

NAG Toolbox

nag_sparse_complex_herm_basic_setup (f11gr)

1 Purpose

`nag_sparse_complex_herm_basic_setup` (f11gr) is a setup function, the first in a suite of three functions for the iterative solution of a complex Hermitian system of simultaneous linear equations. `nag_sparse_complex_herm_basic_setup` (f11gr) must be called before `nag_sparse_complex_herm_basic_solver` (f11gs), the iterative solver. The third function in the suite, `nag_sparse_complex_herm_basic_diag` (f11gt), can be used to return additional information about the computation.

These three functions are suitable for the solution of large sparse complex Hermitian systems of equations.

2 Syntax

```
[lwreq, work, ifail] = nag_sparse_complex_herm_basic_setup(method, precon, n,
tol, maxitn, anorm, sigmax, maxits, monit, 'sigcmp', sigcmp, 'norm_p', norm_p,
'weight', weight, 'iterm', iterm, 'sigtol', sigtol)

[lwreq, work, ifail] = f11gr(method, precon, n, tol, maxitn, anorm, sigmax,
maxits, monit, 'sigcmp', sigcmp, 'norm_p', norm_p, 'weight', weight, 'iterm',
iterm, 'sigtol', sigtol)
```

3 Description

The suite consisting of the functions `nag_sparse_complex_herm_basic_setup` (f11gr), `nag_sparse_complex_herm_basic_solver` (f11gs) and `nag_sparse_complex_herm_basic_diag` (f11gt) is designed to solve the complex Hermitian system of simultaneous linear equations $Ax = b$ of order n , where n is large and the matrix of the coefficients A is sparse.

`nag_sparse_complex_herm_basic_setup` (f11gr) is a setup function which must be called before the iterative solver `nag_sparse_complex_herm_basic_solver` (f11gs). `nag_sparse_complex_herm_basic_diag` (f11gt), the third function in the suite, can be used to return additional information about the computation. Either of two methods can be used:

1. Conjugate Gradient Method (CG)

For this method (see Hestenes and Stiefel (1952), Golub and Van Loan (1996), Barrett *et al.* (1994) and Dias da Cunha and Hopkins (1994)), the matrix A should ideally be positive definite. The application of the Conjugate Gradient method to indefinite matrices may lead to failure or to lack of convergence.

2. Lanczos Method (SYMMLQ)

This method, based upon the algorithm SYMMLQ (see Paige and Saunders (1975) and Barrett *et al.* (1994)), is suitable for both positive definite and indefinite matrices. It is more robust than the Conjugate Gradient method but less efficient when A is positive definite.

Both CG and SYMMLQ methods start from the residual $r_0 = b - Ax_0$, where x_0 is an initial estimate for the solution (often $x_0 = 0$), and generate an orthogonal basis for the Krylov subspace $\text{span}\{A^k r_0\}$, for $k = 0, 1, \dots$, by means of three-term recurrence relations (see Golub and Van Loan (1996)). A sequence of real symmetric tridiagonal matrices $\{T_k\}$ is also generated. Here and in the following, the index k denotes the iteration count. The resulting real symmetric tridiagonal systems of equations are usually more easily solved than the original problem. A sequence of solution iterates $\{x_k\}$ is thus generated such that the sequence of the norms of the residuals $\{\|r_k\|\}$ converges to a required tolerance. Note that, in general, the convergence is not monotonic.

In exact arithmetic, after n iterations, this process is equivalent to an orthogonal reduction of A to real symmetric tridiagonal form, $T_n = Q^H A Q$; the solution x_n would thus achieve exact convergence. In

finite-precision arithmetic, cancellation and round-off errors accumulate causing loss of orthogonality. These methods must therefore be viewed as genuinely iterative methods, able to converge to a solution **within a prescribed tolerance**.

The orthogonal basis is not formed explicitly in either method. The basic difference between the two methods lies in the method of solution of the resulting real symmetric tridiagonal systems of equations: the conjugate gradient method is equivalent to carrying out an LDL^H (Cholesky) factorization whereas the Lanczos method (SYMMLQ) uses an LQ factorization.

Faster convergence can be achieved using a **preconditioner** (see Golub and Van Loan (1996) and Barrett *et al.* (1994)). A preconditioner maps the original system of equations onto a different system, say

$$\bar{A}\bar{x} = \bar{b}, \quad (1)$$

with, hopefully, better characteristics with respect to its speed of convergence: for example, the condition number of the matrix of the coefficients can be improved or eigenvalues in its spectrum can be made to coalesce. An orthogonal basis for the Krylov subspace $\text{span}\{\bar{A}^k\bar{r}_0\}$, for $k = 0, 1, \dots$, is generated and the solution proceeds as outlined above. The algorithms used are such that the solution and residual iterates of the original system are produced, not their preconditioned counterparts. Note that an unsuitable preconditioner or no preconditioning at all may result in a very slow rate, or lack, of convergence. However, preconditioning involves a trade-off between the reduction in the number of iterations required for convergence and the additional computational costs per iteration. Also, setting up a preconditioner may involve non-negligible overheads.

A preconditioner must be **Hermitian and positive definite**, i.e., representable by $M = EE^H$, where M is nonsingular, and such that $\bar{A} = E^{-1}AE^{-H} \sim I_n$ in (1), where I_n is the identity matrix of order n . Also, we can define $\bar{r} = E^{-1}r$ and $\bar{x} = E^Hx$. These are formal definitions, used only in the design of the algorithms; in practice, only the means to compute the matrix-vector products $v = Au$ and to solve the preconditioning equations $Mv = u$ are required, that is, explicit information about M , E or their inverses is not required at any stage.

The first termination criterion

$$\|r_k\|_p \leq \tau \left(\|b\|_p + \|A\|_p \times \|x_k\|_p \right) \quad (2)$$

is available for both conjugate gradient and Lanczos (SYMMLQ) methods. In (2), $p = 1, \infty$ or 2 and τ denotes a user-specified tolerance subject to $\max(10, \sqrt{n})\epsilon \leq \tau < 1$, where ϵ is the **machine precision**. Facilities are provided for the estimation of the norm of the matrix of the coefficients $\|A\|_1 = \|A\|_\infty$, when this is not known in advance, used in (2), by applying Higham's method (see Higham (1988)). Note that $\|A\|_2$ cannot be estimated internally. This criterion uses an error bound derived from **backward** error analysis to ensure that the computed solution is the exact solution of a problem as close to the original as the termination tolerance requires. Termination criteria employing bounds derived from **forward** error analysis could be used, but any such criteria would require information about the condition number $\kappa(A)$ which is not easily obtainable.

The second termination criterion

$$\|\bar{r}_k\|_2 \leq \tau \max(1.0, \|b\|_2/\|r_0\|_2) (\|\bar{r}_0\|_2 + \sigma_1(\bar{A}) \times \|\Delta\bar{x}_k\|_2) \quad (3)$$

is available only for the Lanczos method (SYMMLQ). In (3), $\sigma_1(\bar{A}) = \|\bar{A}\|_2$ is the largest singular value of the (preconditioned) iteration matrix \bar{A} . This termination criterion monitors the progress of the solution of the preconditioned system of equations and is less expensive to apply than criterion (2). When $\sigma_1(\bar{A})$ is not supplied, facilities are provided for its estimation by $\sigma_1(\bar{A}) \sim \max_k \sigma_1(T_k)$. The interlacing property $\sigma_1(T_{k-1}) \leq \sigma_1(T_k)$ and Gerschgorin's theorem provide lower and upper bounds from which $\sigma_1(T_k)$ can be easily computed by bisection. Alternatively, the less expensive estimate $\sigma_1(\bar{A}) \sim \max_k \|T_k\|_1$ can be used, where $\sigma_1(\bar{A}) \leq \|T_k\|_1$ by Gerschgorin's theorem. Note that only order of magnitude estimates are required by the termination criterion.

Termination criterion (2) is the recommended choice, despite its (small) additional costs per iteration when using the Lanczos method (SYMMLQ). Also, if the norm of the initial estimate is much larger

than the norm of the solution, that is, if $\|x_0\| \gg \|x\|$, a dramatic loss of significant digits could result in complete lack of convergence. The use of criterion (2) will enable the detection of such a situation, and the iteration will be restarted at a suitable point. No such restart facilities are provided for criterion (3).

Optionally, a vector w of user-specified weights can be used in the computation of the vector norms in termination criterion (2), i.e., $\|v\|_p^{(w)} = \|v^{(w)}\|_p$, where $(v^{(w)})_i = w_i v_i$, for $i = 1, 2, \dots, n$. Note that the use of weights increases the computational costs.

The sequence of calls to the functions comprising the suite is enforced: first, the setup function `nag_sparse_complex_herm_basic_setup` (f11gr) must be called, followed by the solver `nag_sparse_complex_herm_basic_solver` (f11gs). `nag_sparse_complex_herm_basic_diag` (f11gt) can be called either when `nag_sparse_complex_herm_basic_solver` (f11gs) is carrying out a monitoring step or after `nag_sparse_complex_herm_basic_solver` (f11gs) has completed its tasks. Incorrect sequencing will raise an error condition.

4 References

Barrett R, Berry M, Chan T F, Demmel J, Donato J, Dongarra J, Eijkhout V, Pozo R, Romine C and Van der Vorst H (1994) *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods* SIAM, Philadelphia

Dias da Cunha R and Hopkins T (1994) PIM 1.1 — the parallel iterative method package for systems of linear equations user's guide — Fortran 77 version *Technical Report* Computing Laboratory, University of Kent at Canterbury, Kent, UK

Golub G H and Van Loan C F (1996) *Matrix Computations* (3rd Edition) Johns Hopkins University Press, Baltimore

Hestenes M and Stiefel E (1952) Methods of conjugate gradients for solving linear systems *J. Res. Nat. Bur. Stand.* **49** 409–436

Higham N J (1988) FORTRAN codes for estimating the one-norm of a real or complex matrix, with applications to condition estimation *ACM Trans. Math. Software* **14** 381–396

Paige C C and Saunders M A (1975) Solution of sparse indefinite systems of linear equations *SIAM J. Numer. Anal.* **12** 617–629

5 Parameters

5.1 Compulsory Input Parameters

1: **method** – CHARACTER(*)

The iterative method to be used.

method = 'CG'

Conjugate gradient method.

method = 'SYMMLQ'

Lanczos method (SYMMLQ).

Constraint: **method** = 'CG' or 'SYMMLQ'.

2: **precon** – CHARACTER(1)

Determines whether preconditioning is used.

precon = 'N'

No preconditioning.

precon = 'P'

Preconditioning.

Constraint: **precon** = 'N' or 'P'.

- 3: **n** – INTEGER
 n , the order of the matrix A .
Constraint: **n** > 0.
- 4: **tol** – REAL (KIND=nag_wp)
The tolerance τ for the termination criterion.
If **tol** \leq 0.0, $\tau = \max(\sqrt{\epsilon}, \sqrt{n}\epsilon)$ is used, where ϵ is the *machine precision*.
Otherwise $\tau = \max(\mathbf{tol}, 10\epsilon, \sqrt{n}\epsilon)$ is used.
Constraint: **tol** < 1.0.
- 5: **maxitn** – INTEGER
The maximum number of iterations.
Constraint: **maxitn** > 0.
- 6: **anorm** – REAL (KIND=nag_wp)
If **anorm** > 0.0, the value of $\|A\|_p$ to be used in the termination criterion (2) (**iterm** = 1).
If **anorm** \leq 0.0, **iterm** = 1 and **norm_p** = '1' or 'I', then $\|A\|_1 = \|A\|_\infty$ is estimated internally by nag_sparse_complex_herm_basic_solver (f11gs).
If **iterm** = 2, then **anorm** is not referenced.
Constraint: if **iterm** = 1 and **norm_p** = 2, **anorm** > 0.0.
- 7: **sigmax** – REAL (KIND=nag_wp)
If **sigmax** > 0.0, the value of $\sigma_1(\bar{A}) = \|E^{-1}AE^{-H}\|_2$.
If **sigmax** \leq 0.0, $\sigma_1(\bar{A})$ is estimated by nag_sparse_complex_herm_basic_solver (f11gs) when either **sigcmp** = 'S' or termination criterion (3) (**iterm** = 2) is employed, though it will be used only in the latter case.
Otherwise, **sigmax** is not referenced.
- 8: **maxits** – INTEGER
Suggested value: **maxits** = min(10, n) when **sigtol** is of the order of its default value (0.01).
The maximum iteration number $k = \mathbf{maxits}$ for which $\sigma_1(T_k)$ is computed by bisection (see also Section 3). If **sigcmp** = 'N' or **sigmax** > 0.0, then **maxits** is not referenced.
Constraint: if **sigcmp** = 'S' and **sigmax** \leq 0.0, $1 \leq \mathbf{maxits} \leq \mathbf{maxitn}$.
- 9: **monit** – INTEGER
If **monit** > 0, the frequency at which a monitoring step is executed by nag_sparse_complex_herm_basic_solver (f11gs): the current solution and residual iterates will be returned by nag_sparse_complex_herm_basic_solver (f11gs) and a call to nag_sparse_complex_herm_basic_diag (f11gt) made possible every **monit** iterations, starting from iteration number **monit**. Otherwise, no monitoring takes place. There are some additional computational costs involved in monitoring the solution and residual vectors when the Lanczos method (SYMMLQ) is used.
Constraint: **monit** \leq **maxitn**.

5.2 Optional Input Parameters

- 1: **sigcmp** – CHARACTER(1)
Suggested value: **sigcmp** = 'N'.

Default: 'N'

Determines whether an estimate of $\sigma_1(\bar{A}) = \|E^{-1}AE^{-H}\|_2$, the largest singular value of the preconditioned matrix of the coefficients, is to be computed using the bisection method on the sequence of tridiagonal matrices $\{T_k\}$ generated during the iteration. Note that $\bar{A} = A$ when a preconditioner is not used.

If **sigmax** > 0.0 (see below), i.e., when $\sigma_1(\bar{A})$ is supplied, the value of **sigcmp** is ignored.

sigcmp = 'S'

$\sigma_1(\bar{A})$ is to be computed using the bisection method.

sigcmp = 'N'

The bisection method is not used.

If the termination criterion (3) is used, requiring $\sigma_1(\bar{A})$, an inexpensive estimate is computed and used (see Section 3).

Constraint: **sigcmp** = 'S' or 'N'.

2: **norm_p** – CHARACTER(1)

Suggested value:

if **iterm** = 1, **norm_p** = '1';

if **iterm** = 2, **norm_p** = '2'.

Default:

if **iterm** = 1, '1';

otherwise '2'.

Defines the matrix and vector norm to be used in the termination criteria.

norm_p = '1'

Use the l_1 norm.

norm_p = 'I'

Use the l_∞ norm.

norm_p = '2'

Use the l_2 norm.

Constraints:

if **iterm** = 1, **norm_p** = '1', 'I' or '2';

if **iterm** = 2, **norm_p** = '2'.

3: **weight** – CHARACTER(1)

Suggested value: **weight** = 'N'.

Default: 'N'

Specifies whether a vector w of user-supplied weights is to be used in the vector norms used in the computation of termination criterion (2) (**iterm** = 1): $\|v\|_p^{(w)} = \|v^{(w)}\|_p$, where $v_i^{(w)} = w_i v_i$, for $i = 1, 2, \dots, n$. The suffix $p = 1, 2, \infty$ denotes the vector norm used, as specified by the argument **norm_p**. Note that weights cannot be used when **iterm** = 2, i.e., when criterion (3) is used.

weight = 'W'

User-supplied weights are to be used and must be supplied on initial entry to nag_sparse_complex_herm_basic_solver (f11gs).

weight = 'N'

All weights are implicitly set equal to one. Weights do not need to be supplied on initial entry to nag_sparse_complex_herm_basic_solver (f11gs).

Constraints:

if **iterm** = 1, **weight** = 'W' or 'N';
if **iterm** = 2, **weight** = 'N'.

4: **iterm** – INTEGER

Suggested value: **iterm** = 1.

Default: 1

Defines the termination criterion to be used.

iterm = 1

Use the termination criterion defined in (2) (both conjugate gradient and Lanczos (SYMMLQ) methods).

iterm = 2

Use the termination criterion defined in (3) (Lanczos method (SYMMLQ) only).

Constraints:

if **method** = 'CG', **iterm** = 1;
if **method** = 'SYMMLQ', **iterm** = 1 or 2.

5: **sigtol** – REAL (KIND=nag_wp)

Suggested value: **sigtol** = 0.01 should be sufficient in most cases.

Default: 0.01

The tolerance used in assessing the convergence of the estimate of $\sigma_1(\bar{A}) = \|\bar{A}\|_2$ when the bisection method is used.

If **sigtol** \leq 0.0, the default value **sigtol** = 0.01 is used. The actual value used is $\max(\mathbf{sigtol}, \epsilon)$.

If **sigcmp** = 'N' or **sigmax** $>$ 0.0, then **sigtol** is not referenced.

Constraint: if **sigcmp** = 'S' and **sigmax** \leq 0.0, **sigtol** $<$ 1.0.

5.3 Output Parameters

1: **lwreq** – INTEGER

The minimum amount of workspace required by nag_sparse_complex_herm_basic_solver (f11gs). (See also Section 5 in nag_sparse_complex_herm_basic_solver (f11gs).)

2: **work**(*lwork*) – COMPLEX (KIND=nag_wp) array

lwork = 120.

The array **work** is initialized by nag_sparse_complex_herm_basic_setup (f11gr). It must **not** be modified before calling the next function in the suite, namely nag_sparse_complex_herm_basic_solver (f11gs).

3: **ifail** – INTEGER

ifail = 0 unless the function detects an error (see Section 5).

6 Error Indicators and Warnings

Errors or warnings detected by the function:

ifail = $-i$

If **ifail** = $-i$, parameter i had an illegal value on entry. The parameters are numbered as follows:

1: **method**, 2: **precon**, 3: **sigcmp**, 4: **norm_p**, 5: **weight**, 6: **iterm**, 7: **n**, 8: **tol**, 9: **maxitn**, 10: **anorm**, 11: **sigmax**, 12: **sigtol**, 13: **maxits**, 14: **monit**, 15: **lwreq**, 16: **work**, 17: **lwork**, 18: **ifail**.

It is possible that **ifail** refers to a parameter that is omitted from the MATLAB interface. This usually indicates that an error in one of the other input parameters has caused an incorrect value to be inferred.

ifail = 1

nag_sparse_complex_herm_basic_setup (f11gr) has been called out of sequence.

ifail = -99

An unexpected error has been triggered by this routine. Please contact NAG.

ifail = -399

Your licence key may have expired or may not have been installed correctly.

ifail = -999

Dynamic memory allocation failed.

7 Accuracy

Not applicable.

8 Further Comments

When $\sigma_1(\bar{A})$ is not supplied (**sigmax** ≤ 0.0) but it is required, it is estimated by nag_sparse_complex_herm_basic_solver (f11gs) using either of the two methods described in Section 3, as specified by the argument **sigcmp**. In particular, if **sigcmp** = 'S', then the computation of $\sigma_1(\bar{A})$ is deemed to have converged when the differences between three successive values of $\sigma_1(T_k)$ differ, in a relative sense, by less than the tolerance **sigtol**, i.e., when

$$\max \left(\frac{|\sigma_1^{(k)} - \sigma_1^{(k-1)}|}{\sigma_1^{(k)}}, \frac{|\sigma_1^{(k)} - \sigma_1^{(k-2)}|}{\sigma_1^{(k)}} \right) \leq \mathbf{sigtol}.$$

The computation of $\sigma_1(\bar{A})$ is also terminated when the iteration count exceeds the maximum value allowed, i.e., $k \geq \mathbf{maxits}$.

Bisection is increasingly expensive with increasing iteration count. A reasonably large value of **sigtol**, of the order of the suggested value, is recommended and an excessive value of **maxits** should be avoided. Under these conditions, $\sigma_1(\bar{A})$ usually converges within very few iterations.

9 Example

This example solves a complex Hermitian system of simultaneous linear equations using the conjugate gradient method, where the matrix of the coefficients A , has a random sparsity pattern. An incomplete Cholesky preconditioner is used (nag_sparse_real_symm_precon_ichol (f11ja) and nag_sparse_real_symm_precon_ichol_solve (f11jb)).

9.1 Program Text

```
function f11gr_example
fprintf('f11gr example results\n\n');

% Solve sparse Hermitian system Ax = b using CG method with
% Incomplete Cholesky preconditioning (IC)
```

```

% Define A and b
n = nag_int(9);
nz = nag_int(23);
a = complex(zeros(3*nz,1));
irow = zeros(3*nz, 1, nag_int_name);
icol = zeros(3*nz, 1, nag_int_name);
a(1:nz) = [ 6 + 0i;          -1 + 1i; 6 + 0i;          0 + 1i; 5 + 0i;
           5 + 0i;          2 - 2i; 4 + 0i;          1 + 1i; 2 + 0i; 6 + 0i;
          -4 + 3i; 0 + 1i; -1 + 0i; 6 + 0i;          -1 - 1i; 0 - 1i; 9 + 0i;
           1 + 3i; 1 + 2i; -1 + 0i; 1 + 4i; 9 + 0i];

irow(1:nz) = [1; 2;2; 3;3; 4; 5;5; 6;6;6; 7;7;7;7; 8;8;8; 9;9;9;9;9];
icol(1:nz) = [1; 1;2; 2;3; 4; 1;5; 3;4;6; 2;5;6;7; 4;6;8; 1;5;6;8;9];

b = [ 8 + 54i; -10 - 92i; 25 + 27i;
      26 - 28i; 54 + 12i; 26 - 22i;
      47 + 65i; 71 - 57i; 60 + 70i];

% Setup IC factorization
lfill = nag_int(0);
dtol = 0;
mic = 'N';
dscale = 0;
ipiv = zeros(n, 1, nag_int_name);

[a, irow, icol, ipiv, istr, nnzc, npivm, ifail] = ...
    f11jn( ...
        nz, a, irow, icol, lfill, dtol, mic, dscale, ipiv);

% Iterative method setup
method = 'CG';
precon = 'Preconditioned';
tol = (x02aj)^(3/8);
maxitn = nag_int(20);
anorm = 0;
sigmax = 0;
maxits = nag_int(9);
monit = nag_int(2);

[lwreq, work, ifail] = ...
    f11gr( ...
        method, precon, nag_int(n), tol, maxitn, anorm, sigmax, ...
        maxits, monit, 'sigcmp', 's', 'norm_p', 'l');

% Reverse communication loop calling f11ge
irevcm = nag_int(0);
u = complex(zeros(n,1));
v = b;
wgt = zeros(n,1);

while (irevcm ~= 4)
    [irevcm, u, v, work, ifail] = ...
        f11gs( ...
            irevcm, u, v, wgt, work);

    if (irevcm == 1)
        % v = Au
        [v, ifail] = f11xs( ...
            a(1:nz), irow(1:nz), icol(1:nz), 'N', u);
    elseif (irevcm == 2)
        % Solve (IC)v = u
        [v, ifail] = f11jp( ...
            a, irow, icol, ipiv, istr, 'N', u);
    elseif (irevcm == 3)
        % Monitoring
        [itn, stplhs, strrhs, anorm, sigmax, its, sigerr, ifail] = ...
            f11gt(work);
        fprintf('\nMonitoring at iteration number %2d\n',itn);
        fprintf('residual norm:           %14.4e\n', stplhs);
        fprintf('\n Solution Vector\n');
        disp(u);
    end
end

```

```

    fprintf('\n  Residual Vector\n');
    disp(v);
end
end

% Get information about the computation
[itn, stplhs, stprhs, anorm, sigmax, its, sigerr, ifail] = ...
fllgt(work);

fprintf('\nNumber of iterations for convergence:      %4d\n', itn);
fprintf('Residual norm:                               %14.4e\n', stplhs);
fprintf('Right-hand side of termination criteria: %14.4e\n', stprhs);
fprintf('i-norm of matrix a:                           %14.4e\n', anorm);
fprintf('\n  Solution Vector\n');
disp(u);
fprintf('\n  Residual Vector\n');
disp(v);

```

9.2 Program Results

fllgr example results

```

Monitoring at iteration number 2
residual norm:                1.4937e+01

```

```

Solution Vector
0.2142 + 4.5333i
-1.6589 -12.6722i
2.4101 + 7.4551i
4.4400 - 6.4174i
9.1135 + 3.7812i
4.4419 - 4.0382i
1.4757 + 1.2662i
8.4872 - 3.5347i
5.9948 + 0.9685i

```

```

Residual Vector
-1.8370 + 3.6956i
-0.6501 + 0.2546i
-0.1262 - 0.1362i
-0.1312 + 0.1413i
-1.1471 + 0.7339i
-0.5505 - 1.0535i
1.7165 - 1.4614i
-0.3583 + 0.2876i
-0.3028 - 0.3532i

```

```

Monitoring at iteration number 4
residual norm:                1.4602e+00

```

```

Solution Vector
1.0061 + 8.9847i
1.9637 - 7.9768i
3.0067 + 7.0285i
3.9830 - 5.9636i
5.0390 + 5.0432i
6.0488 - 4.0771i
6.9710 + 3.0168i
8.0118 - 1.9806i
9.0074 + 0.9646i

```

```

Residual Vector
0.0115 - 0.0282i
0.0135 - 0.1734i
0.0182 + 0.0196i
0.0189 - 0.0204i
-0.0909 - 0.1090i
-0.2389 + 0.3244i
0.1903 - 0.0155i
0.0516 - 0.0414i

```

0.0436 + 0.0509i

| | |
|--|------------|
| Number of iterations for convergence: | 5 |
| Residual norm: | 9.0594e-14 |
| Right-hand side of termination criteria: | 2.8433e-03 |
| i-norm of matrix a: | 2.2000e+01 |

Solution Vector

1.0000 + 9.0000i
2.0000 - 8.0000i
3.0000 + 7.0000i
4.0000 - 6.0000i
5.0000 + 5.0000i
6.0000 - 4.0000i
7.0000 + 3.0000i
8.0000 - 2.0000i
9.0000 + 1.0000i

Residual Vector

1.0e-13 *

-0.0178 + 0.0000i
0.0355 - 0.2842i
-0.0355 + 0.0355i
0.0355 - 0.0711i
-0.0711 + 0.0355i
-0.0711 + 0.0000i
0.0000 + 0.0000i
0.0000 - 0.0711i
0.0000 - 0.1421i
