

Adaptive model trust region methods for generalized eigenvalue problems[§]

P.-A. Absil^{*†} C. G. Baker^{*} K. A. Gallivan^{*†} A. Sameh[‡]

Technical Report FSU-CSIT-04-28
December 9, 2004

Abstract

Computing a few eigenpairs of large-scale matrices is a significant problem in science and engineering applications and a very active area of research. In this paper, two methods that compute extreme eigenpairs of positive-definite matrix pencils are combined into a hybrid scheme that inherits the advantages of both constituents. The hybrid algorithm is developed and analyzed in the framework of model-based methods for trace minimization.

1 Introduction

We consider the computation of a few smallest eigenvalues and the corresponding eigenvectors of the generalized eigenvalue problem

$$Ax = \lambda Bx, \tag{1}$$

where A and B are $n \times n$ symmetric positive-definite matrices. Positive definiteness of A and B guarantees that all the eigenvalues are real and positive. This eigenvalue problem appears in particular in the computation of the lower modes of vibration of a mechanical structure, assuming that there are no rigid-body modes and that all the degrees of freedom are mass-supplied: A is the stiffness matrix, B is the mass matrix, x is a mode of vibration and λ is the square of the circular frequency associated with the mode x .

Inverse iteration (INVIT) and Rayleigh quotient iteration (RQI) are the conceptually simplest methods for the generalized eigenvalue problem [Par80, §15-9]. Interestingly, they have complementary properties: unshifted INVIT converges to the smallest (or *leftmost*) eigenpair for almost all initial conditions but with linear local convergence only; RQI has cubic local convergence but it can converge to different eigenpairs depending on the initial condition.

It is therefore quite natural to try to combine these two methods and obtain a hybrid method that enjoys strong global convergence and fast local convergence; see in particular Szyld [Szy88] for the problem of finding an eigenvalue in a given interval. For the computation of the leftmost eigenpair, such a hybrid method would use INVIT until reaching the basin of attraction of the

^{*}School of Computational Science, Florida State University, Tallahassee, FL 32306-4120, USA (<http://www.csit.fsu.edu/{~absil,~cbaker,~gallivan}>).

[†]These authors' work was supported by the National Science Foundation under Grant ACI0324944 and by the School of Computational Science of Florida State University through a postdoctoral fellowship.

[‡]Department of Computer Sciences, Purdue University, West Lafayette, IN 47907-2066, USA.

[§]This is an expanded version of a conference paper submitted to ICCS 2005 on 08 Dec 2004. Please check the first author's web page for the latest electronic version of this paper.

leftmost eigenvector under RQI; then the method would switch to RQI in order to exploit its superlinear convergence. However, to our knowledge, a practical and reliable switching criterion between the two methods has yet to be found that guarantees global convergence of the hybrid method: if the switch is made too early, RQI may not converge to the leftmost eigenspace. A second drawback of *exact* INVIT and RQI is their possibly high computational cost, since large-scale linear systems with system matrix $(A - \sigma B)$ have to be solved exactly.

In this paper, we propose a remedy to these two difficulties. In Phase I, in order to reduce the computational cost, we replace INVIT with the Basic Tracemin algorithm of Sameh *et al.* [SW82, ST00]. Basic Tracemin proceeds by successive unconstrained approximate minimization of inexact quadratic local models of a generalized Rayleigh quotient cost function, using the stiffness matrix A as the model Hessian. This method is closely related to INVIT: if the model minimizations are carried out exactly (which happens in particular when an exact factorization of A is used for preconditioning), then Basic Tracemin is mathematically equivalent to (block) INVIT.

Extending an observation made by Edelman *et al.* [EAS98, §4.4], we point out that the stiffness matrix used as the model Hessian is quite different from the true Hessian of the generalized Rayleigh quotient—this is why the model is called “inexact”. However, we emphasize that the choice of an inexact Hessian in Phase I does not conflict with the findings in [EAS98]: using the exact Hessian is important only when the iteration gets close to the solution, in order to achieve superlinear convergence. Therefore, using an inexact Hessian in Phase I is not necessarily a liability. On the contrary, the stiffness matrix as the model Hessian offers a useful property: as shown in [SW82, ST00], any decrease in the inexact model induces a decrease in the cost function (i.e., the Rayleigh quotient). Therefore, in the presence of an inexact preconditioner, Basic Tracemin can be thought of as an inexact INVIT, that reduces computational cost per step while preserving the global convergence property. Due to its link with INVIT, it quickly purges the eigenvectors whose eigenvalues are well separated from the p leftmost eigenvalues. This is particularly true when a good preconditioner is available, as is often the case in the very sparse problems encountered in structural mechanics. Consequently, the Basic Tracemin iteration is efficient when the iterates are still far away from the solution.

On the other hand, close to the solution, Basic Tracemin suffers from the linear rate of convergence (due to the use of an inexact Hessian), especially when the p leftmost eigenvalues are not well separated from the immediately higher ones. This is why a superlinear method, like RQI, is required in Phase II. However, we want to remedy both drawbacks of convergence failure and high computational cost mentioned above. This motivates the use of the recently proposed Riemannian trust-region (RTR) algorithm [ABG04c, ABG04b, ABG04a]. Unlike Basic Tracemin, this method uses the *true* Hessian as the model Hessian (for superlinear convergence), along with a trust-region safeguard that prevents convergence to non-leftmost eigenpairs that could otherwise occur if switching takes place too early. Moreover, the computational cost is reduced by using a truncated conjugate-gradient algorithm to solve the trust-region problems inexactly while preserving the global and locally superlinear convergence.

However, in spite of its convergence properties (which make it an excellent choice for Phase II), there is a reason not to use the trust-region method in Phase I: far from the solution, the trust-region confinement often makes it difficult to exploit the full power of a good preconditioner—efficient preconditioned steps are likely to be rejected for falling outside of the trust region. A possible remedy, which we are currently investigating, is to relax the trust-region requirement so as to accept efficient preconditioned steps. Another remedy, which we study in this paper, is to use a Basic Tracemin / RTR hybrid.

In summary, we use Basic Tracemin in Phase I (far away from the solution) and the RTR algo-

gorithm with exact Hessian in Phase II (close to the solution). In this work, we develop and analyze the hybrid scheme by unifying the two constituents and their combination using the framework of adaptive model-based algorithms for the minimization of the generalized Rayleigh quotient.

2 Model-based scheme for trace minimization

We want to compute the p leftmost eigenpairs of the generalized eigenproblem (1), where A and B are positive definite $n \times n$ matrices. We denote the eigenpairs by $(\lambda_1, v_1), \dots, (\lambda_n, v_n)$ with $0 < \lambda_1 \leq \dots \leq \lambda_n$ and take v_1, \dots, v_n B -orthonormal; see, e.g., [Ste01] for details. The methods presented here aim at computing the leftmost p -dimensional eigenspace \mathcal{V} of (A, B) , namely, $\mathcal{V} = \text{span}(v_1, \dots, v_p)$. To ensure uniqueness of \mathcal{V} , we assume that $\lambda_p < \lambda_{p+1}$. When p is small, which is inherent in most applications, it is computationally inexpensive to recover the eigenvectors v_1, \dots, v_p from \mathcal{V} by solving a reduced-order generalized eigenvalue problem.

It is well known (see for example [ST00]) that the leftmost eigenspace \mathcal{V} of (A, B) is the column space of any minimizer of the Rayleigh cost function

$$f : \mathbb{R}_*^{n \times p} \rightarrow \mathbb{R} : Y \mapsto \text{trace}((Y^T B Y)^{-1} (Y^T A Y)), \quad (2)$$

where $\mathbb{R}_*^{n \times p}$ denotes the set of full-rank $n \times p$ matrices. It is readily checked that the right-hand side only depends on $\text{colsp}(Y)$. Therefore, f induces a well-defined real-valued function on the set of p -dimensional subspaces of \mathbb{R}^n .

The proposed methods iteratively compute the minimizer of f by (approximately) minimizing successive models of f . The minimization of the models themselves is done via an iterative process, which is referred to as *inner iteration*, to distinguish it with the principal, *outer iteration*. We present here the process in a way that does not require a background in differential geometry; we refer to [ABG04b] for the mathematical foundations of the technique.

Let Y , a full-rank $n \times p$ matrix, be the current iterate. The task of the inner iteration is to produce a correction S of Y such that $f(Y + S) < f(Y)$. A difficulty is that corrections of Y that do not modify its column space do not affect the value of the cost function. This situation leads to unpleasant degeneracy if it is not addressed. Therefore, we require S to satisfy some complementarity condition with respect to the space $\mathcal{V}_Y := \{YM : M \text{ } p \times p \text{ invertible}\}$. Here, in order to simplify later developments, we impose complementarity via B -orthogonality, namely $S \in \mathcal{H}_Y$ where

$$\mathcal{H}_Y := \{Z \in \mathbb{R}^{n \times p} : Y^T B Z = 0\}.$$

Consequently, the inner iteration aims at minimizing the function

$$\widehat{f}_Y(S) := \text{trace} \left(((Y + S)^T B (Y + S))^{-1} ((Y + S)^T A (Y + S)) \right), \quad S \in \mathcal{H}_Y. \quad (3)$$

A Taylor expansion of \widehat{f}_Y around $S = 0$ yields the “exact” quadratic model

$$\begin{aligned} m_Y^{\text{exact}}(S) &= \text{trace}((Y^T B Y)^{-1} (Y^T A Y)) + \text{trace}((Y^T B Y)^{-1} S^T 2 A Y) \\ &\quad + \frac{1}{2} \text{trace}((Y^T B Y)^{-1} S^T 2 (A S - B S (Y^T B Y)^{-1} Y^T A Y)), \quad S \in \mathcal{H}_Y. \end{aligned} \quad (4)$$

Throughout this paper, we let $P := I - B Y (Y^T B^2 Y)^{-1} Y^T B$ denote the orthogonal projector onto \mathcal{H}_Y , where Y is the current iterate. From (4), and using the inner product

$$\langle Z_1, Z_2 \rangle := \text{trace}((Y^T B Y)^{-1} Z_1^T Z_2), \quad Z_1, Z_2 \in \mathcal{H}_Y, \quad (5)$$

we identify $2PAY$ to be the gradient of \widehat{f}_Y at $S = 0$ and the operator $S \mapsto 2P(AS - BS(Y^T BY)^{-1}Y^T AY)$ to be the Hessian of \widehat{f}_Y at $S = 0$; see [ABG04b] for details. In the sequel, we will use the more general form

$$m_Y(S) = \text{trace}((Y^T BY)^{-1}(Y^T AY)) + \text{trace}((Y^T BY)^{-1}S^T 2AY) + \frac{1}{2}\text{trace}((Y^T BY)^{-1}S^T H_Y[S]), \quad S \in \mathcal{H}_Y \quad (6)$$

to allow for the use of inexact quadratic expansions.

The proposed general algorithm is a model trust-region scheme defined as follows.

Algorithm 1 (outer iteration)

Data: symmetric $n \times n$ matrices A and B , with B positive definite.

Parameters: $\bar{\Delta} > 0$, $\Delta_0 \in (0, \bar{\Delta}]$, and $\rho' \in [0, \frac{1}{4})$.

Input: initial iterate Y_0 (full-rank $n \times p$ matrix).

Output: sequence of iterates $\{Y_k\}$.

for $k = 0, 1, 2, \dots$ until an outer stopping criterion is satisfied:

- Using Algorithm 2, obtain S_k that (approximately) solves the trust-region subproblem

$$\min_{S \in H_{Y_k}} m_{Y_k}(S) \quad \text{s.t.} \quad \|S\|_M^2 := \text{trace}((Y_k^T B Y_k)^{-1}S^T S) \leq \Delta_k^2, \quad (7)$$

where m is defined in (6) and M is a preconditioner.

- Evaluate

$$\rho_k := \frac{\widehat{f}_{Y_k}(0) - \widehat{f}_{Y_k}(S_k)}{m_{Y_k}(0) - m_{Y_k}(S_k)} \quad (8)$$

where \widehat{f} is defined in (3).

- Update the trust-region radius:

if $\rho_k < \frac{1}{4}$
 $\Delta_{k+1} = \frac{1}{4}\Delta_k$
else if $\rho_k > \frac{3}{4}$ and $\|S_k\| = \Delta_k$
 $\Delta_{k+1} = \min(2\Delta_k, \bar{\Delta})$
else
 $\Delta_{k+1} = \Delta_k$;

- Update the iterate:

if $\rho_k > \rho'$,
 $Y_{k+1} = \text{orth}(Y_k + S_k)$, where orth denotes an orthonormalization process
which prevents loss of rank;
else
 $Y_{k+1} = Y_k$;

end (for).

The inner iteration (Algorithm 2) employed to generate S_k is a preconditioned truncated conjugate gradient method, directly inspired from the work of Steihaug [Ste83] and Toint [Toi81]. The notation $(PMP)^\dagger R$ below denotes the solution \tilde{R} of the system $PMP\tilde{R} = R$, $P\tilde{R} = \tilde{R}$, $PR = R$, which is given by the Olsen formula $\tilde{R} = M^{-1}R - M^{-1}BY(Y^T BM^{-1}BY)^{-1}Y^T BM^{-1}R$.

Algorithm 2 (inner iteration)

Set $S_0 = 0$, $R_0 = PAY_k = AY_k - BY_k(Y_k^T B^2 Y_k)^{-1} Y_k^T BAY_k$, $\tilde{R}_0 = (PMP)^\dagger R_0$, $\delta_0 = -R_0$;
for $j = 0, 1, 2, \dots$ *until an inner stopping criterion is satisfied, perform the following operations,*
where $\langle \cdot, \cdot \rangle$ denotes the inner product (5) and H_{Y_k} denotes model Hessian in (6).

if $\langle \delta_j, H_{Y_k} \delta_j \rangle \leq 0$
Compute τ such that $S = S_j + \tau \delta_j$ minimizes $m(S)$ in (6)
and satisfies $\|S\|_M = \Delta$;
return S ;
Set $\alpha_j = \langle R_j, \tilde{R}_j \rangle / \langle \delta_j, H_{Y_k} \delta_j \rangle$; Set $S_{j+1} = S_j + \alpha_j \delta_j$;
if $\|S_{j+1}\|_M \geq \Delta$
Compute $\tau \geq 0$ such that $S = S_j + \tau \delta_j$ satisfies $\|S\|_M = \Delta$;
return S ;
Set $R_{j+1} = R_j + \alpha H_{Y_k} \delta_j$; Set $\tilde{R}_{j+1} = (PMP)^\dagger R_{j+1}$;
Set $\beta_{j+1} = \langle R_{j+1}, \tilde{R}_{j+1} \rangle / \langle R_j, \tilde{R}_j \rangle$; Set $\delta_{j+1} = -\tilde{R}_{j+1} + \beta_{j+1} \delta_j$;
end (for).

We refer to [ABG04b] for a convergence analysis of Algorithm 1.

The two-phase scheme outlined in Section 1 can now be formalized.

Algorithm 3 *Phase I: Iterate Basic Tracemin—or, formally, Algorithm 1 with $H[S] := PAPS$, $\Delta_0 := +\infty$ and $\rho' := 0$ —until some switching criterion is satisfied.*

Phase II: Continue using Algorithm 1 with $H[S] := P(AS - BS(Y^T BY)^{-1} Y^T AY)$, some initial Δ and some $\rho' \in (0, \frac{1}{4})$.

In the algorithms above, stopping and switching criteria, as well as the choices of some parameters and preconditioner, were left unspecified: depending on the information available about the structure of the problem and on the accuracy/speed requirements, various choices may be appropriate. We refer to the next section for examples of specific choices.

3 Numerical experiments

In order to illustrate the practical relevance of the scheme proposed in Algorithm 3, we conducted experiments on the Calgary Olympic Saddledome arena matrices ($n = 3562$) available on Matrix Market: $A = \text{BCSSTK24}$ and $B = \text{BCSSTM24}$. The task was to compute the $p = 5$ leftmost eigenpairs of (A, B) , which correspond to the lower modes of vibration.

With a view to achieving superlinear convergence when the exact model (4) is utilized, we used an inner stopping criterion of the form

$$\|R_j\| \leq \|R_0\| \min(\|R_0\|^\theta, \kappa), \quad (9)$$

for some $\theta > 0$ and $\kappa > 0$. We chose $\theta = 1$ (striving for quadratic local convergence) and $\kappa = .5$. Several other choices are possible, notably based on the discussion in [SW82] in the case of Phase I.

In the first step of Phase II, we chose $\Delta := \|S_-\|_M$, where S_- is the last S computed in Phase I, and we chose $\rho' = .1$. The outer stopping criterion was a threshold on the norm of PAY_k . Transition between Phase I and Phase II was forced after various prescribed numbers of outer iterations, to illustrate the effect of different switching points on the behaviour of the algorithm. An initial iterate Y_0 was selected from a normal distribution. In all experiments, the preconditioner was kept constant throughout the iteration.

In the first set of experiments (Figures 1 and 2), the preconditioner M was chosen as an incomplete Cholesky factorization of A with relative drop tolerance set to 10^{-6} . Basic Tracemin converges linearly and (relatively) slowly. RTR (curve ‘0’) is eventually faster than Basic Tracemin but it is initially slower, due to the trust-region constraint. Curves ‘5’ and ‘10’ show that there is a rather large “sweet-spot” for efficient switching from Basic Tracemin to RTR such that the hybrid method performs better than both of its pure components. If switching is done too late (curve ‘20’), then superlinear convergence is still observed and the problem is solved, but with less efficiency than the properly switched versions. Preliminary results suggest that the evolution of the trace function (2) gives valuable information on when switching should take place (Figure 6). Note that all that is at stake in the choice of the switching criterion between Basic Tracemin and RTR is efficiency, and neither success nor accuracy.

In a second set of experiments (Figures 3 and 4), we initially applied a permutation on A and B as returned by the Matlab function `symamd(A)`. The preconditioner was then defined to be the exact Cholesky factorization of A , which is very sparse due to the approximate minimum degree permutation. Consequently, all algorithms converged (much) faster. Note that Phase I is mathematically equivalent to INVIT in this case. We observe a similar sweet-spot for the switching and the superiority of the global superlinear convergence of the trust-region-based method.

Finally, in Figure 5, we conducted experiments where RTR was replaced by a block RQI.¹ Because RQI does not have global convergence to the leftmost eigenspace, convergence to a non-leftmost eigenspace may occur if switching is done too early. This was observed in the experiments. To our knowledge, there is no switching criterion that guarantees convergence to the leftmost eigenspace when RQI is used. This is a major reason for using RTR instead of RQI in Phase II (another reason is to reduce the computational cost with the truncated CG inner iteration).

4 Conclusion

We have shown that an appropriate combination of the Basic Tracemin algorithm of [SW82, ST00] and the Riemannian trust-region algorithm of [ABG04a] yields an efficient method for high precision computation of the smallest eigenpairs of positive-definite generalized eigenproblems. In future work, we will further investigate the choice of the inner stopping criterion and the switching criterion between the two algorithms. For the latter, there is evidence that the decrease of the trace function (2) provides an useful guideline.

References

- [ABG04a] P.-A. Absil, C. G. Baker, and K. A. Gallivan, *A truncated-CG style method for symmetric generalized eigenvalue problems*, submitted, 2004.
- [ABG04b] ———, *Trust-region methods on Riemannian manifolds*, Tech. Report FSU-CSIT-04-13, School of Computational Science, Florida State University, July 2004, <http://www.csit.fsu.edu/~absil/Publi/RTR.htm>.
- [ABG04c] ———, *Trust-region methods on Riemannian manifolds with applications in numerical linear algebra*, Proceedings of the 16th International Symposium on Mathematical Theory of Networks and Systems (MTNS2004), Leuven, Belgium, 5–9 July 2004, 2004.

¹We used the NG algorithm mentioned in [LE02, ASVM04], which reduces to RQI in the $p = 1$ case. This algorithm refines estimates of invariant subspaces, without favoring convergence to the leftmost one.

- [ASVM04] P.-A. Absil, R. Sepulchre, P. Van Dooren, and R. Mahony, *Cubically convergent iterations for invariant subspace computation*, SIAM J. Matrix. Anal. Appl. **26** (2004), no. 1, 70–96.
- [EAS98] A. Edelman, T. A. Arias, and S. T. Smith, *The geometry of algorithms with orthogonality constraints*, SIAM J. Matrix Anal. Appl. **20** (1998), no. 2, 303–353.
- [LE02] E. Lundström and L. Eldén, *Adaptive eigenvalue computations using Newton’s method on the Grassmann manifold*, SIAM J. Matrix Anal. Appl. **23** (2002), no. 3, 819–839.
- [Par80] B. N. Parlett, *The symmetric eigenvalue problem*, Prentice-Hall, Inc., Englewood Cliffs, N.J. 07632, 1980, republished by SIAM, Philadelphia, 1998.
- [ST00] A. Sameh and Z. Tong, *The trace minimization method for the symmetric generalized eigenvalue problem*, J. Comput. Appl. Math. **123** (2000), 155–175.
- [Ste83] T. Steihaug, *The conjugate gradient method and trust regions in large scale optimization*, SIAM J. Numer. Anal. **20** (1983), 626–637.
- [Ste01] G. W. Stewart, *Matrix algorithms, vol II: Eigensystems*, Society for Industrial and Applied Mathematics, Philadelphia, 2001.
- [SW82] A. H. Sameh and J. A. Wisniewski, *A trace minimization algorithm for the generalized eigenvalue problem*, SIAM J. Numer. Anal. **19** (1982), no. 6, 1243–1259.
- [Szy88] Daniel B. Szyld, *Criteria for combining inverse and Rayleigh quotient iteration*, SIAM J. Numer. Anal. **25** (1988), no. 6, 1369–1375. MR MR972460 (90a:65089)
- [Toi81] Ph. L. Toint, *Towards an efficient sparsity exploiting Newton method for minimization*, Sparse Matrices and Their Uses (I. S. Duff, ed.), Academic Press, London, 1981, pp. 57–88.




Figure 1: Experiments on Algorithm 3 with an incomplete Cholesky factorization of A as preconditioner. The distance to the solution is measured as the largest principal angle between the column space of Y and the leftmost p -dimensional eigenspace. Since the numerical cost of inner iterations differs in both phases, the distance is plotted versus an estimation of the number of operations. Note that the formation of the preconditioner does not appear in the operation count. Circles and stars correspond to outer iterations of Basic Tracemin and RTR, respectively. The numbers on the curves indicate after how many steps of Basic Tracemin switching occurred to RTR.

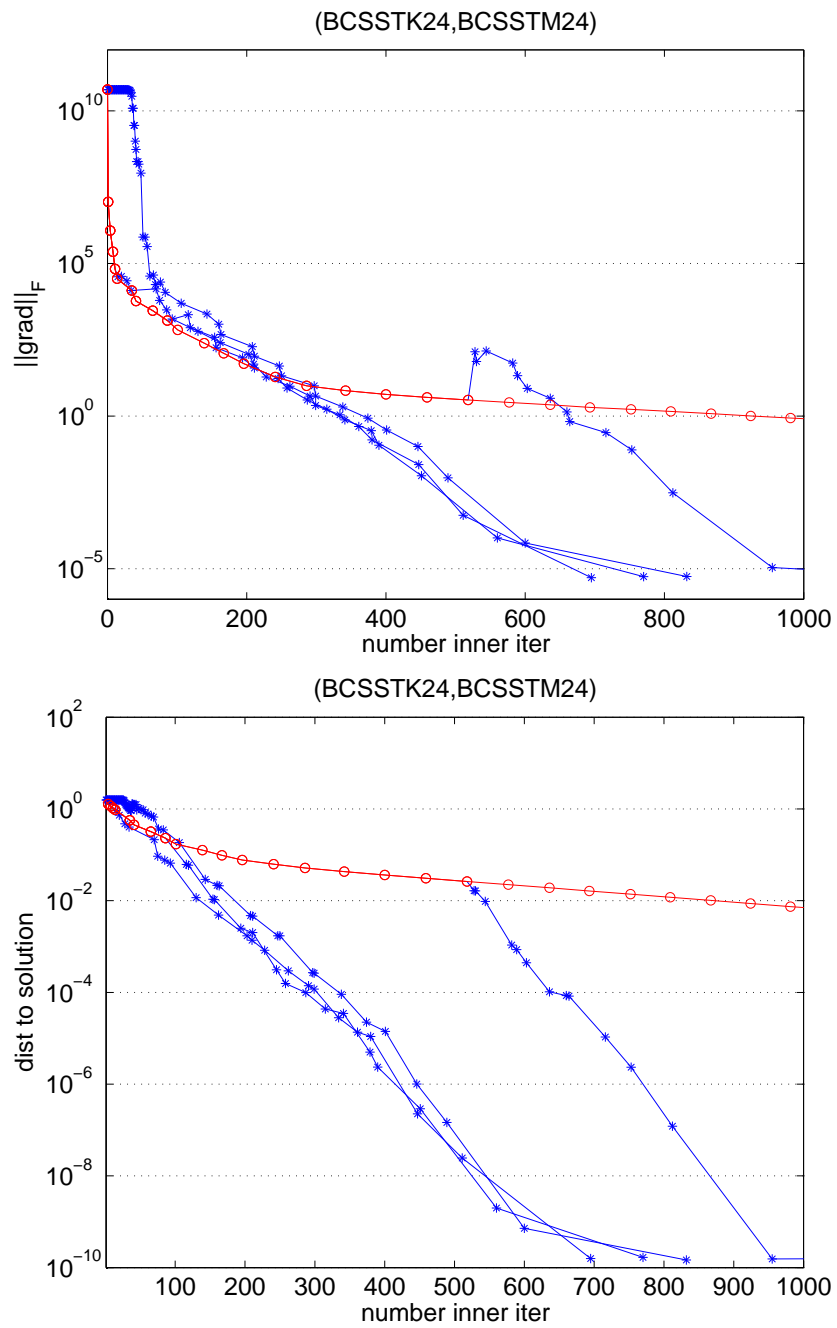


Figure 2: See Figure 1, with number of inner iterations instead operation counts.

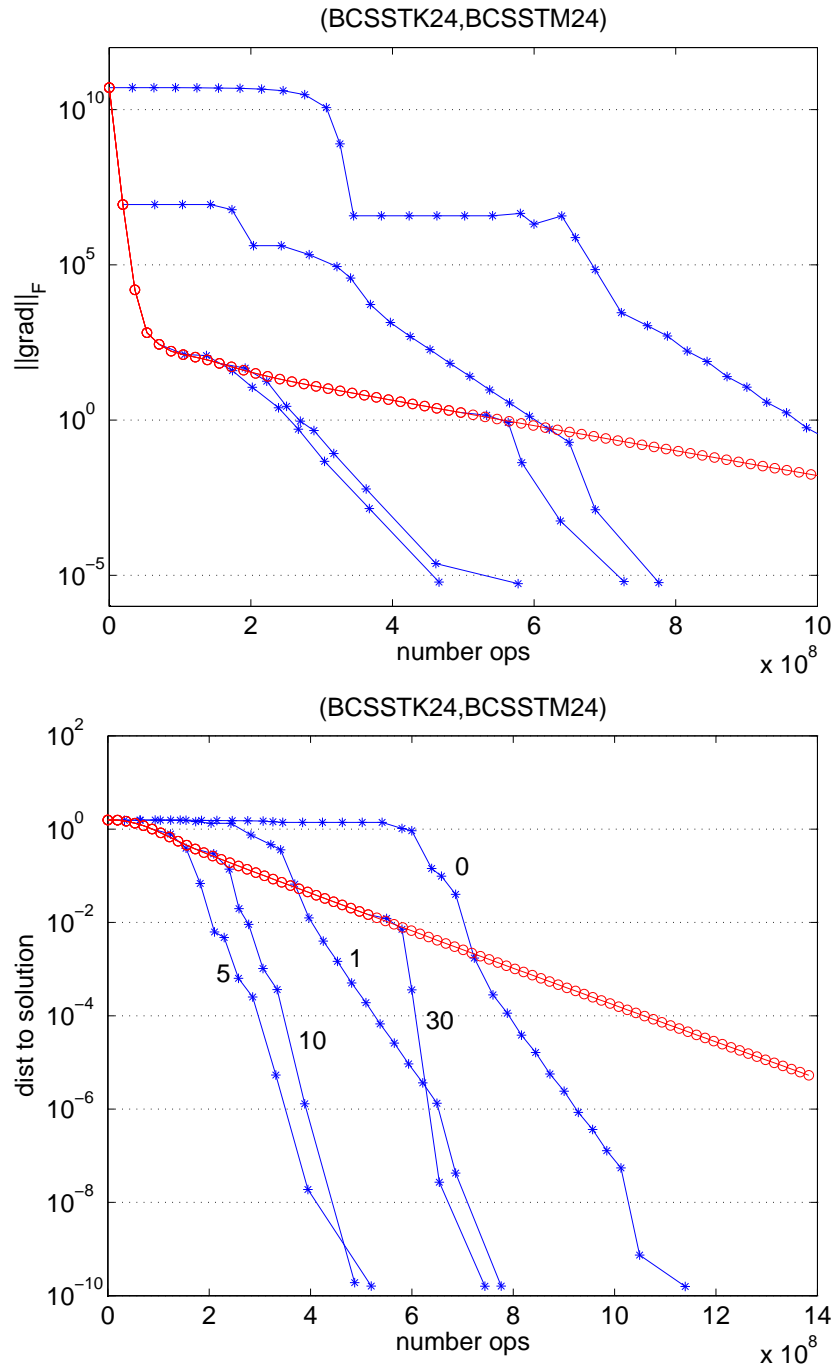


Figure 3: Same as Figure 1, with exact preconditioner after approximate minimum degree permutation.




Figure 4: Same as Figure 3, with number of inner iterations instead of operation counts.

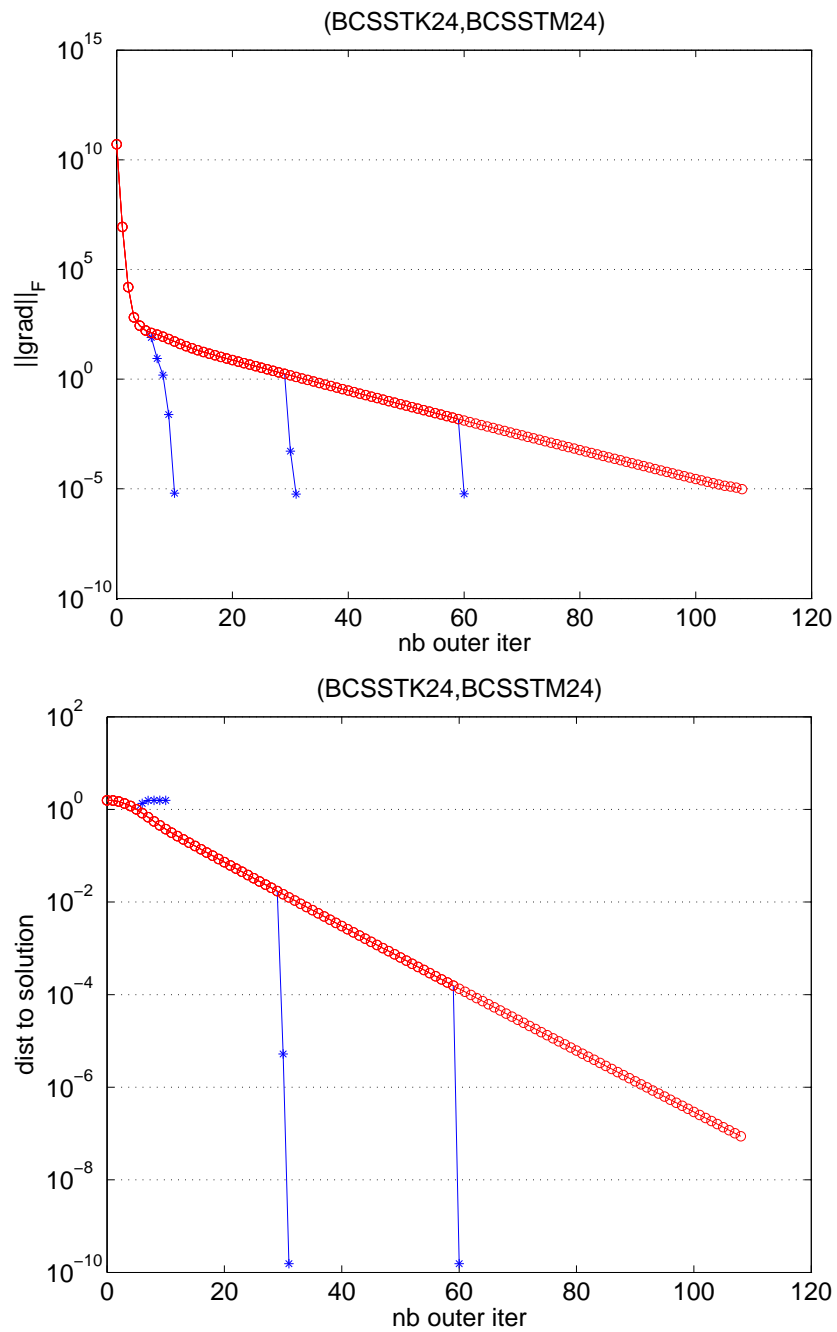


Figure 5: Experiments using a block RQI in Phase II. The graphs show that if switching occurs too early, RQI may converge to an eigenspace different from the leftmost one.




Figure 6: Evolution of the trace function (2) for the experiment in Figure 3. The second plot gives the distance between the current value and the value at the solution.