Large-scale Eigensolvers for SciDAC Applications

Eugene Vecharynski Meiyue Shao Chao Yang Esmond Ng
Lawrence Berkeley National Laboratory

Overview

Eigenvalue problems arise in a number of SciDAC applications. We highlight some recent progress on 1) computing a large number eigenpairs of a Hermitian matrix in the context of density functional theory based electronic structure calculation 2) computing a few selected eigenpairs of a non-Hermitian matrix in the context of equation-of-motion coupled cluster (EOM-CC) calculation and complex scaling configuration interaction 3) computing all or a selected number of eigenpairs of the Bethe–Salpeter and Casida Hamiltonian matrices which have a special structure.

Computing a large invariant subspace of a Hermitian matrix

Motivation:

- Large-scale density functional theory based electronic structure calculations require computing a large number of lowest eigenpairs (10^3 pairs or more).
- Density functional perturbation theory requires many more lowest eigenpairs (10^3-10^5).

The challenge:

- Existing eigensolvers contain repeated calls of the Rayleigh–Ritz procedure that becomes a bottleneck when many eigenpairs are computed on a massively distributed-memory parallel machines.
- Standard computational kernels for solving dense eigenvalue problems (ScaLAPACK) do not scale beyond a certain number of cores.

Our goal:

- Compute many lowest eigenpairs on massively parallel high performance computers.
- Avoid or reduce the amount of the RR computations.

The Projected Preconditioned Conjugate Gradient (PPCG) algorithm

- The new eigensolver for computing large invariant subspaces of Hermitian matrices.
- The standard Rayleigh–Ritz procedure is replaced by a sequence of small dense eigenvalue problems plus a QR factorization of the approximate eigenspace.
- The Rayleigh–Ritz computation is performed only once every 5-10 iterations.
- Takes advantage of the available preconditioning techniques.
- Relatively easy to implement.

[1] E. Vecharynski, C. Yang, J. E. Pask: *A projected preconditioned conjugate gradient algorithm for computing many extreme eigenpairs of a Hermitian matrix*, J. Comp. Phys., Vol. 290, pp. 73–89, 2015

Performance of the PPCG algorithm in Quantum Espresso

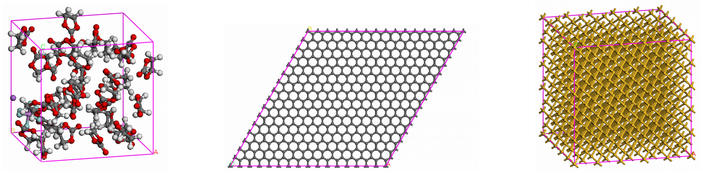

Benchmark systems: the solvation of LiPF6 in ethylene carbonate and propylene carbonate liquids containing 318 atoms (left), the 16 by 16 supercell of graphene containing 512 carbon atoms (center), and 5 by 5 by 5 supercell of bulk silicon containing 1000 silicon atoms (right).

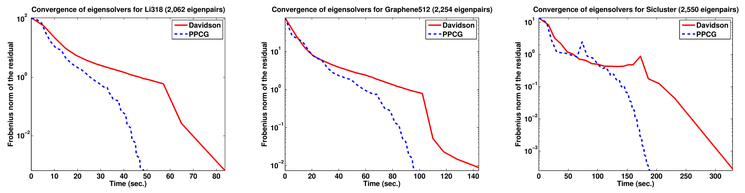

(a) Li318 (480 cores) (b) Graphene512 (576 cores) (c) Silicon1000 (2, 400 cores)

Performance profile

- Li318,
- Nev = 2000
- ncpus=480,
- Tol = 1e-6
- RR every 5 iterations
- Max subspace dimension for Davidson is 2k

component	PPCG	Davidson
HX	10	6.4
RR	10	55
GEMM	15	10
Cholesky QR	6	0
Total	41	71.4

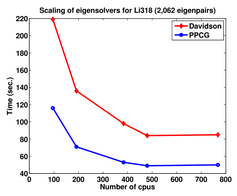

Ongoing work: PPCG in QBox

The ultimate goal is to perform a molecular dynamics simulation for up to 10,000 atoms to study lithium-ion batteries

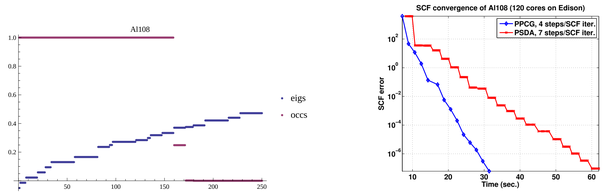

- 2x speedup due to a reduced SCF iteration count compared to a default QBox eigensolver (PSDA)

Non-Hermitian eigenvalue problems: computational challenges

Computing a subset of eigenpairs closest to the given shift σ

- Electronic resonant states (method of complex coordinate rotation).
- Equation-of-motion coupled-cluster (EOM-CC) method.

Difficulties with the existing solution approaches

- Require inverting $A - \sigma I$ ("shift-and-invert").
- Performance issues
  - Limited degree of parallelism ("one-by-one" eigenpair computation).
  - Failure to fully take advantage of BLAS3.
- Robustness issues.

Our goal:

- Develop a novel eigensolver that overcomes the known difficulties.

The Generalized Preconditioned Locally Harmonic Residual (GPLHR) method

- Uses the harmonic Rayleigh–Ritz procedure to extract approximate eigenpairs from low-dimensional search subspaces.
- Performs block iterations, effectively leverages BLAS3 kernels, provides multiple levels of concurrency.
- Takes advantage of the available preconditioning techniques.
- Robust, better convergence if memory is limited/tight.
- Provides an option of switching between the approximate eigenvector and Schur vectors iterations.

[1] E. Vecharynski, C. Yang, F. Xue: *Generalized preconditioned locally harmonic residual method for non-Hermitian eigenproblems*, SIAM J. Sci. Comput., submitted (2015)
[2] D. Zuev, E. Vecharynski, C. Yang, N. Orms, and A. I. Krylov: *New algorithms for iterative matrix-free eigensolvers in quantum chemistry*, J. Comp. Chem., Vol. 36, Issue 5, pp. 273–284, 2015

GPLHR in Q-Chem: EOM-CC benchmark

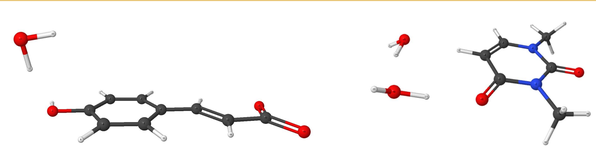

Benchmark systems: hydrated photoactive yellow protein chromophore PYPa-W$_p$ (left) and dihydrated 1,3-dimethyluracil (mU)$_2$-(H$_2$O)$_2$ (right).

PYPa-W$_p$/6-31+G(d,p)
GPLHR (σ = 11 a.u.)

nroots[a]	niters[b]	m	Max. # of stored vectors	# matvec[c]
1	4	1	8	9
2	4	1	16	18
3	4	1	24	27
5	8	1	40	63

[a] The number of requested eigenpairs. [b] The number of iterations to converge all eigenpairs. [c] The total number of matrix-vector multiplications. Davidson failed to deliver the solution.

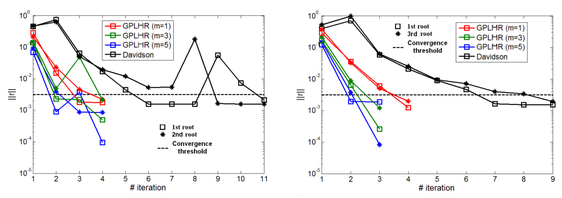

Left: PYPa-W$_p$/6-31+G(d,p) for the pairs with converged energies of 4.11 and 4.20 eV; Right: (mU)$_2$-(H$_2$O)$_2$/6-311+G(d,p) for the pairs with converged energies of 8.89 and 10.04 eV.

GPLHR in other applications

GPLHR exhibits a rapid and reliable convergence for a variety of standard and generalized eigenvalue problems across different applications (e.g., crystal growth simulation or stability analysis of fluid flows)

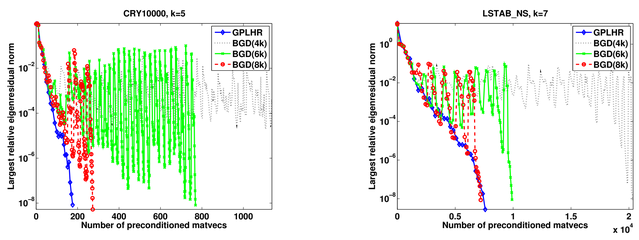

- Often requires significantly less memory to maintain convergence rate similar to state-of-the-art approaches (e.g., block Davidson methods).

Eigenvalue problems with a paired matrix structure

- The Bethe–Salpeter eigenvalue (BSE) problem
- The linear response (LR) eigenvalue problems in time-dependent density functional theory (TDDFT)

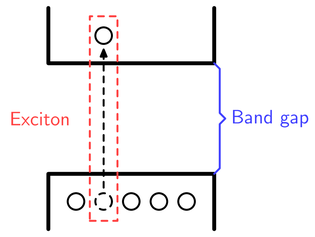

An exciton (electron-hole pair)

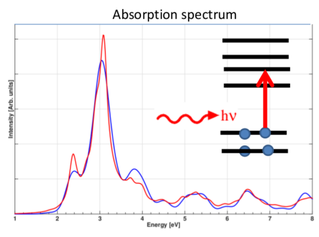

Absorption spectrum

Computing all/selected eigenpairs of the Casida Hamiltonian

Problem setting

- The Casida Hamiltonian matrix is of the form

$$H = \begin{bmatrix} A & B \\ -\overline{B} & -\overline{A} \end{bmatrix} \in \mathbb{C}^{2n\times 2n},$$

where $A = A^*$ is Hermitian, $B = B^T$ is complex symmetric, with

$$\begin{bmatrix} A & B \\ \overline{B} & \overline{A} \end{bmatrix} \succ 0 \quad \text{(i.e., } \begin{bmatrix} A & B \\ \overline{B} & \overline{A} \end{bmatrix} \text{ is Hermitian positive definite.)}$$

- In BSE, often all eigenpairs of H are required.
- In TDDFT, H is real and sparse for molecules, and only several smallest positive eigenpairs are needed.

Our goal:

- Develop structure-preserving parallel eigensolver for BSE.
- Develop and apply block preconditioned eigensolver techniques tailored specifically to the LR problem.

Methodology

- BSE can be reduced to a real positive definite Hamiltonian eigenvalue problem and solved by a skew-symmetric eigensolver.
- When H is real it is equivalent to solving an $n \times n$ product eigenvalue problem $MKx = \lambda^2 x$, where $M = A + B \succ 0$, $K = A - B \succ 0$.
- Cholesky factorization and SVD are used to solve $MKx = \lambda^2 x$ when all eigenpairs are needed.
- Several smallest eigenpairs of MK are computed by properly preconditioned iterative eigensolvers applied to the symmetric generalized eigenvalue problem $KMKx = \lambda^2 Kx$.

Experimental results

Parallel scalability of the dense BSE solver

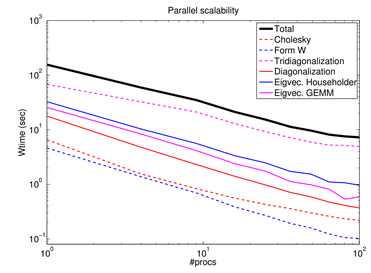

- The parallel solver is built on top of ScaLAPACK.
- The scalability of the solver is comparable to that of GEMM.

Efficiency of new LR eigensolvers

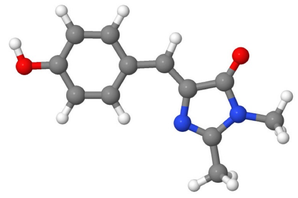

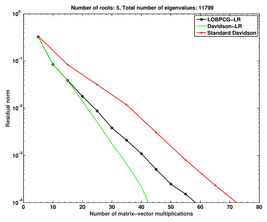

- New eigensolvers (LOBPCG-LR and Davidson-LR) achieve 2x speedup compared to the traditional Davidson approach for the LR eigenproblem.
- The proposed approaches require 2x less matrix-vector products and offer a significant reduction in memory usage.