Eigensolvers for Large Electronic Structure Calculations

Osni Marques

(*OAMarques@lbl.gov*)

Acknowledgments:

A. Canning, J. Dongarra, J. Langou, S. Tomov, C. Voemel and L.-W. Wang

Introduction

Photo luminescence of semi-conducting materials:

1. Electrons in stable initial state
2. Energy ⇒ electron "jumps" to previously unoccupied energy level
3. Electron jumps back ⇒ light

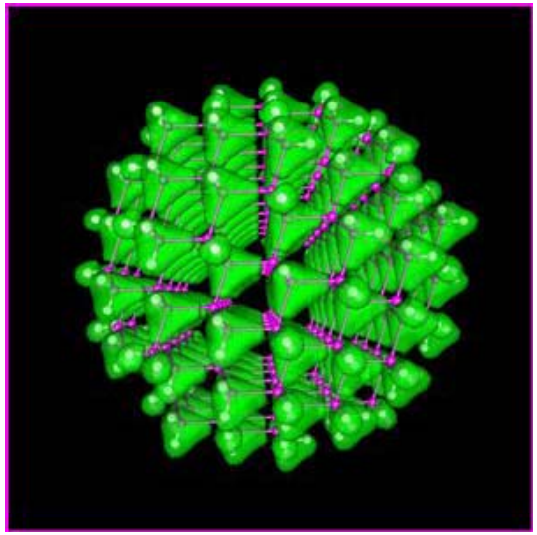

CdSe quantum dot (size)

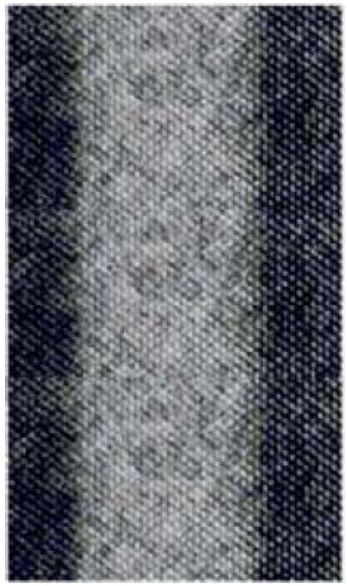

Problem and Physical Interpretation

$$H\Psi_i = \varepsilon_i \Psi_i$$

Schrödinger Equation

- Complex Hamiltonian $H = [-\frac{1}{2}\Delta + V]$
  - Δ is the kinetic energy term
  - V is the potential energy term
  - Implicitly defined by matrix-vector product (via FFT)
- Real eigenvalue ε_i
  - Discrete energy level
  - Can be occupied by electron or unoccupied
  - Clustered, multiplicities
- Complex eigenvector Ψ_i
  - Profile gives probability of finding electron at spatial location

Simulation Code: ESCAN (Energy SCAN)

- Solves single particle problem (*density functional theory*)
- Semi-empirical potential
- Non-selfconsistent calculations
- Plane-waves for larger systems
- Optical of electronic properties of interest
- Interior eigenvalue problem
- Folded spectrum method

- For more info, contact Lin-Wang Wang (*lwwang@lbl.gov*)

Spectral Transformations

- Shift-invert Rayleigh quotient

$$\rho([H - e_{ref} I]^{-1}, w)$$

- Folded spectrum

$$\rho([H - e_{ref} I]^{2}, w)$$

- Harmonic Rayleigh quotient

$$\rho(w) = \frac{w^{*}[H - e_{ref} I]w}{w^{*}[H - e_{ref} I]^{2} w}$$

ESCAN: *Folded Spectrum Approach*

$$[-\frac{1}{2}\nabla^2 + V(r)]\psi_i(r) = \varepsilon_i\psi_i(r)$$

$$H\psi_i = \varepsilon_i\psi_i \quad \Rightarrow \quad (H - \varepsilon_{ref} I)^2\psi_i = (\varepsilon_i - \varepsilon_{ref})^2\psi_i$$

Lowest 200 eigenstates of a CdSe system (n=4241), ε_{ref}=-5.0

⋮	
47	-6.85525E+00
48	-6.85525E+00
49	-6.70916E+00
50	-6.56302E+00
51	-6.38244E+00
52	-6.38244E+00
53	-2.11151E+00
54	-1.28873E+00
55	-9.98501E-01
56	-8.93434E-01
57	-8.93433E-01
58	-7.61729E-01
⋮	

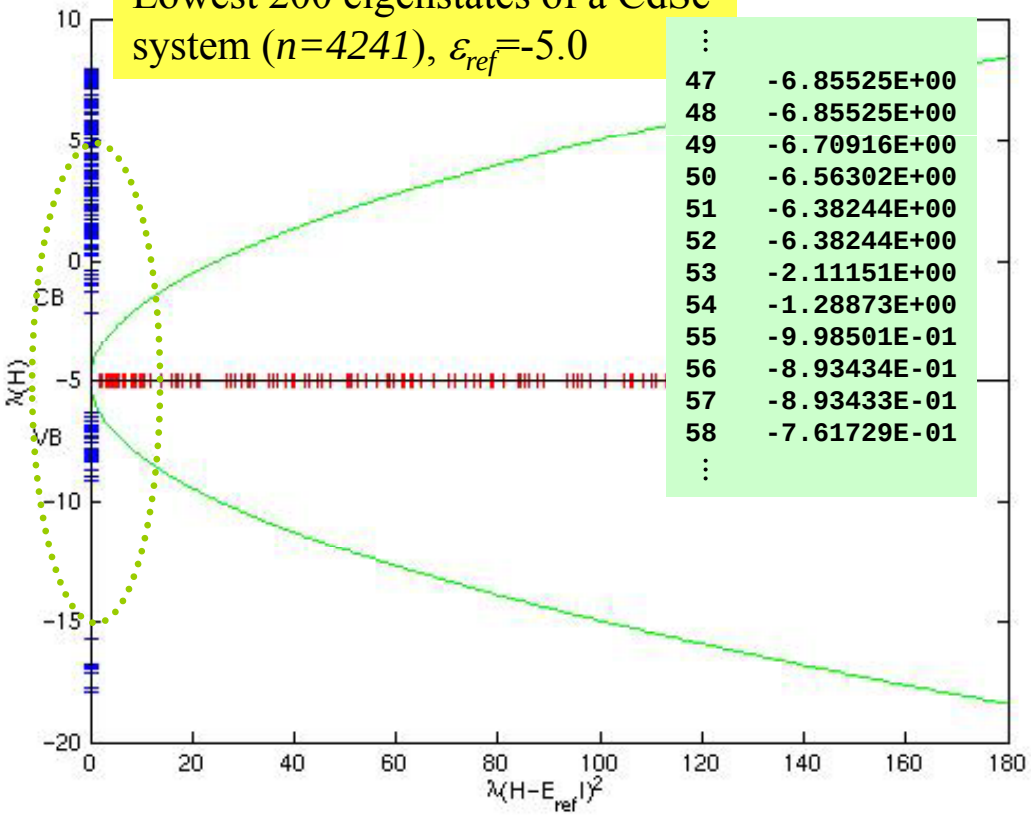

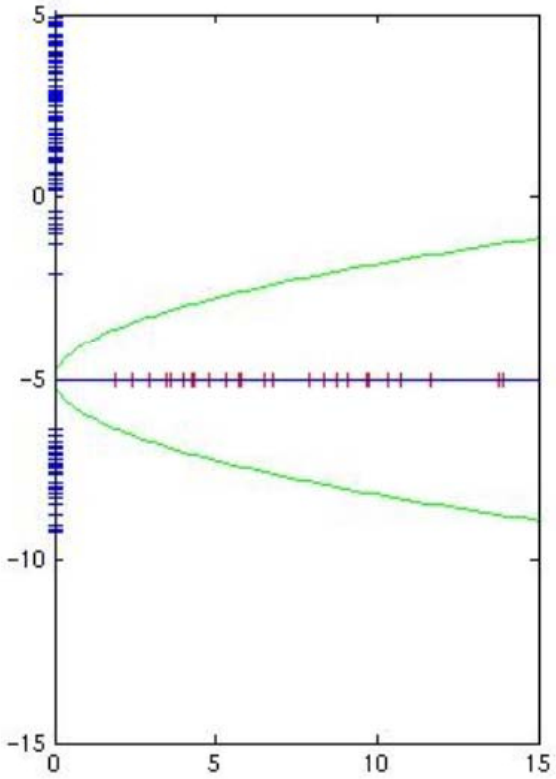

Eigensolvers of Choice

Algorithm	Details	Parameters
Banded PCG	Conjugate-Gradient (CG)-based Rayleigh-Quotient Minimization; implemented by Wang and Zunger.	**nline**
PARPACK	Implicit restarted Arnoldi (IRA); implemented by Lehoucq, Maschoff, Sorensen and Yang.	**ncv**
LOBPCG	Locally Optimal Block-Preconditioned CG; based on A. Knyazev.	-
PRIMME	Jacobi-Davidson, Preconditioned Iterative Multimethod Eigensolver; implemented by A. Stathopoulos and J. Combs.	**max basis size** **min restart size** **max block size** **max prev retain** **max inner iterations** ⋮

PRIMME $\rightarrow$ *tol* on $[(H - \varepsilon_{ref} I)^2 - (\varepsilon_i - \varepsilon_{ref})^2]\psi_i$

Banded PCG / LOBPCG

```
1   do i = 1, niter
2      do m = 1, numEvals
3         orthonormalize X(m) to X(1 : m − 1)
4         ax = A X(m)
5         do j = 1, nline
6            λ(m) = X(m) · ax
7            if (||ax − λ(m) X(m)||_2 < tol .or. j == nline) exit
8            r_{j+1} = (I − X(m) X(m)^*) ax
9            β = (r_{j+1} · Pr_{j+1}) / (r_j · Pr_j)
10           d_{j+1} = −P r_{j+1} + β d_j
11           d_{j+1} = (I − X(m)X(m)^*) d_{j+1}
12           γ = ||d_{j+1}||_2^{-1}
13           θ = 0.5 |atan( 2 γ d_{j+1} · ax / (λ(m) − γ^2 d_{j+1} · A d_{j+1}) )|
14           X(m) = cos(θ) X(m) + sin(θ) γ d_{j+1}
15           ax = cos(θ) ax + sin(θ) γ A d_{j+1}
16        enddo
17     enddo
18     [X, λ] = Rayleigh − Ritz on span{X}
19  enddo
```

```
1   do i = 1, niter
2      R = P (A X_i − λ X_i)
3      check convergence criteria
4      [X_i, λ] = Rayleigh − Ritz on span{X_i, X_{i−1}, R}
5   enddo
```

$$f(x) = \frac{x^* Hx}{x^* x}$$

$$\nabla f(x) = Hx - x\left(\frac{x^* Hx}{x^* x}\right) = r(x)$$

nline iterations of nonlinear CG

CdSe system (*n=4241*), ε_{ref}=-4.25, 30 line minimizations per iteration

iteration	line minimizations
1	30-30-30-30-30-30-30-30
2	30-30-30-30-30-30-30-30
10	26-23-1-30-30-30-30-30
15	1-1-1-1-30-30-30-30
20	1-1-1-1-1-30-30-30
25	1-1-1-1-1-1-1-30
30	1-1-1-1-1-1-1-30

Arnoldi with Implicit Restarts

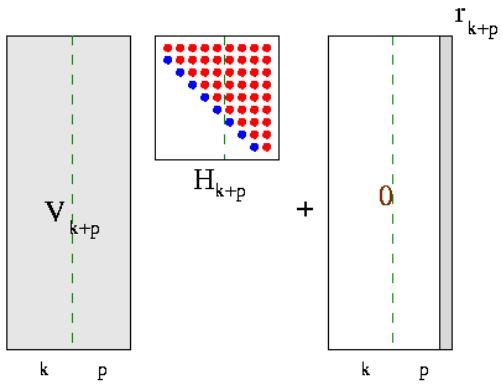

$$\text{Krylov Subspace}: \quad \mathsf{K}(A, q_1, j) = span(q_1, Aq_1, \ldots A^{j-1}q_1)$$

Arnoldi factorization at step *k*+*p*.

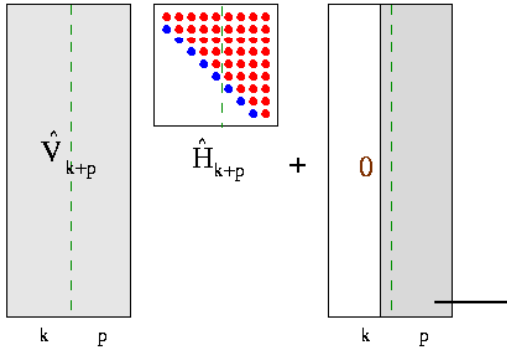

QR iteration on *H* with "special" shifts to promote convergence to

- the *k* eigenvalues with largest real part, or
- the *k* eigenvalues with largest magnitude, or
- the *k* eigenvalues with smallest real part, or
- the *k* eigenvalues with smallest magnitude

becomes non zero

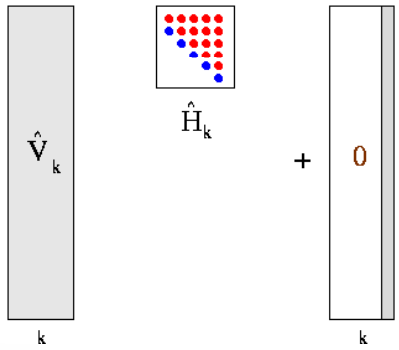

After discarding the last p columns, the final set represents a length *k* Arnoldi factorization.

Davidson / Jacobi-Davidson

$V_1 = [v_1],\ \|v_1\| = 1$

for $j = 1, 2, \cdots p$

a) $W_j = V_j^T A V_j$

b) solve $W_j \hat{y} = \hat{\theta} \hat{y}$

c) choose $(\hat{\theta}_j, \hat{y}_j)$

d) $\hat{x}_j = V_j \hat{y}_j$

e) $r_j = (\hat{\theta}_j I - A)\hat{x}_j$

f) if $\|r_j\| \le$ tol, stop

g) $t_j = (\hat{\theta}_j I - diag(A))^{-1} r_j$

h) $[V_j \ \ t_j] \Rightarrow V_{j+1},\ V_{j+1}^T V_{j+1} = I$

end for

→ (re)starting vector
block strategy: $V_1 = [v_1^{(1)}\ v_1^{(2)}\ \ldots]$

→ projection into subspace

→ min or max eigenvalue/eigenvector

→ Ritz vector

→ residual vector

→
- preconditioning of an auxiliary problem
- depends on diagonal dominancy
- $P = (\theta_1 I - diag(A))$ may be ill conditioned
- Jacobi-Davidson solves approximately $(I - \hat{x}_j \hat{x}_j^T)(A - \hat{\theta}_j I)(I - \hat{x}_j \hat{x}_j^T) t_j = -r_j$ (by QMR for example)

Test Cases

IBM SP3, 16 processors

System	atoms	n	time matvec (s)
Cd20Se19	39	11,331	0.005 (1.0)
Cd83Se81	164	34,143	0.014 (2.8)
Cd232Se235	467	75,645	0.043 (8.6)
Cd534Se527	1071	141,625	0.105 (21.)

- $neig = 10$
- $tol = 10^{-6}$
- $\varepsilon_{ref} = -4.8\text{eV}$
- diagonal preconditioner: $P = (I + (-\frac{1}{2}\nabla^2 + V_{avg} - \varepsilon_{ref}) / E_k)^2$
- IBM SP5 (8 to 32 processors)

Cd20Se19 (*n=11331, 8 procs*)

Folded Spectrum

ALGORITHM	nline	basis size	rest. size	prev. ret.	inner iter.	matvecs	time (s)
PCG	100	-	-	-	-	4956	9.4
LOBPCG	-	-	-	-	-	4756	19.3
PARPACK	-	20	-	-	-	14630	27.2
PARPACK	-	25	-	-	-	9712	18.1
PARPACK	-	30	-	-	-	7474	14.1
PARPACK	-	35	-	-	-	5838	11.1
PRIMME JDQMR	-	16	8	1	-	8546	14.6
PRIMME MIN_MATVECS	-	16	8	2	0	**1750**	**3.9**
PRIMME MIN_TIME	-	16	8	1	-1	4720	8.0

PRIMME JDQMR:
- adaptive stopping criterion for inner QMR

PRIMME MIN_MATVECS: currently GD_Olsen_plusk
- GD+k
- preconditioner applied to ($r+\varepsilon x$)

PRIMME MIN_TIME: currently JDQMR_Etol
- JDQMR
- stops after resid reduces by a 0.1 factor

eigenvalues	PARPACK
-6.19176	-6.43238
-6.19176	-6.43238
-6.34729	-6.60944
-6.38668	-6.60945
-6.43238	-6.71546
-6.43238	-6.71546
-6.60944	-6.88809
-6.60945	-6.91577
-6.71546	-6.98363
-6.71546	-7.08253

Cd20Se19 (*n=11331, 8 procs*)

Folded Spectrum

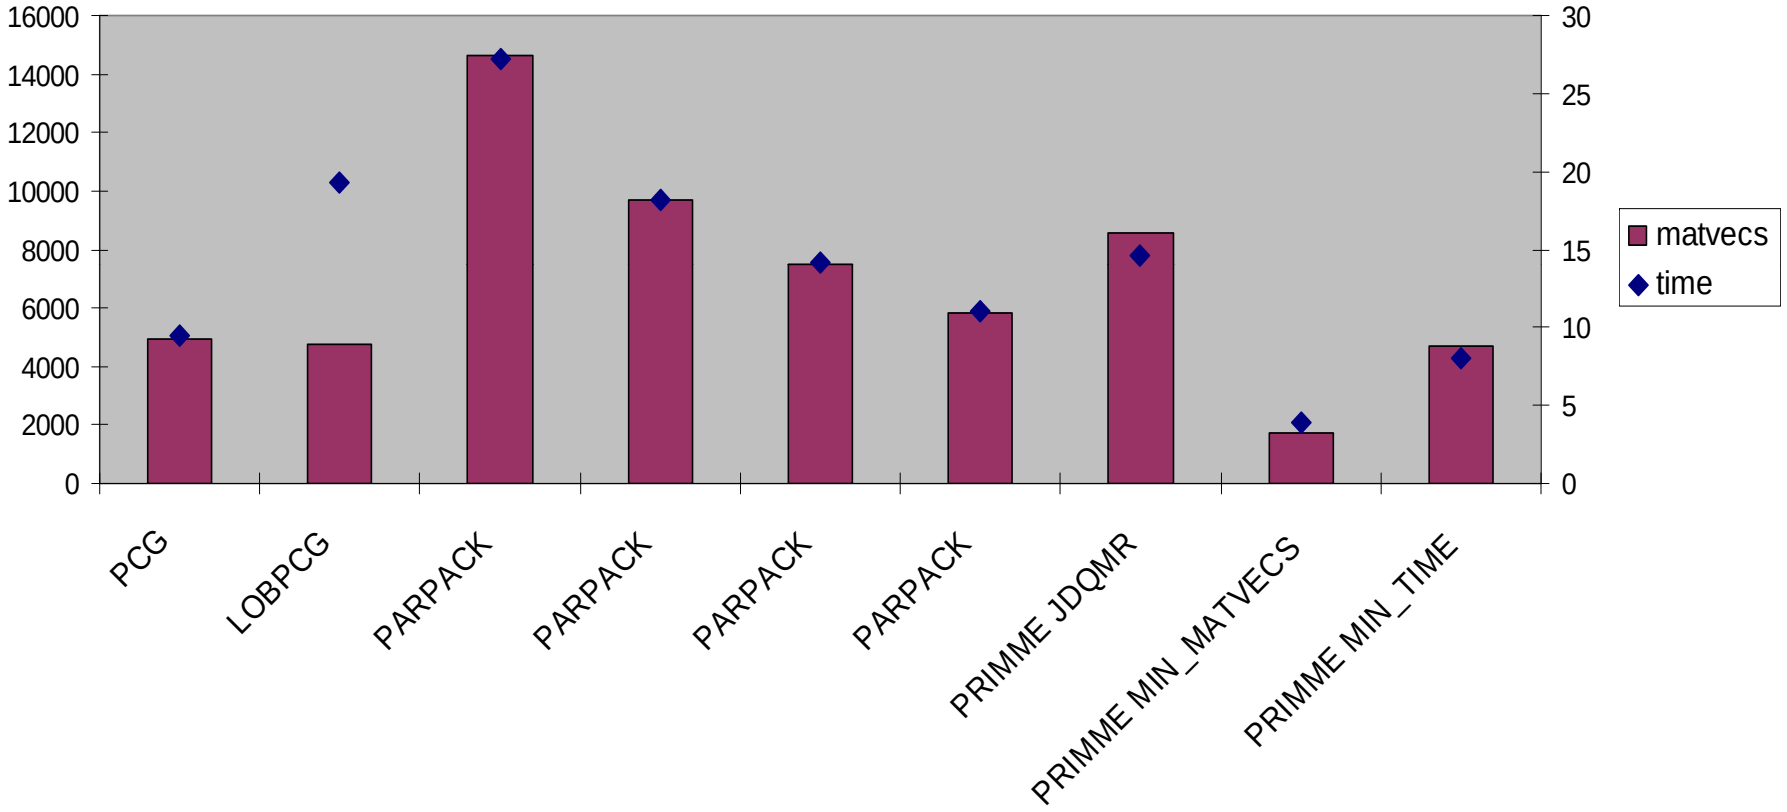

Cd20Se19 (*n=11331, 8 procs*)

Unfolded Spectrum

ALGORITHM	nline	basis size	rest. size	prev. ret.	inner iter.	matvecs	time (s)
PARPACK	-	20	-	-	-	****	****
PARPACK	-	25	-	-	-	**1326**	**2.9**
PARPACK	-	30	-	-	-	**1310**	**2.9**
PARPACK	-	35	-	-	-	**1293**	**2.9**
PRIMME MIN_MATVECS	-	16	8	2	0	4185	10.0
PRIMME MIN_TIME	-	16	8	1	0	**3350**	**7.0**

Implicit restarted Lanczos is (surprisingly) fast but (as in the folded spectrum) misses some eigenvalues.

eigenvalues	PARPACK
-6.19176	-6.19176
-6.19176	-6.34729
-6.34729	-6.38668
-6.38668	-6.43238
-6.43238	-6.43238
-6.43238	-6.60944
-6.60945	-6.60945
-6.60945	-6.71546
-6.71546	-6.71546
-6.71546	-6.88809

Cd83Se81 (*n=34143, 16 procs*)

ALGORITHM	nline	basis size	rest. size	prev. ret.	inner iter.	matvecs	time (s)
PCG	100	-	-	-	-	17920	65.6
PCG	200	-	-	-	-	15096	52.7
LOBPCG	-	-	-	-	-	10688	69.9
PARPACK	-	50	-	-	-	24252	86.7
PARPACK	-	100	-	-	-	15126	60.3
PRIMME MIN_MATVECS	-	30	10	2	0	**3670**	**12.7**
PRIMME MIN_TIME	-	30	10	1	-1	11808	36.7

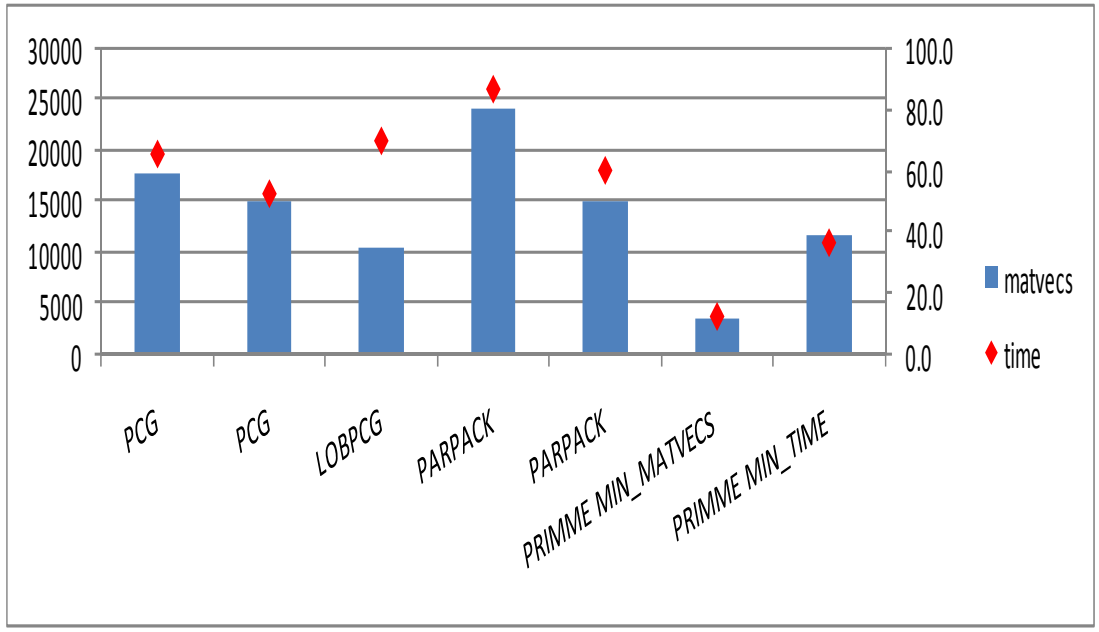

eigenvalues	PARPACK (1)	PARPACK (2)
-5.72654	-5.83003	-5.83003
-5.72654	-5.83003	-5.83003
-5.78686	-5.85207	-5.85207
-5.83003	-5.98438	-5.98438
-5.83003	-6.01278	-6.01278
-5.85207	-6.01278	-6.01278
-5.98438	-6.02422	-6.02422
-6.01278	-6.02751	-6.02422
-6.01278	-6.02751	-6.02751
-6.02422	-6.11332	-6.02751

Cd83Se81: trade-off between matrix-vector and orthogonalization

IBM SP3, 16 processors

Method	time(s)	matvecs	time in matvec (s)	% in matvec
Banded PCG	236.1	15096	201.46	85.3%
LOBPCG	190.4	10688	146.11	76.8%
JDQMR*	75.7	5314	73.20	96.7%
GD+1*	100.8	4084	57.17	56.6%

* Earlier version of PRIMME

Cd232Se235 (*n=75645*) and Cd534Se527 (*n=141625*)

Cd232Se235 (16 processors)

ALGORITHM	nline	basis size	restart size	prev. ret.	inner iter.	matvecs	time (s)
PCG	200	-	-	-	-	15754	106.4
LOBPCG	-	-	-	-	-	11864	121.4
PARPACK	-	30	-	-	-	****	****
PRIMME MIN_MATVECS	-	16	8	2	0	**3742**	**25.0**
PRIMME MIN_TIME	-	16	8	1	-1	11708	73.4

eigenvalues
-5.51570
-5.51570
-5.53926
-5.58286
-5.58286
-5.60869
-5.67889
-5.69688
-5.69688
-5.71672

Cd534Se527 (32 processors)

ALGORITHM	nline	basis size	restart size	prev. ret.	inner iter.	matvecs	time (s)
PCG	100	-	-	-	-	23810	228.0
LOBPCG	-	-	-	-	-	16862	254.7
PARPACK	-	30	-	-	-	20060	190.9
PRIMME MIN_MATVECS	-	16	8	2	0	**4762**	**46.0**
PRIMME MIN_TIME	-	16	8	1	-1	11259	109.1

eigenvalues
-5.39076
-5.39076
-5.40313
-5.44361
-5.44361
-5.48316
-5.49335
-5.51804
-5.51804
-5.52054

Entries of *H*

Cd20Se19

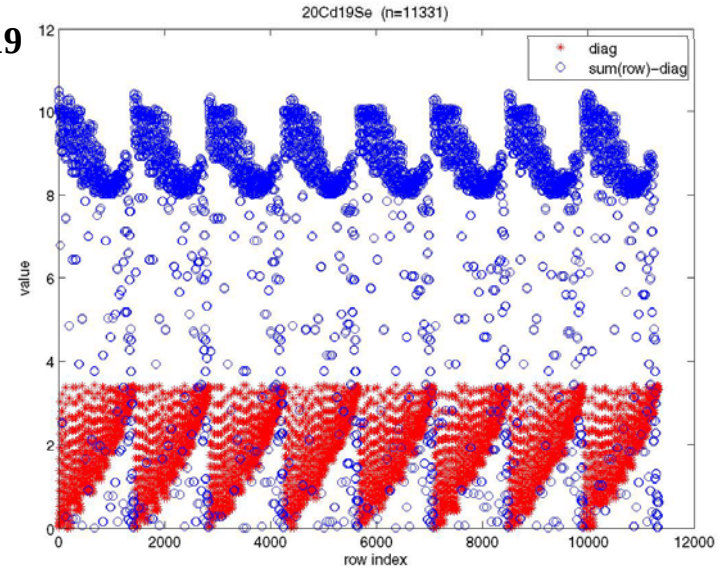

Cd83Se81

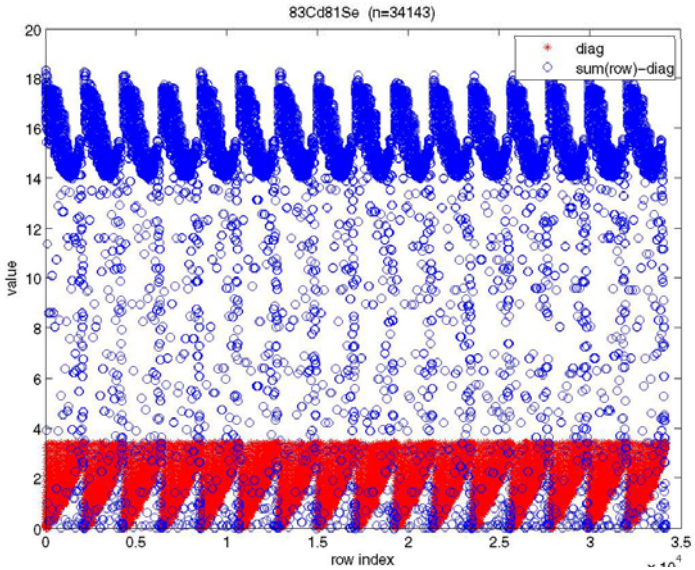

Cd232Se235

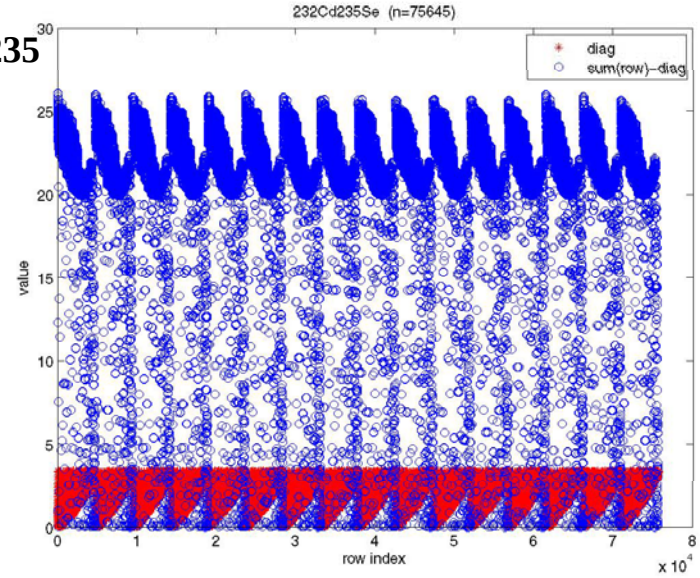

Cd534Se527

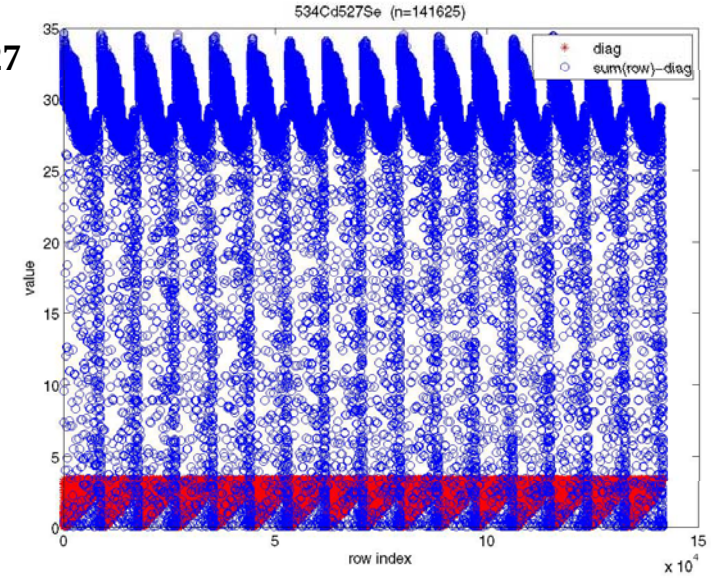

Cd675Se652 (*n=2717000*)

- Energy levels at the valence band maximum (VBM) and at the conductivity band minimum (CBM); $\varepsilon_{ref} = -0.4$eV and 0.6eV.
- *neig*=6, 64 processors, IBM SP5.

ALGORITHM	matvecs	time (s)	
CBM (folded spectrum)			
PCG *	335966	53670	(5.3x)
LOBPCG *	148486	28380	(2.8x)
PRIMME MIN_MATVECS**	**62334**	**10211**	(1.0x)
PRIMME MIN_TIME**	271492	43242	(4.2x)
VBM (folded spectrum)			
PCG *	101904	15671	(1.8x)
LOBPCG *	240030	41400	(4.7x)
PRIMME MIN_MATVECS**	**54362**	**8758**	(1.0x)
PRIMME MIN_TIME**	254810	39112	(4.5x)

VBM	CBM
-0.723983	1.357240
-0.723983	1.646169
-0.723983	1.646169
-0.729462	1.646169
-0.729462	1.923527
-0.729462	1.923535

* *not all eigenvalues satisfy tol = 1.0e-6 with the max number of iterations*
** *tol = 1.0e-10*

Quantum Wire System

- InAs nanowire embedded in bulk InP
- 66,624 atoms, ~2.3x10^6 equations, 64 processors
- CBM: ε_{ref}=-5.1eV, *neig* = 6

ALGORITHM	matvecs	time (s)	req tol
PCG	21931	1072	1.E-06
LOBPCG	20337	1377	1.E-06
PRIMME MIN_MATVECS (1)	**5438**	**292**	**1.E-06**
PRIMME MIN_MATVECS (2)	**8504**	**418**	**1.E-08**
PRIMME MIN_TIME (1)	16490	757	1.E-06
PRIMME MIN_TIME (2)	28076	1392	1.E-08

VBM	CBM
-5.73241	-4.89017
-5.73241	-4.71187
-5.73423	-4.68034
-5.74245	-4.68034
-5.74360	-4.55008

- VBM: ε_{ref}=-5.4eV, *neig* = 6

ALGORITHM	matvecs	time (s)	req tol
PCG	149726	7278	1.E-06
LOBPCG	56207	3690	1.E-06
PRIMME MIN_MATVECS (1)	**12670**	**2572**	**1.E-06**
PRIMME MIN_MATVECS (2)	**26326**	**1424**	**1.E-08**
PRIMME MIN_TIME	36310	1683	1.E-06

missed one eigenvalue; tighter tol fixed the problem

Conclusions and References

- Davidson type algorithms can significantly reduce the time required for eigenvalue calculations.
- Different algorithms (implementations) may require different tolerances.
- More work is needed for the unfolded spectrum (harmonic Ritz values).

- The Use of Bulk States to Accelerate the Band Edge State Calculation of a Semiconductor Quantum Dot, C. Voemel, S. Tomov, L.-W. Wang, O. Marques and J. Dongarra. Journal of Computational Physics, Vol. 223, pp. 774-782, 2007.
- State-of-the-art Eigensolvers for Electronic Structure Calculations of Large Scale Nano-systems, C. Voemel, S. Tomov, L.-W. Wang, O. Marques and J. Dongarra. To appear in Journal of Computational Physics.