Discontinuous methods for massively parallel quantum molecular dynamics: Lithium ion interface dynamics from first principles

John Pask (Lead), **Vincenzo Lordi, Mitchell Ong**
Condensed Matter and Materials Division, Lawrence Livermore National Laboratory

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

Lin Lin
Department of Mathematics, University of California, Berkeley
Computational Research Division, Lawrence Berkeley National Laboratory

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

Valerio Pascucci, Attila Gyulassi**
Scientific Computing and Imaging Institute, University of Utah
***SciDAC SDAV Institute*

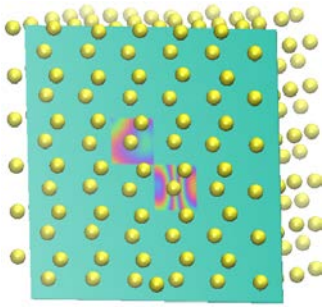

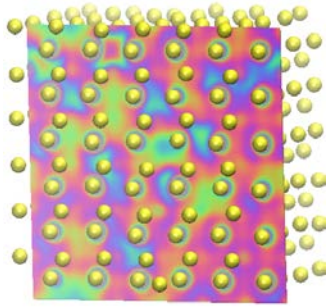

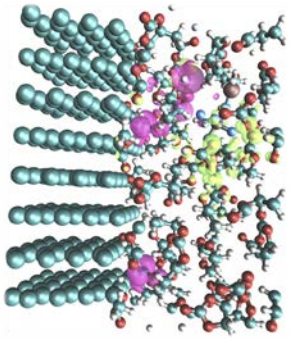

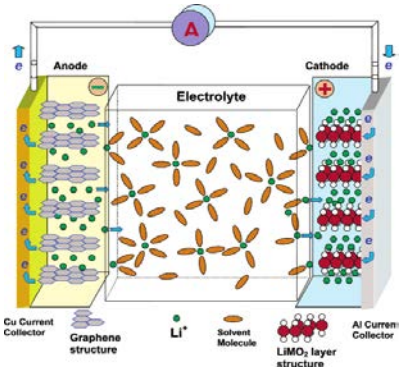

Overview

- 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)** based 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**

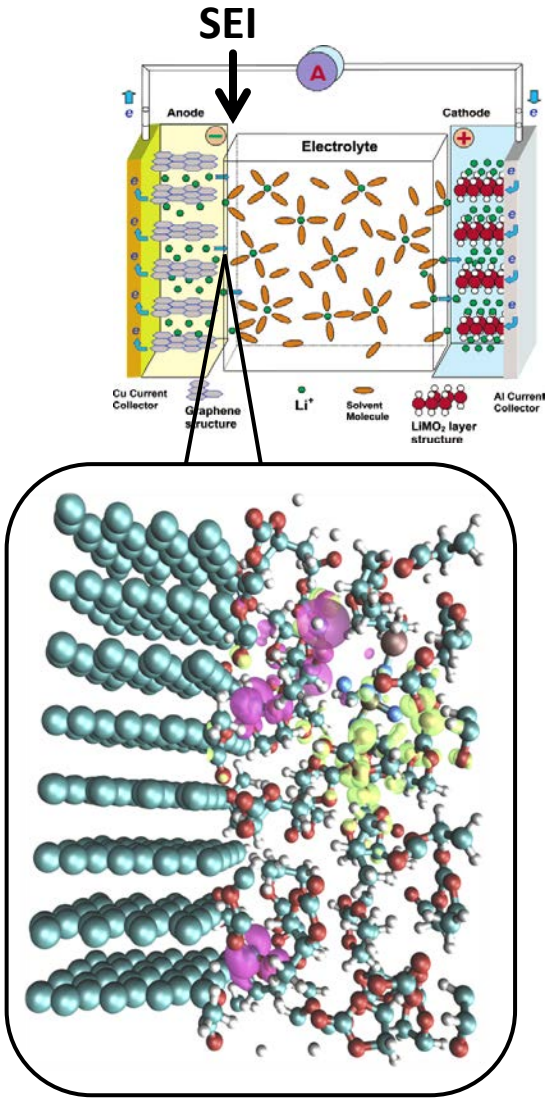

QMD snapshot of SEI layer in Li-ion cell

Simulations

- Initial phase of project, while new DG/PEXSI 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

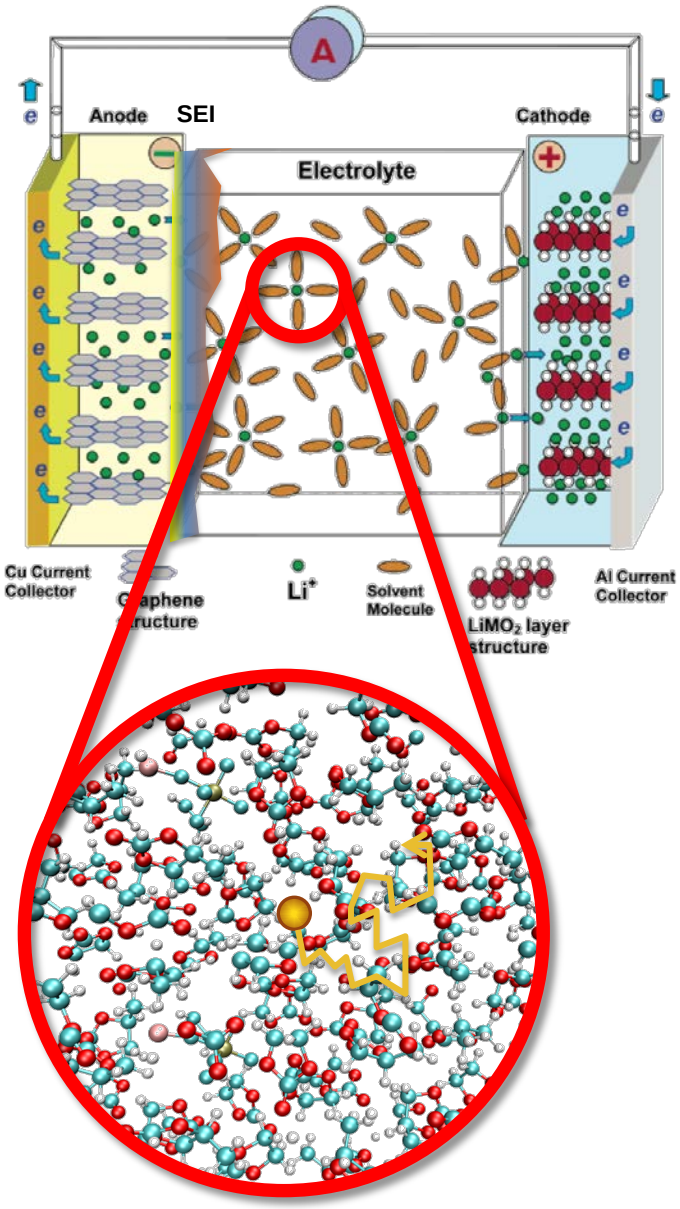

[Adapted from K. Xu, *Chem. Rev.* **104**, 4303 (2004)]

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

Ion solvation

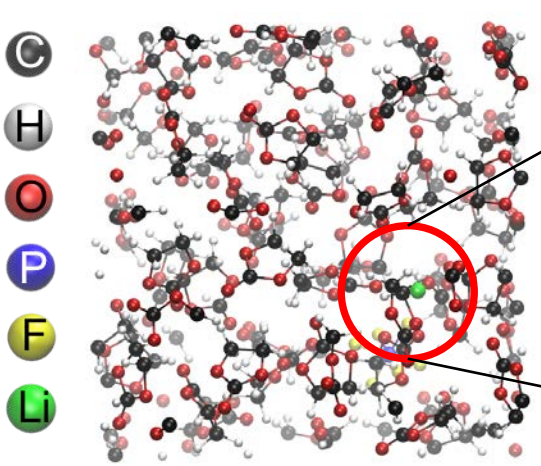

638 atoms
2190 electrons

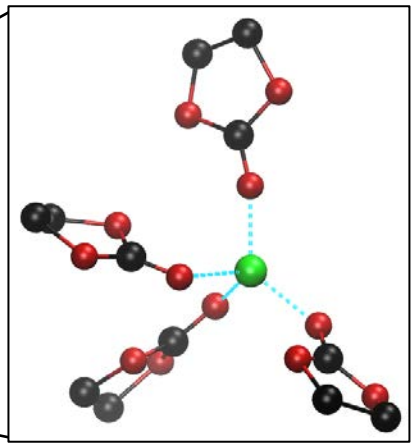

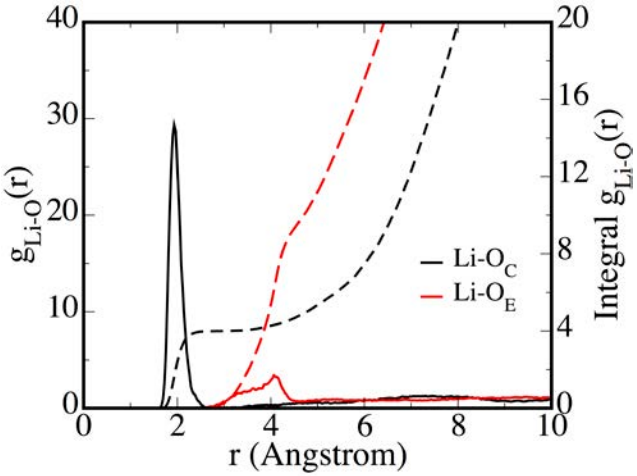

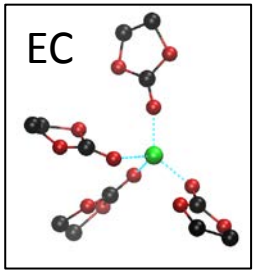

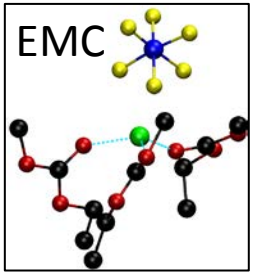

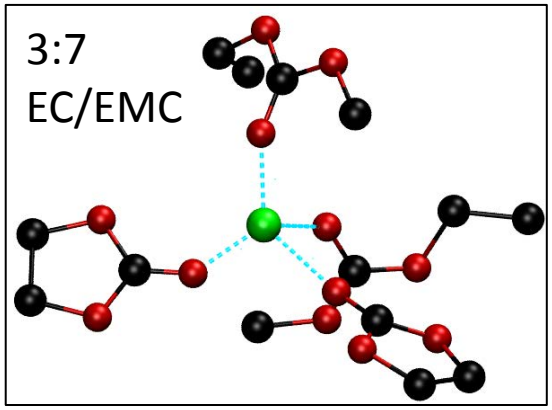

- Li^+ prefers tetrahedrally-coordinated first solvation shell
- Stronger solvation for Li^+ than PF_6^-
- Mixture shows preference for Li^+ to be solvated by EC over EMC
- PF_6^- more mobile than Li^+ due to weaker solvation

Ion diffusion

Ethylene Carbonate (EC)

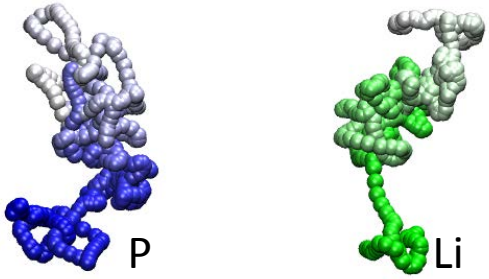

Ethyl Methyl Carbonate (EMC)

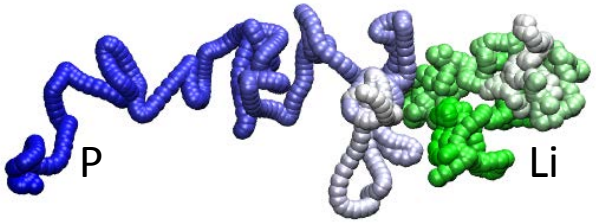

3:7 EC/EMC Mixture

	D_{Li} (10^{-6} cm^2/s)	D_P (10^{-6} cm^2/s)
EC	5.2 ± 0.8	7.1 ± 0.9
EMC	**9.6 ± 1.6**	**30.8 ± 8.8**
3:7 EC/EMC	2.6 ± 1.3	5.7 ± 2.4

- Motion of Li^+ and PF_6^- in EMC less correlated than in EC
- Mixture shows correlated motion similar to EC
- Faster diffusion for Li^+ seen in EMC than EC
- PF_6^- has larger diffusion coefficient than Li^+ since Li^+ is more strongly solvated than PF_6^-
- **Size of coefficient tied to solvation structure**

Simulations

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

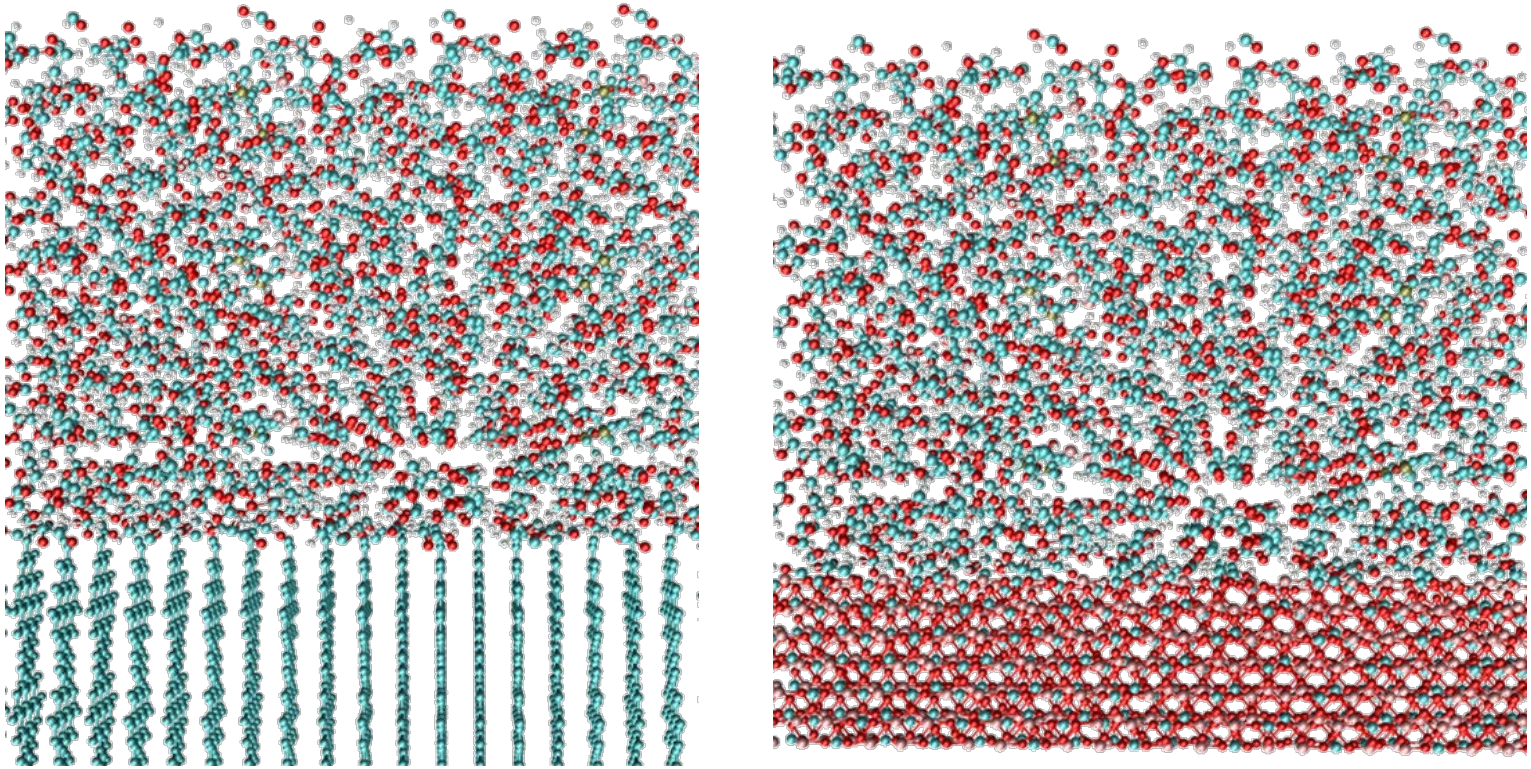

EC/PC mixture (+ $LiPF_6$) on graphite (left) and Li_2CO_3 (right), used to study chemical reactions on the anode surface (for initial SEI formation) and a representative SEI compound (for SEI growth/evolution)

How to reach the needed length and time scales?

Throw off the shackles of continuity.

Rethink the need for eigenfunctions.

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

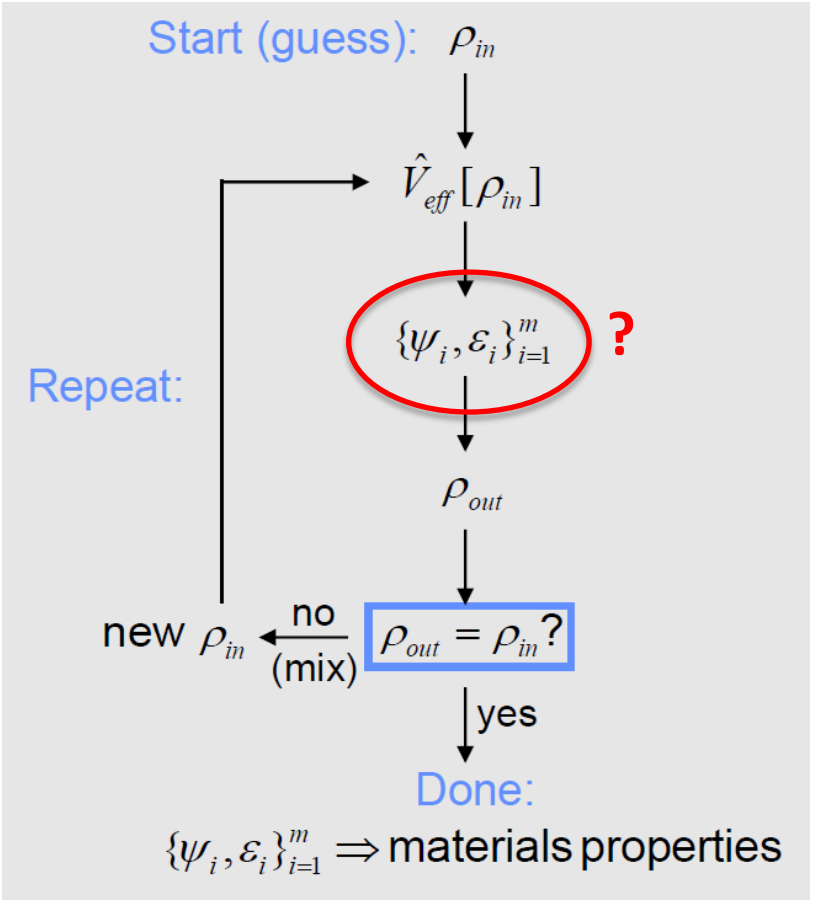

Pushing the current state of the art: Qbox application and development

- Strong scaling to **65,536 cores** on BG/Q
  - Uses hardware threading & SIMD registers on BG/Q
  - Optimized BLAS/ScaLAPACK kernels on BG/Q
  - Preconditioned steepest descent for occupied subspace
- **→ 1,700-atom anode-electrolyte system in 40 sec per QMD step**
- Up to 2,014-atom anode-electrolyte systems presently: 79 sec per MD step on 64k cores

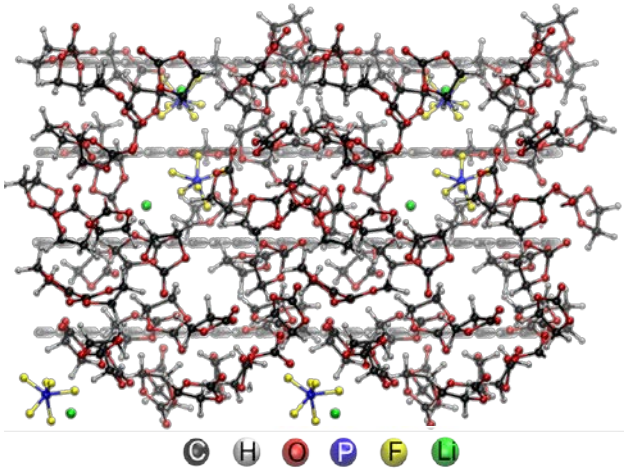

1700-atom anode-electrolyte cell

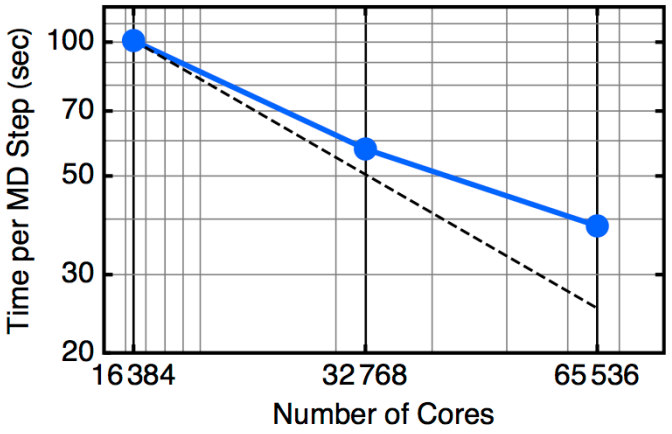

Qbox strong scaling on BG/Q

Current focus: Rayleigh-Ritz

- 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 **[1]**
  - Projected Preconditioned Conjugate Gradient (**PPCG**): Replace 3*m* x 3*m* Rayleigh-Ritz with *m* 3x3 ones **[2]**
  - PPCG now implemented in Quantum Espresso parallel planewave code
  - Initial tests indicate **factor of two speed up (already)** and **potentially superior parallel scalability** relative to current-state-of-the-art Davidson solver
- **Goal: Metallic as fast as insulating, 2,000-atom metallic QMD**

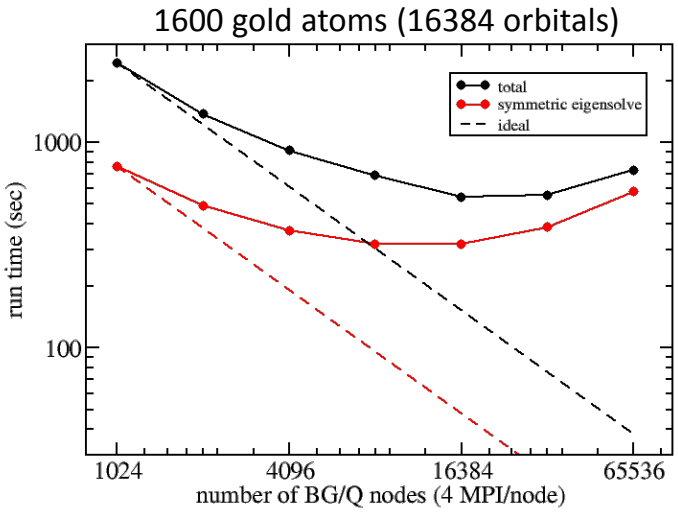

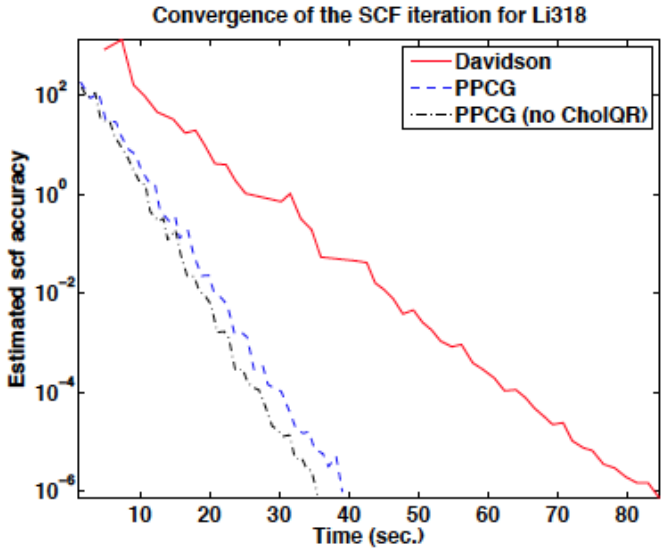

[1] Wen, Yang, Liu, Zhang, J. Sci. Comp., submitted
[2] Vecharynski, Yang, Pask, J. Comp. Phys., submitted

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

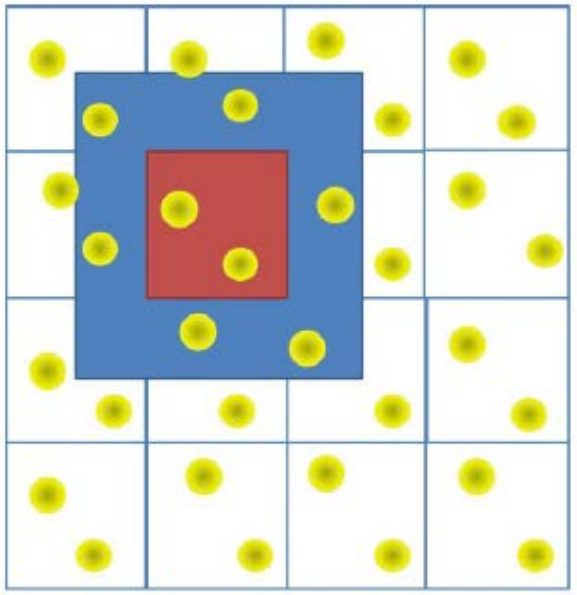

Solve large *N*-atom problem in highly efficient basis of O(*N*) local Kohn-Sham solutions

DG method

- Discontinuity is accommodated by surface ("flux") terms [1 - 4]
- 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; **[2]** Lin, Lu, Ying, E, JCP 2012; **[3]** Lin, Ying, PRB 2012, **[4]** Kaye, Lin, Yang, submitted

Energies, forces, degrees of freedom

- Energies to **< 1e-3 Ha/atom,** forces to **< 1e-3 Ha/au** absolute error with **~10 basis funcs/atom**

Si 1x1x4: random displacements

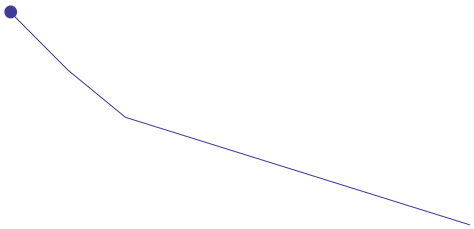

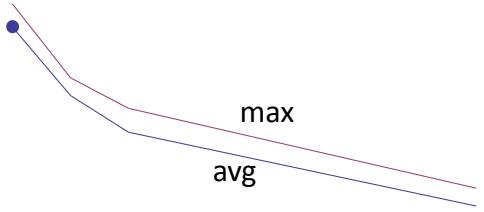

~40 DOF/atom
for 3D

- Hard atoms: Li-ion system – Li, P, F, C, H, O
- Forces to **1e-4 Ha/au** absolute error with **15 basis funcs/atom**

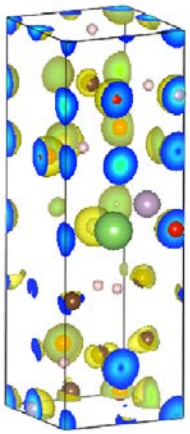

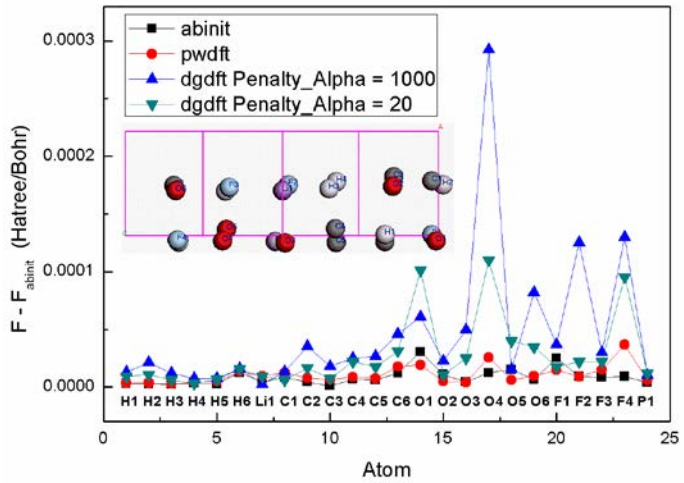

- Largest system so far: **4,392 atoms** on 2,196 CPUs by direct diagonalization

Molecular dynamics

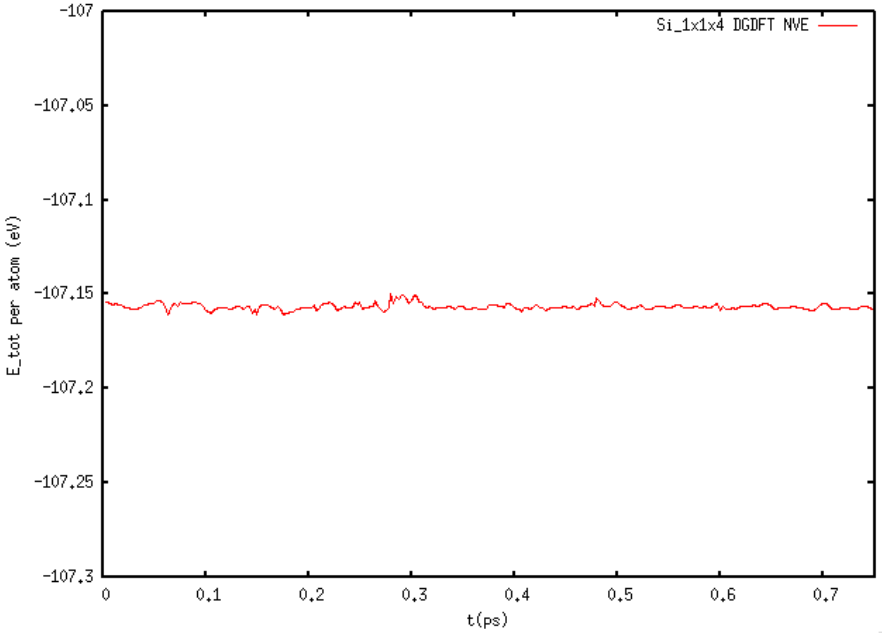

NVE MD simulation of liquid Si at 3000K using DGDFT

Energy drift < 1.5 meV/atom/ps

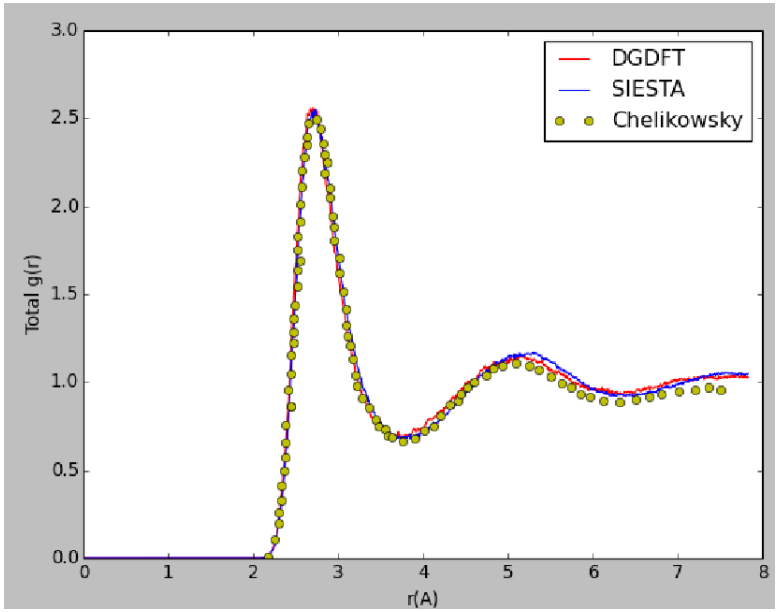

Pair correlation for $Si_{0.12}Al_{0.88}$ alloy with 200 atoms

Agreement with previous literature and independent calculations

Current focus: local K-S solves

- Solution of the local ~50-atom Kohn-Sham problems in each element
- In collaboration with **FASTMath**, we are parallelizing the local K-S solutions to remove this bottleneck
- **OpenMP**: speedup by factor of 6
  - Time per SCF reduced from **180.2 s** (OMP=1) to **29.4 s** (OMP=12) for 24-atom Li-ion system
- **MPI**: PWDFT can use more than **160 cores** with **70-fold speedup** to generate DG basis
  - Time per SCF ~10 s for 20-atom Yb_8O_{12} system
- More to come!

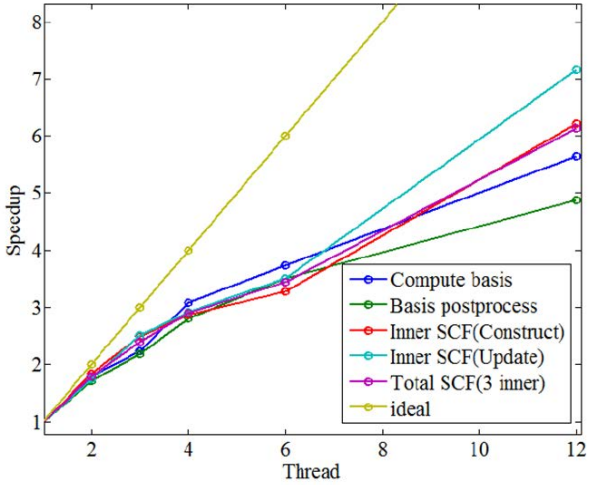

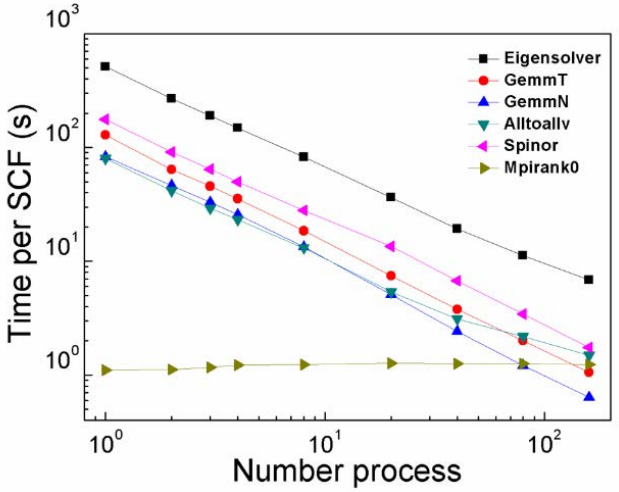

For the largest systems: PEXSI

- 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

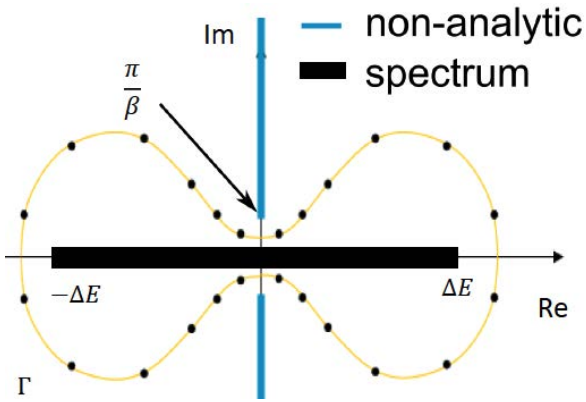

- Need efficient inversion
- Need only diagonal → **Selected Inversion** [2]
- → **Pole Expansion and Selected Inversion (PEXSI)**
- **No need to compute eigenfunctions or eigenvalues**
- Scaling O(*N*) for quasi-1D systems; **O(N^2) for metallic 3D**

[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]

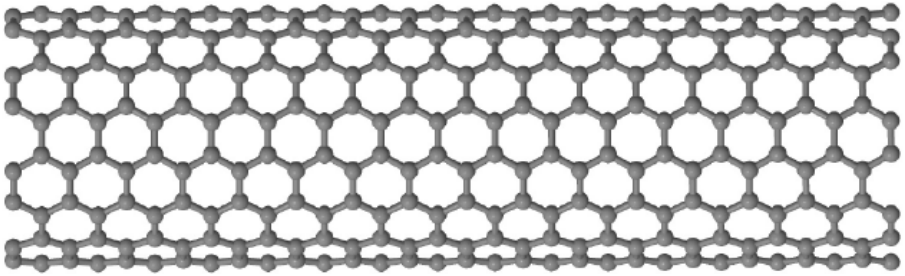

- Accuracy of expansion at T = 300K

# 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

- Largest 3D system so far: **24,000-atom** water using SIESTA DZP basis [2]
  - **314 sec** SCF iteration on 10,240 cores
  - **30-fold speedup** relative to diagonalization

[1] Lin, Chen, Yang, He, J. Phys.: Cond. Mat., 2013
[2] Lin, Garcia, Huhs, Yang, J. Phys.: Cond. Mat., 2014

PEXSI released

The PEXSI library is available online http://pexsi.org/, BSD license

- Integrated with SIESTA for accelerating atomic-orbital based calculations
- Being integrated into CP2K and other electronic structure packages
- Massively parallel to **10,000 – 100,000 processors** on high performance computers

[M. Jacquelin, L. Lin and C. Yang, submitted]

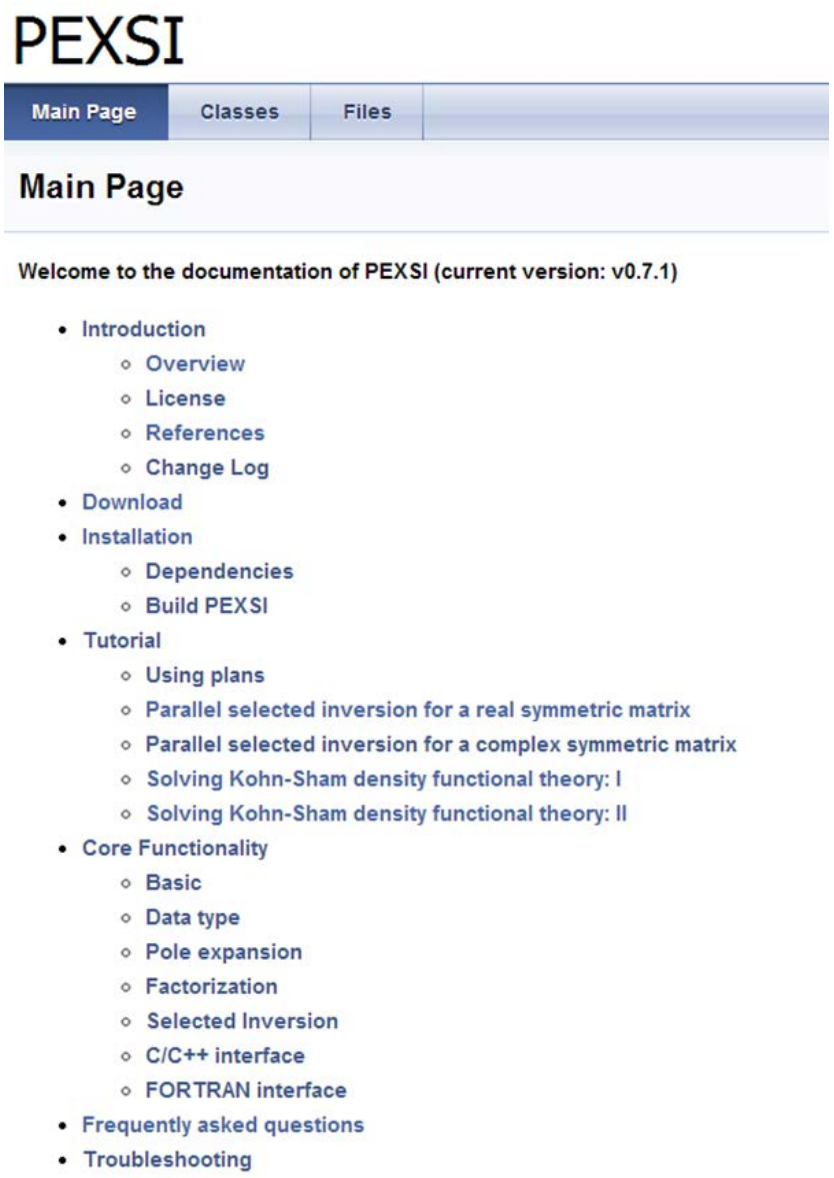

DG + PEXSI

- **2544-atom** Li-ion electrolyte

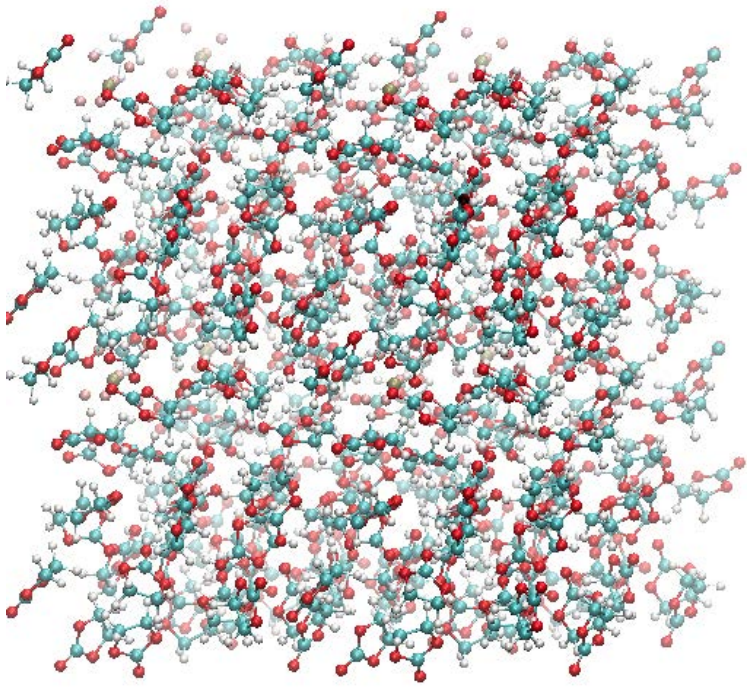

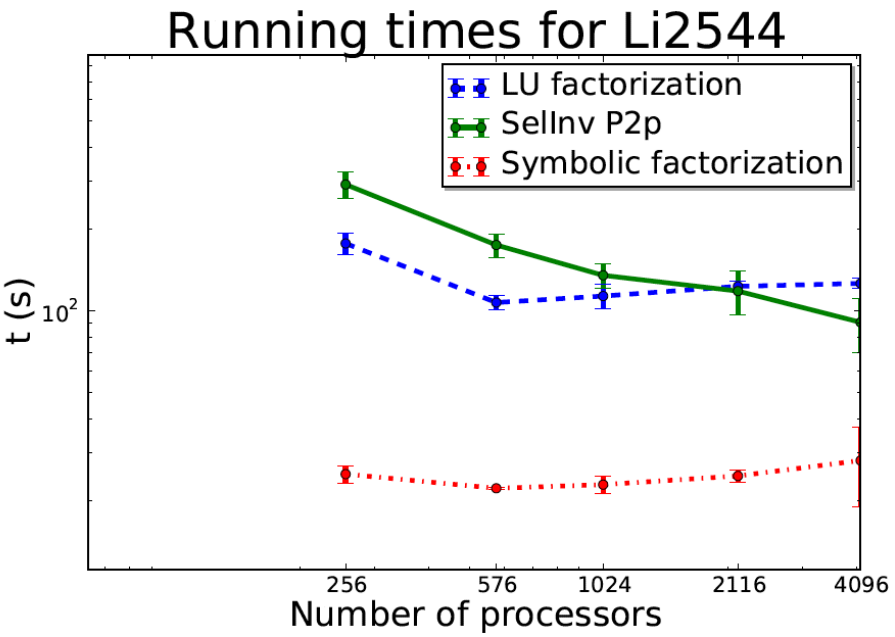

- 12x12x12 element partition, 68 basis functions per atom
- PESXI: ~250 sec per pole per μ iteration on 1024 cores
- High accuracy -> 40 poles, ~2 μ iterations -> **~500 sec/SCF on 40,960 cores**

Current focus: factorization

- Parallel scaling of *LU* factorization
- 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 sparse direct factorization
  - More robust symbolic factorization
  - New symmetric factorization code under development: exploring block fan-out and fan-both methods for better parallel scaling
  - Leverage results of previous SCF iteration: previous *H*, *LU*?
- To get SCF step times down from minutes to seconds for 3D systems may require iterative methods to fully leverage information from previous SCF/MD steps

New collaboration: SDAV

- Valerio Pascucci, Attila Gyulassi (University of Utah), Timo Bremer (LLNL)
- Systematic topological analysis of quantum mechanical data: density, potential, wavefunctions
- Compute Morse-Smale complex: minima, maxima, saddle points → mountains, voids, ridges, valleys, *connectivity*
- Evolution in space and time
- Bond formation and breaking
- Voids/tunnels for Li+ transport
- On the fly → inform/accelerate MD?

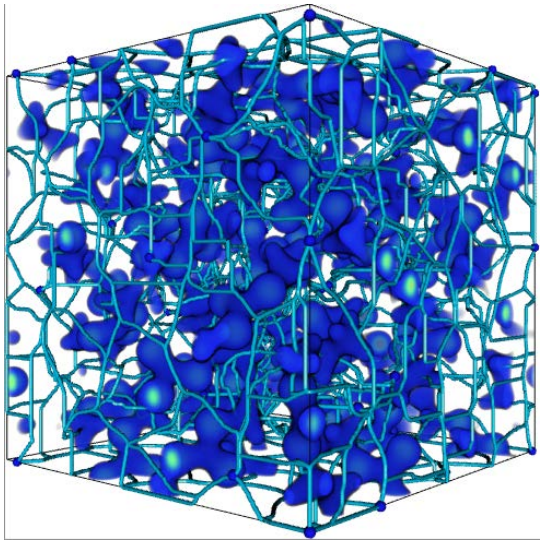

Valley lines connecting voids in QMD charge density

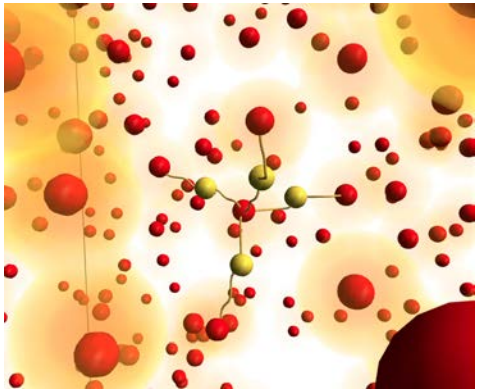

Saddle-max-saddle connections in distance field

Thanks for your attention!

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