## UNIVERSITY <br> OF Minnesota twin cities

Divide and conquer algorithms for large eigenvalue problems

Yousef Saad
Department of Computer Science and Engineering

## University of Minnesota

> PMAA 14
> Lugano, July 4, 2014

## Collaborators:

> Joint work with: Haw-ren Fang and Vassileos Kalantzis
> Grady Schoefield and Jim Chelikowsky [UT Austin] [windowing into PARSEC]
> Work supported in part by NSF (to 2012) and now by DOE

## Introduction

Q. How do you compute eigenvalues in the middle of the spectrum of a large Hermitian matrix?

A:
Common practice: Shift and invert + some projection process (Lanczos, subspace iteration..)

Main

1) Select a shift (or sequence of shifts) $\sigma$;
2) Factor $A-\sigma I: \quad A-\sigma I=L D L^{T}$
3) Apply Lanczos algorithm to $(A-\sigma I)^{-1}$
$>$ Solves with $A-\sigma I$ carried out using factorization
> Limitation: factorization

$Q$What if factoring $\boldsymbol{A}$ is too expensive (e.g., Large 3-D similation)?

A: Obvious answer: Use iterative solvers ...
$>$ But: systems highly indefinite $\rightarrow$ Won’t work well.
> Other common issue:
Need a very large number of eigenvalues and eigenvectors
> Applications: Excited states in quantum physics: TDDFT, GW, ... or just plain Density Functional Theory (DFT)
$>$ Example: in real-space code (PARSEC), Hamiltonian can be of size a few Millions, and number of ev's in the tens of thousands

## I. Polynomial filtered Lanczos

> Possible solution: Use Lanczos with polynomial filtering.
$>$ In short: just replace $(A-\sigma I)^{-1}$ in S.I. Lanczos by $p_{k}(A)$ where $p_{k}(t)=$ polynomial of degree $k$
> Idea not new (and not too popular in the past)

## What is new?

1. Very large problems;
2. (tens of) Thousands of eigenvalues;
3. Parallelism.
> Important application: compute the spectrum by pieces ['spectrum slicing' a term coined by B. Parlett]
> Main attraction: reduce cost of orthogonalization

## Low-pass, high-pass, \& barrier (mid-pass) filters



$>$ See Reference on Lanczos + pol. filtering: Bekas, Kokiopoulou, YS (2008) for motivation, etc.
> H.-r Fang and YS "Filtlan" paper [SISC,2012] and code

## Misconception: High degree polynomials are bad



## 'Spectrum slicing' or 'windowing'

Rationale. Eigenvectors on both ends of wanted spectrum need not be orthogonalized against each other :

> Idea: Get the spectrum by 'slices' or 'windows'
$>$ Can get a few hundreds or thousands of vectors at a time.

> Deceivingly simple looking idea.
> Issues:

- Deal with interfaces : duplicate/missing eigenvalues
- Window size [need estimate of eigenvalues]
- polynomial degree


## Spectrum slicing in PARSEC

$>$ Implemented in our code:
Pseudopotential Algorithm for Real-Space Electronic Calcultions (PARSEC)
$>$ See:
'A Spectrum Slicing Method for the Kohn-Sham Problem', G. Schofield, J. R. Chelikowsky and YS, Computer Physics Comm., vol 183 (2011) pp. 487-505.
$>$ Refer to this paper for details on windowing and 'initial proof of concept'

## Computing the polynomials: Jackson-Chebyshev

Chebyshev-Jackson approximation of a function $f$ :

$$
f(x) \approx \sum_{i=0}^{k} g_{i}^{k} \gamma_{i} \boldsymbol{T}_{i}(x)
$$

$\gamma_{i}=\frac{2-\delta_{i 0}}{\pi} \int_{-1}^{1} \frac{1}{\sqrt{1-x^{2}}} f(x) d x \quad \delta_{i 0}=$ Kronecker symbol $=$ explicitly known
$>$ The $g_{i}^{k}$ 's for $\boldsymbol{k}=50,100,150$
$>$ The $g_{i}^{k}$ 's dampen high order terms
> Alternative: Lanczos $\sigma$-damping


## A mid-pass (barrier) filter-3 damping methods

Mid-pass polynom. filter [-1 . 3 . 6 1]; Degree $=30$


## Tests - Test matrices

> Experiments on two dual-core AMD Opteron(tm) Processors 2214 @ 2.2GHz and 16GB memory.

## Test matrices:

* Five Hamiltonians from electronic structure calculations,
* Andrews matrix $N=60,000, n \boldsymbol{n z} \approx 760 K$, interval $[4,5]$; nev $=1,844$ eigenvalues, ( 3,751 to the left of $\boldsymbol{\eta}$ )
* A discretized Laplacian (FD) $n=10^{6}$, interval $=[1,1.01]$, nev=276, (>17,000 on the left of $\eta$ )
> Here : report only on Andrews and Laplacean


## Results for Andrews - set 1 of stats

| method | degree | \# iter | \# matvecs | memory |
| :---: | :---: | :---: | :---: | :---: |
| filt. Lan. <br> (mid-pass) | $d=20$ | 9,440 | 188,800 | 4,829 |
|  | $d=30$ | 6,040 | 180,120 | 2,799 |
|  | $d=50$ | 3,800 | 190,000 | 1,947 |
|  | $d=100$ | 2,360 | 236,000 | 1,131 |
| filt. Lan. (high-pass) | $d=10$ | 5,990 | 59,900 | 2,799 |
|  | $d=20$ | 4,780 | 95,600 | 2,334 |
|  | $d=30$ | 4,360 | 130,800 | 2,334 |
|  | $d=50$ | 4,690 | 234,500 | 2,334 |
| Part. $\perp$ Lanczos |  | 22,345 | 22,345 | 10,312 |
| ARPACK |  | 30,716 | 30,716 | 6,129 |

## Results for Andrews - CPU times (sec.)

| method | degree | $\rho(A) v$ | reorth | eigvec | total |
| :---: | :---: | :---: | :---: | :---: | :---: |
| filt. Lan. (mid-pass) | $d=20$ | 2,797 | 192 | 4,834 | 9,840 |
|  | $d=30$ | 2,429 | 115 | 2,151 | 5,279 |
|  | $d=50$ | 3,040 | 65 | 521 | 3,810 |
|  | $d=100$ | 3,757 | 93 | 220 | 4,147 |
| filt. Lan. (high-pass) | $d=10$ | 1,152 | 2,911 | 2,391 | 7,050 |
|  | $d=20$ | 1,335 | 1,718 | 1,472 | 4,874 |
|  | $d=30$ | 1,806 | 1,218 | 1,274 | 4,576 |
|  | $d=50$ | 3,187 | 1,032 | 1,383 | 5,918 |
| Part. $\perp$ Lanczos |  | 217 | 30,455 | 64,223 | 112,664 |
| ARPACK |  | 345 | 423,492 | †18,094 | 441,934 |

## Results for Laplacian - Matvecs and Memory

| method | degree | \# iter \# matvecs |  | memory |
| :---: | :---: | ---: | ---: | ---: |
| mid-pass filter | 600 | 1,400 | 840,000 | 10,913 |
|  | 1,000 | 950 | 950,000 | 7,640 |
|  | 1,600 | 710 | $1,136,000$ | 6,358 |

Results for Laplacian - CPU times

| method | degree | $\rho(\boldsymbol{A}) \boldsymbol{v}$ | reorth eigvec | total |  |
| :---: | :---: | :---: | :---: | ---: | ---: |
| mid-pass filter | 600 | 97,817 | 927 | 241 | 99,279 |
|  | 1,000 | 119,242 | 773 | 162 | 120,384 |
|  | 1,600 | 169,741 | 722 | 119 | 170,856 |

## II. Domain decomposition ideas

> Main idea: Cauchy integral-based method [e.g. FEAST]
> ... within a domain-decomposition framework:

$\leftarrow$ edgeseparators
vertex-
separators $\rightarrow$

Two classical ways of partitioning a graph.
> We use edge-separators (vertex-based partitioning)

Distributed graph and its matrix representation

$>$ Stack all interior variables $u_{1}, u_{2}, \cdots, u_{p}$ into a vector $\boldsymbol{u}$, then interface variables $\boldsymbol{y}$

Result:

$$
\underbrace{\left(\begin{array}{ccccc}
B_{1} & & & \ldots & E_{1} \\
& B_{2} & & \ldots & E_{2} \\
\vdots & & \ddots & & \vdots \\
& & & B_{p} & E_{p} \\
E_{1}^{T} & E_{2}^{T} & \ldots & E_{p}^{T} & C
\end{array}\right)}_{P A P^{T}}\left(\begin{array}{c}
u_{1} \\
u_{2} \\
\vdots \\
u_{p} \\
y
\end{array}\right)=\lambda\left(\begin{array}{c}
u_{1} \\
u_{2} \\
\vdots \\
u_{p} \\
y
\end{array}\right)
$$

## Notation:



Write as:
$\boldsymbol{A}=\left(\begin{array}{cc}\boldsymbol{B} & \boldsymbol{E} \\ \boldsymbol{E}^{T} & \boldsymbol{C}\end{array}\right)$

## First idea: Schur complement techniques (On-going work)

$>$ Eliminate interior variables $\boldsymbol{u}_{i}$ - Result:

$$
\underbrace{\left(\begin{array}{cccc}
\boldsymbol{S}_{1}(\lambda) & \boldsymbol{E}_{12} & \ldots & \boldsymbol{E}_{1 p} \\
\boldsymbol{E}_{21} & S_{2}(\lambda) & \ldots & \boldsymbol{E}_{2 p} \\
\vdots & & \ddots & \vdots \\
\boldsymbol{E}_{p 1} & \boldsymbol{E}_{p, 2} & \ldots & \boldsymbol{S}_{p}(\lambda)
\end{array}\right)}_{S(\lambda)} \underbrace{\left(\begin{array}{c}
\boldsymbol{y}_{1} \\
\boldsymbol{y}_{2} \\
\vdots \\
\boldsymbol{y}_{p}
\end{array}\right)}_{y}=0
$$

$S_{i}(\lambda)=C_{i}-E_{i}^{T}(B-\lambda I)^{-1} E_{i} \equiv$ Local Schur Complement
> Nonlinear eigenvalue problem.

$$
S(\lambda) y=0
$$

Involves only interface variables.
> Related to AMLS - see also Bekas and YS (2005)

## Next: Schur complements + FEAST

$$
\begin{aligned}
A-s I & =\left(\begin{array}{cc}
\boldsymbol{B}-s \boldsymbol{I} & \boldsymbol{E} \\
\boldsymbol{E}^{T} & C-s \boldsymbol{I}
\end{array}\right) \rightarrow \\
(\boldsymbol{A}-s \boldsymbol{I})^{-1} & =\left[\begin{array}{cc}
* & -(\boldsymbol{B}-s \boldsymbol{I})^{-1} \boldsymbol{E} \boldsymbol{S}(s)^{-1} \\
\hline * & \boldsymbol{S}(s)^{-1}
\end{array}\right]
\end{aligned}
$$

> Then, Cauchy integral formula for spectral projector yields:

$$
\begin{gathered}
P=\frac{-1}{2 i \pi} \int_{\Gamma} R(s) d s \equiv\left[\begin{array}{c|c}
* & -W \\
\hline * & G
\end{array}\right] \text { with } \\
G=\frac{-1}{2 i \pi} \int_{\Gamma} S(s)^{-1} d s, \quad W=\frac{-1}{2 i \pi} \int_{\Gamma}(B-s I)^{-1} \boldsymbol{E} S(s)^{-1} d s
\end{gathered}
$$

> Advantage: Does not involve inverse of whole matrix
$>$ Let

$$
\boldsymbol{P}=\left[\boldsymbol{P}_{1}, \boldsymbol{P}_{2}\right] \equiv\left[\begin{array}{c|c}
* & -\boldsymbol{W} \\
\hline * & \boldsymbol{G}
\end{array}\right]
$$

$>$ We know how to compute $\boldsymbol{P}_{2}$ or $\boldsymbol{P}_{2} \times$ randn $(\boldsymbol{s}, \boldsymbol{n} \boldsymbol{s})$
Q: How can we recover eigenvectors of $\boldsymbol{A}$ from $P_{2}$ ?
A: Write $\boldsymbol{P}$ as $\boldsymbol{P}=\boldsymbol{V} \boldsymbol{V}^{\boldsymbol{T}}$, and $\boldsymbol{V}=\binom{V_{u}}{V_{s}}$ then note:

$$
P_{2}=V V_{s}^{T}
$$

$>$ Just capture the range of $\boldsymbol{P}_{\mathbf{2}}$
$>$ Can use Lanczos on $P_{2} P_{2}^{T}$ or just a random $X \in \mathbb{R}^{s \times n s}$
> Advantage of Lanczos: stops when dimension reached
$>$ Drawbacks: 1) sequential; 2) $\approx$ Doubles the work
$>$ So far: Both idea tested in matlab

## Need better poles

> Approach is a one-shot method [no easy way to iterate]
Q: How can we improve accuracy?
A: Select poles carefully.
> Current choices: trapezoidal rule, Gauss, Zolotarev,...
> None of these allows for repeated ('multiple') poles e.g.,

$$
r(z)=\frac{\alpha_{1}}{z-\sigma}+\frac{\alpha_{2}}{(z-\sigma)^{2}}+\cdots+\frac{\alpha_{k}}{(z-\sigma)^{k}}
$$

$>$ This can be useful for any rational filtering approach
$>$ Next: See what we can do with one double pole



## Two double poles + comparison with compounding

Sigma $=+/-0.6+/-1 i ; \quad$ pow $=22$


Who needs a circle? Two poles ${ }^{2}$ far from the origin


## Conclusion

## Part I:

>FiltLan is appealing when number of eigenvectors to be computed is large and when Matvecs are not too expensise
> Will not work too well for generalized eigenvalue problem
> Code available here www.cs.umn.edu/~saad/software/
filtlan

## Part II:

> Many ideas still to explore in Domain Decomposition for interior eigenvalue problems
> Viewpoint: look at rational filtering from angle of approx. theory

