The Use of a Preconditioner in Iterative Methods for Solving Large-scale Eigenvalue Problems

Chao Yang

Computational Research Division

Lawerence Berkeley National Laboratory
Berkeley, California, USA

Limitation of a Krylov Subspace

- May require a high degree polynomial $\phi(\lambda)$ to produce an accurate approximation $z = \phi(A)v_0$;
  - Subspace of large dimension
  - Many restarts
- Spectral transformation may be prohibitively costly (sometime impossible)
- Not easy to introduce a preconditioner

Preconditioner for $Ax = b$

- Solve $K^{-1}Ax = K^{-1}b$;
- Choose K such that
  - $\kappa(K^{-1}A) \ll \kappa(A)$;
  - Eigenvalue of $\kappa(K^{-1}A) \ll \kappa(A)$ are clustered;
  - K is easy to construct, and solving $Ky = z$ is more efficient than solving $Ax = b$.

Preconditioner for $Ax = \lambda x$?

- Eigenvectors are not preserved under $K^{-1}A$;
- Cannot extract correct spectral info from $\mathcal{K}(K^{-1}A, b; m)$ directly;
- However, preconditioning does make sense if we treat an eigenvalue problem as
  - a system of nonlinear equations (JD)
  - an optimization problem (LOBPCG)

Nonlinear Equation Point of View

- Because $\lambda(x) = x^T Ax/x^T x$,

$$Ax = \lambda(x)x,$$

  is a nonlinear equation in x;
- Alternative formulation

$$\begin{aligned} Ax &= (x^T Ax)x \\ x^T x &= 1 \end{aligned}$$

- Many solutions;

Solve by Newton's Correction

- Given a starting guess u such that $u^T u = 1$;
- Let $\theta = u^T A u$;
- Seek (z, δ) pair such that

$$A(u + z) = (\theta + \delta)(u + z)$$

- Ignore the 2nd order term δz (Newton correction) and impose

$$u^T z = 0$$

The Correction Equation

- Augmented form

$$\begin{pmatrix} A - \theta I & u \\ u^T & 0 \end{pmatrix} \begin{pmatrix} z \\ -\delta \end{pmatrix} = \begin{pmatrix} -r \\ 0 \end{pmatrix},$$

  where $r = Au - \theta u$;

- Projected form

$$(I - uu^T)(A - \theta I)(I - uu^T)z = -r$$

  where $u^T z = 0$

Solving the Correction Eq (Direct)

- Assume θ not converged yet, block elimination yields

$$\begin{pmatrix} I & 0 \\ u^T(A-\theta I)^{-1} & 1 \end{pmatrix} \begin{pmatrix} A-\theta I & u \\ 0 & \gamma \end{pmatrix} \begin{pmatrix} z \\ -\delta \end{pmatrix} = \begin{pmatrix} -r \\ 0 \end{pmatrix},$$

where $\gamma = u^T(A-\theta)^{-1}u$

- $$\delta = \frac{u^T(A-\theta I)^{-1}r}{u^T(A-\theta I)^{-1}u}.$$

- Back substitution yields

$$z = \delta(A-\theta I)^{-1}u - u$$

,

Connection with Inverse Iteration

- Adding correction z to u directly

$$x = u + z = u + \delta(A - \theta I)^{-1}u - u = \delta(A - \theta I)^{-1}u$$

- Quadratic convergence in general
- Cubic convergence for symmetric problems
- But requires solving

$$(A - \theta I)x = u$$

accurately

Jacobi-Davidson (JD)

- Solving the correction equation iteratively

$$(I - uu^T)(A - \theta I)(I - uu^T)z = -r$$

  where $u^T z = 0$
- Allows the use of a preconditioner;
- Instead of adding z to u, construct a search space $\mathcal{S} = \{u, z\}$;
- Extract Ritz pairs from $\mathcal{S}$ through Rayleigh-Ritz

JD Inner-outer Iteration

Input: A, v_0, tol;

Output: (u, θ) such that $\|Au - \theta u\|$ is small

1. $u \leftarrow v_0/\|v_0\|$, $V \leftarrow (u)$, $\theta = u^T Au$, $r \leftarrow Au - \theta u$;
2. while ($\|r\| > tol$)
   (a) Iteratively solve the correction equation $(I - VV^T)(A - \theta I)(I - VV^T)z = -r$ approximately;
   (b) $z \leftarrow (I - VV^T)z$;
   (c) $V \leftarrow (V, z)$, $H = V^T AV$;
   (d) Solve $Hy = \theta y$ and select the desired (y, θ);
   (e) $u \leftarrow Vy$, $r \leftarrow Au - \theta u$;

Practical Issues

- Choose an iterative solver and a preconditioner (for the correction equation);
- Set tolerance for the inner iteration;
- Shift selection;
- Restart (set a limit on the dimension of V);
- Compute more than one eigenpair (JDQR, JDQZ);

Preconditioner for the Correction Ec

- If $K \approx A$, then $\hat{K} \approx \hat{A}$, where
  $\hat{K} = (I - VV^T)K(I - VV^T)$, $\hat{A} = (I - VV^T)A(I - VV^T)$;
- Precondition: must solve $\hat{K}x = b$;
- Use augmented formulation and block elimination
  $\hat{K}^{\dagger} = (I - YG^{-1}V^T)K^{-1} = K^{-1}(I - VG^{-1}U^T)$,
  where $Y = K^{-1}V$, $G = V^TY$, $U = K^{-T}Y$;
- MATVEC:

$$y \leftarrow (I - VG^{-1}Y^T)K^{-1}A(I - YG^{-1}V^T)x$$

Termination Criterion - Inner Iteratio

- Fixed tolerance (larger than that required for the eigen-residual)
- Dynamic tolerance (tighter when eigen-residual small)
  - Can estimate eigen-residual from the correction equation residual for certain solvers (CG, sQMR)
  - See Notay (2002), Stathopoulos (2005)

Example

- normal mode vibration analysis for macromolecules
- 3000-atom, $n = 9000$, interested in low frequency modes (small eigenvalues)

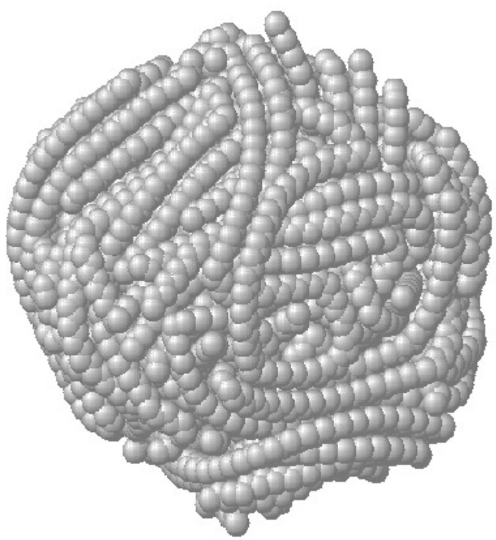

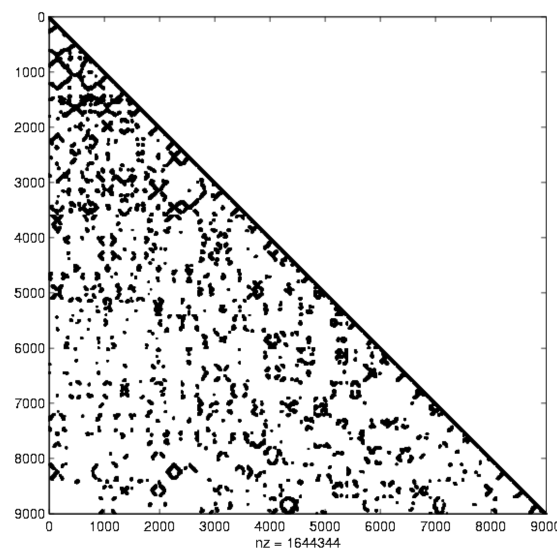

Effect of Preconditioner

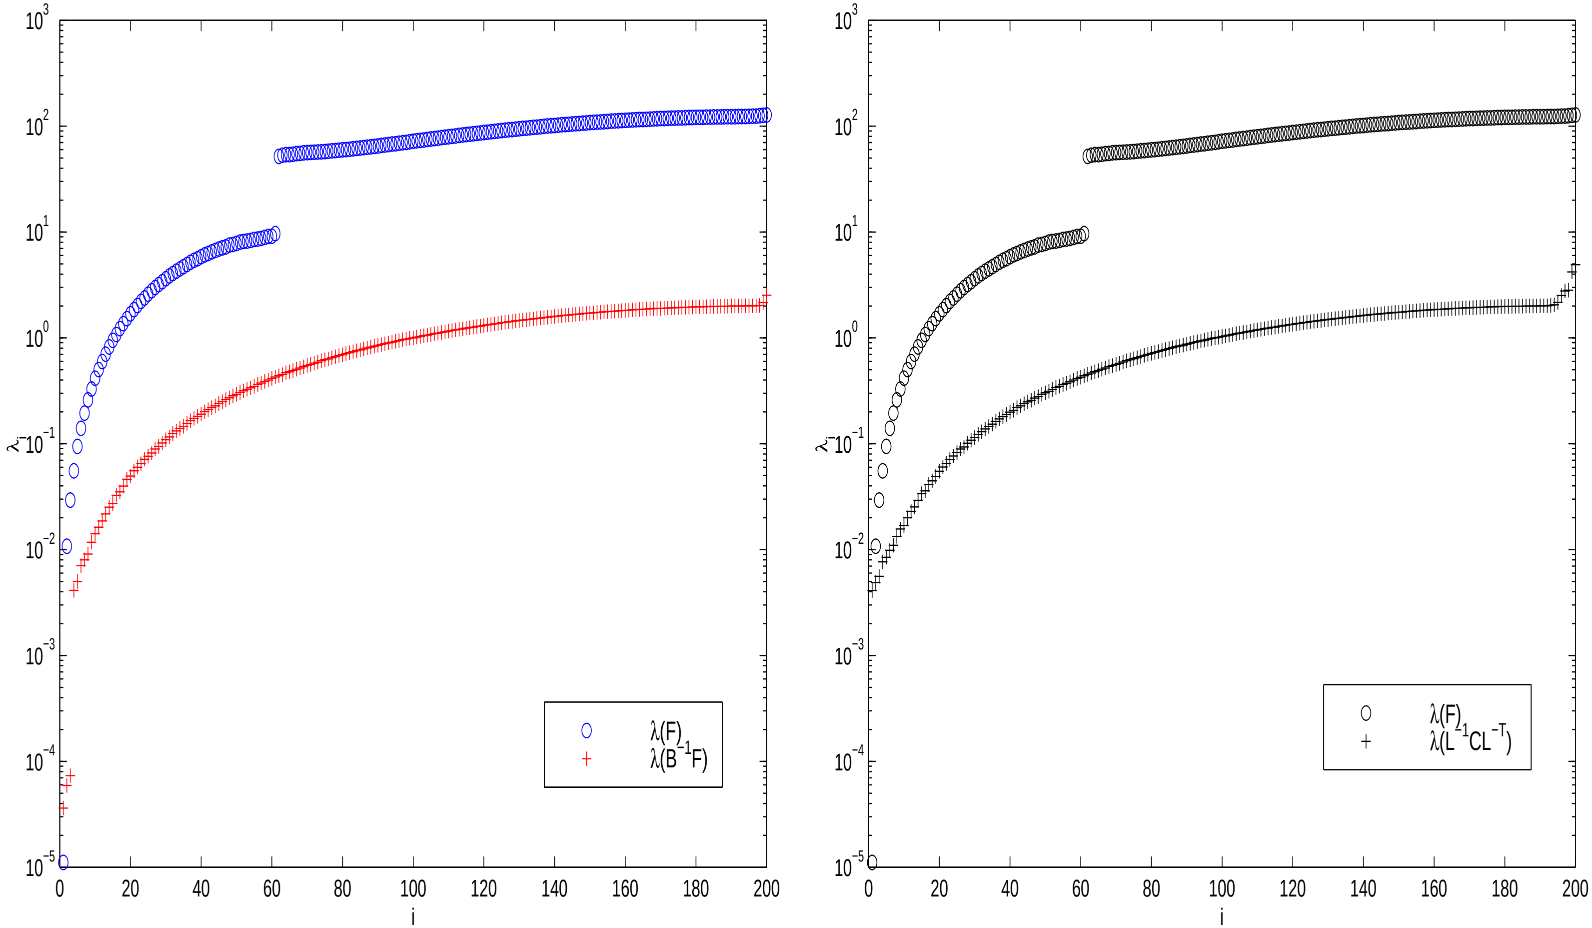

Convergence History

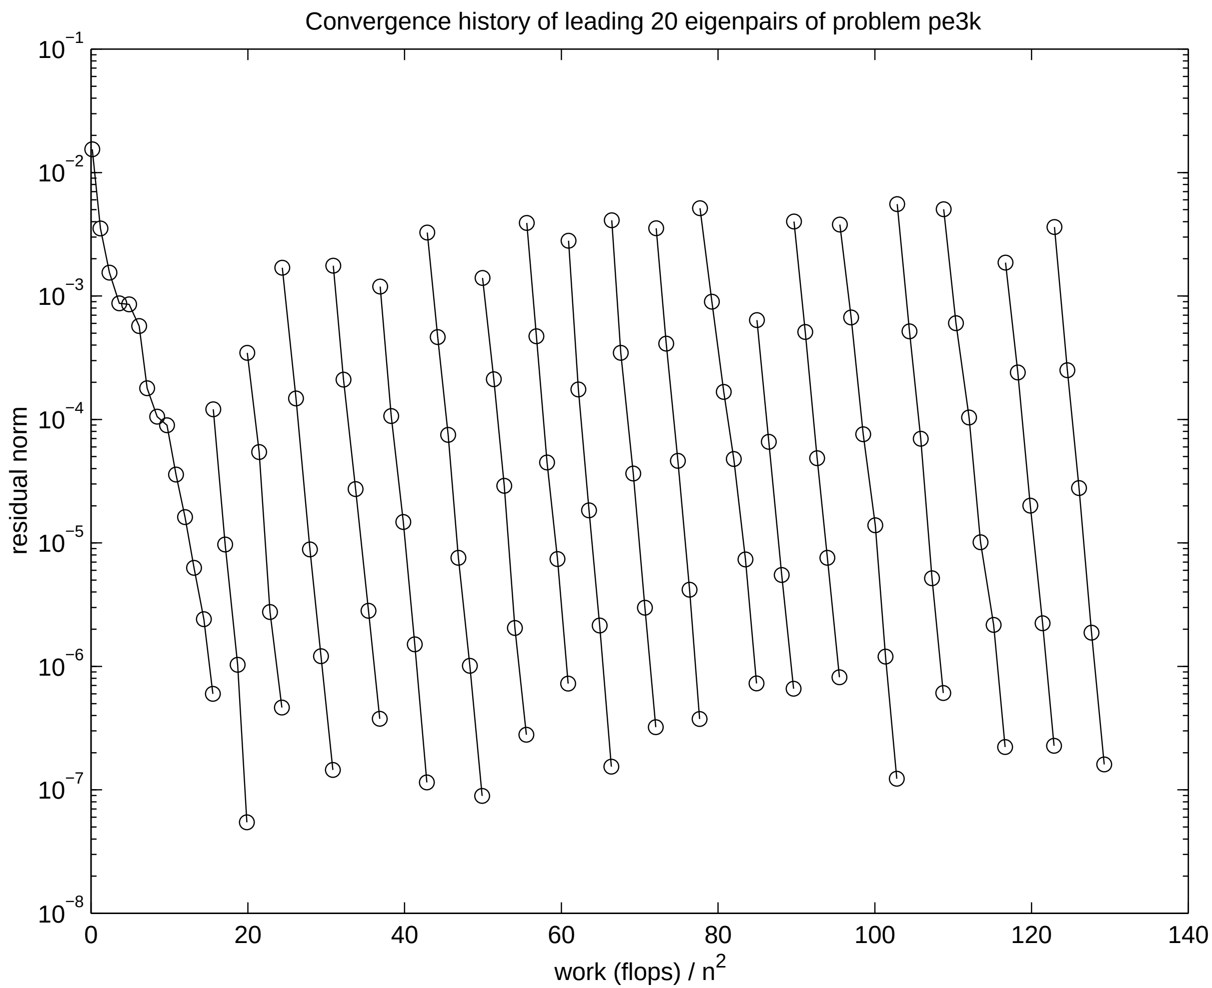

C. Yang et al. SIAM J. Sci. Comp. 2001

Other Issues

- Block version (not trivial)
- Can extend JD to polynomial or rational eigenvalue problems (Van der Vorst, Voss, Lin)
- Automatic parameter tuning
- Missing eigenvalues?

The Optimization View

- Only valid for symmetric problems, extreme eigenvalues
- Constrained optimization

$$\min_{x^T x = 1} x^T A x$$

- Lagrangian $\mathcal{L}(x, \lambda) = x^T A x - \lambda(x^T x - 1)$
- KKT condition

$$\begin{aligned} Ax - \lambda x &= 0 \\ x^T x &= 1 \end{aligned}$$

Geometry

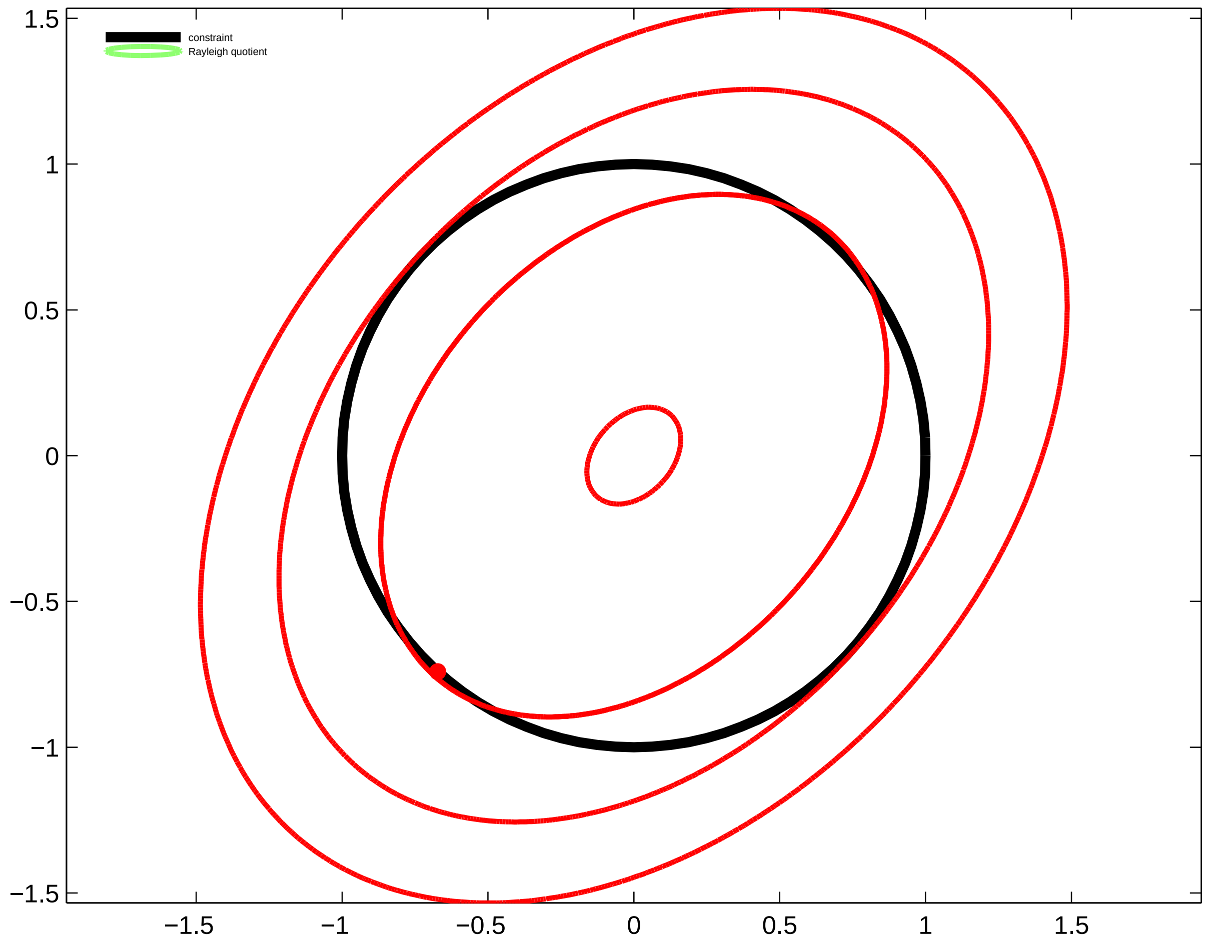

Direct Constrained Minimization

- Assume x_k is current approximation;
- Update by

$$x_{k+1} = \alpha x_k + \beta p_k$$

  - p_k is a descent (search) direction;
  - α, β are chosen so that
    - $x_{k+1}^T x_{k+1} = 1$;
    - $\rho(x_{k+1}) < \rho(x_k)$, where $\rho(x) = x^T A x$;

Search Direction

- Steepest descent

$$r_k = -\nabla_x \mathcal{L}(x_k, \theta_k) = -(Ax_k - \theta_k x_k)$$

- Conjugate gradient

$$p_k = p_{k-1} + \gamma r_k$$

Choose γ so that $p_k^T A p_{k-1} = 0$

- Maintaining the orthonormality constraint?

$$x_{k+} = \alpha x_k + \beta p_{k-1} + \gamma r_k$$

,

Subspace Minimization

- Let $V = (x_k, p_{k-1}, r_k)$, then $x_{k+1} = Vy_k$, for some $y_k \in \mathbb{R}^3$;
- Must solve

$$\min_{y_k^T V^T V y_k = 1} y_k^T V^T A V y_k$$

- Equivalent to solving

$$\begin{aligned} G y_k &= \lambda B y_k \\ y_k^T B y_k &= 1 \end{aligned}$$

where $B = V^T V$ and $G = V^T A V$;

Compute More Eigenpairs

- Trace minimization

$$\min_{X^T X = I_m} \frac{1}{2}\text{trace}(X^T A X)$$

where $X \in \mathbb{R}^{n\times m}$;

- Gradient

$$R_k = \nabla_X \mathcal{L}(X_k, \Lambda_k) = AX_k - X_k\Lambda_k,$$

where $\Lambda_k = X_k^T A X_k$;

LOBPCG (Knyazev)

Input: A, K, $X_0 \in \mathbb{R}^{n\times m}$, tol;

Output: (X, Λ) such that $\|AX - X\Lambda\|$ is small

1. Orthonormalize columns of X_0, $\Lambda = X_0^* A X_0$, $i = 1$, $P_0 = [\,]$,$R_1 = AX_0 - X_0\Lambda$,;
2. while ($\|R_i\| > tol$)
   (a) Set $V = (X_{i-1}, P_{i-1}, K^{-1}R_i)$;
   (b) Compute the eigenvectors G corresponding to the m smallest eigenvalues of (H, B), where $B = V^T V$ and $H = V^T A V$;
   (c) $X_i = VG(1:m,:)$; $\Lambda_i = X_i^T A X_i$, $P_i = VG(m+1:3m,:)$;
   (d) $R_{i+1} = AX_i - X_i\Lambda$, $i \leftarrow i+1$;

Practical Issues

- Choice of preconditioner
- Linear dependency between columns of V;
- Deflation (not all eigenpairs converge at the same rate)
- Extension to (symmetric) generalized eigenvalue problem (straightforward)

Example

- Small accelerator model ($n = 750$)
- Generalized problem
- Interested in the smallest eigenvalue

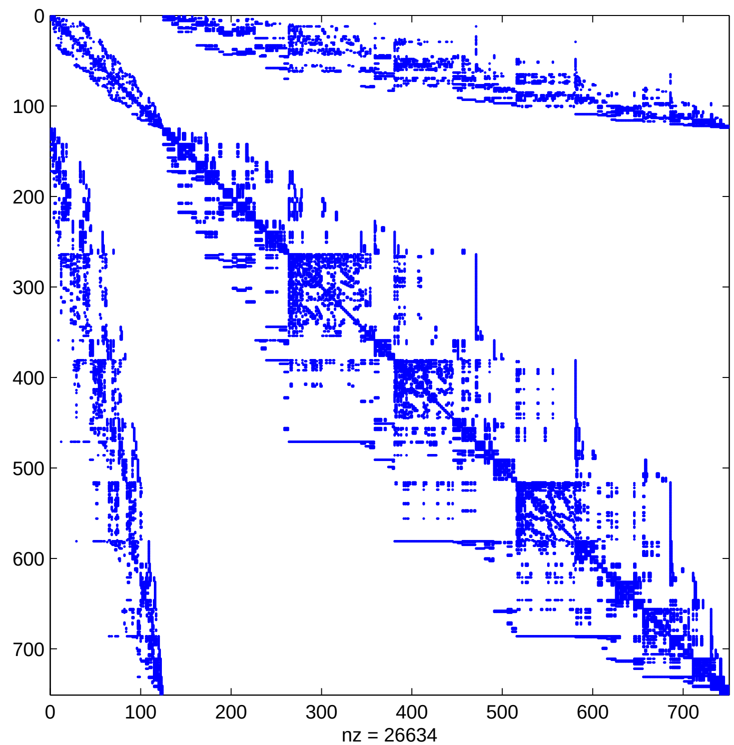

Change in residual

Diagonal Preconditioner

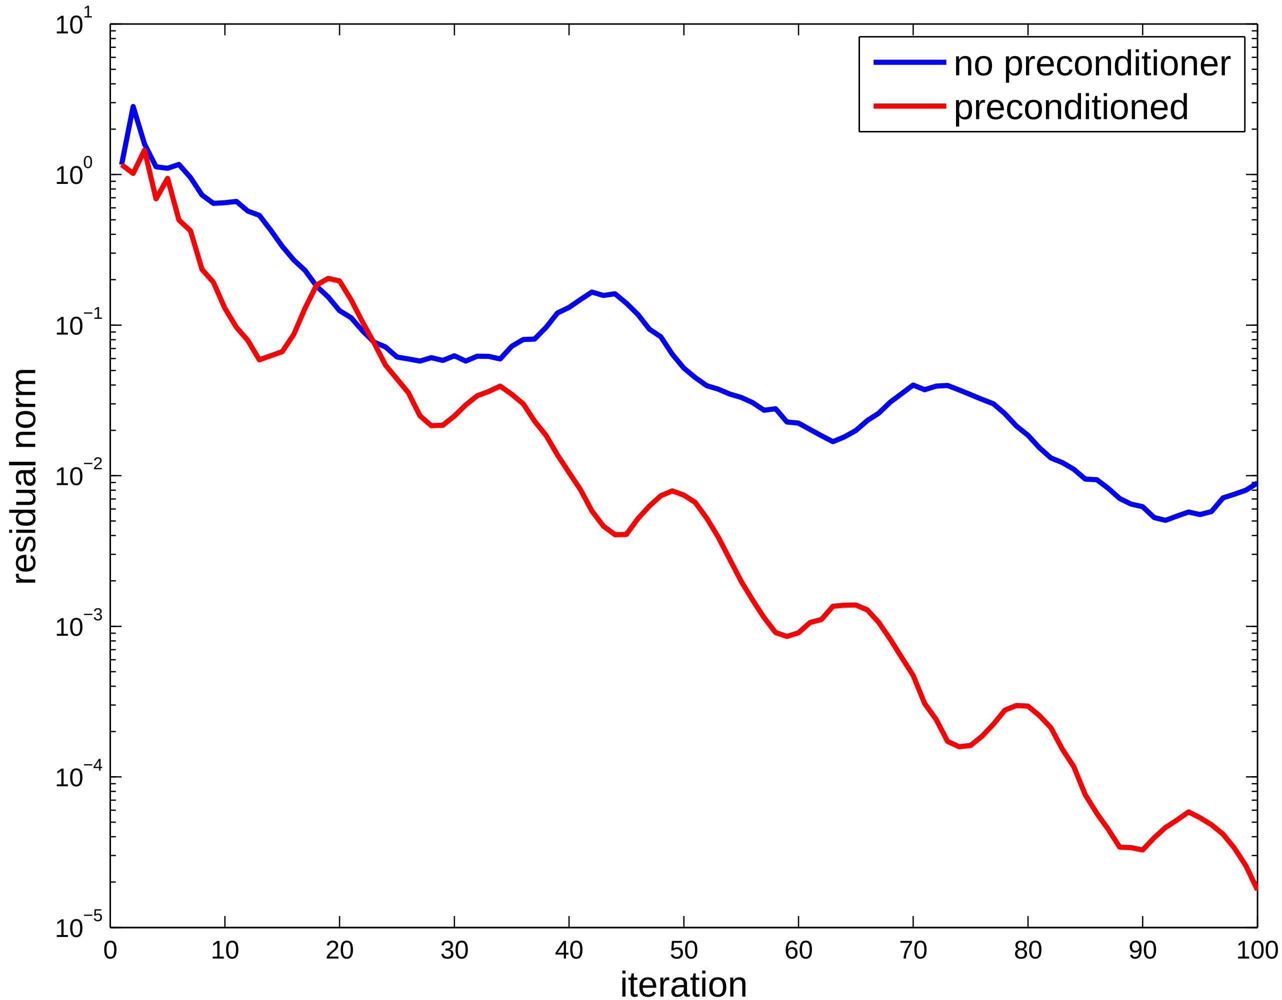

Extension to a Nonlinear EV Proble[m]

- Quantum many-body problem reduced to single-particle wavefunctions through DFT;
- single particle wavefunctions (orbitals)

  $X = (x_1, x_2, ..., x_k), \quad X^* X = I_k, \quad x_i \in \mathbb{C}^n$
  - n - real space grid size;
  - k - number of occupied states;
- Charge density $\rho(X) = \text{diag}(XX^*)$;
- Kohn-Sham total energy

$$E_{tot}(X) = E_{kinetic}(X) + E_{ion}(X) + E_{Hartree}(X) + E_{xc}(X),$$

KS Total Energy Minimization

$$\min_{X^*X=I_k} E_{tot}(X) \equiv E_{kinetic}(X) + E_{ion}(X) + E_{Hartree}(X) + E_{xc}(X),$$

where

$$\begin{aligned}
E_{kinetic} &= \frac{1}{2}\text{trace}(X^*LX) \\
E_{ionic} &= \frac{1}{2}\left[\text{trace}(XD_{ion}X^*) + \sum_i \sum_\ell (x^* w_\ell)^2\right] \\
E_{Hartree} &= \frac{1}{4}\rho(X)^T S \rho(X) \\
E_{xc} &= e^T(f_{xc}[\rho(X)])
\end{aligned}$$

First order (KKT) Condition

- KKT condition

$$\begin{aligned} \nabla_X \mathcal{L}(X,\Lambda) &= 0; \\ X^*X &= I. \end{aligned}$$

- Kohn-Sham equation

$$\begin{aligned} H(X)X &= X\Lambda, \\ X^*X &= I. \end{aligned}$$

- Kohn-Sham Hamiltonian

$$H(X) = L + D_{ion} + \sum_{\ell} w_\ell w_\ell^* + \text{diag}(S\rho(X)) + \text{diag} g_{xc}(\rho(X))$$

;

Two Approaches

- Work with the KS equation
  - Self-Consistent Field Iteration
- Minimize the total energy directly
  - Direct Constrained Minimization (extension of LOBPCG)

The SCF Iteration

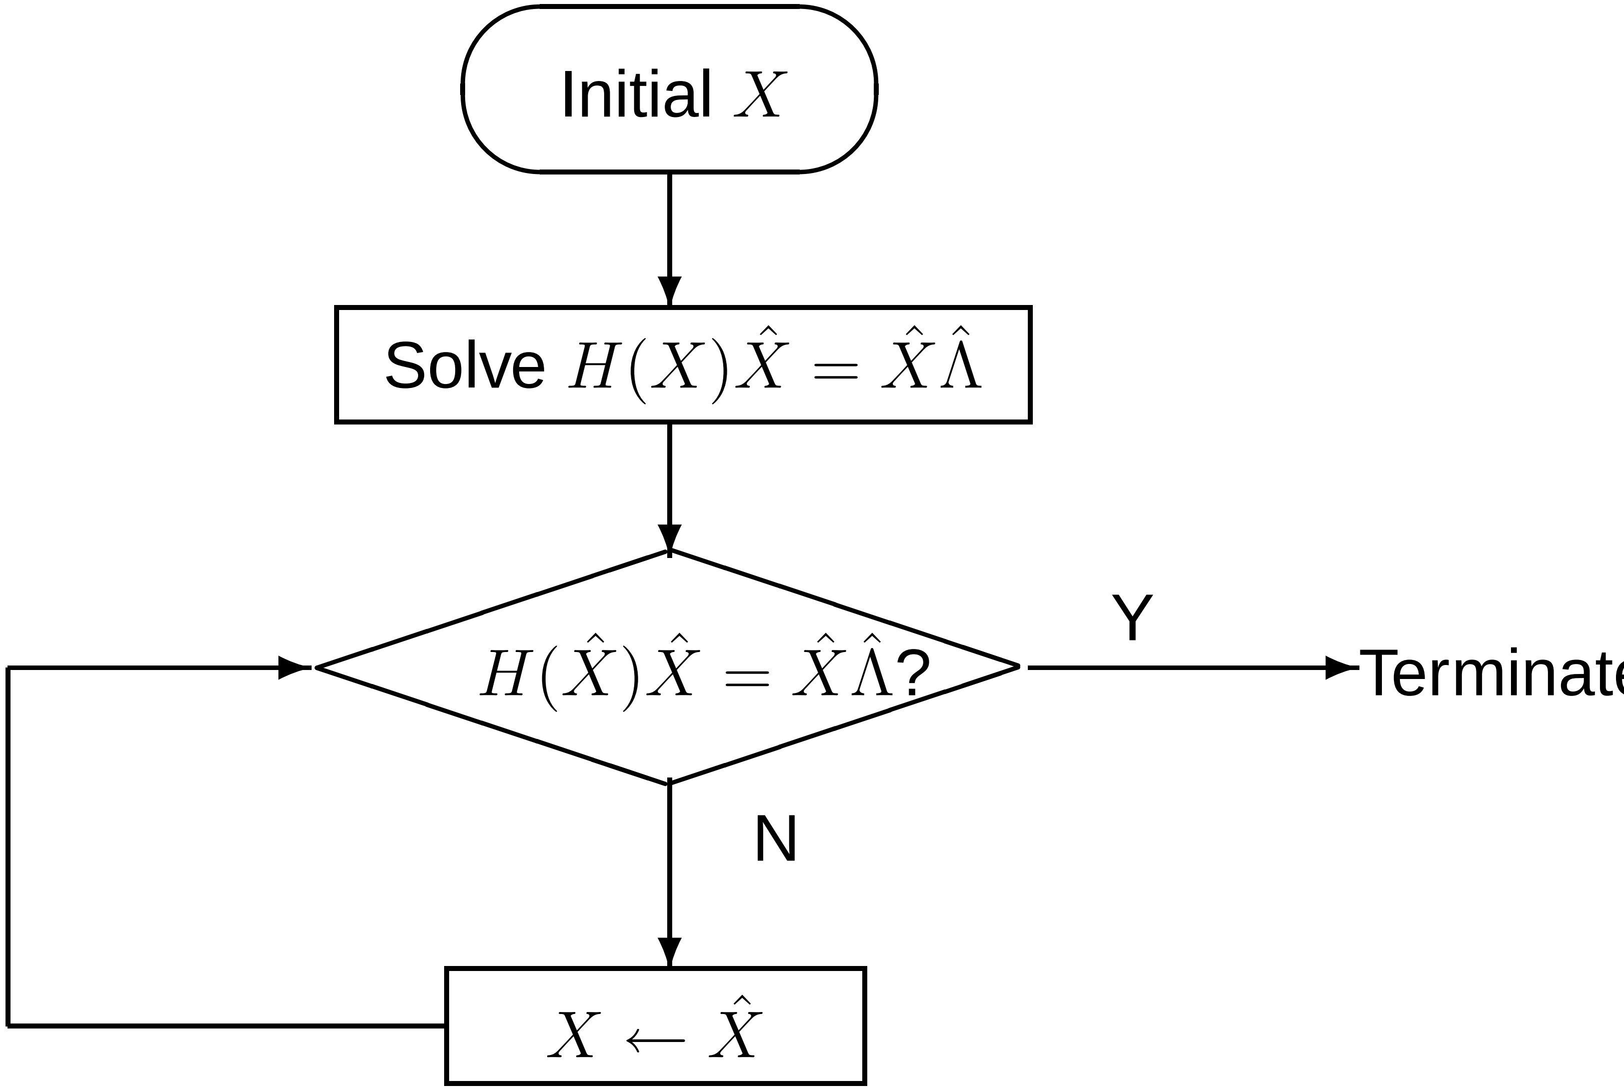

Problems with SCF

- Does not always converge, $E_{tot}(X)$ may increase;
- Little convergence theory
  - Fixed-point iteration $X^{(i+)} = F(X^{(i)})$. But what is $F(\cdot)$?
- Solve large-scale linear eigenvalue at each iteration?
  - LAPACK routines (expensive)
  - Iterative solvers (PCG, stopping criterion, accuracy)

Direct Constrained Minimization

- Minimize the total energy directly;
- Block method
- Wavefunction update similar to LOBPCG

$$X^{(i+1)} = X^{(i)}G_x + P^{(i-1)}G_p + R^{(i)}G_r,$$

$$R^{(i)} = K^{-1}(H(X^{(i)})X^{(i)} - X^{(i)}\Theta^{(i)}),$$
$$\Theta^{(i)} = X^{(i)^*}H(X^{(i)})X^{(i)}.$$

- Choose G_x, G_p, G_r to
  - Minimize E_{tot} in $Y = (X^{(i)}, P^{(i-1)}, R^{(i)})$;
  - Maintain $X^{(i+1)^*}X^{(i+1)} = I$.

Minimization within a Subspace

- Let $Y = (X^{(i)}, P^{(i-1)}, R^{(i)})$;
- Solve

$$min_G E_{tot}(YG)$$
$$\text{s.t.} G^* Y^* Y G = I_k$$

- Equivalent to solving

$$\begin{aligned} \hat{H}(G)G &= BG\Omega_k \\ G^* BG &= I_k, \end{aligned}$$

$$\hat{H}(G) = Y^* H(YG)Y, \;\; B = Y^* Y,$$

$$H(YG) =$$
$$L + D_{ion} + \sum_\ell w_\ell w_\ell^T + \text{Diag}(S\rho(YG)) + \text{Diag}(g_{xc}(\rho(YG)))$$

Solving the Projected Problem

- A smaller nonlinear (generalized) eigenvalue problem

$$\begin{aligned} \hat{H}(G)G &= BG\Omega_k \\ G^*BG &= I_k, \end{aligned}$$

where $H(G) \in \mathbb{C}^{3k\times 3k}$, $G \in \mathbb{C}^{3k\times k}$.

- Modify SCF by introducting a trust region (TRSCF)

The Optimization View of SCF

- SCF minimizes a sequence of surrogate models
- Objective:
  - $E_{tot}(X) = E_{kinetic}(X) + E_{ion}(X) + E_{Hartree}(X) + E_{xc}(X)$
  - $E_{surrogate}(X) = \frac{1}{2}\text{trace}(X^* H(X^{(i)}) X)$
- Gradient:
  - $\nabla E_{tot}(X) = H(X)X$
  - $\nabla E_{surrogate}(X) = H(X^{(i)})X$

$$\boxed{\nabla E_{tot}(X^{(i)}) = \nabla E_{surrogate}(X^{(i)})}$$

Toy Example

$$E_{tot}(x) = \frac{1}{2}x^T L x + \frac{\alpha}{4}\rho(x)^T L^{-1}\rho(x)$$

$$L = \begin{pmatrix} 2 & -1 \\ -1 & 2 \end{pmatrix}, x = \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}, \rho(x) = \begin{pmatrix} x_1^2 \\ x_2^2 \end{pmatrix}$$

$$\boxed{\begin{array}{c} \min E_{tot}(x) \\ \text{s.t.}\, x_1^2 + x_2^2 = 1 \end{array}}$$

- When does SCF work?
- How can it fail?

SCF Works ($\alpha = 2.0$)

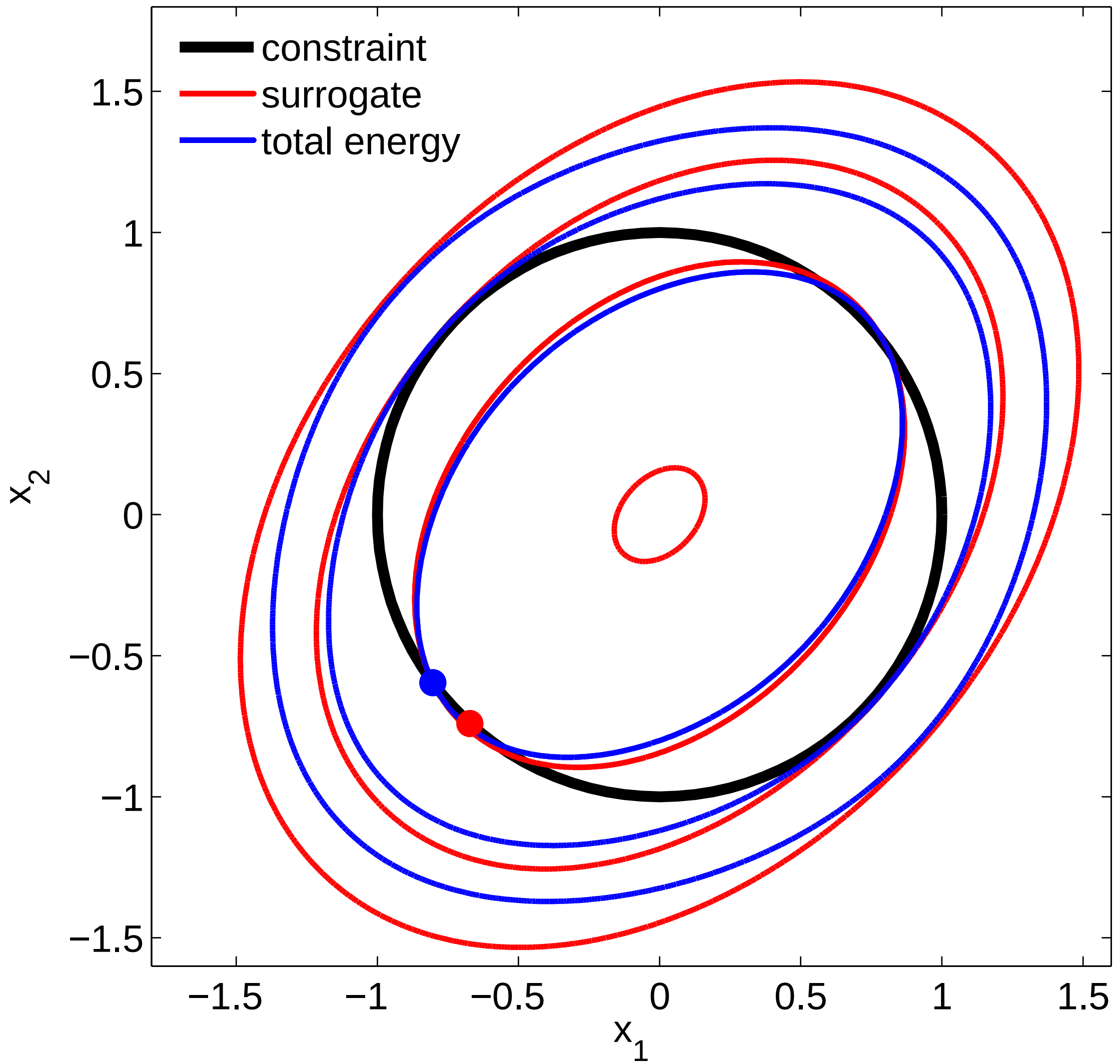

SCF Works

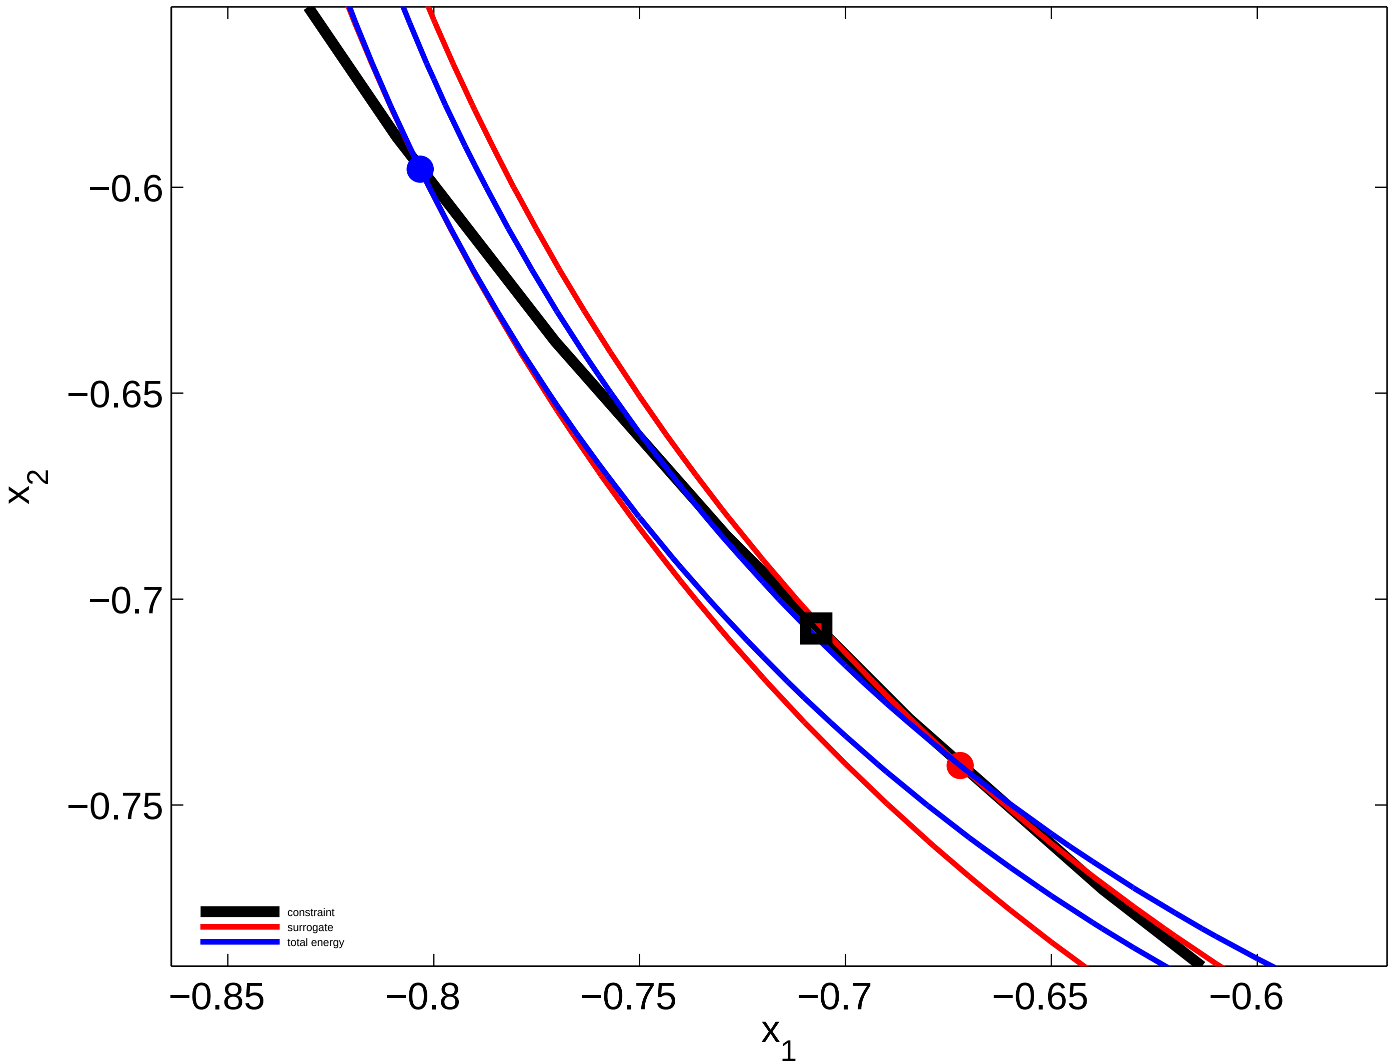

−0.6
−0.65
−0.7
−0.75
x_2
−0.85
−0.8
−0.75
−0.7
−0.65
−0.6
x_1
constraint
surrogate
total energy

SCF Fails ($\alpha = 12.0$)

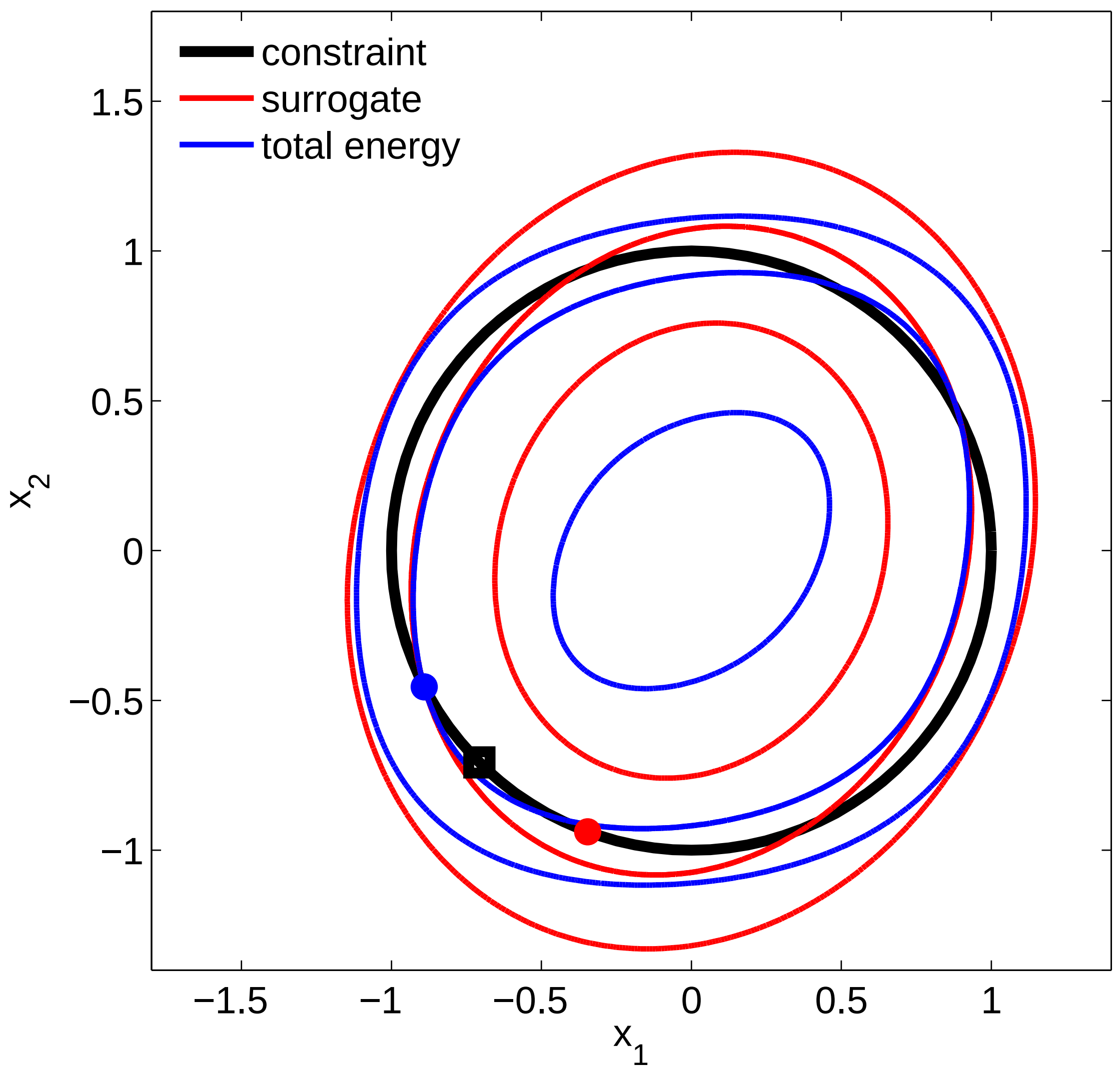

SCF Fails

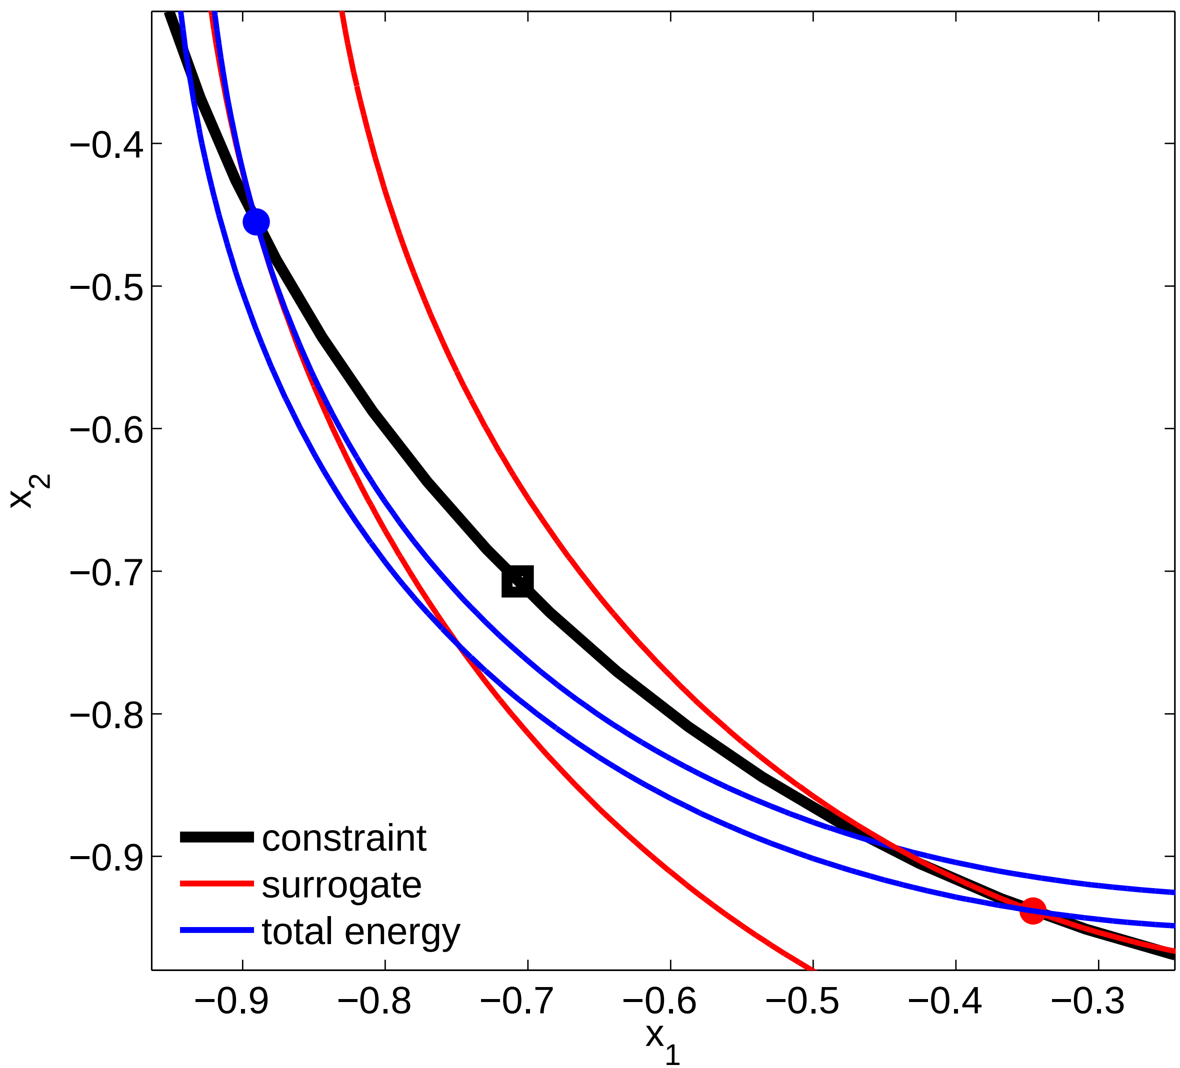

−0.4
−0.5
−0.6
−0.7
−0.8
−0.9
x_2
−0.9
−0.8
−0.7
−0.6
−0.5
−0.4
−0.3
x_1
constraint
surrogate
total energy

Improving SCF

- Construct better surrogate
  - Cannot afford to use local quadratic approx (Hessian too expensive)
  - Charge-mixing
- Use Trust Region to restrict the update of the wavefunction in a small neighborhood of the gradient matching point
  - TRSCF (Thogersen, Olsen, Yeager & Jorgensen 2004)

Trust Region

- Must be defined on quantities that are "rotationally invariant"
  - Density matrix $D(X) = D(XQ) = XX^*$;
  - Charge density $\rho(X) = \rho(XQ) = \text{diag}(D(X))$;

$$\boxed{\text{Use } \|D(X) - D(X^{(0)})\|_F \leq \Delta}$$

- The trust region subproblem must be easy to solve

Trust Region Subproblem

- Instead of solving

$$\min_{X^*X=I} q(X)$$
$$s.t. \|D(X) - D(X^{(0)})\|_F \leq \Delta$$

- We solve the penalized problem

$$\min_{X^*X-I} q(X) + \frac{\sigma}{2}\|D(X) - D(X^{(0)}\|_F^2$$

or equivalently

$$\min_{X^*X-I} q(X) - \frac{\sigma}{2}\text{trace}\left[X^* X^{(0)} X^{(0)^*} X\right]$$

TRSCF

- In each TRSCF iteration, we solve

$$\min_{X^*X=I} \frac{1}{2}\text{trace}\left[X^* H(X^{(0)}X - \sigma X^* X^{(0)} X^{(0)^*} X\right]$$

- First order (KKT) condition

$$\begin{aligned}\left[H(X^{(0)}) - \sigma X^{(0)} X^{(0)^*}\right] X &= X\Lambda \\ X^* X &= I\end{aligned}$$

- At convergence

$$\lambda_1 - \sigma, \lambda_2 - \sigma, ..., \lambda_k - \sigma, \lambda_{k+1}, ..., \lambda_n$$

- How to pick σ?

The Effect of Trust Region

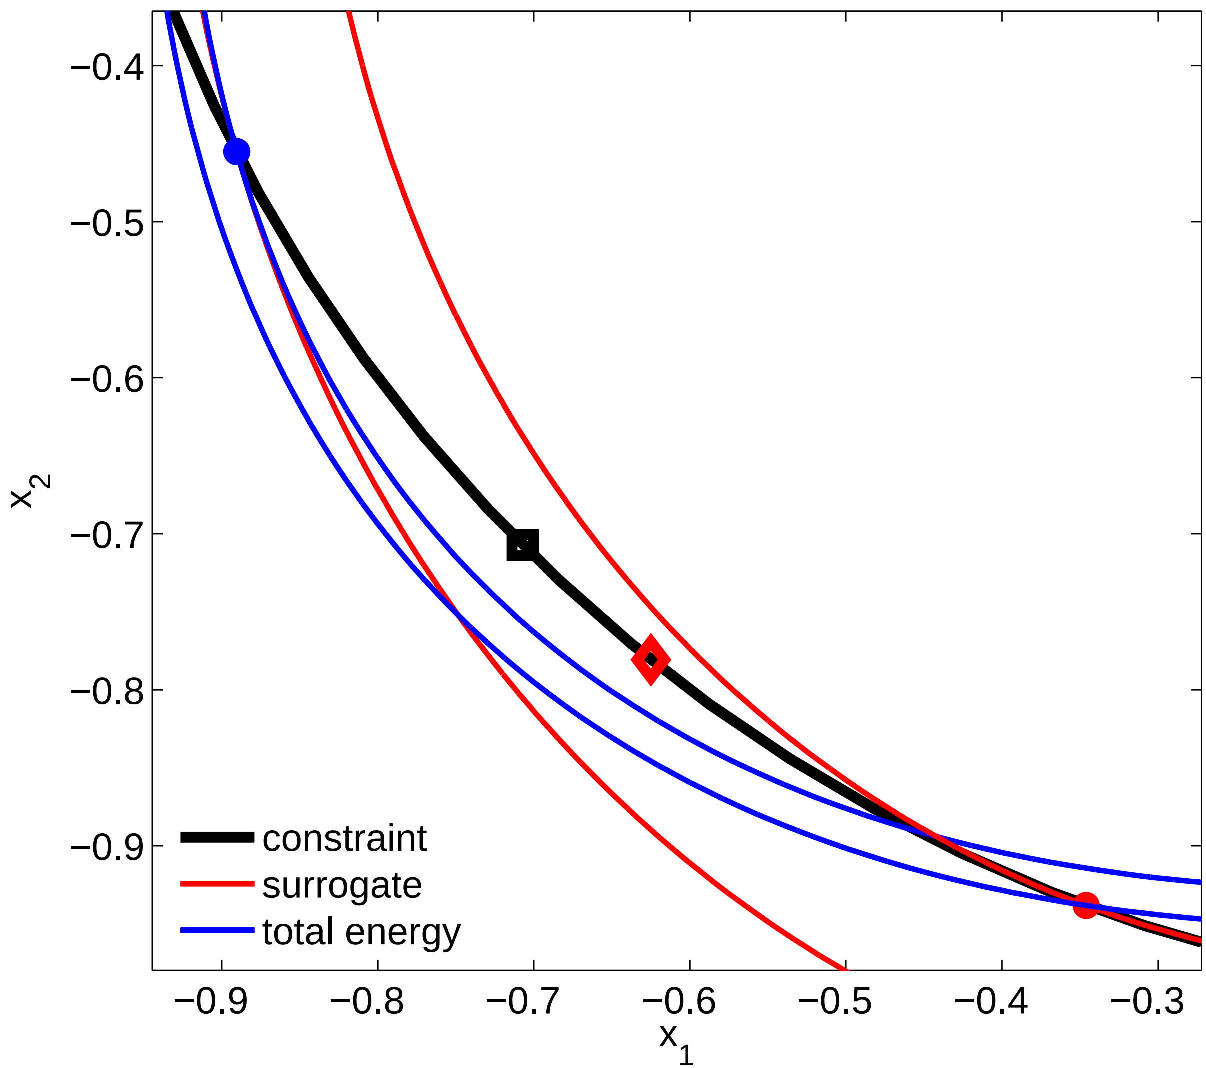

DCM

Input: L,D_{ion},w_ℓ, $X_0 \in \mathbb{R}^{n\times m}$;

Output: X such that $E_{tot}(X)$ is minimized

1. Orthonormalize columns of X_0, $\Theta = X_0^* H(X^{(0)}) X_0$, $i = 1$, $P_0 = []$;
2. while (not converged)
   (a) $R_i = H(X_i)X_i - X_i\Lambda$,
   (b) Set $Y = (X_{i-1}, P_{i-1}, L^{-1}R_i)$;
   (c) Solve $\min_{G^*Y^*YG=I_k} E_{tot}(YG)$;
   (d) $X_i = YG(1:m,:)$; $\Lambda_i = X_i^T A X_i$, $P_i = YG(m+1:3m,:)$,$i \leftarrow i+1$;

Numerical Example

- Atomistic sytem: NiPtO
- Discretization: spectral method with plane wave basis $n = 96 \times 48 \times 48$ in real space, $N = 15179$ (number of basis functions) in frequency space
- Number of occupied states $k = 43$
- PETOT version of SCF does 10 PCG step per outer iteration
- DCM 5 inner iteration
- Compare the change of total energy with respect to timing
- Parallelized using MPI on 64 IBM Power3 CPUs

Convergence

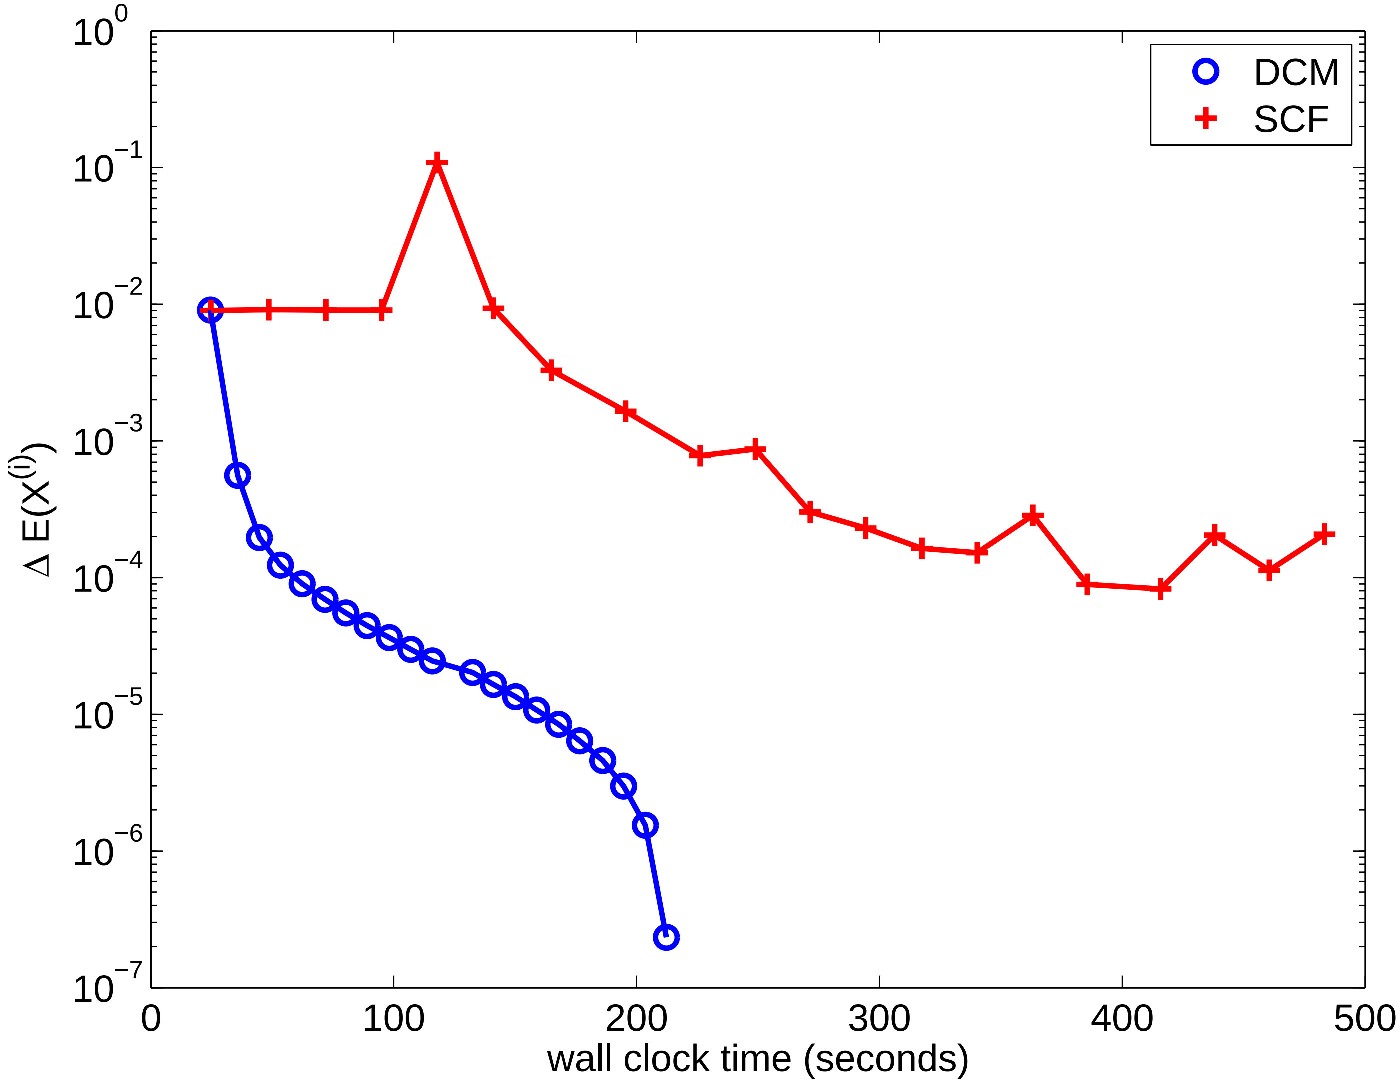

Conclusion

- The use of a preconditioner is natural if we treat an eigenvalue problem (both linear and nonlinear) as
  - a nonlinear system of equations
  - a constrained nonlinear optimization problem (for certain class of problems)
- The choice of preconditioner is application dependant
- A preconditioner must be deflated in Jacobi-Davidson