Exercise 7

The main goal of this assignment is to experiment with stationary iterative system solvers.

1. Matlab provides a large number of interesting test matrices that are worth exploring. An overview can be obtained by typing `help gallery` in the command window. Today we are interested in one particular example from finite elements, see the excerpt below.

```
>> help private/wathen
WATHEN Wathen matrix (sparse).
A = GALLERY('WATHEN',NX,NY) is a sparse matrix with random
elements. It is an N-by-N finite element matrix, where
       N = 3*NX*NY + 2*NX + 2*NY + 1.
A is precisely the "consistent mass matrix" for a regular NX-by-NY
grid of 8-node elements in two space dimensions. A is symmetric
positive definite for any (positive) values of the "density",
RHO(NX,NY), which is chosen randomly. In particular, if
D = DIAG(DIAG(A)), then 0.25 <= EIG(INV(D)*A) <= 4.5, for any
positive integers NX and NY and any densities RHO(NX,NY).
B = GALLERY('WATHEN',NX,NY,1) returns B = DIAG(DIAG(A))\A,
which is A row-scaled to have unit diagonal.
```

   (a) Generate a sample of Wathen's matrix with `A=gallery('wathen',8,8)`. Visualize its sparsity pattern with the `spy` command. Factorize the SPD matrix using `chol` and `plot` the sparsity pattern of the triangular factor. (You can use the `subplot` command to have the plots appear next to each other).

   (b) Compare the results with a reverse Cuthill–McKee-permuted matrix, using `symrcm`.

   (c) Compare the results with an approximate minimum degree permutation of the matrix, using `symamd`.

   (d) Compare the elimination trees of the three permuted versions of $\boldsymbol{A}$ (original, permuted with `symrcm` / `symamd`). Use the Matlab function `etreeplot` to generate the trees. Which permutation generates the shortest tree? By comparing heights of three elimination trees, can you predict which version of $\boldsymbol{A}$ best minimizes fill-in?

   (e) Now generate samples of Wathen matrix with `A = gallery('wathen', N, N)` , for $N = 16, 32, 64, 128, 256$. For all the three permuted versions of $\boldsymbol{A}$ (original, permuted with `symrcm` / `symamd`), measure the times for the Cholesky factorization $\boldsymbol{A} = \boldsymbol{L}\boldsymbol{L}^T$ and for the forward/backward substitution, i.e, $\boldsymbol{L}\boldsymbol{y} = \boldsymbol{b}$, $\boldsymbol{L}^T\boldsymbol{x} = \boldsymbol{y}$. Use `tic` and `toc` to measure the elapsed time. Also, compute the number of non-zero elements in the Cholesky factor $\boldsymbol{L}$. Make a table or plot your results. What do you observe? How do you explain your results?

Please submit your solution via e-mail to Peter Arbenz (`arbenz@inf.ethz.ch`) by November 9, 2017. (12:00). Please specify the tag **FEM17** in the subject field.

Remark: You could conduct similar experiments with examples from the SuiteSparse Matrix Collection (formerly the University of Florida Sparse Matrix Collection) at `https://sparse.tamu.edu/` which contains many interesting sparse matrices.

2. Consider a 2×2 symmetric positive definite matrix $\boldsymbol{A}$. Show that both the Jacobi and the Gauss–Seidel iterations converge. Compare the spectral radii of the respective iteration matrices. Which method converges faster?

3. Show that if A is singular and a stationary method with $\boldsymbol{G} = \boldsymbol{M}^{-1}\boldsymbol{N}$ exists, then

$$\rho(\boldsymbol{G}) \geq 1.$$

Does this mean that the method never converges for any starting vector $\boldsymbol{x}_0$?

4. [Saad] Consider ***Richardson's iteration*** given by,

$$\boldsymbol{x}_{k+1} = \boldsymbol{x}_k + \alpha(\boldsymbol{b} - \boldsymbol{A}\boldsymbol{x}_k)$$

where α is a positive scalar, and the preconditioner is $\boldsymbol{M} = (1/\alpha)\boldsymbol{I}$.

   (a) Rewrite this iteration in fixed point form.

$$\boldsymbol{x}_{k+1} = \boldsymbol{G}_\alpha \boldsymbol{x}_k + \boldsymbol{c}.$$

   (b) Assume that the eigenvalues λ_i, $i = 1, \ldots, n$, of the matrix $\boldsymbol{A}$, are all real such that $\lambda_{\min} \leq \lambda_i \leq \lambda_{\max}$. What is the range of eigenvalues of the matrix $\boldsymbol{G}_\alpha$?

   (c) Show that if $\lambda_{\min} < 0$ and $\lambda_{\max} > 0$, then Richardson's iteration will always diverge for *some* initial starting vector $\boldsymbol{x}_0$.

   (d) Let us now assume that all eigenvalues are positive, $0 < \lambda_{\min} \leq \lambda_{\max}$. Under what conditions does Richardson's iteration certainly converge?

   What is the optimal value of α? In other words, what is the value of α that minimizes the spectral radius $\rho(\boldsymbol{G}_\alpha)$?

5. Determine experimentally the optimal SSOR relaxation parameter ω for the tridiagonal matrix `tridiag`$([-1, 2, -1])$. Is it the same for all matrix sizes?

Please submit your solution via e-mail to Peter Arbenz (`arbenz@inf.ethz.ch`) by November 9, 2017. (12:00). Please specify the tag **FEM17** in the subject field.