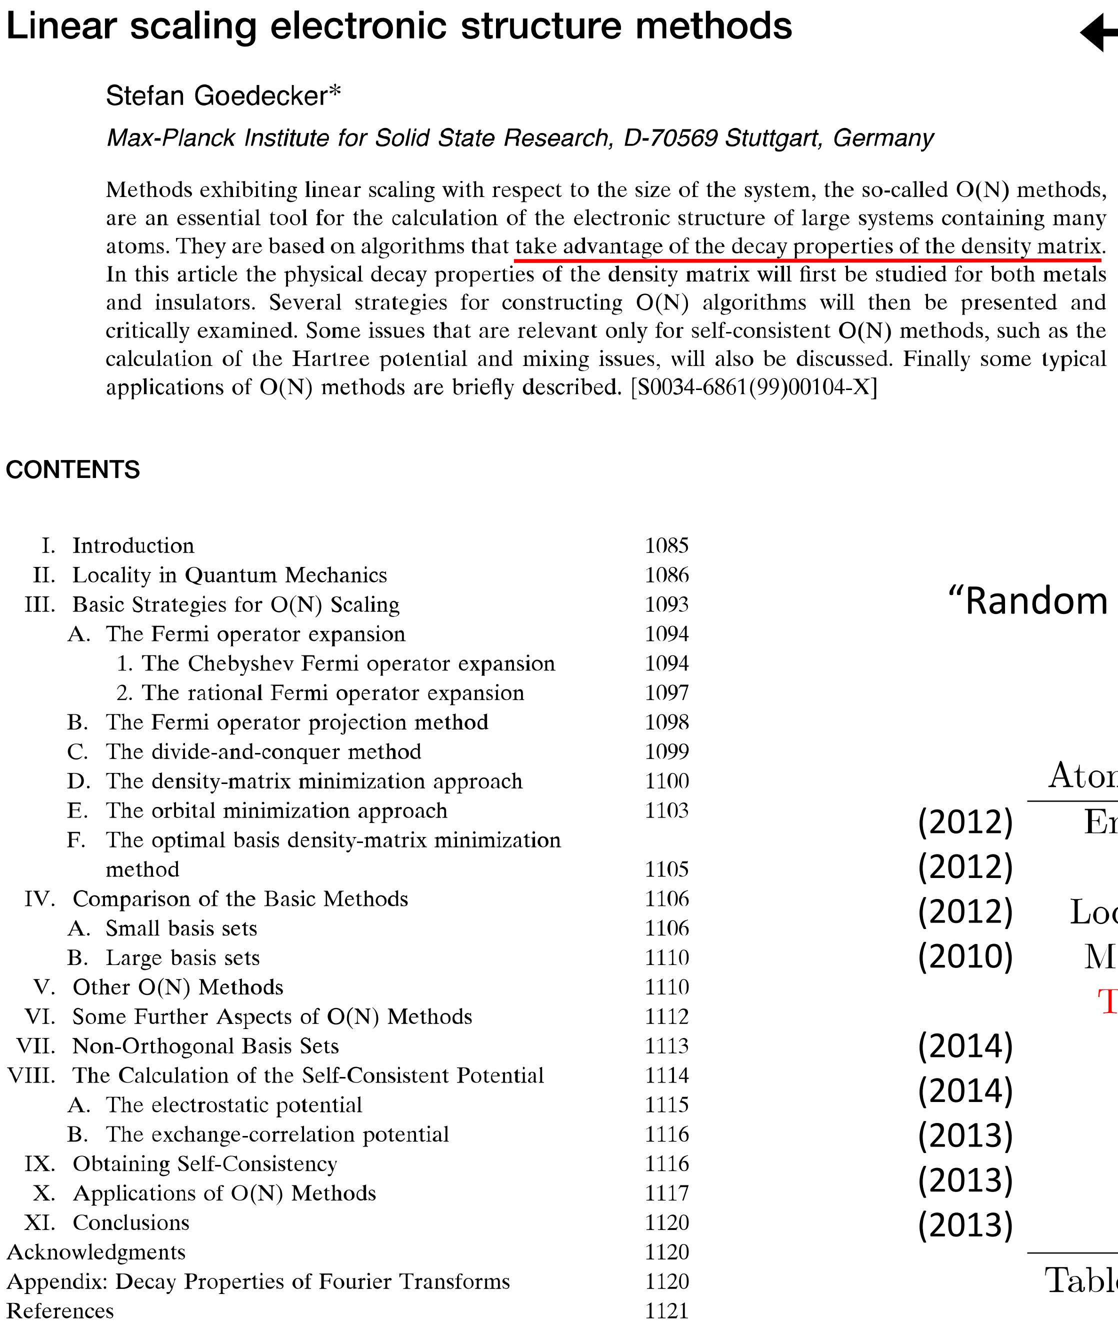

Linear scaling electronic structure methods

Stefan Goedecker*

Max-Planck Institute for Solid State Research, D-70569 Stuttgart, Germany

Methods exhibiting linear scaling with respect to the size of the system, the so-called O(N) methods, are an essential tool for the calculation of the electronic structure of large systems containing many atoms. They are based on algorithms that take advantage of the decay properties of the density matrix. In this article the physical decay properties of the density matrix will first be studied for both metals and insulators. Several strategies for constructing O(N) algorithms will then be presented and critically examined. Some issues that are relevant only for self-consistent O(N) methods, such as the calculation of the Hartree potential and mixing issues, will also be discussed. Finally some typical applications of O(N) methods are briefly described. [S0034-6861(99)00104-X]

CONTENTS

← [Rev. Mod. Phys. **71**, 1085 (1999)]

post-HF methods:

[Saebo & Pulay, Annu. Rev. Phys. Chem. 44, 213 (1993)]

stochastic methods:

[Silver & Roder, Int. J. Mod. Phys. C **5**, 735 (1994)]

[Baer, Neuhauser, & Rabani, Phys. Rev. Lett. **111**, 106402 (2013)]

available linear-scaling DFT codes: ONETEP

“Random Phase Approximation: Models and Algorithms for Electron Correlation”
[my 2014 BES early career proposal]

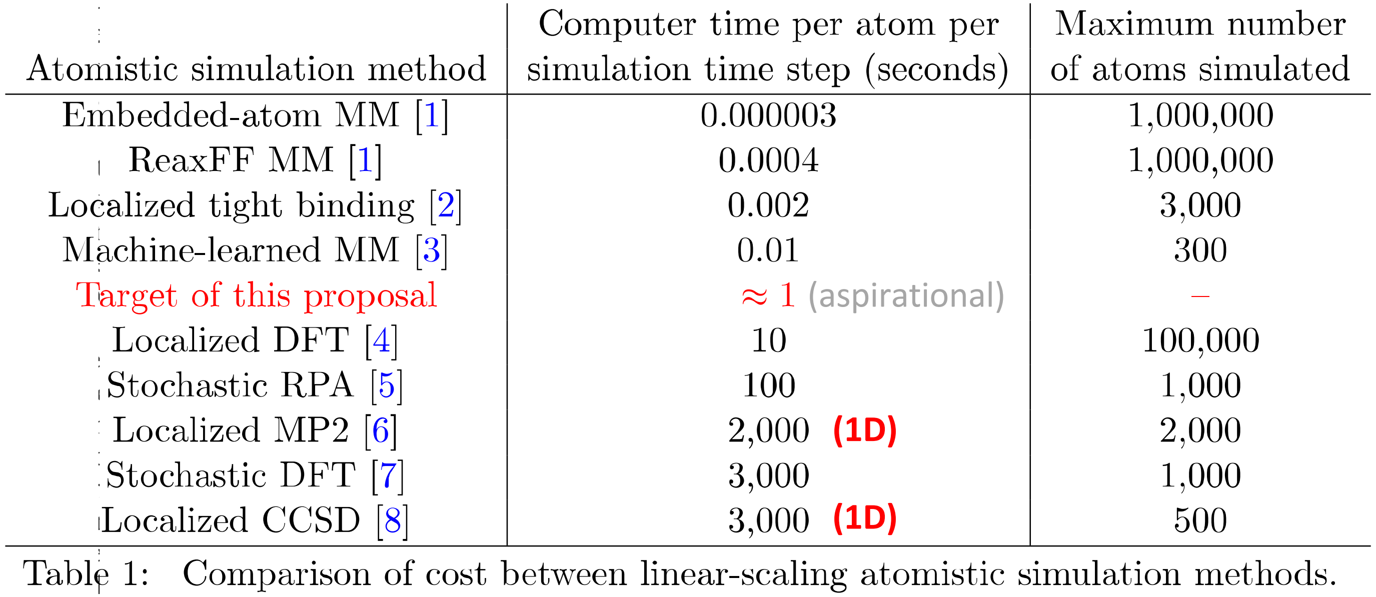

	Atomistic simulation method	Computer time per atom per simulation time step (seconds)	Maximum number of atoms simulated
(2012)	Embedded-atom MM [1]	0.000003	1,000,000
(2012)	ReaxFF MM [1]	0.0004	1,000,000
(2012)	Localized tight binding [2]	0.002	3,000
(2010)	Machine-learned MM [3]	0.01	300
	Target of this proposal	≈ 1 (aspirational)	–
(2014)	Localized DFT [4]	10	100,000
(2014)	Stochastic RPA [5]	100	1,000
(2013)	Localized MP2 [6]	2,000 **(1D)**	2,000
(2013)	Stochastic DFT [7]	3,000	1,000
(2013)	Localized CCSD [8]	3,000 **(1D)**	500

Table 1: Comparison of cost between linear-scaling atomistic simulation methods.

“Realistic” test problems

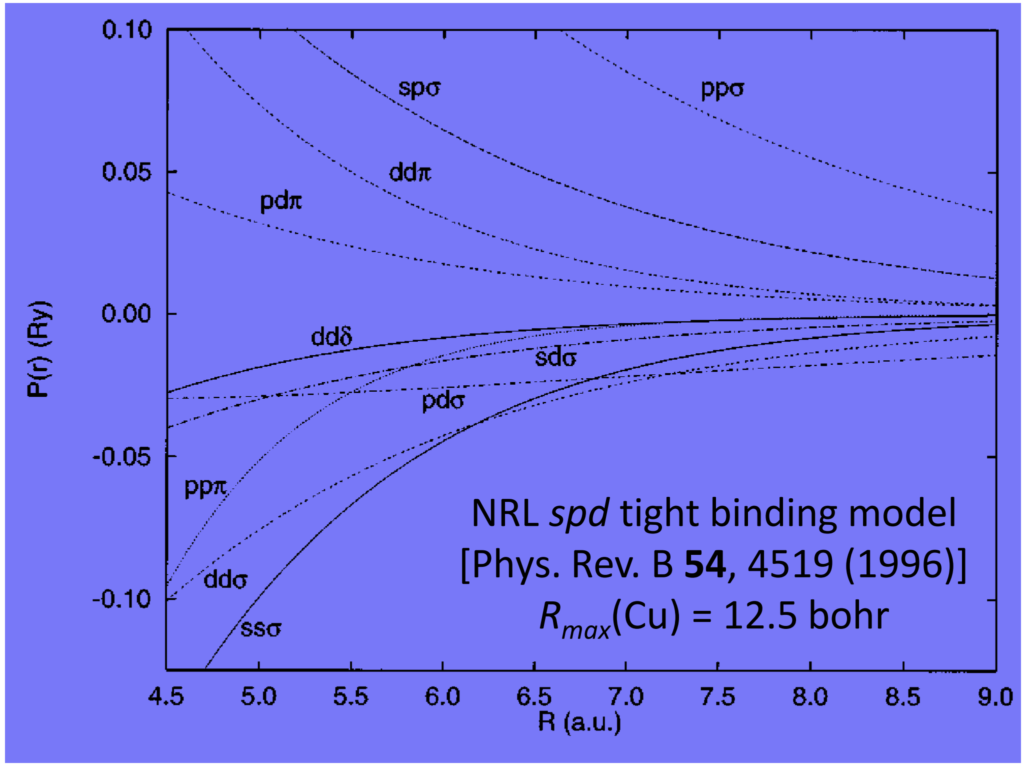

NRL *spd* tight binding model
[Phys. Rev. B **54**, 4519 (1996)]
R_{max}(Cu) = 12.5 bohr

target temperatures: 0.03 eV (ambient)
3.0 eV (warm dense matter)

target force accuracy: 0.01 eV/A

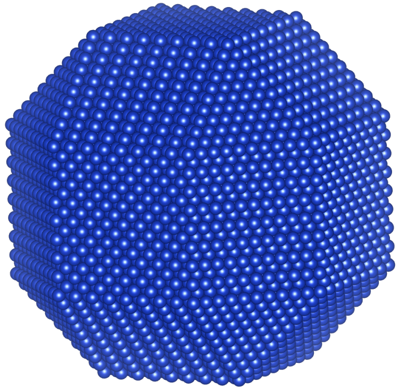

rhombicuboctahedral
copper clusters

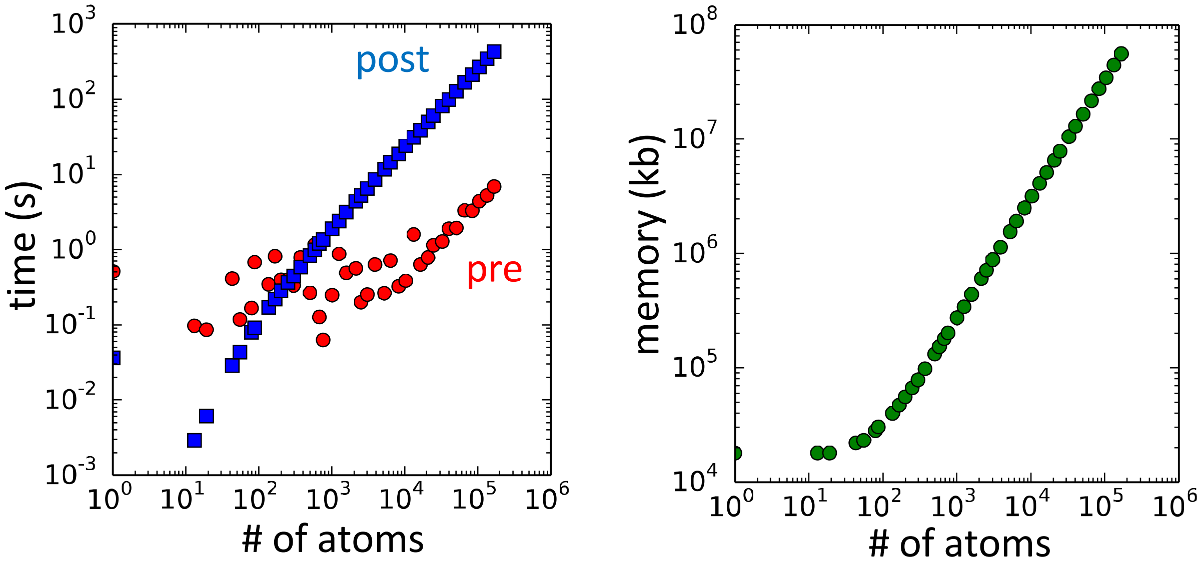

pre/post-processing benchmarks

Basic numerical ingredients

- NRL tight-binding model
- pseudorandom number generation
- block-sparse matrix operations
- geometric localization / partitioning
- polynomial Fermi-Dirac approximation
- rational Fermi-Dirac approximation
- iterative linear solvers (CG & CR)
- PEXSI solver interface

3k lines of C w/ OpenMP & MPI (PEXSI only)

input: Fermi energy & *.xyz
output: total energy & forces

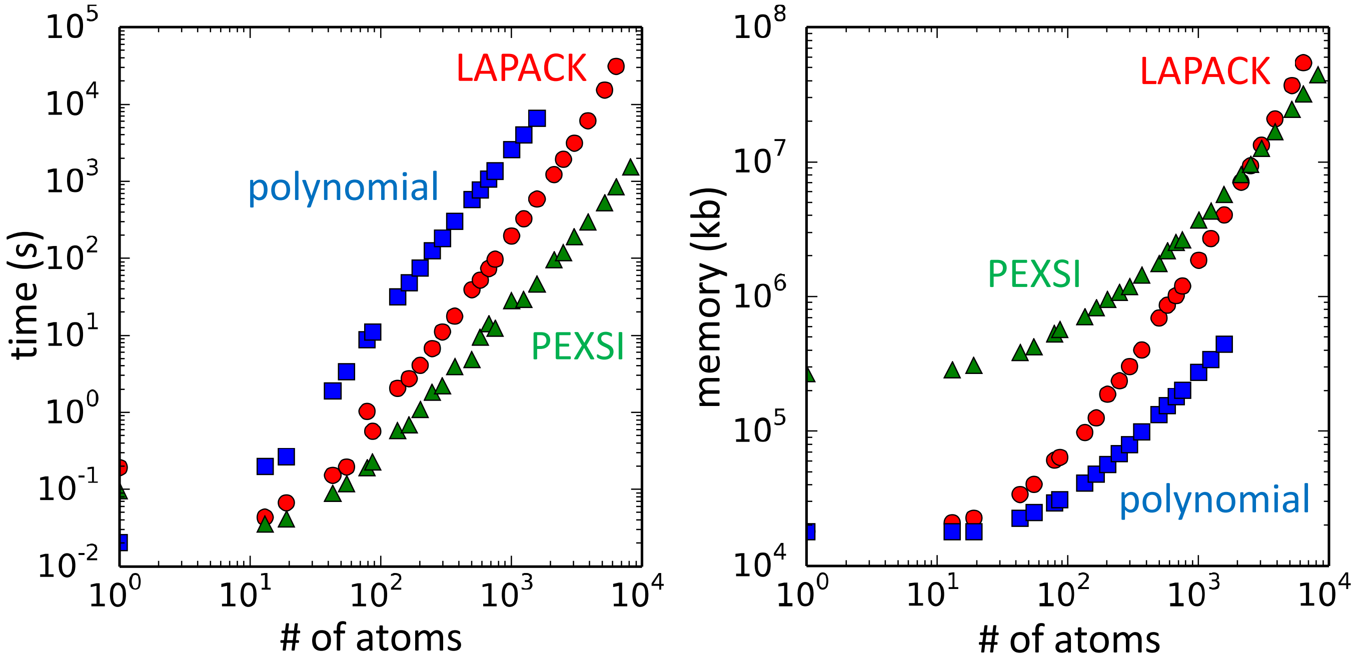

6-pole rational *or* degree-16 polynomial @ $T = 3.0$ eV

Localization *a priori*

$$f(H)e_i \approx f(H_i)e_i$$

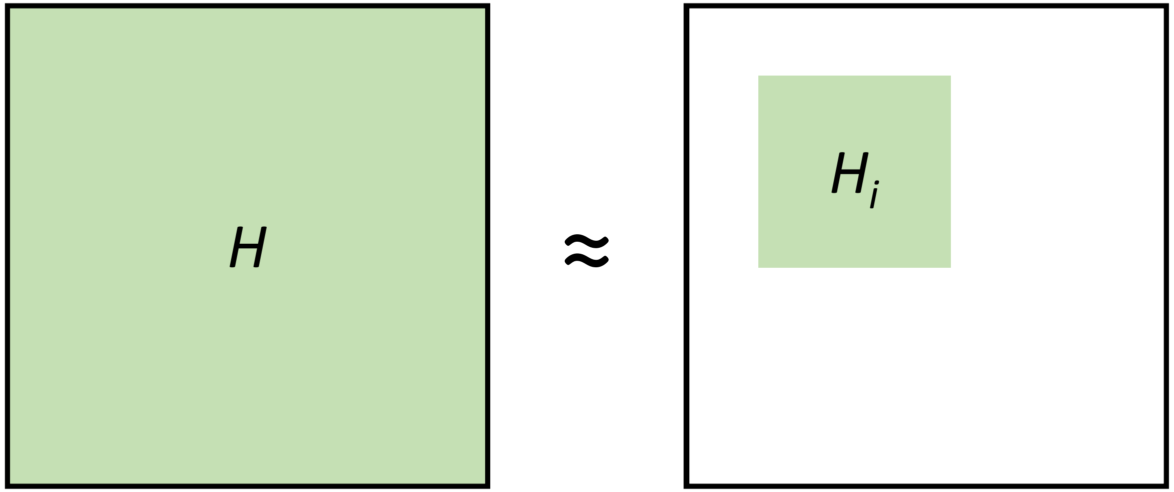

error depends on *T*, gap, & sparsity

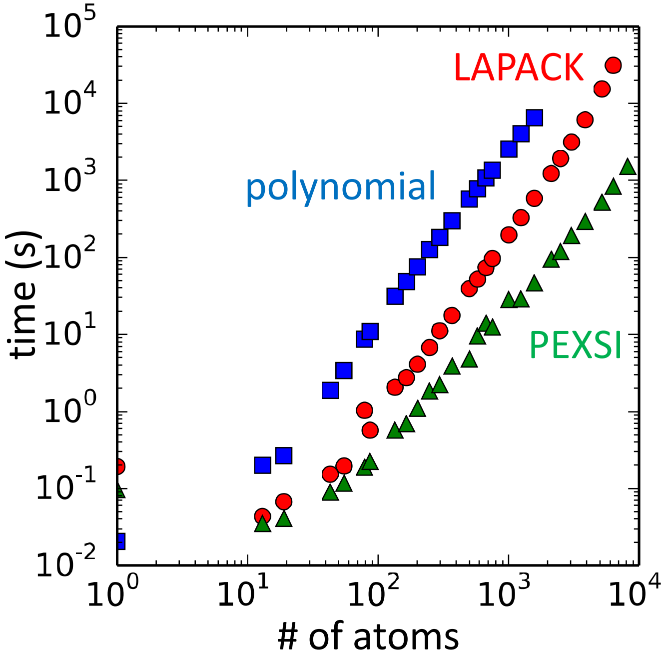

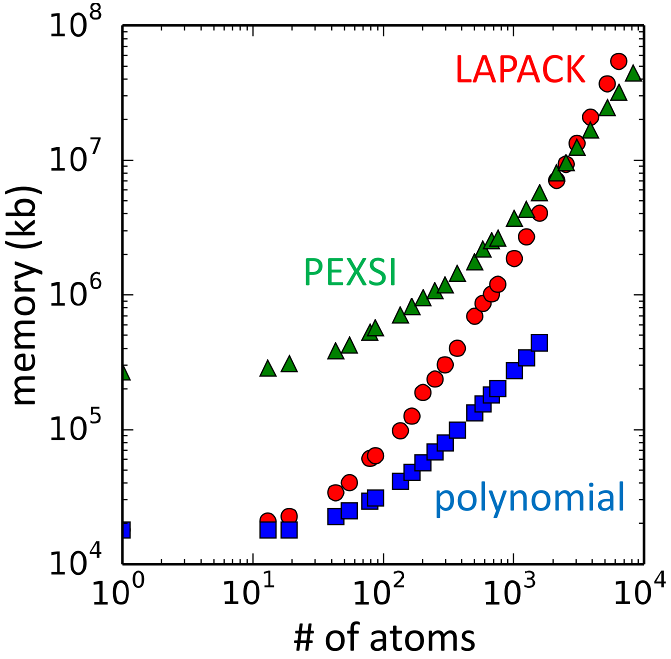

Randomization

$$\mathrm{tr}[f(H)] \approx \frac{1}{s}\sum_{i=1}^{s} x_i^T f(H) x_i$$

stochastic matrix trace estimation

[Avron & Toledo, J. ACM **58**, 8 (2011)]

error ∝ $s^{-1/2}$

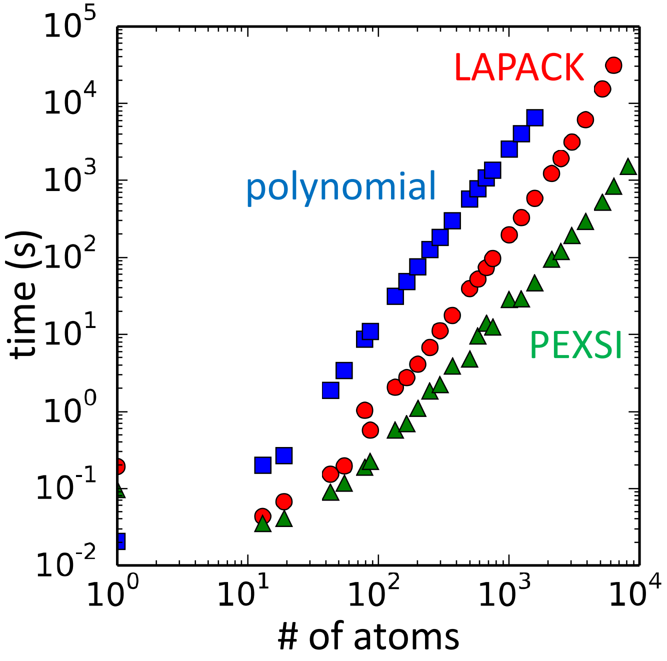

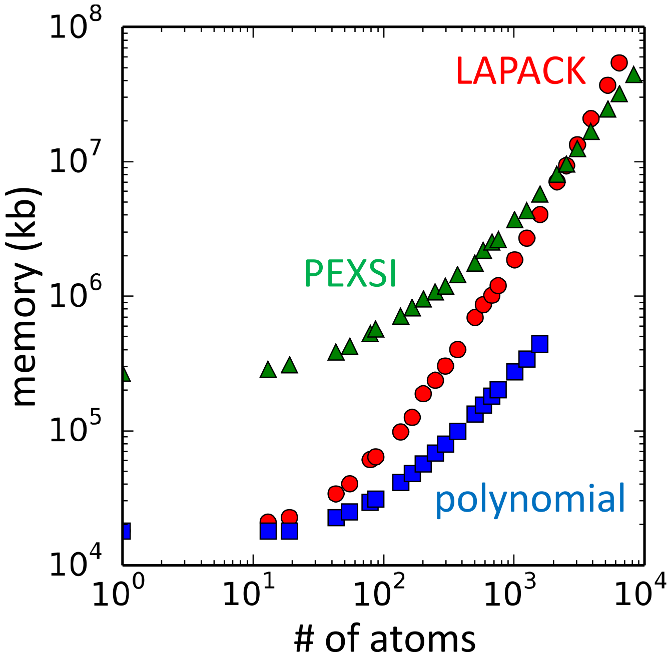

Localization *a posteriori*

$$\min_{P} \sum_{i=1}^{s} \| [P - f(H)] x_i \|_2^2$$

fixed-sparsity analog of
compressed sensing

[Candes, Romberg, & Tao, Comm. Pure Appl. Math. **59**, 1207 (2006)]

error ∝ [localization error] * $s^{-1/2}$

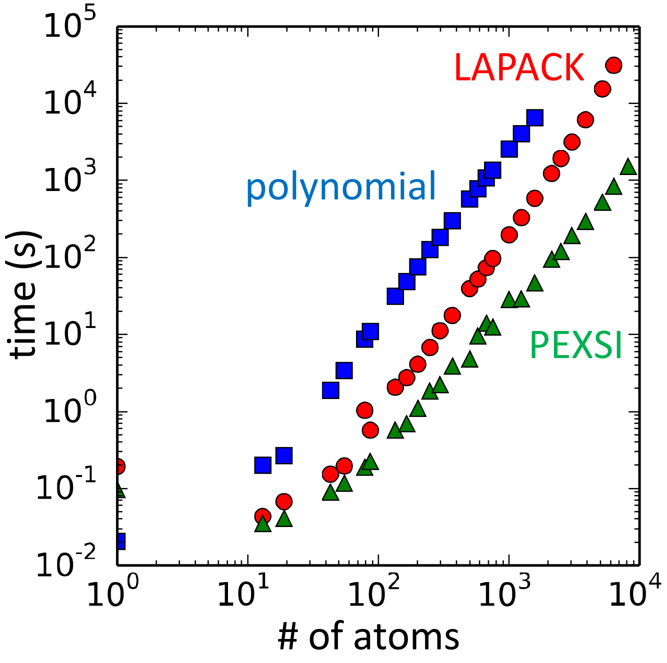

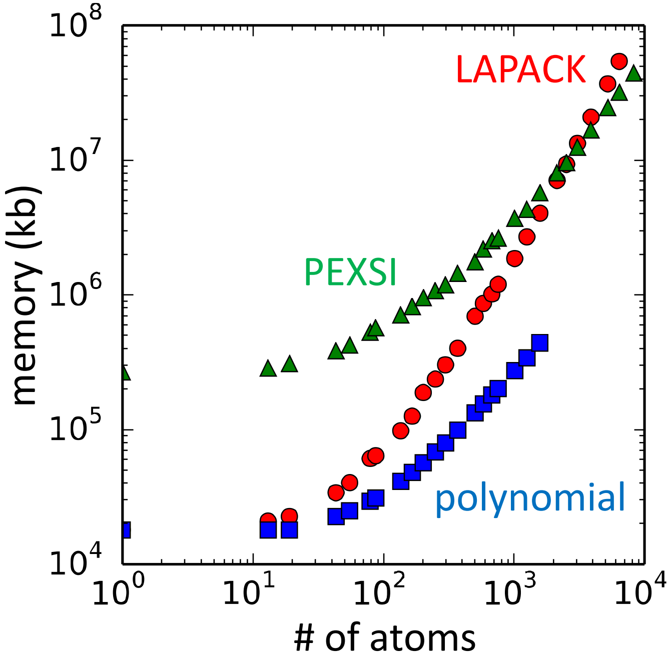

Rational approximation & preconditioning

$$f(H)x_i = \sum_j c_j (H - z_j S)^{-1} x_i$$

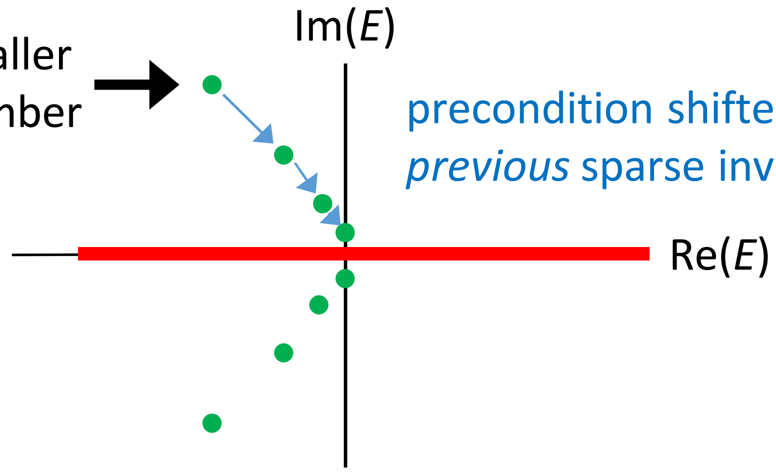

sparser & smaller condition number

precondition shifted linear solver w/ *previous* sparse inverse from larger shift

@ 2018 APS March meeting (3PM, March 7th)

Does this have a future?

- Yes, if moderate numerical accuracy suffices
- Efficient log(1/ε) multiscale methods are *long-term* research
- Holistic design: electron correlation model from fast ingredients
- Compatible w/ fast MP2 & RPA methods
  [**Moussa**, J. Chem. Phys. **140**, 014107 (2014)]

Future of electronic structure?

- Method/software development has always been technically difficult & will only get more difficult as sophistication increases.

- Electronic structure (e.g.@Sandia: QMCPACK, SeqQuest, Socorro, PyQuante, MPQC, LMTO-ASA, MECCA) has struggled to achieve as much practical impact as classical MD (e.g.@Sandia: **LAMMPS**).

My research plan:

1. linear-scaling solver
2. RPA correlation model

3a. semi-empirical basis (stopgap)
3b. all-electron basis

"Big-science" electronic structure:	**At Sandia:**
Are there enough people?	No.
Is there enough money?	No.
Is there a design that *can* succeed?	No.
Is there enough consensus?	No.