Discontinuous methods for large-scale quantum molecular dynamics: challenges and outlook

John Pask (Director)
Condensed Matter and Materials Division, Lawrence Livermore National Laboratory

Vincenzo Lordi, Mitchell Ong
Condensed Matter and Materials Division, Lawrence Livermore National Laboratory

Chao Yang*, Lin Lin, Mathias Jacquelin, Gaigong Zhang
Computational Research Division, Lawrence Berkeley National Laboratory
**SciDAC FASTMath Institute*

Erik Draeger
Center for Applied Scientific Computing, Lawrence Livermore National Laboratory

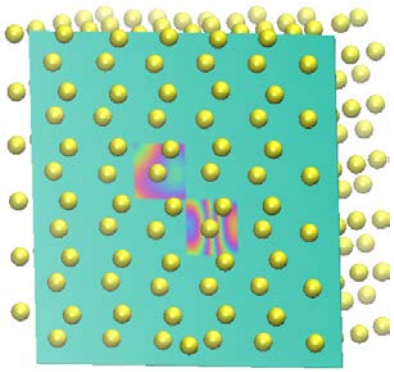

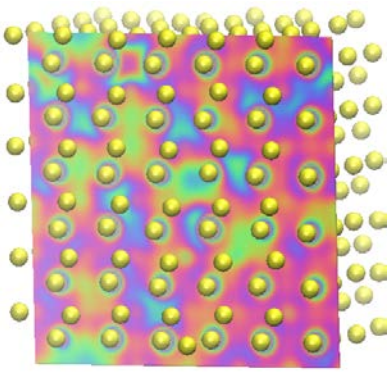

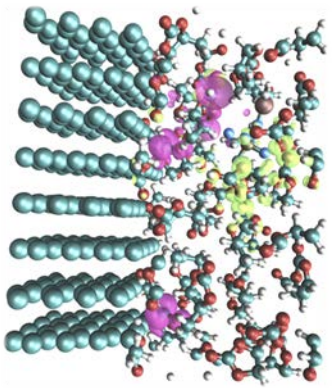

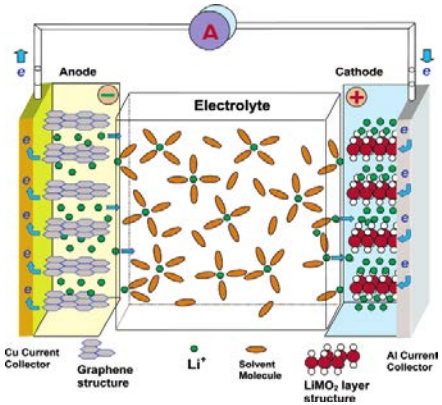

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Overview

- Li-ion batteries have revolutionized consumer electronics and have the potential to do the same for transportation (e.g., plug-in hybrids, all-electrics, aircraft) and electrical distribution (e.g., load leveling)
- To do so, energy/power density, lifetime, safety must be increased
- Key issue: **solid-electrolyte interphase** (**SEI**) layer at electrolyte-anode interface, product of electrolyte decomposition
- Understanding has been hindered by need for both quantum mechanical description and sufficiently large length/time scales to capture necessary complexity
- In this work, we:
  - *Develop* new **Discontinuous Galerkin (DG)** electronic structure method to accomplish quantum molecular dynamics (QMD) on an unprecedented scale
  - *Apply* new method to advance understanding of the **chemistry & dynamics of electrolyte/SEI/anode systems**

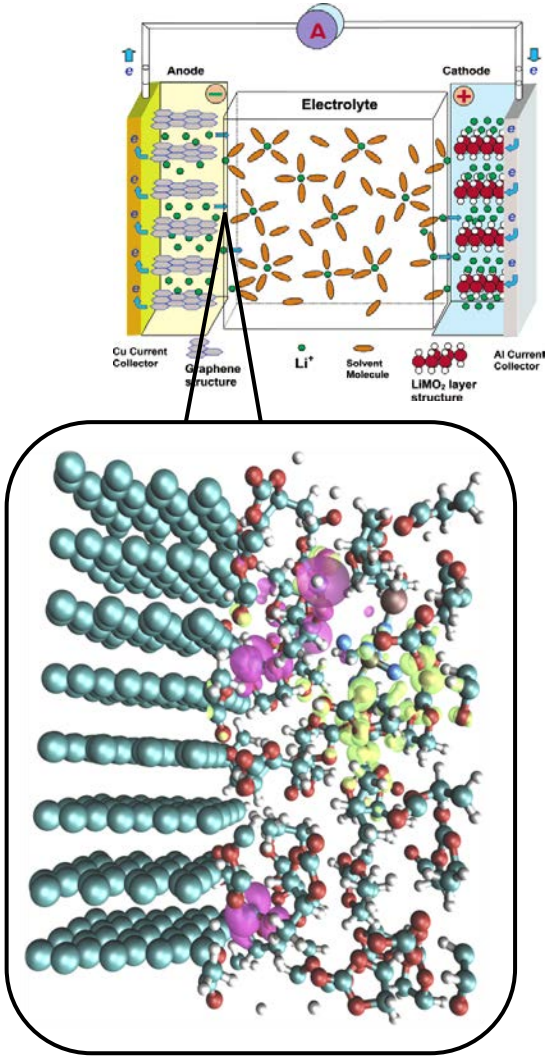

QMD snapshot of SEI layer in Li-ion cell

Management

- All postdocs up and running
- Monthly meetings, alternating between LLNL and LBL: proximity has proved a significant advantage
- Skype, GotoMeeting, phone, e-mail between

Simulations

- Initial phase of project, while new DG code is developed and optimized: **Qbox** [1] for systems of < 2,000 atoms
- Li^+ solvation and diffusion: determine diffusion coefficients, effect of counter-ion, differences in bulk vs near interface

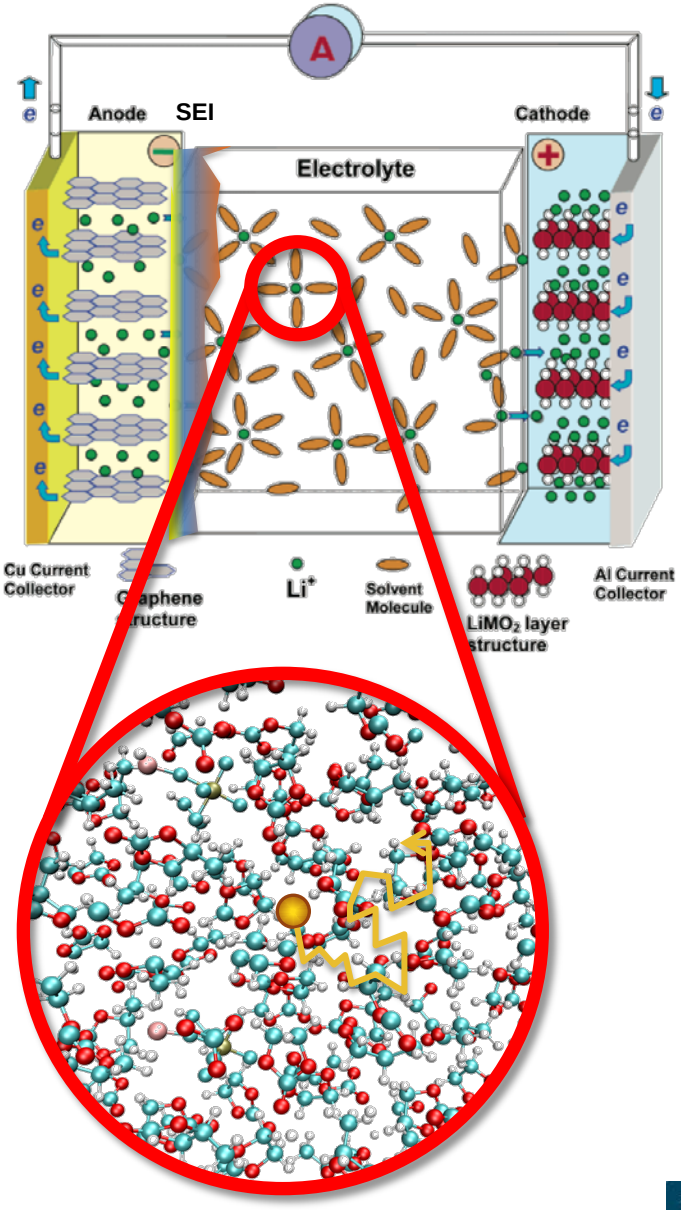

[Adapted from K. Xu, *Chem. Rev.* **104**, 4303 (2004)]

Molecular dynamics simulation of 50/50 ethylene carbonate/propylene carbonate electrolyte

[1] Gygi, Draeger, et al., Proc. ACM/IEEE Supercomputing '06; Gygi, IBM J. Res. Dev., 2008

Simulations

- As the new **DGDFT** method and code ramp up, we transition to it for larger scale simulations, up to 10,000 atoms and more
- Full electrolyte-anode and electrolyte-SEI systems

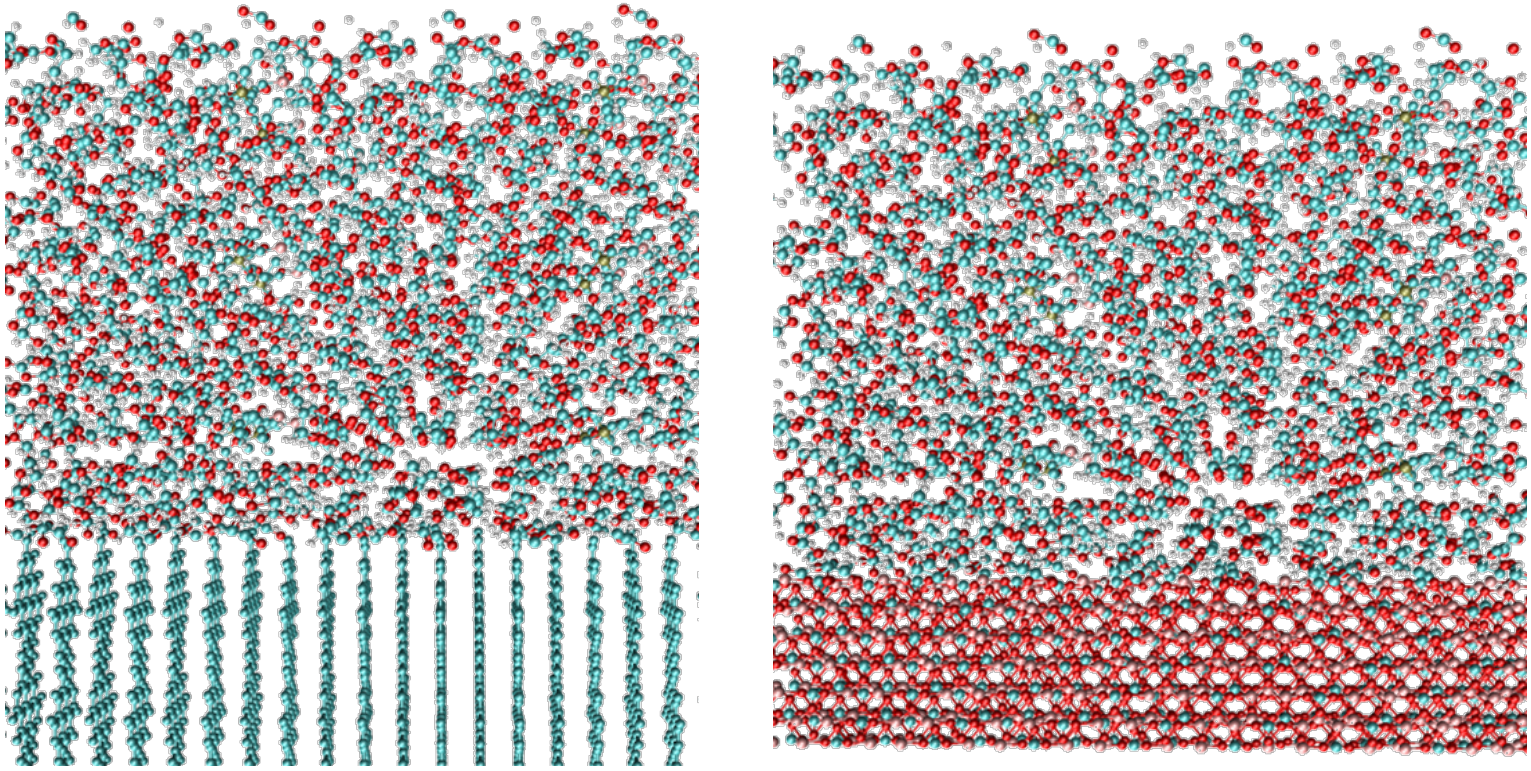

EC/PC mixture (+ $LiPF_6$) on graphite (left) and Li_2CO_3 (right), used to study chemical reactions on the anode surface (for initial SEI formation) and a representative SEI compound (for SEI growth/evolution)

Quantum molecular dynamics (QMD)

- Solve Kohn-Sham equations for electronic structure, compute quantum mechanical forces, move atoms, repeat – thousands to hundreds of thousand of times

Kohn-Sham equations

$$-\tfrac{1}{2}\nabla^2\psi_i(\mathbf{x}) + \hat{V}_{eff}\psi_i(\mathbf{x}) = \varepsilon_i\psi_i(\mathbf{x}), \quad \text{(Schrödinger)}$$

$$\hat{V}_{eff} = V_I^{\ell} + \hat{V}_I^{n\ell} + V_H + V_{xc},$$

$$V_I^{\ell} = \sum_a V_{I,a}(\mathbf{x}),$$

$$\hat{V}_I^{n\ell}\psi_i = \sum_a \int d\mathbf{x}' V_{I,a}^{n\ell}(\mathbf{x},\mathbf{x}')\psi_i(\mathbf{x}'),$$

$$V_H = -\int d\mathbf{x}' \frac{\rho_e(\mathbf{x}')}{|\mathbf{x}-\mathbf{x}'|}, \quad \text{(Poisson)}$$

$$V_{xc} = V_{xc}(\mathbf{x};\rho_e),$$

$$\rho_e = -\sum_i f_i\psi_i^*(\mathbf{x})\psi_i(\mathbf{x}),$$

~ 10^4 atoms, more eigenfunctions

Self-consistent field (SCF) solution process

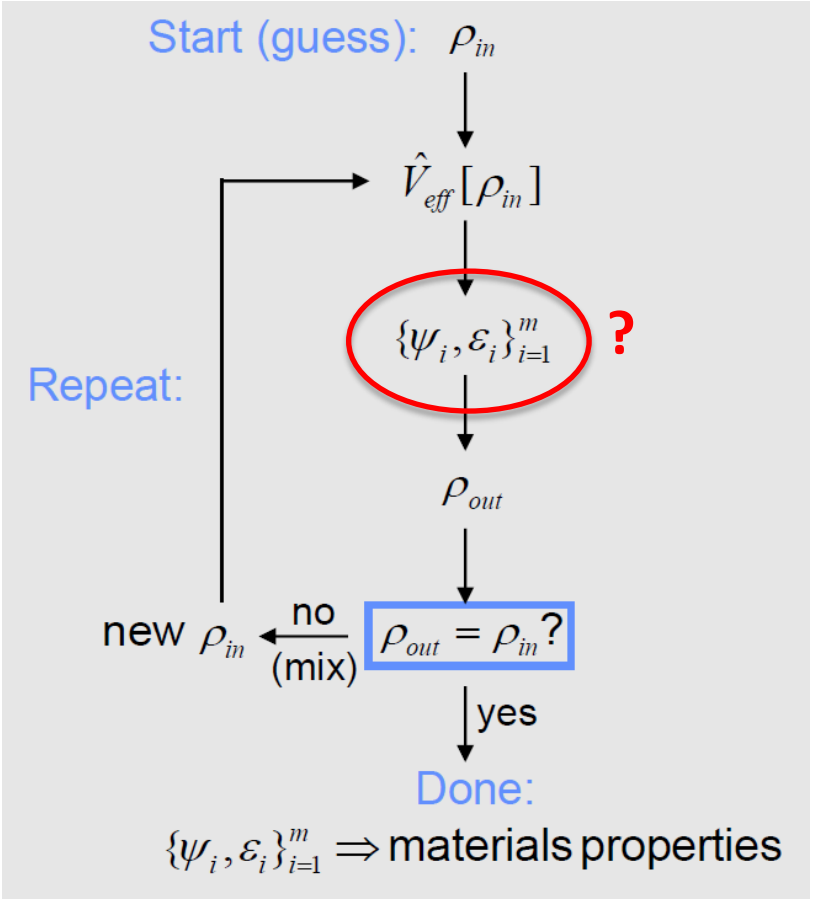

Pushing the current state of the art: Qbox application and development

- Sped up by **factor of three** in metallic calculations by implementing Harris-Foulkes estimator
- Strong scaling to **65,536 cores** on BG/Q
  - Uses hardware threading & SIMD registers on BG/Q
  - Preconditioned steepest descent for occupied subspace
- **→ 1,700-atom anode-electrolyte system in 40 sec per QMD step → 10-40 ps per month**
- Year 3 milestone

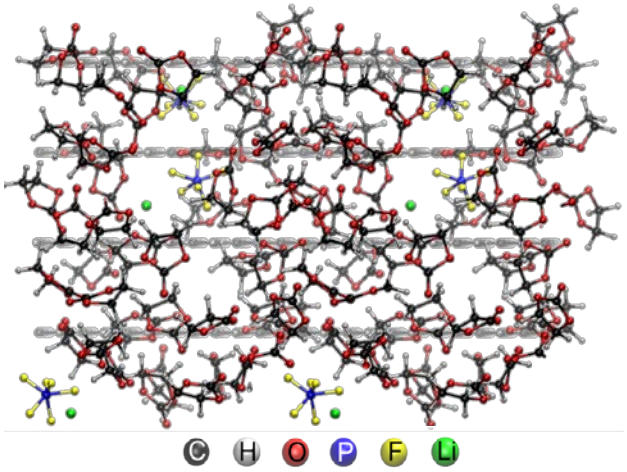

1700-atom anode-electrolyte cell

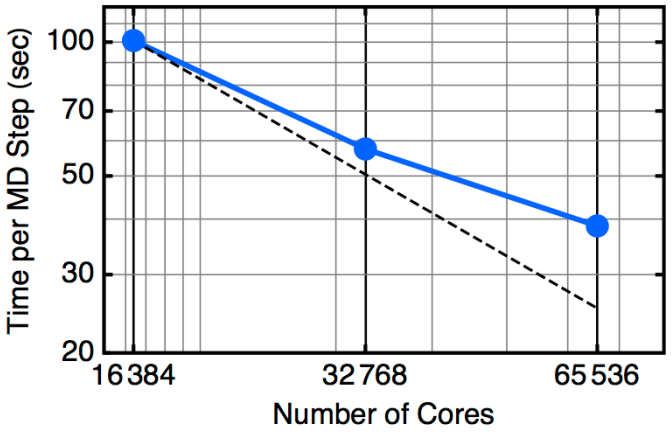

Qbox strong scaling on BG/Q

Issue

- Metallic calculations rely on diagonalization in the occupied subspace (Rayleigh-Ritz)
- Limits both efficiency (N^3) and parallel scaling
- In collaboration with **FASTMath**, we are investigating alternatives to minimize or eliminate Rayleigh-Ritz entirely
  - Trace penalty minimization: minimize trace of Rayleigh quotient and penalty term to enforce orthogonality
  - Vector-update formulation of Locally Optimal Preconditioned Conjugate Gradient method
  - Chebyshev filtered subspace iteration?
  - Fermi Operator Expansion in subspace
- **Goal: Metallic as fast as insulating, 2,000-atom metallic QMD routine**

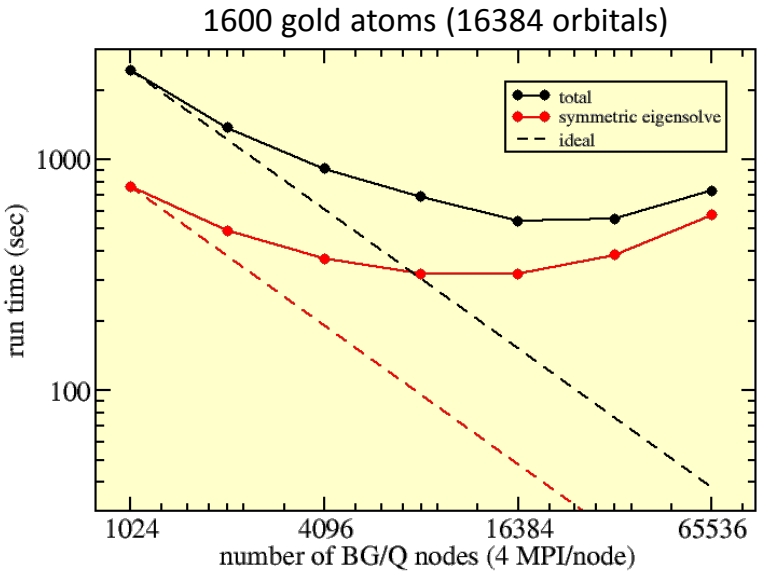

Moving beyond the current state of the art: DGDFT

- DG framework allows solving the Kohn-Sham equations in a discontinuous basis
- Because basis can be discontinuous, can possess number of desirable properties simultaneously:
  - **Efficient** (few tens of DOF/atom)
  - **Systematically improvable**
  - **Strictly local:** identically zero outside prescribed subdomain, zero overlap across subdomains
  - **Orthonormal:** standard eigenproblem, well-cond.
- How?
  - Partition domain into subdomains (elements)
  - Solve Kohn-Sham equations in each element
  - Basis is union of local Kohn-Sham solutions

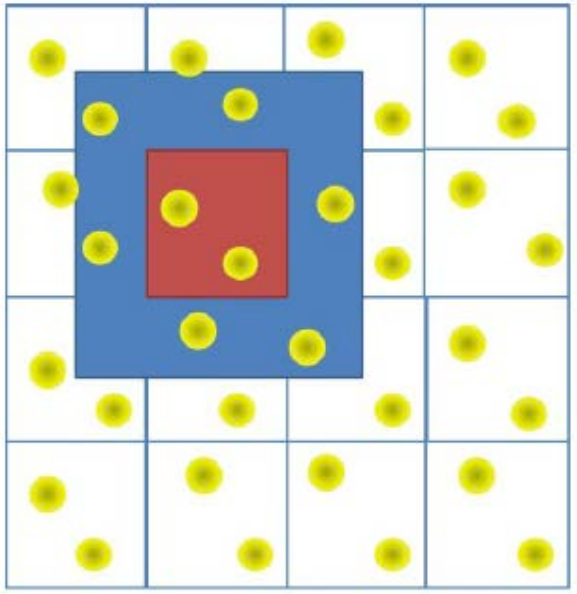

Solve large *N*-atom problem in highly efficient basis of O(*N*) local Kohn-Sham solutions

DG formulation

- Discontinuity is accommodated by surface ("flux") terms [1]
- Kohn-Sham Hamiltonian becomes

$$H_{DG}(k',j';k,j) = \frac{1}{2}\left\langle \nabla u_{k',j'}, \nabla u_{k,j}\right\rangle_{\mathcal{T}} + \alpha\left\langle [[u_{k',j'}]], [[u_{k,j}]]\right\rangle_{\mathcal{S}} - \frac{1}{2}\left\langle [[u_{k',j'}]], \{\{\nabla u_{k,j}\}\}\right\rangle_{\mathcal{S}}$$
$$- \frac{1}{2}\left\langle \{\{\nabla u_{k',j'}\}\}, [[u_{k,j}]]\right\rangle_{\mathcal{S}} + \left\langle u_{k',j'}, V_{\text{eff}}\, u_{k,j}\right\rangle_{\mathcal{T}} + \sum_{\ell}\gamma_\ell \left\langle u_{k',j'}, b_\ell\right\rangle_{\mathcal{T}}\left\langle b_\ell, u_{k,j}\right\rangle_{\mathcal{T}}$$

$\mathcal{T}$ = elements
$\mathcal{S}$ = element surfaces
$u_{k,j}$ = jth basis function in kth element
$\{\{\cdot\}\}$ and $[[\cdot]]$ = average and jump operators across surfaces

- Kohn-Sham equations: $H_{DG}c_i = \varepsilon_i c_i$
- Wavefunctions: $\psi_i = \sum_{E_k\in\mathcal{T}}\sum_{j=1}^{J_k} c_{i;k,j}u_{k,j}$
- Density: $\rho = \sum_{E_k\in\mathcal{T}}\sum_{i=1}^{N}\left|\sum_{j=1}^{J_k} c_{i;k,j}u_{k,j}\right|^2$
- Energy: $E_{\text{tot}} = \sum_{i=1}^{N}\varepsilon_i - \frac{1}{2}\iint\frac{\rho(x)\rho(y)}{|x-y|}\,dx\,dy + \int \epsilon_{\text{xc}}[\rho(x)]\,dx - \int \epsilon'_{\text{xc}}[\rho(x)]\rho(x)\,dx$

[1] Lin, Lu, Ying, E, JCP 2012

Energies, forces, degrees of freedom

- Total energies converged to **< 1e-3 Ha/atom** absolute error with **15 basis funcs/atom**

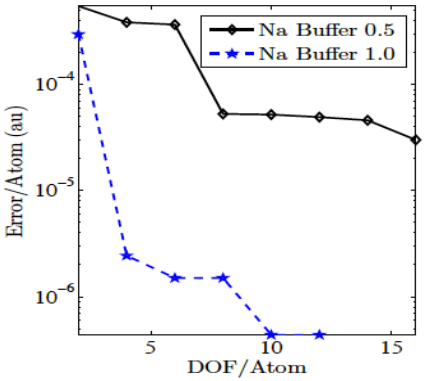

(a) Quasi-1D Na

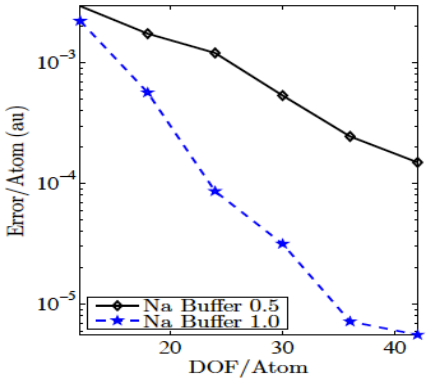

(b) Bulk 3D Na

- Forces converged to **< 1e-4 Ha/au** absolute error with **15 basis funcs/atom**

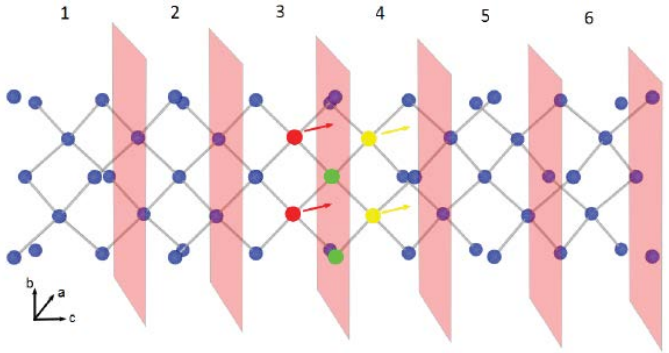

Si: random displacements, series of translations

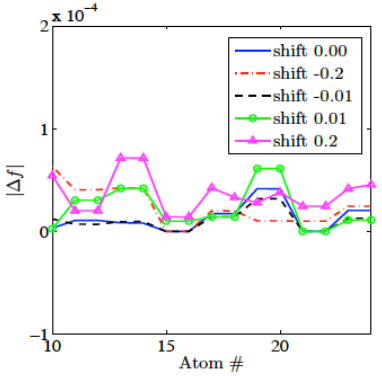

(a) Error of force

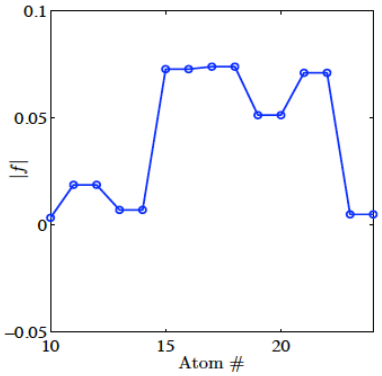

(b) Magnitude of force

- Largest system so far: **4,392 atoms** on 2,196 CPUs by direct diagonalization [1]
- New, parallel C++ code written: modular, extensible, nonlocal potentials, …

[1] Lin, Ying, PRB 2012

Issue

- Solution of the local ~ 50-atom Kohn-Sham problems (!)
- DG basis is so small and straightforward to evaluate that solution of the local K-S problems has become the bottleneck
- In collaboration with **FASTMath**, we are currently parallelizing the local K-S solutions to remove this bottleneck, and enable scaling of the code as a whole to thousands of times more cores
  - Harvesting massively parallel **Qbox** planewave code to accomplish as optimally and scalably as possible
  - Considering alternative spectral approaches to accommodate non-periodic potential in extended elements

For the largest systems: PEXSI

- Solving for Kohn-Sham wavefunctions of *N* atom system scales as O(N^3)
- Solve for density directly instead

$$\rho(x) = \mathrm{diag}\left(f_\beta(\hat{H}[\rho(x)] - \mu\delta(x,x'))\right)$$

$\hat{H}$ = Hamiltonian, μ = chemical potential, $f_\beta(x) = 2/(1+e^{\beta x})$
$\beta = 1/k_B T$, k_B = Boltzmann constant, T = temperature

- Need efficient approximation of Fermi function → **Pole expansion** [1]

$$f_\beta(\varepsilon - \mu) \approx \mathfrak{Im} \sum_{l=1}^{P} \frac{\omega_l^\rho}{\varepsilon - (z_l + \mu)}$$

$z_l, \omega_l^\rho \in \mathbb{C}$ are complex shifts and weights

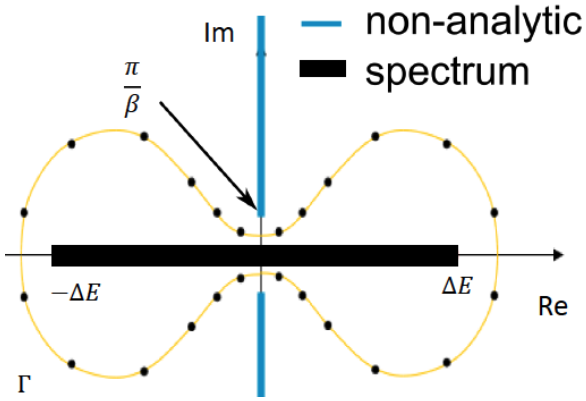

- Need efficient inversion
- Need only diagonal → **Selected Inversion** [2]
- → **Pole Expansion and Selected Inversion (PEXSI)**
- **No need to compute eigenfunctions or eigenvalues**
- Scaling O(*N*) for quasi-1D systems; **O(N^2) for metallic 3D**

[1] Lin, Lu, Ying, E, 2009; [2] Lin, Yang, Meza, Lu, Ying, E, 2011

Energies, forces, poles

- Metallic carbon nanotube, CNT (8,8), 512 atoms, atomic orbital basis [1]

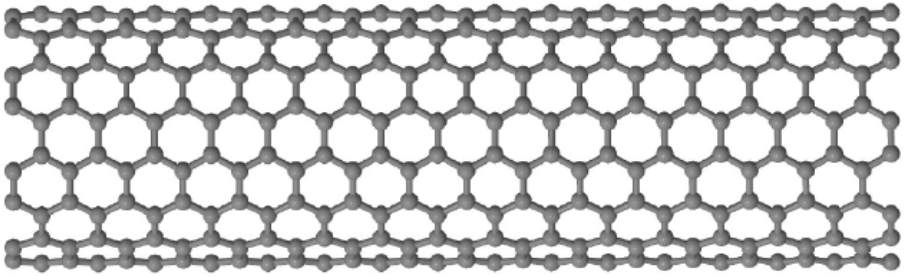

- Accuracy of expansion at T = 300K

# Poles	$E_{\mathrm{PEpSI}} - E_{\mathrm{ref}}$ (eV)	MAE Force (eV/Angstrom)
20	5.868351108	0.400431
40	0.007370583	0.001142
60	0.000110382	0.000026
80	0.000000360	0.000002

- New parallel PEXSI code in development
- Largest system so far: **20,256 atoms** on 256 CPUs

[1] Lin, Chen, Yang, He, J. Phys.: Cond. Mat., 2013

Issue

- Parallel scaling of *LU* factorization and SelInv (selected inversion)
- By pipelining and overlapping communication with computation, SelInv now faster and better scaling than SuperLU_DIST
- SuperLU_DIST scales to only ~1000 CPU
- In collaboration with **FASTMath**, we are exploring alternatives for better scaling LU
  - More robust symbolic factorization
  - Symmetric factorization: PARDISO (block fan-out), CLIQUE (multi-frontal)
  - Incorporating ideas from new pipelined SelInv
  - Incomplete factorizations
  - Leveraging results of previous SCF iteration

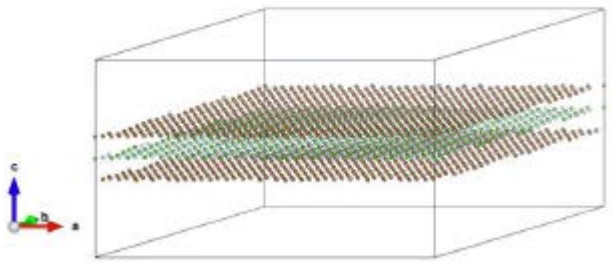

Scalability

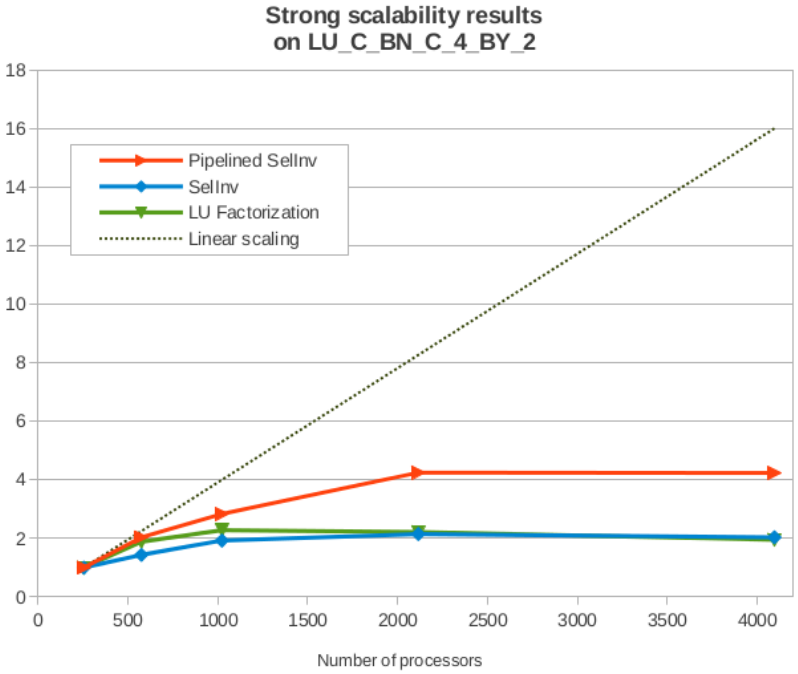

Thanks for your attention!

Please visit us at http://www.dgdft-scidac.org