CS 550: Iterative and Multigrid Methods
Spring 2007

Homework, Set 1 **Due Thursday February 22, 2007**

Write descriptive solutions. Comment your code!

The goal of this homework set is to investigate the projection methods of chapters 4, 5, and part of 6. You should build off of the splat package developed in HW1. note: assume sorted CSR format.

1. `basicsolve.h` Implement efficiently weigthed-Jacobi and Gauss-Seidel with the following prototypes

```
wJacobi(CSRMat *A, DVec *b, DVec *x, const int iter, const double omega);
GaussSeidel(CSRMat *A, DVec *b, DVec *x, const int iter);
// A x = b
// x     = initial guess and return vector
// A,b   = CSR matrix and rhs
// iter  = number of sweeps to perform
// omega = weigthed Jacobi parameter
```

2. Construct an example problem from FDIFF that compares the convergence of `wJacobi` versus `GaussSeidel`. A plot of iterations versus the (log) relative residual is useful.

3. `basicsolve.h` The recurrence

$$x_{k+1} = x_k + \alpha r_k$$

where $\alpha \in \mathbb{R}$ is known as Richardson iteration. Implement Richardson with the following prototype

```
Richardson(CSRMat *A, DVec *b, DVec *x, const int iter, const double alpha
   );
```

4. What polynomial $p(A)$ at step m does Richardson Iteration correspond to?

5. (extra credit) Consider $A = 2I_n + R$, where $R_{i,j} = 0.5/\sqrt{n} * \texttt{drand48}()$. What choice of α should be used? Consider $B = A + D$ where D is diagonal and $D_{kk} = (-2 + 2\sin\theta_k) + \imath\cos\theta_k$ and $\theta_k = \frac{k\pi}{m-1}$, for $0 \leq k \leq n-1$. What choice of α is needed here? Confirm your suspicions numerically.

6. `krylovsolve.h` Implement efficiently the GMRES algorithm outlined in class. Specifically, use the Arnoldi-MGS version of orthogonalization (versus Householder) and use Givens Rotations as described in section 6.5.3. Use the following prototype

```
GMRES(CSRMat *A, DVec *b, DVec *x,
      const int m, const doulbe rtol, const int maxit,
      int &iter, DVec* res);
// A x = b
// x     = initial guess and return vector
// A,b   = CSR matrix and rhs
// m     = restart
// rtol  = relative residual stopping criteria
// maxit = total number of iterations allowed (outer and inner)
// iter  = number of iterations needed to converge
// res   = residual history if not NULL
```

7. In class we attempted to bound the size of $\|r_k\|_2$ and found a strong dependence on $\kappa(T)$ (the deviation of A's normality) and on the polynomial selected from the Krylov space. Consider the matrix $A = TDT^{-1}$ with

$$D = \begin{bmatrix} \lambda_1 & & & \\ & \lambda_2 & & \\ & & \ddots & \\ & & & \lambda_n \end{bmatrix} \quad T = \begin{bmatrix} 1 & \beta & & \\ & \ddots & \ddots & \\ & & 1 & \beta \\ & & & 1 \end{bmatrix}$$

. For $\beta = 0$, $\kappa(T) = 1.0$ since the A is normal. Set $\beta = 0.9$. Investigate the following cases and describe your results. It will be helpful to generate log-linear plots of both residual history and residual reduction $\|r_k\|_2/\|r_{k-1}\|_2$ versus iteration. Also, the Ritz values, $\theta_i \in \Lambda(H_m)$ are useful but not necessary. Use an initial guess of 0 and $b = Ax$, where $x = S\mathbf{1}$.

   (i) Initialize $\lambda_j = j$, $j = 1, \ldots, n$. Consider $D(2,2) = 1.0 + \alpha$ as $\alpha \to 0$

   (ii) Initialize $\lambda_j = j$, $j = 1, \ldots, n$. Consider $D_{1:j,1:} = J_j$, where J_j a Jordan block of size j with a multiple eigenvalue of 1.0. Look at j increasing slightly.

8. Implement CG as described in algorithm 6.18. Consider a diagonal system ($n = 100$) with the minimum diagonal entry to be 0.0005 and the rest to be evenly distributed in $[0.08, 1.21]$. Use $x_0 = 0$ and $b = 1$. Investigate the behavior of CG. You may want to look at both the residual and the A-norm of the error.

```
CG(CSRMat *A, DVec *b, DVec *x,
      const doulbe rtol, const int maxit, int &iter, DVec* res);
// A x = b
// x     = initial guess and return vector
// A,b   = CSR matrix and rhs
// rtol  = relative residual stopping criteria
// maxit = total number of iterations allowed (outer and inner)
// iter  = number of iterations needed to converge
// res   = residual history if not NULL
```

9. Suppose A is a dense symmetric positive definite 1000×1000 matrix with $\kappa(A) = 100$. Estimate the flops required to solve $Ax = b$ to ten digits with Cholesky, Richardson with α_{opt}, and with CG.

Please hand in a tarred-gziped file with the following name:
`CS550_hw2_FirstInitialLastName.tgz`
For example:
`CS550_hw2_LOlson.tgz`
This should unpack in the directory
`CS550_hw2_LOlson`
with the command
`tar xvfz CS550_hw2_LOlson.tgz`
Comments or homework write-up should be included in this directory under `hw2.pdf`.