



Science and
Technology
Facilities Council



GALAHAD

LLST

USER DOCUMENTATION

GALAHAD Optimization Library version 4.0

1 SUMMARY

Given a real m by n matrix \mathbf{A} , a real n by n symmetric diagonally dominant matrix \mathbf{S} a real m vector \mathbf{b} and a scalar $\Delta > 0$, this package finds a **minimizer of the linear least-squares objective function** $\|\mathbf{Ax} - \mathbf{b}\|_2$, **where the vector \mathbf{x} is required to satisfy the constraint** $\|\mathbf{x}\|_{\mathbf{S}} \leq \Delta$, where the \mathbf{S} -norm of \mathbf{x} is $\|\mathbf{x}\|_{\mathbf{S}} = \sqrt{\mathbf{x}^T \mathbf{S} \mathbf{x}}$. This problem commonly occurs as a trust-region subproblem in nonlinear least-squares calculations. The package may also be used to solve the related problem in which \mathbf{x} is instead required to satisfy the **equality constraint** $\|\mathbf{x}\|_{\mathbf{S}} = \Delta$. The matrix \mathbf{S} need not be provided in the commonly-occurring ℓ_2 -trust-region case for which $\mathbf{S} = \mathbf{I}$, the n by n identity matrix.

Factorization of matrices of the form

$$\begin{pmatrix} \lambda \mathbf{S} & \mathbf{A}^T \\ \mathbf{A} & -\mathbf{I} \end{pmatrix} \quad (1.1)$$

for a succession of scalars λ will be required, so this package is most suited for the case where such a factorization may be found efficiently. If this is not the case, the package GALAHAD_LSTR may be preferred.

ATTRIBUTES — Versions: GALAHAD_LLST_single, GALAHAD_LLST_double. **Uses:** GALAHAD_CLOCK, GALAHAD_SYMBOLS, GALAHAD_SPACE, GALAHAD_RAND, GALAHAD_NORMS, GALAHAD_ROOTS, GALAHAD_SPECFILE, GALAHAD_SBLs, GALAHAD_SLS, GALAHAD_IR, GALAHAD_MOP **Date:** October 2008. **Origin:** N. I. M. Gould, Rutherford Appleton Laboratory. **Language:** Fortran 95 + TR 15581 or Fortran 2003. **Parallelism:** Some options may use OpenMP and its runtime library.

2 HOW TO USE THE PACKAGE

Access to the package requires a USE statement such as

Single precision version

```
USE GALAHAD_LLST_single
```

Double precision version

```
USE GALAHAD_LLST_double
```

If it is required to use both modules at the same time, the derived types SMT_TYPE, LLST_control_type, LLST_history_type, LLST_inform_type, LLST_data_type, (Section 2.4) and the subroutines LLST_initialize, LLST_solve, LLST_terminate (Section 2.5) and LLST_read_specfile (Section 2.7) must be renamed on one of the USE statements.

2.1 Matrix storage formats

The matrices \mathbf{A} and (if required) \mathbf{S} may be stored in a variety of input formats.

2.1.1 Dense storage format

The matrix \mathbf{S} is stored as a compact dense matrix by rows, that is, the values of the entries of each row in turn are stored in order within an appropriate real one-dimensional array. Since \mathbf{S} is symmetric, only the lower triangular part (that is the part h_{ij} for $1 \leq j \leq i \leq n$) need be held. In this case the lower triangle should be stored by rows, that is component $i*(i-1)/2 + j$ of the storage array H%val will hold the value h_{ij} (and, by symmetry, h_{ji}) for $1 \leq j \leq i \leq n$. The same is true for \mathbf{S} if it is used. If \mathbf{A} is used, the entire matrix must be supplied, and component $n*(i-1) + j$ of the storage array A%val will hold the value a_{ij} for $i = 1, \dots, m, j = 1, \dots, n$.

All use is subject to the conditions of the GNU Lesser General Public License version 3.
See <http://galahad.rl.ac.uk/galahad-www/cou.html> for full details.

2.1.2 Sparse co-ordinate storage format

Only the nonzero entries of the matrices are stored. For the l -th entry of \mathbf{S} , $1 \leq j \leq i \leq n$, its row index i , column index j and value h_{ij} are stored in the l -th components of the integer arrays `H%row`, `H%col` and real array `H%val`, respectively. Note that only the entries in the lower triangle should be stored. The same scheme may be used for \mathbf{S} if it is required. If \mathbf{A} is used, the entire matrix must be supplied using the same scheme in the integer arrays `A%row`, `A%col` and real array `A%val`.

2.1.3 Sparse row-wise storage format

Again only the nonzero entries are stored, but this time they are ordered so that those in row i appear directly before those in row $i + 1$. For the i -th row of \mathbf{S} , the i -th component of the integer array `A%ptr` holds the position of the first entry in this row, while `A%ptr(m + 1)` holds the total number of entries plus one. The column indices j , $1 \leq j \leq i$, and values h_{ij} of the entries in the i -th row are stored in components $l = \text{H\%ptr}(i), \dots, \text{H\%ptr}(i + 1) - 1$ of the integer array `H%col`, and real array `H%val`, respectively. Note that as before only the entries in the lower triangle should be stored. For sparse matrices, this scheme almost always requires less storage than its predecessor. This scheme may also be used for \mathbf{S} and \mathbf{A} if they are required, excepting that for \mathbf{A} the whole matrix must be stored.

2.1.4 Diagonal storage format

If \mathbf{S} is diagonal (i.e., $h_{ij} = 0$ for all $1 \leq i \neq j \leq n$) only the diagonal entries h_{ii} , $1 \leq i \leq n$, need be stored, and the first n components of the array `H%val` may be used for the purpose. The same applies to \mathbf{S} if it is required. This scheme is inappropriate and thus unavailable for \mathbf{A} .

2.2 Integer kinds

We use the term long `INTEGER` to denote `INTEGER(kind=long)`, where `long = selected_int_kind(18)`.

2.3 OpenMP

OpenMP may be used by the `GALAHAD_LLST` package to provide parallelism for some solver options in shared memory environments. See the documentation for the `GALAHAD` package `SLS` for more details. To run in parallel, OpenMP must be enabled at compilation time by using the correct compiler flag (usually some variant of `-openmp`). The number of threads may be controlled at runtime by setting the environment variable `OMP_NUM_THREADS`.

The code may be compiled and run in serial mode.

2.4 The derived data types

Six derived data types are accessible from the package.

2.4.1 The derived data type for holding matrices

The derived data type `SMT_TYPE` is used to hold the matrices \mathbf{A} and perhaps \mathbf{S} . The components of `SMT_TYPE` used here are:

- `m` is a scalar component of type default `INTEGER`, that holds the number of rows in the matrix.
- `n` is a scalar component of type default `INTEGER`, that holds the number of columns in the matrix.
- `ne` is a scalar variable of type default `INTEGER`, that holds the number of matrix entries.

All use is subject to the conditions of the GNU Lesser General Public License version 3.
See <http://galahad.rl.ac.uk/galahad-www/cou.html> for full details.

- `type` is a rank-one allocatable array of type default `CHARACTER`, that is used to indicate the matrix storage scheme used. Its precise length and content depends on the type of matrix to be stored.
- `val` is a rank-one allocatable array of type default `REAL` (double precision in `GALAHAD_LLST_double`) and dimension at least `ne`, that holds the values of the entries. Each pair of off-diagonal entries $s_{ij} = s_{ji}$ of the *symmetric* matrix **S** is represented as a single entry (see §2.1.1–2.1.3). Any duplicated entries that appear in the sparse co-ordinate or row-wise schemes will be summed.
- `row` is a rank-one allocatable array of type default `INTEGER`, and dimension at least `ne`, that may hold the row indices of the entries. (see §2.1.2).
- `col` is a rank-one allocatable array of type default `INTEGER`, and dimension at least `ne`, that may hold the column indices of the entries (see §2.1.2–2.1.3).
- `ptr` is a rank-one allocatable array of type default `INTEGER`, and dimension at least `n + 1`, that may holds the pointers to the first entry in each row (see §2.1.3).

2.4.2 The derived data type for holding control parameters

The derived data type `LLST_control_type` is used to hold controlling data. Default values may be obtained by calling `LLST_initialize` (see Section 2.5.1). The components of `LLST_control_type` are:

- `error` is a scalar variable of type default `INTEGER`, that holds the stream number for error messages. Printing of error messages in `LLST_solve` and `LLST_terminate` is suppressed if `error` ≤ 0 . The default is `error` = 6.
- `out` is a scalar variable of type default `INTEGER`, that holds the stream number for informational messages. Printing of informational messages in `LLST_solve` is suppressed if `out` < 0 . The default is `out` = 6.
- `print_level` is a scalar variable of type default `INTEGER`, that is used to control the amount of informational output which is required. No informational output will occur if `print_level` ≤ 0 . If `print_level` = 1 a single line of output will be produced for each iteration of the process. If `print_level` ≥ 2 this output will be increased to provide significant detail of each iteration. The default is `print_level` = 0.
- `new_a` is a scalar variable of type default `INTEGER`, that is used to indicate how **A** has changed (if at all) since the previous call to `LLST_solve`. Possible values are:
- 0 **A** is unchanged.
 - 1 the values in **A** have changed, but its nonzero structure is as before.
 - 2 both the values and structure of **A** have changed.

The default is `new_a` = 2.

- `new_s` is a scalar variable of type default `INTEGER`, that is used to indicate how **S** (if required) has changed (if at all) since the previous call to `LLST_solve`. Possible values are:

- 0 **S** is unchanged.
- 1 the values in **S** have changed, but its nonzero structure is as before.
- 2 both the values and structure of **S** have changed.

The default is `new_s` = 2.

- `max_factorizations` is a scalar variable of type default `INTEGER`, that holds the maximum number of factorizations which will be permitted. If `max_factorizations` is set to a negative number, there will be no limit on the number of factorizations allowed. The default is `max_factorizations` = -1.

All use is subject to the conditions of the GNU Lesser General Public License version 3.
See <http://galahad.rl.ac.uk/galahad-www/cou.html> for full details.

`taylor_max_degree` is a scalar variable of type default `INTEGER`, that specifies the maximum degree of Taylor approximant that will be used to approximate the secular function when trying to improve λ ; a first-degree approximant results in Newton's method. The higher the degree, the better in general the improvement, but the larger the cost. Thus there is a balance between many cheap low-degree approximants and a few more expensive higher-degree ones. Our experience favours higher-degree approximants. The default is `taylor_max_degree = 3`, which is the highest degree currently supported.

`initial_multiplier` is a scalar variables of type `REAL` (double precision in `GALAHAD_LLST_double`), that should be set to an initial estimate of the required multiplier λ_* (see Section 4). The algorithm will only use this value if `%use_initial_multiplier` is set `.TRUE.` (see below), and otherwise will be reset by `LLST_solve`. A good initial estimate may sometimes dramatically improve the performance of the package. The default is `initial_multiplier = 0.0`.

`lower` is a scalar variables of type `REAL` (double precision in `GALAHAD_LLST_double`), that holds the value of any known lower bound on the required multiplier λ_* . A good lower bound may sometimes dramatically improve the performance of the package, but an incorrect value might cause the method to fail. Thus resetting `lower` from its default should be used with caution. The default is `lower = - HUGE(1.0)` (`-HUGE(1.0D0)` in `GALAHAD_LLST_double`).

`upper` is a scalar variables of type `REAL` (double precision in `GALAHAD_LLST_double`), that holds the value of any known upper bound on the required multiplier λ_* . A good upper bound may sometimes dramatically improve the performance of the package, but an incorrect value might cause the method to fail. Thus resetting `upper` from its default should be used with caution. The default is `upper = HUGE(1.0)` (`HUGE(1.0D0)` in `GALAHAD_LLST_double`).

`stop_normal` is a scalar variable of type `REAL` (double precision in `GALAHAD_LLST_double`), that hold values for the standard convergence tolerances of the method (see Section 4). In particular, the method is deemed to have converged when the computed solution \mathbf{x} and its multiplier λ satisfy either $\lambda = 0$ and $\|\mathbf{x}\|_S < \Delta$ or $\|\mathbf{x}\|_S - \Delta \leq \text{stop_normal} * \max(1, \Delta)$. The defaults is `stop_normal = $u^{0.75}$` , where u is `EPSILON(1.0)` (`EPSILON(1.0D0)` in `GALAHAD_LLST_double`).

`equality_problem` is a scalar variable of type default `LOGICAL`, that may be set `.TRUE.` if the user requires that the solution occur on the constraint boundary (i.e., that the inequality constraint be replaced by $\|\mathbf{x}\|_S = \Delta$), and `.FALSE.` otherwise. The default is `equality_problem = .FALSE..`

`use_initial_multiplier` is a scalar variable of type default `LOGICAL`, that may be set `.TRUE.` if the user wishes to use the value of initial multiplier supplied in `%initial_multiplier`, and `.FALSE.` if the initial value will be chosen automatically. The default is `use_initial_multiplier = .FALSE..`

`space_critical` is a scalar variable of type default `LOGICAL`, that may be set `.TRUE.` if the user wishes the package to allocate as little internal storage as possible, and `.FALSE.` otherwise. The package may be more efficient if `space_critical` is set `.FALSE..` The default is `space_critical = .FALSE..`

`deallocate_error_fatal` is a scalar variable of type default `LOGICAL`, that may be set `.TRUE.` if the user wishes the package to return to the user in the unlikely event that an internal array deallocation fails, and `.FALSE.` if the package should be allowed to try to continue. The default is `deallocate_error_fatal = .FALSE..`

`definite_linear_solver` is a scalar variable of type default `CHARACTER` and length 30, that specifies the external package to be used to solve any symmetric positive-definite linear system that might arise. Current possible choices are `'sils'`, `'ma27'`, `'ma57'`, `'ma77'`, `'ma86'`, `'ma87'`, `'ma97'`, `'ssids'`, `'pardiso'` and `'wsmp'`. See the documentation for the GALAHAD package SLS for further details. The default is `definite_linear_solver = 'sils'`.

All use is subject to the conditions of the GNU Lesser General Public License version 3.
See <http://galahad.rl.ac.uk/galahad-www/cou.html> for full details.

`prefix` is a scalar variable of type default CHARACTER and length 30, that may be used to provide a user-selected character string to preface every line of printed output. Specifically, each line of output will be prefaced by the string `prefix(2:LEN(TRIM(prefix))-1)`, thus ignoring the first and last non-null components of the supplied string. If the user does not want to preface lines by such a string, they may use the default `prefix = ""`.

`SBLS_control` is a scalar variable of type `SBLS_control_type` that is used to control various aspects of the factorization package SBLS. See the documentation for GALAHAD_SBLS for more details.

`SLS_control` is a scalar variable of type `SLS_control_type` that is used to control various aspects of the factorization package SLS. See the documentation for GALAHAD_SLS for more details.

`IR_control` is a scalar variable of type `IR_control_type` that is used to control various aspects of the iterative refinement package IR. See the documentation for GALAHAD_IR for more details.

2.4.3 The derived data type for holding history information

The derived data type `LLST_history_type` is used to hold the value of $\|\mathbf{x}(\lambda)\|_{\mathbf{S}}$, where $\mathbf{x}(\lambda)$ satisfies $(\mathbf{H} + \lambda\mathbf{S})\mathbf{x}(\lambda) = -\mathbf{c}$ and $\mathbf{A}\mathbf{x}(\lambda) = \mathbf{0}$ for a specific λ arising during the computation. The components of `LLST_history_type` are:

`lambda` is a scalar variable of type default REAL (double precision in GALAHAD_LLST_double), that gives the value λ .

`x_norm` is a scalar variable of type default REAL, that gives the corresponding value $\|\mathbf{x}(\lambda)\|_{\mathbf{S}}$.

2.4.4 The derived data type for holding timing information

The derived data type `LLST_time_type` is used to hold elapsed CPU and system clock times for the various parts of the calculation. The components of `LLST_time_type` are:

`total` is a scalar variable of type default REAL (double precision in GALAHAD_LLST_double), that gives the total CPU time spent in the package.

`assemble` is a scalar variable of type default REAL (double precision in GALAHAD_LLST_double), that gives the CPU time spent assembling the matrix (1.1) from its constituent parts.

`analyse` is a scalar variable of type default REAL (double precision in GALAHAD_LLST_double), that gives the CPU time spent analysing required matrices prior to factorization.

`factorize` is a scalar variable of type default REAL (double precision in GALAHAD_LLST_double), that gives the CPU time spent factorizing the required matrices.

`solve` is a scalar variable of type default REAL (double precision in GALAHAD_LLST_double), that gives the CPU time spent using the factors to solve relevant linear equations.

`clock_total` is a scalar variable of type default REAL (double precision in GALAHAD_LLST_double), that gives the total elapsed system clock time spent in the package.

`clock_assemble` is a scalar variable of type default REAL (double precision in GALAHAD_LLST_double), that gives the elapsed system clock time spent assembling the matrix (1.1) from its constituent parts.

`clock_analyse` is a scalar variable of type default REAL (double precision in GALAHAD_LLST_double), that gives the elapsed system clock time spent analysing required matrices prior to factorization.

`clock_factorize` is a scalar variable of type default REAL (double precision in GALAHAD_LLST_double), that gives the elapsed system clock time spent factorizing the required matrices.

`clock_solve` is a scalar variable of type default REAL (double precision in GALAHAD_LLST_double), that gives the elapsed system clock time spent using the factors to solve relevant linear equations.

All use is subject to the conditions of the GNU Lesser General Public License version 3.
See <http://galahad.rl.ac.uk/galahad-www/cou.html> for full details.

2.4.5 The derived data type for holding informational parameters

The derived data type `LLST_inform_type` is used to hold parameters that give information about the progress and needs of the algorithm. The components of `LLST_inform_type` are:

`status` is a scalar variable of type default `INTEGER`, that gives the current status of the algorithm. See Section 2.6 for details.

`alloc_status` is a scalar variable of type default `INTEGER`, that gives the status of the last internal array allocation or deallocation. This will be 0 if `status = 0`.

`bad_alloc` is a scalar variable of type default `CHARACTER` and length 80, that gives the name of the last internal array for which there were allocation or deallocation errors. This will be the null string if `status = 0`.

`factorizations` is a scalar variable of type default `INTEGER`, that gives the number of factorizations of the matrix (1.1) for different λ , performed during the calculation.

`max_entries_factors` is a scalar variable of type default `INTEGER`, that gives the maximum number of entries in any of the matrix factorizations performed during the calculation.

`len_history` is a scalar variable of type default `INTEGER`, that gives the number of $(\lambda, \|\mathbf{x}(\lambda)\|_{\mathbf{s}})$ pairs encountered during the calculation.

`r_normj` is a scalar variable of type default `REAL` (double precision in `GALAHAD_LLST_double`), that holds the value of the norm of the residual $\|\mathbf{Ax} - \mathbf{b}\|_2$.

`x_norm` is a scalar variable of type default `REAL` (double precision in `GALAHAD_LLST_double`), that holds the value of $\|\mathbf{x}\|_{\mathbf{s}}$.

`multiplier` is a scalar variable of type default `REAL` (double precision in `GALAHAD_LLST_double`), that holds the value of the Lagrange multiplier λ associated with the constraint.

`time` is a scalar variable of type `LLST_time_type` whose components are used to hold elapsed CPU and system clock times for the various parts of the calculation (see Section 2.4.4).

`history` is an array argument of dimension `len_history` and type `LLST_history_type` that contains a list of pairs $(\lambda, \|\mathbf{x}(\lambda)\|_{\mathbf{s}})$ encountered during the calculation (see Section 2.4.3).

`SBLS_inform` is a scalar variable of type `SBLS_inform_type`, that holds informational parameters concerning the analysis, factorization and solution phases performed by the GALAHAD sparse matrix factorization package SBLS. See the documentation for the package SBLS for details of the derived type `SBLS_inform_type`.

`SLS_inform` is a scalar variable of type `SLS_inform_type`, that holds informational parameters concerning the analysis, factorization and solution phases performed by the GALAHAD sparse matrix factorization package SLS. See the documentation for the package SLS for details of the derived type `SLS_inform_type`.

`IR_inform` is a scalar variable of type `IR_inform_type`, that holds informational parameters concerning the iterative refinement subroutine contained in the GALAHAD refinement package IR. See the documentation for the package IR for details of the derived type `IR_inform_type`.

2.4.6 The derived data type for holding problem data

The derived data type `LLST_data_type` is used to hold all the data for a particular problem between calls of LLST procedures. This data should be preserved, untouched, from the initial call to `LLST_initialize` to the final call to `LLST_terminate`.

All use is subject to the conditions of the GNU Lesser General Public License version 3.
See <http://galahad.rl.ac.uk/galahad-www/cou.html> for full details.

2.5 Argument lists and calling sequences

There are three procedures for user calls (see Section 2.7 for further features):

1. The subroutine `LLST_initialize` is used to set default values and initialize private data.
2. The subroutine `LLST_solve` is called to solve the problem.
3. The subroutine `LLST_terminate` is provided to allow the user to automatically deallocate array components of the private data, allocated by `LLST_solve`, at the end of the solution process.

We use square brackets [] to indicate OPTIONAL arguments.

2.5.1 The initialization subroutine

Default values are provided as follows:

```
CALL LLST_initialize( data, control, inform )
```

`data` is a scalar INTENT(INOUT) argument of type `LLST_data_type` (see Section 2.4.6). It is used to hold data about the problem being solved.

`control` is a scalar INTENT(OUT) argument of type `LLST_control_type` (see Section 2.4.2). On exit, `control` contains default values for the components as described in Section 2.4.2. These values should only be changed after calling `LLST_initialize`.

`inform` is a scalar INTENT(OUT) argument of type `LLST_inform_type` (see Section 2.4.5). A successful call to `LLST_initialize` is indicated when the component `status` has the value 0. For other return values of `status`, see Section 2.6.

2.5.2 The optimization problem solution subroutine

The optimization problem solution algorithm is called as follows:

```
CALL LLST_solve( m, n, radius, f, C, H, X, data, control, inform[, M, A] )
```

`m` is a scalar INTENT(IN) argument of type default INTEGER, that must be set to the number of residuals, m .
Restriction: $m > 0$.

`n` is a scalar INTENT(IN) argument of type default INTEGER, that must be set to the number of unknowns, n .
Restriction: $n > 0$.

`radius` is a scalar INTENT(IN) variable of type default REAL (double precision in `GALAHAD_LLST_double`), that must be set on initial entry to the value of the radius of the trust-region constraint, Δ .
Restriction: $\Delta > 0$.

`A` is a scalar INTENT(IN) argument of type `SMT_TYPE` that holds the constraint matrix **A**. The following components are used:

`A%type` is an allocatable array of rank one and type default CHARACTER, that is used to indicate the storage scheme used. If the dense storage scheme (see Section 2.1.1) is used, the first five components of `A%type` must contain the string DENSE. For the sparse co-ordinate scheme (see Section 2.1.2), the first ten components of `A%type` must contain the string COORDINATE, for the sparse row-wise storage scheme (see Section 2.1.3), the first fourteen components of `A%type` must contain the string SPARSE_BY_ROWS.

For convenience, the procedure `SMT_put` may be used to allocate sufficient space and insert the required keyword into `A%type`. For example, if we wish to store **A** using the co-ordinate scheme, we may simply

All use is subject to the conditions of the GNU Lesser General Public License version 3.
See <http://galahad.rl.ac.uk/galahad-www/cou.html> for full details.

```
CALL SMT_put( A%type, 'COORDINATE' )
```

See the documentation for the GALAHAD package SMT for further details on the use of SMT_put.

A%*m* is a scalar variable of type default INTEGER, that holds the number of rows of **A**.

A%*ne* is a scalar variable of type default INTEGER, that holds the number of entries in **A** in the sparse co-ordinate storage scheme (see Section 2.1.2). It need not be set for either of the other schemes.

A%*val* is a rank-one allocatable array of type default REAL (double precision in GALAHAD_LLST_double), that holds the values of the entries of the constraint matrix **A** in any of the storage schemes discussed in Section 2.1.

A%*row* is a rank-one allocatable array of type default INTEGER, that holds the row indices of **A** in the sparse co-ordinate storage scheme (see Section 2.1.2). It need not be allocated for either of the other schemes.

A%*col* is a rank-one allocatable array variable of type default INTEGER, that holds the column indices of **A** in either the sparse co-ordinate (see Section 2.1.2), or the sparse row-wise (see Section 2.1.3) storage scheme. It need not be allocated when the dense storage scheme is used.

A%*ptr* is a rank-one allocatable array of dimension A%*m*+1 and type default INTEGER, that holds the starting position of each row of **A**, as well as the total number of entries plus one, in the sparse row-wise storage scheme (see Section 2.1.3). It need not be allocated when the other schemes are used.

B is an array INTENT(IN) argument of dimension *m* and type default REAL (double precision in GALAHAD_LLST_double), whose *i*-th entry holds the component b_i of the vector **b** in the residual.

X is an array INTENT(OUT) argument of dimension *n* and type default REAL (double precision in GALAHAD_LLST_double), that holds an estimate of the solution **x** of the problem on exit.

data is a scalar INTENT(INOUT) argument of type LLST_data_type (see Section 2.4.6). It is used to hold data about the problem being solved. It must not have been altered **by the user** since the last call to LLST_initialize.

control is a scalar INTENT(IN) argument of type LLST_control_type. (see Section 2.4.2). Default values may be assigned by calling LLST_initialize prior to the first call to LLST_solve.

inform is a scalar INTENT(INOUT) argument of type LLST_inform_type (see Section 2.4.5). On initial entry, the component status must be set to 1. The remaining components need not be set. A successful call to LLST_solve is indicated when the component status has the value 0. For other return values of status, see Section 2.6.

S is an OPTIONAL scalar INTENT(IN) argument of type SMT_TYPE that holds the diagonally dominant scaling matrix **S**. It need only be set if $\mathbf{S} \neq \mathbf{I}$ and in this case the following components are used:

S%*type* is an allocatable array of rank one and type default CHARACTER, that is used to indicate the storage scheme used. If the dense storage scheme (see Section 2.1.1) is used, the first five components of S%*type* must contain the string DENSE. For the sparse co-ordinate scheme (see Section 2.1.2), the first ten components of S%*type* must contain the string COORDINATE, for the sparse row-wise storage scheme (see Section 2.1.3), the first fourteen components of S%*type* must contain the string SPARSE_BY_ROWS, and for the diagonal storage scheme (see Section 2.1.4), the first eight components of S%*type* must contain the string DIAGONAL.

For convenience, the procedure SMT_put may be used to allocate sufficient space and insert the required keyword into S%*type*. For example, if we wish to store **S** using the co-ordinate scheme, we may simply

```
CALL SMT_put( M%type, 'COORDINATE' )
```

See the documentation for the GALAHAD package SMT for further details on the use of SMT_put.

S%*ne* is a scalar variable of type default INTEGER, that holds the number of entries in the **lower triangular** part of **S** in the sparse co-ordinate storage scheme (see Section 2.1.2). It need not be set for any of the other three schemes.

All use is subject to the conditions of the GNU Lesser General Public License version 3.
See <http://galahad.rl.ac.uk/galahad-www/cou.html> for full details.

`S%val` is a rank-one allocatable array of type default `REAL` (double precision in `GALAHAD_LLST_double`), that holds the values of the entries of the **lower triangular** part of the scaling matrix **S** in any of the storage schemes discussed in Section 2.1.

`S%row` is a rank-one allocatable array of type default `INTEGER`, that holds the row indices of the **lower triangular** part of **S** in the sparse co-ordinate storage scheme (see Section 2.1.2). It need not be allocated for any of the other three schemes.

`S%col` is a rank-one allocatable array variable of type default `INTEGER`, that holds the column indices of the **lower triangular** part of **S** in either the sparse co-ordinate (see Section 2.1.2), or the sparse row-wise (see Section 2.1.3) storage scheme. It need not be allocated when the dense or diagonal storage schemes are used.

`S%ptr` is a rank-one allocatable array of dimension $n+1$ and type default `INTEGER`, that holds the starting position of each row of the **lower triangular** part of **S**, as well as the total number of entries plus one, in the sparse row-wise storage scheme (see Section 2.1.3). It need not be allocated when the other schemes are used.

If `S` is absent, the ℓ_2 -norm, $\|\mathbf{x}\|_2 = \sqrt{\mathbf{x}^T \mathbf{x}}$, will be employed.

2.5.3 The termination subroutine

All previously allocated arrays are deallocated as follows:

```
CALL LLST_terminate( data, control, inform )
```

`data` is a scalar `INTENT(INOUT)` argument of type `LLST_data_type` exactly as for `LLST_solve` that must not have been altered **by the user** since the last call to `LLST_initialize`. On exit, array components will have been deallocated.

`control` is a scalar `INTENT(IN)` argument of type `LLST_control_type` exactly as for `LLST_solve`.

`inform` is a scalar `INTENT(OUT)` argument of type `LLST_inform_type` exactly as for `LLST_solve`. Only the component status will be set on exit, and a successful call to `LLST_terminate` is indicated when this component status has the value 0. For other return values of status, see Section 2.6.

2.6 Warning and error messages

A negative value of `inform%status` on exit from `LLST_solve` or `LLST_terminate` indicates that an error might have occurred. No further calls should be made until the error has been corrected. Possible values are:

- 1. An allocation error occurred. A message indicating the offending array is written on unit `control%error`, and the returned allocation status and a string containing the name of the offending array are held in `inform%alloc_status` and `inform%bad_alloc`, respectively.
- 2. A deallocation error occurred. A message indicating the offending array is written on unit `control%error` and the returned allocation status and a string containing the name of the offending array are held in `inform%alloc_status` and `inform%bad_alloc`, respectively.
- 3. (`LLST_solve` only) One of the restrictions $n > 0$, $m > 0$ or $\text{radius} > 0$ has been violated.
- 9. (`LLST_solve` only) The analysis phase of the factorization of the matrix (1.1) failed.
- 10. (`LLST_solve` only) The factorization of the matrix (1.1) failed.
- 15. (`LLST_solve` only) The matrix **S** appears not to be diagonally dominant.

All use is subject to the conditions of the GNU Lesser General Public License version 3.
See <http://galahad.rl.ac.uk/galahad-www/cou.html> for full details.

- 16. (LLST_solve only) The problem is so ill-conditioned that further progress is impossible.
- 18. (LLST_solve only) Too many factorizations have been required. This may happen if `control%max_factorizations` is too small, but may also be symptomatic of a badly scaled problem.

2.7 Further features

In this section, we describe an alternative means of setting control parameters, that is components of the variable `control` of type `LLST_control_type` (see Section 2.4.2), by reading an appropriate data specification file using the subroutine `LLST_read_specfile`. This facility is useful as it allows a user to change LLST control parameters without editing and recompiling programs that call LLST.

A specification file, or `specfile`, is a data file containing a number of "specification commands". Each command occurs on a separate line, and comprises a "keyword", which is a string (in a close-to-natural language) used to identify a control parameter, and an (optional) "value", which defines the value to be assigned to the given control parameter. All keywords and values are case insensitive, keywords may be preceded by one or more blanks but values must not contain blanks, and each value must be separated from its keyword by at least one blank. Values must not contain more than 30 characters, and each line of the `specfile` is limited to 80 characters, including the blanks separating keyword and value.

The portion of the specification file used by `LLST_read_specfile` must start with a "BEGIN LLST" command and end with an "END" command. The syntax of the `specfile` is thus defined as follows:

```
( .. lines ignored by LLST_read_specfile .. )
BEGIN LLST
  keyword    value
  .....     .....
  keyword    value
END
( .. lines ignored by LLST_read_specfile .. )
```

where `keyword` and `value` are two strings separated by (at least) one blank. The "BEGIN LLST" and "END" delimiter command lines may contain additional (trailing) strings so long as such strings are separated by one or more blanks, so that lines such as

```
BEGIN LLST SPECIFICATION
```

and

```
END LLST SPECIFICATION
```

are acceptable. Furthermore, between the "BEGIN LLST" and "END" delimiters, specification commands may occur in any order. Blank lines and lines whose first non-blank character is `!` or `*` are ignored. The content of a line after a `!` or `*` character is also ignored (as is the `!` or `*` character itself). This provides an easy manner to "comment out" some specification commands, or to comment specific values of certain control parameters.

The value of a control parameters may be of three different types, namely integer, logical or real. Integer and real values may be expressed in any relevant Fortran integer and floating-point formats (respectively). Permitted values for logical parameters are "ON", "TRUE", ".TRUE.", "T", "YES", "Y", or "OFF", "NO", "N", "FALSE", ".FALSE." and "F". Empty values are also allowed for logical control parameters, and are interpreted as "TRUE".

The specification file must be open for input when `LLST_read_specfile` is called, and the associated device number passed to the routine in `device` (see below). Note that the corresponding file is REWINDed, which makes it possible to combine the specifications for more than one program/routine. For the same reason, the file is not closed by `LLST_read_specfile`.

All use is subject to the conditions of the GNU Lesser General Public License version 3.
See <http://galahad.rl.ac.uk/galahad-www/cou.html> for full details.

Control parameters corresponding to the components `SLS_control` and `IR_control` may be changed by including additional sections enclosed by “BEGIN SLS” and “END SLS”, and “BEGIN IR” and “END IR”, respectively. See the specification sheets for the packages `GALAHAD_SLS` and `GALAHAD_IR` for further details.

2.7.1 To read control parameters from a specification file

Control parameters may be read from a file as follows:

```
CALL LLST_read_specfile( control, device )
```

`control` is a scalar `INTENT(INOUT)` argument of type `LLST_control_type` (see Section 2.4.2). Default values should have already been set, perhaps by calling `LLST_initialize`. On exit, individual components of `control` may have been changed according to the commands found in the specfile. Specfile commands and the component (see Section 2.4.2) of `control` that each affects are given in Table 2.1.

command	component of control	value type
error-printout-device	%error	integer
printout-device	%out	integer
print-level	%print_level	integer
use-dense-factorization	%dense_factorization	integer
has-a-changed	%new_a	integer
has-s-changed	%new_s	integer
factorization-limit	%max_factorizations	integer
max-degree-taylor-approximant	%taylor_max_degree	integer
initial-multiplier	%initial_multiplier	real
lower-bound-on-multiplier	%lower	real
upper-bound-on-multiplier	%upper	real
stop-normal-case	%stop_normal	real
use-initial-multiplier	%use_initial_multiplier	logical
initialize-approximate-eigenvector	%initialize_approx_eigenvector	real
equality-problem	%equality_problem	logical
space-critical	%space_critical	logical
deallocate-error-fatal	%deallocate_error_fatal	logical
definite-linear-equation-solver	%definite_linear_solver	character
output-line-prefix	%prefix	character

Table 2.1: Specfile commands and associated components of `control`.

`device` is a scalar `INTENT(IN)` argument of type `default_INTEGER`, that must be set to the unit number on which the specfile has been opened. If `device` is not open, `control` will not be altered and execution will continue, but an error message will be printed on unit `control%error`.

2.8 Information printed

If `control%print_level` is positive, information about the progress of the algorithm will be printed on unit `control%out`. If `control%print_level = 1`, a single line of output will be produced for each iteration of the process. In the first phase of the algorithm, this will include the current estimate of the multiplier and known brackets on its optimal value. In the second phase, the residual $\|x\|_S - \Delta$, the current estimate of the multiplier and the size of the correction will be printed. If `control%print_level ≥ 2` , this output will be increased to provide significant detail of each iteration. This extra output includes times for various phases.

All use is subject to the conditions of the GNU Lesser General Public License version 3.
See <http://galahad.rl.ac.uk/galahad-www/cou.html> for full details.

3 GENERAL INFORMATION

Use of common: None.

Workspace: Provided automatically by the module.

Other routines called directly: None.

Other modules used directly: LLST_solve calls the GALAHAD packages GALAHAD_CLOCK, GALAHAD_SYMBOLS, GALAHAD_SPACE, GALAHAD_RAND, GALAHAD_NORMS, GALAHAD_ROOTS, GALAHAD_SPECFILE, GALAHAD_SBLs, GALAHAD_SLS, GALAHAD_IR and GALAHAD_MOP.

Input/output: Output is under control of the arguments control%error, control%out and control%print_level.

Restrictions: $n > 0$, $\Delta > 0$.

Portability: ISO Fortran 95 + TR 15581 or Fortran 2003. The package is thread-safe.

4 METHOD

The required solution \mathbf{x}_* necessarily satisfies the optimality condition $\mathbf{A}^T \mathbf{x}_* + \lambda_* \mathbf{S} \mathbf{x}_* = \mathbf{A}^T \mathbf{b}$, where $\lambda_* \geq 0$ is a Lagrange multiplier corresponding to the constraint $\|\mathbf{x}\|_{\mathbf{S}} \leq \Delta$; for the equality-constrained problem $\|\mathbf{x}\|_{\mathbf{S}} = \Delta$ and the multiplier is unconstrained.

The method is iterative, and proceeds in two phases. Firstly, lower and upper bounds, λ_L and λ_U , on λ_* are computed using Gershgorin's theorems and other eigenvalue bounds, including those that may involve the Cholesky factorization of \mathbf{S} . The first phase of the computation proceeds by progressively shrinking the bound interval $[\lambda_L, \lambda_U]$ until a value λ for which $\|\mathbf{x}(\lambda)\|_{\mathbf{S}} \geq \Delta$ is found. Here $\mathbf{x}(\lambda)$ and its companion $\mathbf{y}(\lambda)$ are defined to be a solution of

$$(\mathbf{A}^T \mathbf{A} + \lambda \mathbf{S}) \mathbf{x}(\lambda) = \mathbf{A}^T \mathbf{b}; \quad (4.1)$$

along the way the possibility that $\|\mathbf{x}(0)\|_{\mathbf{S}} \leq \Delta$ is examined, and if this transpires the process is terminated with $\mathbf{x}_* = \mathbf{x}(0)$. Once the terminating λ from the first phase has been discovered, the second phase consists of applying Newton or higher-order iterations to the nonlinear "secular" equation $\|\mathbf{x}(\lambda)\|_{\mathbf{S}} = \Delta$ with the knowledge that such iterations are both globally and ultimately rapidly convergent.

The dominant cost is the requirement that we solve a sequence of linear systems (4.1). This may be rewritten as

$$\begin{pmatrix} \lambda \mathbf{S} & \mathbf{A}^T \\ \mathbf{A} & -\mathbf{I} \end{pmatrix} \begin{pmatrix} \mathbf{x}(\lambda) \\ \mathbf{y}(\lambda) \end{pmatrix} = \begin{pmatrix} \mathbf{A}^T \mathbf{b} \\ \mathbf{0} \end{pmatrix}, \quad (4.2)$$

for some auxiliary vector $\mathbf{y}(\lambda)$. In general a sparse symmetric, indefinite factorization of the coefficient matrix of (4.2) is often preferred to a Cholesky factorization of that of (1.1).

Reference: The method is the obvious specification to the linear least-squares problem of that described in detail in

H. S. Dollar, N. I. M. Gould and D. P. Robinson. On solving trust-region and other regularised subproblems in optimization. *Mathematical Programming Computation* **2(1)** (2010) 21–57.

All use is subject to the conditions of the GNU Lesser General Public License version 3.
See <http://galahad.rl.ac.uk/galahad-www/cou.html> for full details.


```
WRITE( 6, "( ' LLST_solve exit status = ', IO ) " ) inform%status
END IF
CALL LLST_terminate( data, control, inform ) ! delete workspace
DEALLOCATE( A%row, A%col, A%val, S%val, A%type, S%type )
END PROGRAM GALAHAD_LLST_EXAMPLE
```

This produces the following output:

```
it      ||x||_S-radius      lambda      d_lambda
L  1  2.041447629184912E+05  0.000000000000000E+00  0.000000000000000E+00
L  2  1.146581773441024E+02  1.360692972795873E-03  1.360692972795873E-03
L  3  5.390346878359741E+00  2.221094114131669E-02  2.085024816852081E-02
L  4  1.658471170973428E-01  3.862126136675395E-02  1.641032022543727E-02
L  5  2.374899743529113E-07  3.947801461809867E-02  8.567532513447176E-04
G  6  4.973799150320701E-14  3.947801586871794E-02  1.250619267867403E-09
Normal stopping criteria satisfied
||x||_S recurred and calculated =  1.00000000E+01  1.00000000E+01
||Ax-b||_2 recurred and calculated =  7.06303867E+01  7.06303867E+01
```

All use is subject to the conditions of the GNU Lesser General Public License version 3.
See <http://galahad.rl.ac.uk/galahad-www/cou.html> for full details.