



THE CHASE LIBRARY

Parameter elimination and software architecture

June 28, 2018 | **E. Di Napoli**, J. Winkelmann, P. Springer |

OUTLINE

Chebyshev Accelerated Subspace Iteration Eigensolver (ChASE): the main algorithm

ChASE: parameter selection

ChASE library

OUTLINE

Chebyshev Accelerated Subspace Iteration Eigensolver (ChASE): the main algorithm

ChASE: parameter selection

ChASE library

SUBSPACE ITERATION METHOD

Properties and algorithm evolution

Basic properties and history

- eigenproblem needs to be in **standard form** $A' = L^{-1}AL^{-T}$ with $B = LL^T$
- it efficiently use **Mat-Mat** linear algebra kernels (e.g. `xHEMM`);
- it explores the whole subspace at once so avoiding stalling when facing small **clusters of eigenvalues**;
- it was first introduced by Rutishauser (1969).

Chebyshev Accelerated Subspace Eigensolver (**ChASE**)

- 1 currently only for dense eigenproblems;
- 2 locking mechanism with values ordering;
- 3 **accurate estimates** of eigenspectrum bounds;
- 4 polynomial **degrees optimized**.

CHASE PSEUDOCODE

INPUT: Hermitian A , tol, deg — **OPTIONAL:** approx. vectors W , approx. values $\{\mu_1 \dots \mu_{\text{nev}}\}$.
OUTPUT: nev wanted eigenpairs (Λ, Y) .

1 **Lanczos step.** Computes **spectral bound** $b_{\text{sup}} > \lambda_N$.

REPEAT UNTIL CONVERGENCE:

2 **Chebyshev filter.** **Filter** a block of vectors W .

3 **Re-orthogonalize** the vectors outputted by the filter; $W = QR$.

4 **Compute** the Rayleigh quotient $G = Q^\dagger A Q$.

5 **Compute** the primitive Ritz pairs $(\tilde{\Lambda}, Z)$ by solving for $GZ = Z\tilde{\Lambda}$.

6 **Compute** the approximate Ritz pairs $(\tilde{\Lambda}, W \leftarrow QZ)$.

7 **Compute** Eigenpair residuals $\text{Res}(w_a, \tilde{\lambda}_a)$ and check for **convergence**.

8 **Deflate** and **lock** the converged vectors and values in (Λ, Y) .

END REPEAT

THE CORE OF THE ALGORITHM: CHEBYSHEV FILTER

The basic principle

Theorem

Let $|\gamma| > 1$ and \mathbb{P}_m denote the set of polynomials of degree smaller or equal to m . Then the extremum

$$\min_{p \in \mathbb{P}_m, p(\gamma)=1} \max_{t \in [-1, 1]} |p(t)|$$

is reached by

$$p_m(t) \doteq \frac{C_m(t)}{C_m(\gamma)}.$$

where C_m is the Chebyshev polynomial of the first kind of order m , defined as



$$C_m(t) = \begin{cases} \cos(m \arccos(t)), & t \in [-1, 1]; \\ \cosh(m \operatorname{arccosh}(t)), & |t| > 1. \end{cases}$$

THE CORE OF THE ALGORITHM: CHEBYSHEV FILTER

Chebyshev polynomials

A generic vector $v = \sum_{i=1}^n s_i x_i$ is very quickly aligned in the direction of the eigenvector x_1

$$v^m = p_m(A)v = \sum_{i=1}^n s_i p_m(A)x_i = \sum_{i=1}^n s_i p_m(\lambda_i)x_i = s_1 x_1 + \sum_{i=2}^n s_i \frac{C_m(\frac{\lambda_i - c}{e})}{C_m(\frac{\lambda_1 - c}{e})} x_i \sim \boxed{s_1 x_1}$$



THE CORE OF THE ALGORITHM: CHEBYSHEV FILTER

In practice

Three-terms recurrence relation

$$C_{m+1}(t) = 2xC_m(t) - C_{m-1}(t); \quad m \in \mathbb{N}, \quad C_0(t) = 1, \quad C_1(t) = t$$

$$Z_m \doteq p_m(\tilde{A}) Z_0 \quad \text{with} \quad \tilde{A} = A - cI_N \quad \text{and} \quad c = \frac{b_{\text{sup}} + \mu_{\text{nev}}}{2} \quad e = \frac{b_{\text{sup}} - \mu_{\text{nev}}}{2}$$

FOR: $i = 1 \rightarrow \text{deg} - 1$

$$Z_{i+1} \leftarrow 2 \frac{\sigma_{i+1}}{e} \tilde{A} Z_i - \sigma_{i+1} \sigma_i Z_{i-1} \quad \boxed{\times \text{HEMM}}$$

END FOR.

OUTLINE

Chebyshev Accelerated Subspace Iteration Eigensolver (ChASE): the main algorithm

ChASE: parameter selection

ChASE library

CHASE PARAMETERS

A typical iterative methods depends on a number of tunable parameters. ChASE is no exception.

General input parameters

- N – Size of eigenproblem
- nev – Number of desired eigenpairs
- nex – Size of search space augmentation
- tol – Required threshold tolerance

Filter parameters

- deg – Polynomial degree
- μ_1 – Estimate for lowest eigenvalue
- b_{sup} – Bound for largest eigenvalue
- $\mu_{nev+nex}$ – Estimate for eigenvalue bounding search space

ESTIMATING THE SPECTRAL PARAMETERS

Estimate of μ_1 and b_{sup} are obtained by the simple repetition of few Lanczos steps¹

1. Compute k Lanczos steps

$$AU = UT_k + f_k e_k^\top \quad T_k = Z^H \tilde{\Lambda}_k Z \quad \tilde{\Lambda}_k = \text{diag}[\tilde{\lambda}_1, \dots, \tilde{\lambda}_k]$$

2. Compute upper bound

$$b_{\text{sup}} = \|f_k\|_2 + \max[\tilde{\lambda}_1, \dots, \tilde{\lambda}_k]$$

3. Estimate lower eigenvalue

$$\mu_1 = \max[\tilde{\lambda}_1, \dots, \tilde{\lambda}_k]$$

$k \sim 25$ is usually sufficient

¹Based on work by Zhou and Li (2011)

ESTIMATING THE SPECTRAL PARAMETERS

Estimating $\mu_{\text{nev}+\text{nex}}$ and requires additional Lanczos steps to build a spectral density².

1. Compute n_{vec} times k Lanczos steps

$$AU^{[j]} = U^{[j]}T_k^{[j]} + f_k^{[j]}e_k^\top \quad T_k^{[j]} = (Z^{[j]})^H \tilde{\Lambda}_k^{[j]} Z^{[j]} \quad \tilde{\Lambda}_k^{[j]} = \text{diag}[\tilde{\lambda}_1^{[j]}, \dots, \tilde{\lambda}_k^{[j]}]$$

2. Compute the spectral density

$$\tilde{\Phi}(t) = \frac{1}{n_{\text{vec}}} \sum_{j=1}^{n_{\text{vec}}} \sum_{i=0}^k (Z_{1,i}^{[j]})^2 g_\sigma(t - \tilde{\lambda}_i^{[j]})$$

3. Find $\bar{t} = \mu_{\text{nev}+\text{nex}}$ such that

$$\int_{-\infty}^{\bar{t}} \tilde{\Phi}(t) dt \approx \frac{\text{nev}+\text{nex}}{N}$$

Width of the Gaussian $\sigma = 0.25$

Number of random vectors $n_{\text{vec}} = 3 \div 5$

²Based on work by Lin et al. (2016)

POLYNOMIAL DEGREE OPTIMIZATION

Residuals vs convergence ratio

Definition

The **convergence ratio** for the eigenvector x_a corresponding to eigenvalue $\lambda_a \notin [\mu_{\text{nev}+\text{nex}}, b_{\text{sup}}]$ is defined as

$$|\rho_a|^{-1} = \min_{\pm} \left| \frac{\lambda_a - c}{e} \pm \sqrt{\left(\frac{\lambda_a - c}{e} \right)^2 - 1} \right|.$$

The further away λ_a is from the interval $[\mu_{\text{nev}+\text{nex}}, b_{\text{sup}}]$ the smaller is $|\rho_a|^{-1}$ and the faster the convergence to x_a is.

Residuals are a function of m and $|\rho|$

$$\text{Res}(w_a^m, \lambda_a) = \text{Const} \times \left| \frac{1}{\rho_a} \right|^m \quad 1 \leq a \leq \text{nev}+\text{nex}.$$

“Const” is independent of m and ρ

POLYNOMIAL DEGREE OPTIMIZATION

Tailoring polynomial degree for each eigenpair

1. Filter with an initial degree $m^{(0)}$ and compute residuals

$$\text{Res}(w_a^{m_a^{(0)}}, \tilde{\lambda}_a) \sim \text{Const} \times |\rho_a|^{-m_a^{(0)}} \quad 1 \leq a \leq \text{nev} + \text{nex}.$$

2. The residual of eigenpair after the next filtering step **would be**

$$\text{Res}(w_a^{m_a^{(1)}}, \tilde{\lambda}_a) \approx \text{Res}(w_a^{m_a^{(0)}}, \tilde{\lambda}_a) \times |\rho_a|^{-m_a^{(1)}} \quad 1 \leq a \leq \text{nev} + \text{nex}.$$

3. Compute the optimal minimal degree such that $\text{Res}(w_a^{m_a^{(0)}}, \tilde{\lambda}_a) \leq \text{tol}$

$$m_a^{(1)} \geq \ln \left| \frac{\text{Res}(w_a^{m_a^{(0)}}, \tilde{\lambda}_a)}{\text{tol}} \right| (\ln |\rho_a|)^{-1} \quad 1 \leq a \leq \text{nev} + \text{nex}.$$

CHASE PSEUDOCODE (OPTIMIZED)

INPUT: Hermitian A , tol , deg — **OPTIONAL:** approx. vectors W , approx. values $\{\mu_1 \dots \mu_{\text{nev}}\}$.
OUTPUT: nev wanted eigenpairs (Λ, Y) .

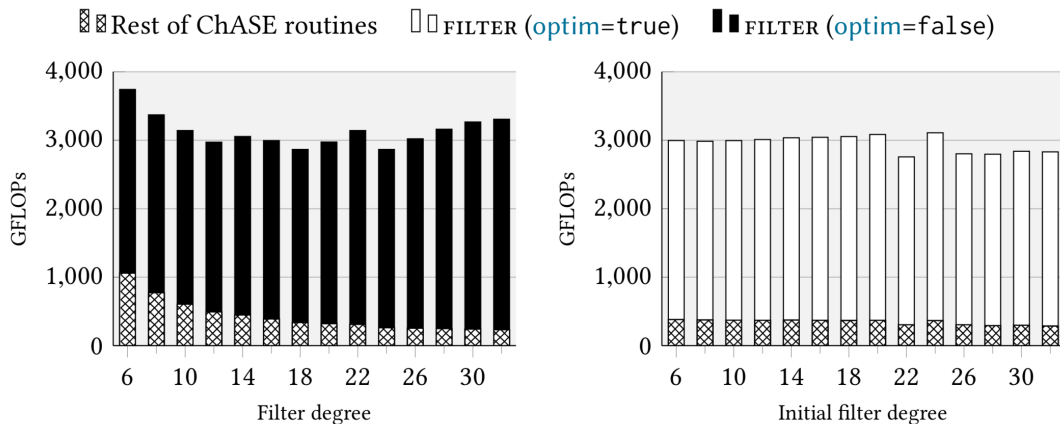
- 1 **Lanczos DoS step.** Computes **spectral estimates** $\mu_1, \mu_{\text{nev}+\text{nex}}$, and $b_{\text{sup}} > \lambda_N$.
- 2 Set the initial vector of degrees $m_a = \text{deg}$

REPEAT UNTIL CONVERGENCE:

- 3 **Chebyshev filter. Filter** a block of vectors W with a vector of degrees m_a .
- 4 **Re-orthogonalize** $W = QR$ & compute the Rayleigh quotient $G = Q^\dagger A Q$.
- 5 **Solve** the reduced problem $GZ = Z\tilde{\Lambda}$ and **compute** the approximate Ritz pairs $(\tilde{\Lambda}, W \leftarrow QZ)$.
- 6 **Compute and store** Ritz vectors residuals $\text{Res}(w_a, \tilde{\lambda}_a)$ and check for **convergence**.
- 7 **Deflate** and **lock** the converged vectors in (Λ, Y) .
- 8 **Optimizer. Compute vector** of polynomial degrees $m_a \geq \ln \left| \frac{\text{Res}(w_a, \tilde{\lambda}_a)}{\text{tol}} \right| / \ln |\rho_a|$.

END REPEAT

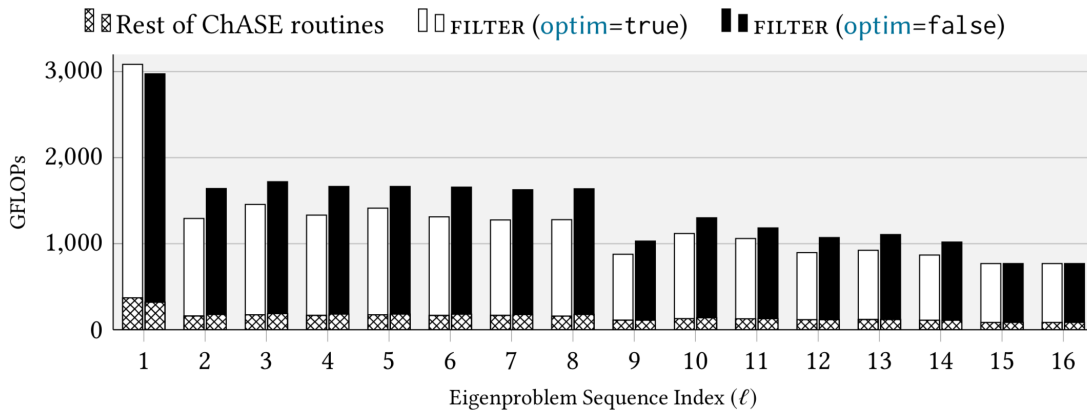
INDEPENDENCE FROM INITIAL DEGREE



(a) *Non* degree-optimized ChASE

(b) Degree-optimized ChASE

DEGREE OPTIMIZATION \Rightarrow FLOPS REDUCTION



CHASE PARAMETERS

In ChASE, filter parameters have been practically eliminated.

General input parameters

- N – Size of eigenproblem
- nev – Number of desired eigenpairs
- nex – Size of search space augmentation
- tol – Required threshold tolerance

Filter parameters

- deg – Polynomial degree
- μ_l – Estimate for lowest eigenvalue
- b_{sup} – Bound for largest eigenvalue
- $\mu_{nev+nex}$ – Estimate for eigenvalue bounding search space

OUTLINE

Chebyshev Accelerated Subspace Iteration Eigensolver (ChASE): the main algorithm

ChASE: parameter selection

ChASE library

CHASE LIBRARY RELEASED



- ChASE is open source (BSD 2.0 license) and available at
- <https://github.com/SimLabQuantumMaterials/ChASE>
- <https://arxiv.org/abs/1805.10121> (Submitted to ACM TOMS)

Highlights

- Modern C++ interface: easy-to-integrate in application codes.
- Multiple parallel implementations: performance portability.
- Excellent strong- and weak-scale performance.

CHASE ABSTRACT CLASS

Listing 1 Class interface that abstracts the ChASE algorithm from the numerical kernels

```
using T = std::complex<double>;
class Chase {
public:
    virtual void Start() = 0; // Alg. 1 line 1
    virtual void End() = 0; // Alg. 1 line 16
    virtual void Resd(double *ritzv, double *resd, // Alg. 1 line 8
                     size_t fixednev) = 0;
    virtual void Lock(size_t new_converged) = 0; // Alg. 3 line 7
    virtual void QR(size_t fixednev) = 0; // Alg. 1 line 5
    virtual void RR(double *ritzv, size_t block) = 0; // Alg. 2
    virtual void HEMM(size_t nev, T alpha, T beta, size_t s) = 0; // Alg. 4 line 9
    virtual void Lanczos(size_t k, double *upperb) = 0; // Alg. 6 line 3
    virtual void Lanczos(size_t M, size_t j, double *upperb, // Alg. 6 line 8
                         double *ritzv, double *Tau,
                         double *ritzV) = 0;
    virtual void LanczosDos(size_t idx, size_t m, T *ritzVc) = 0; // Alg. 6 line 13
    virtual void Shift(T c, bool isunshift = false) = 0; //  $A - cI_n$ 
    virtual void Swap(size_t i, size_t j) = 0; // Swap  $\hat{V}_{:,i}$  and  $\hat{V}_{:,j}$ 

    /* omitted Getters */
};
```

Listing 2 ChaseBLAS: an implementation of ChASE's kernels with BLAS+LAPACK

```
using T = complex<double>;
class ChaseBLAS : public chase::Chase {
public:
    ChaseBLAS(size_t N, size_t nev, size_t nex, T *H, T *V, T *W, double *ritzv,
              double *resid)
        : N_(N), nev_(nev), nex_(nex), locked_(0), H_(H), approxV_(V),
          workspace_(W), ritzv_(ritzv), resid_(resid), config_(N_, nev_, nex_) {}

    void HEMM(size_t block, T alpha, T beta, size_t s) override {
        zhemm('L', 'L', N_, block, &alpha, H_, N_, approxV_ + N_ * (locked_ + s),
              N_, &beta, workspace_ + N_ * (locked_ + s), N_);
        swap(approxV_, workspace_);
    };

    void Lock(size_t new_converged) override {
        memcpy(workspace_ + locked_ * N_, approxV_ + locked_ * N_,
              N_ * (new_converged) * sizeof(T));
        locked_ += new_converged;
    };


    /* Additional functions from Listing 1 omitted for conciseness */
private:
    int N_, nev_, nex_, locked_;
    T *A_, *approxV_, *workspace_;
    double *ritzv_, resid_;
    ChaseConfig<T> config_;
};
```

EXPERIMENTAL SETUP


Table: Matrices used in scaling experiments

	n	n_{ev}	n_{ex}	# Nodes	JURECA		BLUEWATERS	
					# Cores	$\frac{n^2}{\# \text{ Cores}}$	# Cores	$\frac{n^2}{\# \text{ Cores}}$
NaCl	3893	256	51	4	96	N/A	N/A	N/A
	9273	256	51	25	600	N/A	N/A	N/A
AuAg	13,379	972	194	25	600	N/A	N/A	N/A
BSE	22,360	100	20	9	216	2,314,674	64	7,812,025
	32,976	100	20	16	384	2,831,814	128	8,495,442
	47,349	100	20	36	864	2,594,823	288	7,784,471
	62,681	100	20	64	1536	2,557,882	512	7,673,648
	76,674	100	20	100	2400	2,449,542	800	7,348,628

STRONG SCALING



(a) Runtimes



(b) Parallel efficiency

WEAK SCALING ON JURECA

Table: Weak scaling experimental results on JURECA

# Cores	Iterations		Matvecs		Runtime		
	ChASE-BLAS	ChASE-Elemental	ChASE-BLAS	ChASE-Elemental	ChASE-BLAS	ChASE-Elemental	Direct
216	11	11	19,990	20,192	25.1 s	26.0 s	81.5 s
384	10	9	16,778	16,100	23.7 s	24.0 s	141.2 s
864	17	11	23,424	27,506	39.8 s	45.2 s	211.1 s
1536	13	12	23,268	21,940	36.4 s	41.4 s	367.8 s
2400	10	13	22,614	21,720	38.4 s	40.8 s	380.1 s

Tests were performed on the JURECA cluster.

- 2 Intel Xeon E5-2680 v3 Haswell – Up to $0.96 \div 1.92$ TFLOPS DP/SP;
- 2 x NVIDIA K80 (four devices) – Up to $2.91 \div 8.74$ TFLOPS DP/SP.

WEAK SCALING WITH GPUS (BLUE WATERS)

Table: Weak scaling experiment results for ChASE-MPI on BLUEWATERS

# Cores	Iterations	Matvecs	Filter Runtime			Total Runtime		
			CPU	GPU	Speedup	CPU	GPU	Speedup
64	11	20,106	176.4 s	22.8 s	7.7	228.0 s	43.5 s	5.2
128	9	16,856	175.9 s	27.5 s	6.4	236.1 s	52.6 s	4.5
288	12	23,610	231.5 s	30.2 s	7.7	306.8 s	70.3 s	4.4
512	14	23,080	225.5 s	30.1 s	7.5	316.5 s	87.3 s	3.6
800	12	22,868	209.1 s	30.8 s	6.8	299.2 s	89.9 s	3.3

Tests were performed on the BLUE WATERS cluster.

- AMD 6276 Interlagos – Up to 156 GFLOPS DP;
- NVIDIA K20 – Up to 1310 GFLOPS DP.

CONCLUSIONS AND OUTLOOK

Conclusions

- ✓ Modern library based on C++ STL with clear separation b/w algorithm and implementation;
- ✓ ChASE is templates for SP, DP, Real and Complex;
- ✓ ChASE has a pure MPI and a distributed MPI+X implemetation for inner kernels;
- ✓ Eliminated dependence on Chebyshev filter's parameters;
- ✓ Minimized filter's FLOPs count.

Outlook

- Refine ChASE node-level parallelism → multiple GPUs;
- Add support for sparse matrices;
- Modify filter to extend to generalized eigenproblems;
- Eliminate dependence on nex.

THANK YOU



- <https://github.com/SimLabQuantumMaterials/ChASE>
- <https://arxiv.org/abs/1805.10121>



- e.di.napoli@fz-juelich.de
- <http://www.fz-juelich.de/ias/jsc/slqm>