

Program 5 Foundations of Computational Math 2 Spring 2019

Due date: 11:59PM on 15 Monday April 2019

Part 1: Codes

For symmetric positive definite matrix problems, implement

- Symmetric Gauss-Seidel (SGS)
- Preconditioned Minimal Residual (PMR)
- Preconditioned Nonstationary Descent (PD)
- Preconditioned Steepest Descent (PSD)
- Preconditioned Conjugate Gradient (PCG)

PMR should be derived from the application of a two-sided symmetric preconditioner to $Ax = b$. Specifically,

$$\tilde{A} = C^{-1}AC^{-T}, \quad \tilde{x} = C^T x, \quad \tilde{b} = C^{-1}b$$

$$\tilde{x}_{k+1} = \tilde{x}_k + \tilde{\alpha}_k \tilde{r}_k$$

and you must derive a PMR iteration in terms of x_k , A , b , and C and explain how you implement it in a manner similar to all of the other codes. Note that the notes and homework problems have already dealt with this for SD and CG. Of course, PMR should also work on nonsymmetric matrices with a symmetric preconditioner. This, in general may not work that well, i.e., you probably should use a nonsymmetric preconditioner in that case. Your SGS and PMR codes should also include a version where the symmetric banded matrix-vector multiplication is replaced by a dense possibly nonsymmetric matrix for running with the symmetric preconditioner or no preconditioner. SGS should be used with no additional preconditioner, i.e., take $C = I$,

The Nonstationary PD family includes PSD when the direction vector for each iteration is taken as the residual, i.e., $p_k = r_k$. For controlling the experiments, you may wish to impose, without loss of generality, $p_k^T r_k > 0$. Recall that the convergence analysis of PD required that $p_k^T r_k$ does not tend to 0. The codes should be able to operate without preconditioning, i.e., with the preconditioner set to I . The PSD code should be able to operate in a stationary mode, i.e., with a fixed α . The codes must be able to operate in both single and double precision. The matrix-vector product routine should assume that the matrix is symmetric and banded with the nonzeros constrained to the main diagonal, k subdiagonals and k superdiagonals. It should be able to operate with $k = 0$, i.e., a diagonal matrix, up to

$k = 6$. You should also use an implementation that assumes A is symmetric and dense (you should exploit the symmetry in your matrix vector multiplication in this case). This will make it easier to generate matrices with particular spectrum values without worrying about imposing banded structure.

Use problems with known and unknown solutions to provide sufficient evidence of the correctness of your codes presented in an appropriate and convincing manner. For problems with a known solution, x^* , you can monitor the relative error $\|x_k - x^*\|/\|x^*\|$. This is, of course, not a real convergence check but is useful for empirical evaluations. For problems with known or unknown solutions monitoring $\|r_k\|/\|b\|$ is a reasonable assessment of how well the current estimate of the solution solves the system.

Part 2: Influence of the Spectrum

Let $\Lambda \in \mathbb{R}^{n \times n}$ be a diagonal matrix with positive elements and consider solving systems of the form $\Lambda x = b$. Recall that the behavior of CG and SD depends only on the spectrum in the sense that if $A = Q\Lambda Q^T$ is transformed into the diagonal coordinate system by using the eigenvectors then the solution iterates in the two coordinate systems are related via Q and Q^T and the errors at each step have the same 2-norm, i.e., convergence behavior is identical in both coordinate systems.

- (1.1) We have seen several convergence results for the nonstationary descent family, SD, CG and stationary Richardson's family. Design a set of experiments that choose Λ to influence the behavior of these methods. All methods should be run without preconditioning.
- (1.2) Clearly identify the convergence theorems that you are testing with each of the sets of experiments, state the predictions of behavior the theorems predict, and relate the observed behavior to those predictions. Make appropriate use of details of the behavior of a representative iteration as well as statistics and distributions of errors, residuals, and iteration counts from larger sets of runs.
- (1.3) Run both single and double precision experiments and determine the general effect of precision on the ability of the algorithms to solve systems to a particular accuracy. Specifically, how does precision limit the accuracy in terms of relative error or size of the residual relative to b .
- (1.4) Of course, the nonstationary descent family PD is not a specific method, i.e., SD is a particular member of the family. So for the experiments using this family you must manipulate your choice of p_k to investigate its influence on performance. You can intuitively think of SD and CG as two extremes of this family. SD is purely local in its choice of p_k and CG is an optimal choice of p_k .

While P_3 is by definition symmetric, it does not have to be positive definite and therefore, since we are requiring that the preconditioner be positive definite, P_3 may not be usable on all matrices in the test. The Cholesky factorization will fail, i.e., it will not produce lower triangular L with positive diagonal elements such that $P_3 = LL^T$. For example, the following A is positive definite but the corresponding P_3 is not.

$$A = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 \\ 1 & 2 & 2 & 2 & 2 \\ 1 & 2 & 3 & 3 & 3 \\ 1 & 2 & 3 & 4 & 4 \\ 1 & 2 & 3 & 4 & 5 \end{pmatrix}, \quad P_3 = \begin{pmatrix} 1 & 1 & & & \\ 1 & 2 & 2 & & \\ & 2 & 3 & 3 & \\ & & 3 & 4 & 4 \\ & & & 4 & 5 \end{pmatrix}$$

Note that L when it exists for P_3 must be banded with only a main diagonal and a single subdiagonal. This must be exploited in your Cholesky routine and in your preconditioned iteration. This is also true of the Cholesky factor of A ; if A has k subdiagonals then so does L in $A = LL^T$. This can be used to generate test matrices.

Part 4: Nonsymmetric Matrices

You should use the dense nonsymmetric matrix-vector product version of SGS and PMR to perform a small number of experiments that demonstrate the convergence (or lack thereof) results for SGS and PMR discussed in the notes and homeworks. You can generate a nonsymmetric matrix (dense) with a specific spectrum using the Real Schur Decomposition. Recall, that any complex matrix can be written as $A = QRQ^H$ where $Q^H Q = Q Q^H = I$, i.e., Q is a unitary matrix, and R is upper triangular with the eigenvalues (possibly complex) of A on the main diagonal. If A is real then any complex eigenvalues occur in complex conjugate pairs. For real A the Real Schur Decomposition has the form $A = QRQ^T$ where Q is real, $Q^T Q = Q Q^T = I$, and R is real but possibly block upper triangular. In particular, real eigenvalues appear as a 1×1 block (a single element) on the diagonal of R while a pair of complex conjugate eigenvalues appear as a 2×2 real block on the block diagonal of R where the pair of eigenvalues of A are also the eigenvalues of the 2×2 block.

Requirements of Code

- The techniques used to gain an efficient implementation in space and computation must be described clearly and concisely.
- For the banded symmetric matrix-vector multiplication, your code must use $O(n)$ storage for the matrix A and the preconditioner for any size n .
- The computation of the banded symmetric matrix vector product $Av \rightarrow w$ must be done efficiently in $O(n)$ computations using the efficient storage scheme.

- The solution of systems $Pz_k = r_k$ required in each iteration must be done efficiently in $O(n)$ computations using the efficient storage scheme for the appropriate preconditioners. For any other preconditioners you attempt you may use a dense storage approach.
- You must produce single and double precision versions of the codes.
- Your code should collect information at each iteration on
 - $\|e^{(k)}\|_2/\|x_{true}\|_2 = \|x_k - x_{true}\|_2/\|x_{true}\|_2$ assuming $Ax_{true} = b$ and that x_{true} is known and not small in norm.
 - $\|r_k\|_2/\|b\|_2$
- The information collected should be displayed in an appropriate organized manner to support your answers to the questions above. For representative problem instances, when considering the trends involved in the evolution of the values $\|r_k\|_2/\|b\|_2$ and $\|e^{(k)}\|_2/\|x_{true}\|_2$, in addition to plotting the quantities vs. the iteration number k , you should plot the logs vs. the iteration number k . Recall the asymptotic behavior discussion where one would expect linear behavior in the log of the error and residual norms. The observed behavior should be consistent with the theory in so far as roundoff allows. In addition to representative examples you should present global summaries using statistics and distributions of errors (histograms), residuals, and iteration counts from larger sets of runs.

Test Matrices

For Part 2 the key issue is a careful and organized characterization of the attributes of the spectrum of A and how you can generate such spectra in a systematic manner to test the predictions made by the convergence theory. In Part 2 it is not important to do massive amounts of testing but make sure that you clearly describe the experiments and their purpose.

You have seen many ways to derive test matrices

- nonsingular matrices from factorizations (possibly structured to impose structure on A)
- nonsingularity and controlling the spectrum diagonal dominance (strict and irreducible)
- Gershgorin theorem implications for controlling the spectrum
- modifications to random matrices using some of the techniques above
- generating orthogonal matrices and combining with diagonal and other matrices to generate A
- random matrices with after the fact guarantees of nonsingularity and other constraints

The only complications here are the requirements of symmetric and positive definite and you should describe clearly how you are generating your matrices for Part 3.

In addition to these techniques you can make use of matrix families with known eigenvalues and/or eigenvectors. For example, Toeplitz tridiagonal matrices parameterized by a single parameter

$$T_\alpha = \begin{pmatrix} \alpha & -1 & 0 & \dots & \dots & \dots & 0 \\ -1 & \alpha & -1 & 0 & \dots & \dots & 0 \\ 0 & -1 & \alpha & -1 & 0 & \dots & 0 \\ & \ddots & \ddots & \ddots & \ddots & \ddots & \\ 0 & \dots & 0 & -1 & \alpha & -1 & 0 \\ 0 & \dots & \dots & 0 & -1 & \alpha & -1 \\ 0 & \dots & \dots & \dots & 0 & -1 & \alpha \end{pmatrix}$$

have eigenvalues

$$\lambda_j = \alpha - 2 \cos j\theta, \quad \theta = \frac{\pi}{n+1} q_j = (\sin(j\theta), \sin(2j\theta), \dots, \sin(nj\theta))^T$$

Feel free to consult the literature for other such families.

You may use environments like MATLAB or libraries like LAPACK to generate matrices for testing and, more importantly, to analyze the spectrum of the preconditioned matrices to see if the preconditioners have altered the spectrum in such a way as to accelerate convergence. Recall, that the *PCG* and *PSD* iterations can be shown to be equivalent to solving a system with $\tilde{A} = C^{-1}AC^{-1}$ or $\tilde{A} = L^{-1}AL^{-T}$ where the preconditioner $P = C^2$ or $P = LL^T$. This equivalent symmetric form should be when computing the spectrum of the preconditioned matrix since codes for symmetric eigenvalue problems are plentiful and reliable. Clearly, this type of numerical evaluation of the spectrum for both A and \tilde{A} should be done only for problem sizes where the computation is reasonable. **Be sure to cite the libraries/routines or environments/commands used in you solutions.**

Submission of Results

Expected results comprise:

- A document describing your solutions as prescribed in the notes on writing up a programming solution posted on the class website.
- The source code, makefiles, and instructions on how to compile and execute your code including the Math Department machine used, if applicable.
- Code documentation should be included in each routine.
- All text files that do not contain code or makefiles must be PDF files. **Do not send Microsoft word files of any type.**

These results should be submitted by 11:59 PM on the due date. Submission of results is to be done via FSU Dropbox at <https://dropbox.fsu.edu>. Drop the files off for Zhifeng Deng using his MyFSU email zd16d@my.fsu.edu. (Note: **do not use zdeng@math.fsu.edu to drop off the files**) You should login to FSU Dropbox using your MyFSU login. If for some reason you cannot use FSU Dropbox please email the files to Zhifeng at the email address above.