



Science and  
Technology  
Facilities Council



# GALAHAD

# SLLS

USER DOCUMENTATION

GALAHAD Optimization Library version 5.5

## 1 SUMMARY

This package uses a preconditioned, projected-gradient method to solve the **simplex-constrained regularized linear least-squares problem**

$$\text{minimize } q(\mathbf{x}) = \frac{1}{2} \|\mathbf{A}_o \mathbf{x} - \mathbf{b}\|_W^2 + \frac{1}{2} \sigma \|\mathbf{x} - \mathbf{x}_s\|_2^2,$$

where the variables  $\mathbf{x}$  are required to lie within the **regular simplex**

$$\mathbf{e}^T \mathbf{x} = 1 \text{ and } \mathbf{x} \geq 0,$$

or an intersection of **multiple non-overlapping regular simplices**

$$\mathbf{e}_{C_i}^T \mathbf{x}_{C_i} = 1 \text{ and } \mathbf{x}_{C_i} \geq 0 \text{ for } i = 1, \dots, m,$$

where the  $o$  by  $n$  real **design matrix**  $\mathbf{A}_o$ , the vector  $\mathbf{b}$  of **observations**, the **shifts**  $\mathbf{x}_s$ , and the non-negative **weights**  $\mathbf{w}$  and  $\sigma$  are given,  $\mathbf{e}$  is the vector of ones, the vector  $\mathbf{v}_C$  is made up of those entries of  $\mathbf{v}$  indexed by the set  $C$ , the index sets of **cohorts**  $C_i \subseteq \{1, \dots, n\}$  for which  $C_i \cap C_j = \emptyset$  for  $1 \leq i, j \leq m$ , and where the Euclidean and weighted-Euclidean norms are given by  $\|\mathbf{v}\|_2^2 = \mathbf{v}^T \mathbf{v}$  and  $\|\mathbf{v}\|_W^2 = \mathbf{v}^T \mathbf{W} \mathbf{v}$ , respectively, with  $\mathbf{W} = \text{diag}(\mathbf{w})$ . Full advantage is taken of any zero coefficients of the Jacobian matrix  $\mathbf{A}_o$  of the **residuals**  $\mathbf{r}(\mathbf{x}) = \mathbf{A}_o \mathbf{x} - \mathbf{b}$ ; the matrix need not be provided as there are options to obtain matrix-vector products involving  $\mathbf{A}_o$  and its transpose either by reverse communication or from a user-provided subroutine.

**ATTRIBUTES — Versions:** GALAHAD\_SLLS\_single, GALAHAD\_SLLS\_double. **Uses:** GALAHAD\_CPU\_time, GALAHAD\_SYMBOLS, GALAHAD\_SPACE, GALAHAD\_STRING, GALAHAD\_SORT, GALAHAD\_NORMS, GALAHAD\_CONVERT, GALAHAD\_SBLS, GALAHAD\_QPT, GALAHAD\_QPD, GALAHAD\_USERDATA, GALAHAD\_REVERSE, GALAHAD\_SPECFILE. **Date:** July 2022. **Origin:** N. I. M. Gould, STFC-Rutherford Appleton Laboratory. **Language:** Fortran 95 + TR 15581 or Fortran 2003.

## 2 HOW TO USE THE PACKAGE

The package is available with single, double and (if available) quadruple precision reals, and either 32-bit or 64-bit integers. Access to the 32-bit integer, single precision version requires the USE statement

```
USE GALAHAD_SLLS_single
```

with the obvious substitution GALAHAD\_SLLS\_double, GALAHAD\_SLLS\_quadruple, GALAHAD\_SLLS\_single\_64, GALAHAD\_SLLS\_double\_64 and GALAHAD\_SLLS\_quadruple\_64 for the other variants.

If it is required to use more than one of the modules at the same time, the derived types SMT\_type, QPT\_problem\_type, USERDATA\_type, SLLS\_time\_type, SLLS\_control\_type, SLLS\_inform\_type and SLLS\_data\_type (Section 2.3) and the subroutines SLLS\_initialize, SLLS\_solve, SLLS\_terminate, (Section 2.4) and SLLS\_read\_specfile (Section 2.8) must be renamed on one of the USE statements.

### 2.1 Matrix storage formats

When it is explicitly available, an  $m$  by  $n$  matrix  $\mathbf{A}$  (in our case  $\mathbf{A}_o$ ) may be stored in a variety of input formats.

**All use is subject to the conditions of a BSD-3-Clause License.**  
See <https://www.galahad.rl.ac.uk/download/> for full details.

### 2.1.1 Dense row-wise storage format

The matrix  $\mathbf{A}$  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. Component  $n * (i - 1) + j$  of the storage array `A%val` will hold the value  $\mathbf{A}_{i,j}$  for  $i = 1, \dots, m, j = 1, \dots, n$ .

### 2.1.2 Dense column-wise storage format

The matrix  $\mathbf{A}$  is stored as a compact dense matrix by columns, that is, the values of the entries of each column in turn are stored in order within an appropriate real one-dimensional array. Component  $m * (j - 1) + i$  of the storage array `A%val` will hold the value  $\mathbf{A}_{i,j}$  for  $i = 1, \dots, m, j = 1, \dots, n$ .

### 2.1.3 Sparse coordinate storage format

Only the nonzero entries of the matrices are stored. For the  $l$ -th entry of  $\mathbf{A}$ , its row index  $i$ , column index  $j$  and value  $\mathbf{A}_{i,j}$  are stored in the  $l$ -th components of the integer arrays `A%row`, `A%col` and real array `A%val`. The order is unimportant, but the total number of entries `A%ne` is required.

### 2.1.4 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  $\mathbf{A}$ , 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  $j$  and values  $\mathbf{A}_{i,j}$  of the entries in the  $i$ -th row are stored in components  $l = \text{A\%ptr}(i), \dots, \text{A\%ptr}(i + 1) - 1$  of the integer array `A%col`, and real array `A%val`, respectively.

### 2.1.5 Sparse column-wise storage format

Yet again only the nonzero entries are stored, but this time they are ordered so that those in column  $j$  appear directly before those in column  $j + 1$ . For the  $j$ -th column of  $\mathbf{A}$ , the  $j$ -th component of the integer array `A%ptr` holds the position of the first entry in this column, while `A%ptr(n + 1)` holds the total number of entries plus one. The row indices  $i$  and values  $\mathbf{A}_{i,j}$  of the entries in the  $j$ -th column are stored in components  $l = \text{A\%ptr}(j), \dots, \text{A\%ptr}(j + 1) - 1$  of the integer array `A%row`, and real array `A%val`, respectively.

For sparse matrices, the row- and column-wise storage schemes almost always requires less storage than their predecessor.

## 2.2 Real and integer kinds

We use the terms integer and real to refer to the fortran keywords `REAL(rp_)` and `INTEGER(ip_)`, where `rp_` and `ip_` are the relevant kind values for the real and integer types employed by the particular module in use. The former are equivalent to default `REAL` for the single precision versions, `DOUBLE PRECISION` for the double precision cases and quadruple-precision if 128-bit reals are available, and correspond to `rp_ = real32`, `rp_ = real64` and `rp_ = real128` respectively as defined by the fortran `iso_fortran_env` module. The latter are default (32-bit) and long (64-bit) integers, and correspond to `ip_ = int32` and `ip_ = int64`, respectively, again from the `iso_fortran_env` module.

## 2.3 The derived data types

Ten derived data types are accessible from the package.

---

**All use is subject to the conditions of a BSD-3-Clause License.**  
**See <https://www.galahad.rl.ac.uk/download/> for full details.**

### 2.3.1 The derived data type for holding matrices

The derived data type `SMT_TYPE` is used to hold the matrix  $\mathbf{A}$ . The components of `SMT_TYPE` used here are:

- `m` is a scalar component of type `INTEGER(ip_)`, that holds the number of rows in the matrix.
- `n` is a scalar component of type `INTEGER(ip_)`, that holds the number of columns in the matrix.
- `ne` is a scalar variable of type `INTEGER(ip_)`, that holds the number of matrix entries.
- `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 (see §2.3.2).
- `val` is a rank-one allocatable array of type `REAL(rp_)` and dimension at least `ne`, that holds the values of the entries. Any duplicated entries that appear in the sparse coordinate or row-wise schemes will be summed.
- `row` is a rank-one allocatable array of type `INTEGER(ip_)`, and dimension at least `ne`, that may hold the row indices of the entries. (see §2.1.3 and §2.1.5).
- `col` is a rank-one allocatable array of type `INTEGER(ip_)`, and dimension at least `ne`, that may hold the column indices of the entries (see §2.1.3–2.1.4).
- `ptr` is a rank-one allocatable array of type `INTEGER(ip_)`, and dimension at least `m + 1`, that may hold the pointers to the first entry in each row (see §2.1.4). or dimension at least `n + 1`, that may hold the pointers to the first entry in each column (see §2.1.5).

### 2.3.2 The derived data type for holding the problem

The derived data type `QPT_problem_type` is used to hold the problem. The components of `QPT_problem_type` are:

- `n` is a scalar variable of type `INTEGER(ip_)`, that holds the number of optimization variables,  $n$ .
- `o` is a scalar variable of type `INTEGER(ip_)`, that holds the number of residuals,  $o$ .
- `m` is a scalar variable of type `INTEGER(ip_)`, that holds the number of cohorts,  $m$ .
- `regularization_weight` is a scalar variable of type `REAL(rp_)`, that holds the value of the regularization weight,  $\sigma$ , if required. By default `regularization_weight` is set to zero.
- `Ao` is scalar variable of type `SMT_TYPE` that holds the design matrix  $\mathbf{A}_o$  (if it is available). The following components are used here:

`Ao%type` is an allocatable array of rank one and type default `CHARACTER`, that is used to indicate the storage scheme used. If the dense row-wise storage scheme (see Section 2.1.1) is used, the first twelve components of `Ao%type` must contain the string `DENSE_BY_ROWS`, while if the column-wise scheme (see Section 2.1.2) is used, the fifteen components of `Ao%type` must contain the string `DENSE_BY_COLUMNS`. For the sparse coordinate scheme (see Section 2.1.3), the first ten components of `Ao%type` must contain the string `COORDINATE`, for the sparse row-wise storage scheme (see Section 2.1.4), the first fourteen components of `Ao%type` must contain the string `SPARSE_BY_ROWS`. and for the sparse column-wise storage scheme (see Section 2.1.5), the first seventeen components of `Ao%type` must contain the string `SPARSE_BY_COLUMNS`.

For convenience, the procedure `SMT_put` may be used to allocate sufficient space and insert the required keyword into `Ao%type`. For example, if `nlp` is of derived type `SLLS_problem_type` and involves a design matrix we wish to store using the coordinate scheme, we may simply

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

---

**All use is subject to the conditions of a BSD-3-Clause License.**  
**See <https://www.galahad.rl.ac.uk/download/> for full details.**

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

Ao%type should **not** be allocated if  $\mathbf{A}_o$  is unavailable (and access provided to  $\mathbf{A}_o$  provided by other means, see later sections), and in this case the remaining components of  $\mathbf{A}$  need not be set.

Ao%ne is a scalar variable of type INTEGER(ip\_), that holds the number of entries in  $\mathbf{A}_o$  in the sparse coordinate storage scheme (see Section 2.1.3). It need not be set for any of the other three schemes.

Ao%val is a rank-one allocatable array of type REAL(rp\_), that holds the values of the entries of the design matrix  $\mathbf{A}_o$  in any of the storage schemes discussed in Section 2.1.

Ao%row is a rank-one allocatable array of type INTEGER(ip\_), that holds the row indices of  $\mathbf{A}_o$  in the sparse coordinate storage scheme (see Section 2.1.3). It need not be allocated for any of the other three schemes.

Ao%col is a rank-one allocatable array variable of type INTEGER(ip\_), that holds the column indices of  $\mathbf{A}_o$  in either the sparse coordinate (see Section 2.1.3), or the sparse row-wise (see Section 2.1.4) storage scheme. It need not be allocated when the dense or column-wise storage schemes are used.

Ao%row is a rank-one allocatable array variable of type INTEGER(ip\_), that holds the row indices of  $\mathbf{A}_o$  in either the sparse coordinate (see Section 2.1.3), or the sparse column-wise (see Section 2.1.5) storage scheme. It need not be allocated when the dense or row-wise storage schemes are used.

Ao%ptr is a rank-one allocatable array and type INTEGER(ip\_), that must be of dimension o+1 and hold the starting position of each row of  $\mathbf{A}_o$ , as well as the total number of entries plus one, in the sparse row-wise storage scheme (see Section 2.1.4). If the sparse column-wise storage scheme (see Section 2.1.5) is used, it must instead be of dimension n+1 and hold the starting position of each column of  $\mathbf{A}_o$ , as well as the total number of entries plus one. It need not be allocated when the other schemes are used.

B is a rank-one allocatable array of dimension o and type REAL(rp\_), that holds the vector of observations  $\mathbf{b}$  in the residuals. The  $i$ -th component of B,  $i = 1, \dots, o$ , contains  $b_i$ .

COHORT is a rank-one allocatable array of dimension n and type INTEGER(ip\_), that specifies with which cohort variable  $x_j$ ,  $j = 1, \dots, n$ , is associated. If variable  $x_j$  is associated with cohort  $C_i$ ,  $1 \leq i \leq m$ , COHORT(j) should be set to i, while if  $x_j$  is unconstrained COHORT(j) = 0 should be assigned. At least one value COHORT(j) for  $j = 1, \dots, n$  is expected to take the value i for every  $1 \leq i \leq m$ , that is no empty cohorts are allowed. If COHORT is not allocated, a single cohort containing all of the variables will be used instead, i.e., a single unit simplex is employed.

W is a rank-one allocatable array of dimension o and type default REAL(rp\_), that holds the values of the weights,  $\mathbf{w}$ . The  $i$ -th component of W,  $i = 1, \dots, o$ , contains  $w_i$ . If W is not allocated, weights of one will be assumed.

X\_s is a rank-one allocatable array of dimension n and type default REAL(rp\_), that holds the values of the shifts,  $\mathbf{x}_s$ . The  $j$ -th component of X\_s,  $j = 1, \dots, n$ , contains  $x_{sj}$ . If X\_s is not allocated, shifts of zero will be assumed.

X is a rank-one allocatable array of dimension n and type REAL(rp\_), that holds the values  $\mathbf{x}$  of the optimization variables. The  $j$ -th component of X,  $j = 1, \dots, n$ , contains  $x_j$ .

Y is a rank-one allocatable array of dimension m (or 1 if a single unit simplex is implied) and type default REAL(rp\_), that holds the values  $\mathbf{y}$  of estimates of the Lagrange multipliers corresponding to the equality constraint for each cohort (see Section 4). The  $i$ -th component of Y,  $i = 1, \dots, m$ , contains  $y_i$ .

Z is a rank-one allocatable array of dimension n and type default REAL(rp\_), that holds the values  $\mathbf{z}$  of estimates of the dual variables corresponding to the non-negativity constraints (see Section 4). The  $j$ -th component of Z,  $j = 1, \dots, n$ , contains  $z_j$ .

R is a rank-one allocatable array of dimension o and type REAL(rp\_), that holds the residuals,  $\mathbf{r}(\mathbf{x}) = \mathbf{A}_o\mathbf{x} - \mathbf{b}$ , at the point  $\mathbf{x}$ . The  $i$ -th component of R,  $i = 1, \dots, o$ , contains  $r_i(\mathbf{x})$ .

---

**All use is subject to the conditions of a BSD-3-Clause License.**  
**See <https://www.galahad.rl.ac.uk/download/> for full details.**

**G** is a rank-one allocatable array of dimension  $n$  and type `REAL(rp_)`, that holds the gradient of the objective  $\mathbf{g}(\mathbf{x}) = \mathbf{A}_o^T \mathbf{W} \mathbf{r}(\mathbf{x})$ , at the point  $\mathbf{x}$ . The  $j$ -th component of **G**,  $j = 1, \dots, n$ , contains  $\mathbf{g}_j(\mathbf{x})$ .

**X\_status** is a rank-one allocatable array argument of dimension  $n$  and type `INTEGER(ip_)`, that indicates which of the non-negativity constraints are in the current active set. Possible values for **X\_status**( $j$ ),  $j = 1, \dots, n$ , and their meanings are

<0 the  $j$ -th non-negativity constraint is in the active set, at the value zero, and

0 the  $j$ -th non-negativity constraint is not in the active set, or the associated variable is unconstrained.

When `control%crossover` is `.TRUE.`, more specific values are

-1 the  $j$ -th non-negativity constraint is both independent and active with the value zero,

-2 the  $j$ -th non-negativity constraint is at the value zero but linearly dependent on others,

0 the  $j$ -th non-negativity constraint is not in the active set, or the associated variable is unconstrained.

### 2.3.3 The derived data type for holding control parameters

The derived data type `SLLS_control_type` is used to hold controlling data. Default values may be obtained by calling `SLLS_initialize` (see Section 2.4.1), while components may also be changed by calling `SLLS_read_specfile` (see Section 2.8.1). The components of `SLLS_control_type` are:

**error** is a scalar variable of type `INTEGER(ip_)`, that holds the stream number for error messages. Printing of error messages in `SLLS_solve` and `SLLS_terminate` is suppressed if `error`  $\leq 0$ . The default is `error` = 6.

**out** is a scalar variable of type `INTEGER(ip_)`, that holds the stream number for informational messages. Printing of informational messages in `SLLS_solve` is suppressed if `out` < 0. The default is `out` = 6.

**print\_level** is a scalar variable of type `INTEGER(ip_)`, that is used to control the amount of informational output which is required. No informational output will occur if `print_level`  $\leq 0$ . If `print_level` = 1, a single line of output will be produced for each iteration of the process. If `print_level`  $\geq 2$ , this output will be increased to provide significant detail of each iteration. The default is `print_level` = 0.

**start\_print** is a scalar variable of type `INTEGER(ip_)`, that specifies the first iteration for which printing will be permitted in `GALAHAD_SLLS_solve`. If `start_print` is negative, printing will be permitted from the outset. The default is `start_print` = -1.

**stop\_print** is a scalar variable of type `INTEGER(ip_)`, that specifies the last iteration for which printing will be permitted in `GALAHAD_SLLS_solve`. If `stop_print` is negative, printing will be permitted once it has been started by `start_print`. The default is `stop_print` = -1.

**print\_gap** is a scalar variable of type `INTEGER(ip_)`. Once printing has been started, output will occur once every `print_gap` iterations. If `print_gap` is no larger than 1, printing will be permitted on every iteration. The default is `print_gap` = 1.

**maxit** is a scalar variable of type `INTEGER(ip_)`, that holds the maximum number of iterations which will be allowed in `GALAHAD_SLLS_solve`. The default is `maxit` = 1000.

**cold\_start** is a scalar variable of type `INTEGER(ip_)`, that should be set to 0 if a warm start is required (with variables assigned according to `X_stat`, see below), and to any other value if the values given in `prob%X` suffice. The default is `cold_start` = 1.

---

**All use is subject to the conditions of a BSD-3-Clause License.**  
**See <https://www.galahad.rl.ac.uk/download/> for full details.**

`preconditioner` is a scalar variable of type `INTEGER(ip_)`, that specifies the type of preconditioner (scaling) used when computing the search direction. Currently this can be 0 (no preconditioner), 1 (a diagonal preconditioner that normalises the rows of  $A_o$ ) or 2 (a preconditioner supplied by the user either via a subroutine call of `eval_PREC`, see §2.5.4, or via reverse communication, see §2.6), although other values may be introduced in future. The default is `preconditioner = 1`.

`change_max` is a scalar variable of type `INTEGER(ip_)`, that specifies the maximum number of per-iteration changes in the working set permitted when allowing subspace solution rather than steepest descent (see §4). The default is `change_max = 2`.

`cg_maxit` is a scalar variable of type `INTEGER(ip_)`, that holds the maximum number of conjugate-gradient iterations which will be allowed per main iteration in `GALAHAD_SLLS_solve`. The default is `cg_maxit = 1000`, and any negative value will be interpreted as  $n + 1$ .

`arcsearch_max_steps` is a scalar variable of type `INTEGER(ip_)`, that specifies the maximum number of steps allowed in an individual piecewise arc search. The default is `arcsearch_max_steps = 1000`, and a negative value removes the limit.

`stop_d` is a scalar variable of type default `REAL(rp_)`, that holds the required accuracy for the dual infeasibility (see Section 4). The default is `stop_d =  $u^{1/3}$` , where  $u$  is `EPSILON(1.0)` (`EPSILON(1.0D0)` in `GALAHAD_SLLS_double`).

`stop_cg_relative` and `stop_cg_absolute` are scalar variables of type `REAL(rp_)`, that hold the relative and absolute convergence tolerances for the conjugate-gradient iteration that occurs in the face of currently-active constraints when constructing the search direction. `stop_cg_relative = 0.01` and `stop_cg_absolute =  $\sqrt{u}$` , where  $u$  is `EPSILON(1.0)` (`EPSILON(1.0D0)` in `GALAHAD_SLLS_double`).

`alpha_max` is a scalar variable of type `REAL(rp_)`, that specifies The default is `alpha_max =  $10^{20}$` .

`alpha_initial` is a scalar variable of type `REAL(rp_)`, that specifies the initial arc length during the inexact piecewise arc search. The default is `alpha_initial = 1.0`.

`alpha_reduction` is a scalar variable of type `REAL(rp_)`, that specifies the arc length reduction factor for the inexact piecewise arc search. The default is `alpha_reduction = 0.5`.

`arcsearch_acceptance_tol` is a scalar variable of type `REAL(rp_)`, that specifies the required relative reduction during the inexact arc search The default is `arcsearch_acceptance_tol = 0.01`.

`stabilisation_weight` is a scalar variable of type `REAL(rp_)`, that specifies the weight added to the search-direction subproblem to stabilise the computation. The default is `stabilisation_weight =  $10^{-12}$` .

`cpu_time_limit` is a scalar variable of type `REAL(rp_)`, that is used to specify the maximum permitted CPU time. Any negative value indicates no limit will be imposed. The default is `cpu_time_limit = - 1.0`.

`direct_subproblem_solve` is a scalar variable of type default `LOGICAL`, that should be set `.TRUE.` if the search over a promising subspace is carried out using matrix factorization, and `.FALSE.` if the conjugate-gradient least-squares (CGLS) method is to be preferred. The former is generally more effective so long as the cost of factorization isn't exorbitant. The default is `direct_subproblem_solve = .TRUE.`, but the package will override this if the Jacobian is not explicitly available.

`exact_arc_search` is a scalar variable of type default `LOGICAL`, that should be set `.TRUE.` if the exact minimizer along the search arc is required, and `.FALSE.` if an approximation found by backtracking or advancing suffices. The default is `exact_arc_search = .TRUE.`.

---

**All use is subject to the conditions of a BSD-3-Clause License.**  
**See <https://www.galahad.rl.ac.uk/download/> for full details.**

`space_critical` is a scalar variable of type default LOGICAL, that must be set `.TRUE.` if space is critical when allocating arrays and `.FALSE.` otherwise. The package may run faster if `space_critical` is `.FALSE.` but at the possible expense of a larger storage requirement. The default is `space_critical = .FALSE..`

`deallocate_error_fatal` is a scalar variable of type default LOGICAL, that must be set `.TRUE.` if the user wishes to terminate execution if a deallocation fails, and `.FALSE.` if an attempt to continue will be made. The default is `deallocate_error_fatal = .FALSE..`

`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` whose components are used to control the factorization and/or preconditioner used, performed by the package `GALAHAD_SBLs`. See the documentation for `GALAHAD_SBLs` for further details.

### 2.3.4 The derived data type for holding timing information

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

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

`clock_total` is a scalar variable of type default REAL, that gives the total clock time spent in the package.

### 2.3.5 The derived data type for holding informational parameters

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

`status` is a scalar variable of type `INTEGER(ip_)`, that gives the exit status of the algorithm. See Section 2.7 for details.

`alloc_status` is a scalar variable of type `INTEGER(ip_)`, that gives the status of the last attempted 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`.

`factorization_status` is a scalar variable of type `INTEGER(ip_)`, that gives the return status from the matrix factorization.

`iter` is a scalar variable of type `INTEGER(ip_)`, that gives the number of iterations performed.

`obj` is a scalar variable of type `REAL(rp_)`, that holds the value of the regularized objective function  $q(\mathbf{x})$  at the best estimate of the solution found.

`ls_obj` is a scalar variable of type `REAL(rp_)`, that holds the value of the least-squares function  $\frac{1}{2}\|\mathbf{A}_o\mathbf{x} - \mathbf{b}\|_W^2$  at the best estimate of the solution found.

`time` is a scalar variable of type `SLLs_time_type` whose components are used to hold elapsed CPU times for the various parts of the calculation (see Section 2.3.4).

---

**All use is subject to the conditions of a BSD-3-Clause License.**  
**See <https://www.galahad.rl.ac.uk/download/> for full details.**

`SBLS_inform` is a scalar variable of type `SBLS_inform_type` whose components provide information about the progress and needs of the factorization/preconditioner performed by the package `GALAHAD_SBLS`. See the documentation for `GALAHAD_SBLS` for further details.

`lapack_error` is a scalar variable of type `INTEGER(ip_)`, that corresponds to the output value `info` returned from the LAPACK routines `S/D/QPOTRF/S`, `S/D/QSYTRF/S` and `S/D/QPBTRF/S`. See the documentation for LAPACK for further details.

### 2.3.6 The derived data type for holding problem data

The derived data type `SLLS_data_type` is used to hold all the data for a particular problem, or sequences of problems with the same structure, between calls of `SLLS` procedures. This data should be preserved, untouched, from the initial call to `SLLS_initialize` to the final call to `SLLS_terminate`.

### 2.3.7 The derived data type for holding user data

The derived data type `USERDATA_type` is available to allow the user to pass data to and from user-supplied subroutines for function and derivative calculations (see Section 2.5). Components of variables of type `USERDATA_type` may be allocated as necessary. The following components are available:

`integer` is a rank-one allocatable array of type `INTEGER(ip_)`.

`real` is a rank-one allocatable array of type default `REAL(rp_)`

`complex` is a rank-one allocatable array of type default `COMPLEX` (double precision complex in `GALAHAD_SLLS_-double`).

`character` is a rank-one allocatable array of type default `CHARACTER`.

`logical` is a rank-one allocatable array of type default `LOGICAL`.

`integer_pointer` is a rank-one pointer array of type `INTEGER(ip_)`.

`real_pointer` is a rank-one pointer array of type default `REAL(rp_)`

`complex_pointer` is a rank-one pointer array of type default `COMPLEX` (double precision complex in `GALAHAD_SLLS_-double`).

`character_pointer` is a rank-one pointer array of type default `CHARACTER`.

`logical_pointer` is a rank-one pointer array of type default `LOGICAL`.

### 2.3.8 The derived data type for holding reverse-communication data

The derived data type `REVERSE_type` is used to hold data needed for reverse communication when this is required. The components of `REVERSE_type` are:

`lv1` is a scalar variable of type `INTEGER(ip_)`, that may be used to hold the starting position in  $V$  (see below) of the list of indices of nonzero components of  $v$ .

`lvu` is a scalar variable of type `INTEGER(ip_)`, that may be used to hold the finishing position in  $V$  (see below) of the list of indices of nonzero components of  $v$ .

`IV` is a rank-one allocatable array of dimension  $n$  and type `INTEGER(ip_)`, that may be used to hold the indices of the nonzero components of  $v$ . If used, components `IV(lv1:lvu)` of  $V$  (see below) will be nonzero.

---

**All use is subject to the conditions of a BSD-3-Clause License.**  
**See <https://www.galahad.rl.ac.uk/download/> for full details.**

$V$  is a rank-one allocatable array of dimension  $n$  and type `REAL(rp_)`, that is used to hold the components of an input vector  $v$ .

`index` is a scalar variable of type `INTEGER(ip_)`, that may be used to hold a column index of  $A_o$  if required.

`lp` is a scalar variable of type `INTEGER(ip_)`, that is used to record the finishing position in  $P$  and  $IP$  (see below) of the list of indices of nonzero components a requested operation on  $v$ .

$P$  is a rank-one allocatable array of dimension  $o$  and type `REAL(rp_)`, that is used to record the components of the vector that results from a requested operation on  $v$ .

$IP$  is a rank-one allocatable array of dimension  $n$  and type `INTEGER(ip_)`, that is used to record the list of indices of nonzero components of a requested operation on  $v$ .

## 2.4 Argument lists and calling sequences

There are three procedures for user calls (see Section 2.8 for further features):

1. The subroutine `SLLS_initialize` is used to set default values, and initialize private data, before solving one or more problems with the same sparsity and bound structure.
2. The subroutine `SLLS_solve` is called to solve the problem.
3. The subroutine `SLLS_terminate` is provided to allow the user to automatically deallocate array components of the private data, allocated by `SLLS_solve`, at the end of the solution process.

We use square brackets `[ ]` to indicate OPTIONAL arguments.

### 2.4.1 The initialization subroutine

Default values are provided as follows:

```
CALL SLLS_initialize( data, control )
```

`data` is a scalar `INTENT(INOUT)` argument of type `SLLS_data_type` (see Section 2.3.6). It is used to hold data about the problem being solved.

`control` is a scalar `INTENT(OUT)` argument of type `SLLS_control_type` (see Section 2.3.3). On exit, `control` contains default values for the components as described in Section 2.3.3. These values should only be changed after calling `SLLS_initialize`.

### 2.4.2 The simplex-constrained linear least-squares subroutine

The simplex-constrained linear least-squares solution algorithm is called as follows:

```
CALL SLLS_solve( prob, data, control, inform, userdata[, reverse, &
                eval_APROD, eval_ASCOL, eval_AFPROD, eval_PREC] )
```

`prob` is a scalar `INTENT(INOUT)` argument of type `QPT_problem_type` (see Section 2.3.2). It is used to hold data about the problem being solved. The user must allocate and set values for the array components, and set values for the component `prob%B`. Additionally, the user can provide  $A_o$  by allocating the relevant array components and setting values for `prob%Ao` using whichever of the matrix formats described in Section 2.1 is appropriate for the user's application; if the effect of  $A_o$  and its transpose are only available to form products via reverse communication (see `reverse` below) or with a set of user-supplied subroutines (see `eval_APROD`, `eval_ASCOL` and `eval_AFPROD` below), `prob%Ao` is not needed.

---

**All use is subject to the conditions of a BSD-3-Clause License.**  
**See <https://www.galahad.rl.ac.uk/download/> for full details.**

The components `prob%X`, `prob%Y` and `prob%Z` must be set to initial estimates of the primal variables,  $\mathbf{x}$ , Lagrange multipliers,  $\mathbf{y}$ , for the simplex equality constraints and dual variables for the bound constraints,  $\mathbf{z}$ , respectively. Inappropriate initial values will be altered, so the user should not be overly concerned if suitable values are not apparent, and may be content with merely setting `prob%X=0.0`, `prob%Y=0.0` and `prob%Z=0.0`. The components `prob%G` and `prob%R` need not be set or allocated on entry.

On exit, the components `prob%X`, `prob%Y` and `prob%Z` will contain the best estimates of the primal variables  $\mathbf{x}$ , Lagrange multipliers for the simplex equality constraints  $\mathbf{y}$  and dual variables for the bound constraints  $\mathbf{z}$ , respectively. The components `prob%R` and `prob%G` will contain the residuals  $\mathbf{r}(\mathbf{x})$  and gradients  $\mathbf{g}(\mathbf{x})$  at  $\mathbf{x}$ , respectively. The component `X_status` will indicate the status of each variable.

**Restrictions:** `prob%n > 0`, `prob%o > 0` and (if  $\mathbf{A}_o$  is provided) `prob%Ao%ne ≥ 0`. `prob%Ao_type`  $\in \{ \text{'DENSE\_BY\_ROWS', 'DENSE\_BY\_COLUMNS', 'COORDINATE', 'SPARSE\_BY\_ROWS', 'SPARSE\_BY\_COLUMNS'} \}$ .

`data` is a scalar `INTENT(INOUT)` argument of type `SLLS_data_type` (see Section 2.3.6). 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 `SLLS_initialize`.

`control` is a scalar `INTENT(IN)` argument of type `SLLS_control_type` (see Section 2.3.3). Default values may be assigned by calling `SLLS_initialize` prior to the first call to `SLLS_solve`.

`inform` is a scalar `INTENT(INOUT)` argument of type `SLLS_inform_type` (see Section 2.3.5). On initial entry, the component `status` must be set to 1, while other components need not be set. A successful call to `SLLS_solve` is indicated when the component `status` has the value 0. For other return values of `status`, see Sections 2.6 and 2.7.

`userdata` is a scalar `INTENT(INOUT)` argument of type `USERDATA_type` whose components may be used to communicate user-supplied data (see Section 2.3.7) to and from the `OPTIONAL` subroutines `eval_APROD` and `eval_ASCOL` (see below).

`reverse` is an `OPTIONAL` scalar `INTENT(INOUT)` argument of type `REVERSE_type` (see Section 2.3.8). It is used to communicate reverse-communication data between the subroutine and calling program. If `reverse` is `PRESENT` and `eval_APROD` or `eval_ASCOL` (see below) are absent, the user should monitor `inform%status` on exit (see Section 2.6).

`eval_APROD` is an `OPTIONAL` user-supplied subroutine whose purpose is to evaluate the value of the product  $\mathbf{p} + \mathbf{A}_o \mathbf{v}$  or  $\mathbf{p} + \mathbf{A}_o^T \mathbf{v}$  involving the Jacobian (and its transpose) and given vectors  $\mathbf{v}$  and  $\mathbf{p}$ . See Section 2.5.1 for details. If `eval_APROD` is present, it must be declared `EXTERNAL` in the calling program. If `eval_APROD` is absent, `GALAHAD_SLLS_solve` will use reverse communication (see Section 2.6) to obtain Jacobian-vector products if `reverse` is `PRESENT` or otherwise require that the user has provided all relevant components of `prob%A`.

`eval_ASCOL` is an `OPTIONAL` user-supplied subroutine whose purpose is to evaluate the value of a specified column of the Jacobian  $\mathbf{A}_o$ . See Section 2.5.2 for details. If `eval_ASCOL` is present, it must be declared `EXTERNAL` in the calling program. If `eval_ASCOL` is absent, `GALAHAD_SLLS_solve` will use reverse communication (see Section 2.6) to obtain columns of the Jacobian if `reverse` is `PRESENT` or otherwise require that the user has provided all relevant components of `prob%A`.

`eval_AFPD` is an `OPTIONAL` user-supplied subroutine whose purpose is to evaluate the value of the product  $\mathbf{A}_o \mathbf{v}$  or  $\mathbf{A}_o^T \mathbf{v}$  involving the Jacobian (and its transpose) and a given vector  $\mathbf{v}$  in which either only some components of  $\mathbf{v}$  or the resulting product are set/required. See Section 2.5.3 for details. If `eval_AFPD` is present, it must be declared `EXTERNAL` in the calling program. If `eval_AFPD` is absent, `GALAHAD_SLLS_solve` will use reverse communication (see Section 2.6) to obtain Jacobian-sparse-vector products if `reverse` is `PRESENT` or otherwise require that the user has provided all relevant components of `prob%A`.

`eval_PREC` is an `OPTIONAL` user-supplied subroutine whose purpose is to evaluate the value of the product  $\mathbf{P}^{-1} \mathbf{v}$  involving a symmetric, positive definite preconditioner  $\mathbf{P}$  and a given vector  $\mathbf{v}$ . See Section 2.5.4 for details. If

---

**All use is subject to the conditions of a BSD-3-Clause License.**  
See <https://www.galahad.rl.ac.uk/download/> for full details.

`eval_PREC` is present, it must be declared `EXTERNAL` in the calling program. If `eval_PREC` is absent, `GALAHAD_SLLS_solve` will use reverse communication (see Section 2.6) to obtain preconditioning products so long as `reverse` is `PRESENT`; if `reverse` is not `PRESENT`, no preconditioning will be performed.

### 2.4.3 The termination subroutine

All previously allocated arrays are deallocated as follows:

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

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

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

`inform` is a scalar `INTENT(OUT)` argument of type `SLLS_inform_type` exactly as for `SLLS_solve`. Only the component `status` will be set on exit, and a successful call to `SLLS_terminate` is indicated when this component `status` has the value 0. For other return values of `status`, see Section 2.7.

## 2.5 Matrix-vector operations

### 2.5.1 Jacobian-vector products via internal evaluation

If the argument `eval_APROD` is present when calling `GALAHAD_SLLS_solve`, the user is expected to provide a subroutine of that name to evaluate the sum  $\mathbf{p} + \mathbf{A}_o \mathbf{v}$  or  $\mathbf{p} + \mathbf{A}_o^T \mathbf{v}$  involving the product of the residual Jacobian  $\mathbf{A}_o$  or its transpose  $\mathbf{A}_o^T$  and a given vector  $\mathbf{v}$ . The routine must be specified as

```
SUBROUTINE eval_APROD( status, userdata, transpose, V, P )
```

whose arguments are as follows:

`status` is a scalar `INTENT(OUT)` argument of type `INTEGER(ip_)`, that should be set to 0 if the routine has been able to evaluate the sum  $\mathbf{p} + \mathbf{A}_o \mathbf{v}$  or  $\mathbf{p} + \mathbf{A}_o^T \mathbf{v}$  and to a non-zero value if the evaluation has not been possible.

`userdata` is a scalar `INTENT(INOUT)` argument of type `USERDATA_type` whose components may be used to communicate user-supplied data to and from the subroutines `eval_APROD` (see Section 2.3.7).

`transpose` is a scalar `INTENT(IN)` array argument of type `default` that will be set `.TRUE.` if the product involves the transpose of the Jacobian  $\mathbf{A}_o^T$  and `.FALSE.` if the product involves the Jacobian  $\mathbf{A}_o$  itself.

`V` is a rank-one `INTENT(IN)` array argument of type `REAL(rp_)` whose components contain the vector  $\mathbf{v}$ .

`P` is a rank-one `INTENT(INOUT)` array argument of type `REAL(rp_)` whose components on input contain the vector  $\mathbf{p}$  and on output the sum  $\mathbf{p} + \mathbf{A}_o \mathbf{v}$  when `%transpose` is `.FALSE.` or  $\mathbf{p} + \mathbf{A}_o^T \mathbf{v}$  when `%transpose` is `.TRUE.`.

### 2.5.2 Columns of the Jacobian via internal evaluation

If the argument `eval_ASCOL` is present when calling `GALAHAD_SLLS_solve`, the user is expected to provide a subroutine of that name to evaluate the a specified column of the Jacobian  $\mathbf{A}_o$ . The routine must be specified as

```
SUBROUTINE eval_ASCOL( status, userdata, index, P, IP, lp )
```

whose arguments are as follows:

---

**All use is subject to the conditions of a BSD-3-Clause License.**  
**See <https://www.galahad.rl.ac.uk/download/> for full details.**

`status` is a scalar `INTENT (OUT)` argument of type `INTEGER(ip_)`, that should be set to 0 if the routine has been able to evaluate the desired columns of  $\mathbf{A}_o$  and to a non-zero value if the evaluation has not been possible.

`userdata` is a scalar `INTENT (INOUT)` argument of type `USERDATA_type` whose components may be used to communicate user-supplied data to and from the subroutine (see Section 2.3.7).

`index` is a scalar `INTENT (OUT)` argument of type `INTEGER(ip_)`, that should be set to the index of the required column of  $\mathbf{A}_o$ .

`P` is a rank-one `INTENT (OUT)` array argument of type `REAL(rp_)` whose first `lp` components on output contain the nonzeros of the `index`-th column of  $\mathbf{A}_o$ .

`IP` is an `INTENT (INOUT)` rank-one array of dimension at least `lp` and type `INTEGER(ip_)`, that must be set to record the list of indices of nonzero components of the `index`-th column of  $\mathbf{A}_o$ , corresponding to the values stored in `P`.

`lp` is an `INTENT (INOUT)` scalar variable of type `INTEGER(ip_)`, that must be set to record the number of non-zeros in the `col`-th column of  $\mathbf{A}_o$ .

### 2.5.3 Jacobian-vector sub-products via internal evaluation

If the argument `eval_AFFPROD` is present when calling `GALAHAD_SLLS_solve`, the user is expected to provide a subroutine of that name to evaluate the product of the Jacobian  $\mathbf{A}_o$ , or its transpose, with a given vector  $\mathbf{v}$ . Here, either only a subset of the components of the vector  $\mathbf{v}$  are nonzero, or only a subset of the components product  $\mathbf{A}_o^T \mathbf{v}$  are to required. The routine must be specified as

```
SUBROUTINE eval_AFFPROD( status, userdata, transpose, V, P, FREE, n_free )
```

whose arguments are as follows:

`status` is a scalar `INTENT (OUT)` argument of type `INTEGER(ip_)`, that should be set to 0 if the routine has been able to evaluate the product  $\mathbf{p} = \mathbf{A}_o \mathbf{v}$  or  $\mathbf{p} = \mathbf{A}_o^T \mathbf{v}$  and to a non-zero value if the evaluation has not been possible.

`userdata` is a scalar `INTENT (INOUT)` argument of type `USERDATA_type` whose components may be used to communicate user-supplied data to and from the subroutines `eval_APROD` (see Section 2.3.7).

`transpose` is a scalar `INTENT (IN)` array argument of type default that will be set `.TRUE.` if the product involves the transpose of the Jacobian  $\mathbf{A}_o^T$  and `.FALSE.` if the product involves the Jacobian  $\mathbf{A}_o$  itself.

`V` is a rank-one `INTENT (IN)` array argument of type `REAL(rp_)` whose components contain the vector  $\mathbf{v}$ . If `transpose` is `.FALSE.` only those components whose indices `FREE(:n_free)` (see below) will by set, and the remainder will be presumed to be zero.

`P` is a rank-one `INTENT (OUT)` array argument of type `REAL(rp_)` whose components on output must be set to the product  $\mathbf{A}_o \mathbf{v}$  when `%transpose` is `.FALSE.`. If `%transpose` is `.TRUE.`, components with indices `FREE(:n_free)` (see below) should be set to the corresponding components of the product  $\mathbf{A}_o^T \mathbf{v}$ , and the remaining components ignored.

`FREE` is a rank-one `INTENT (IN)` array argument of type `INTEGER(ip_)` that flags the input components of  $\mathbf{v}$  that are set (when `transpose` is `.FALSE.`) or output components of  $\mathbf{A}_o^T \mathbf{v}$  that are required (when `transpose` is `.TRUE.`). Specifically, only indices `FREE(:n_free)` of the relevant vector is set or required, and the remainder should be treated as zero (if `transpose` is `.FALSE.`) or ignored (if `transpose` is `.TRUE.`).

`n_free` is a scalar `INTENT (IN)` argument of type `INTEGER(ip_)`, that specifies the number of components of `FREE` that need be considered.

---

**All use is subject to the conditions of a BSD-3-Clause License.**  
 See <https://www.galahad.rl.ac.uk/download/> for full details.

#### 2.5.4 Application of a preconditioner via internal evaluation

If the argument `eval_PREC` is present when calling `GALAHAD_SLLS_solve`, the user is expected to provide a subroutine of that name to evaluate the product  $\mathbf{P}^{-1}\mathbf{v}$  involving a symmetric, positive definite preconditioner  $\mathbf{P}$  and a given  $n$ -vector  $\mathbf{v}$ . Ideally the  $n$  by  $n$  matrix  $\mathbf{P}$  should be chosen so that the of the product is inexpensive to compute, but also so that  $\mathbf{P}$  is a good approximation of  $\mathbf{A}_o^T\mathbf{A}_o$  in the sense that the eigenvalues of  $\mathbf{P}^{-1}\mathbf{A}_o^T\mathbf{A}_o$  are clustered around a small number of values (preferably all around 1.0). The routine must be specified as

```
SUBROUTINE eval_PREC( status, userdata, V, P )
```

whose arguments are as follows:

`status` is a scalar `INTENT(OUT)` argument of type `INTEGER(ip_)`, that should be set to 0 if the routine has been able to evaluate the product  $\mathbf{P}^{-1}\mathbf{v}$ , and to a non-zero value if the evaluation has not been possible.

`userdata` is a scalar `INTENT(INOUT)` argument of type `USERDATA_type` whose components may be used to communicate user-supplied data to and from the subroutines `eval_APROD` (see Section 2.3.7).

`V` is a rank-one `INTENT(IN)` array argument of type `REAL(rp_)` whose components contain the vector  $\mathbf{v}$ .

`P` is a rank-one `INTENT(INOUT)` array argument of type `REAL(rp_)` whose components on output contain the vector  $\mathbf{p} = \mathbf{P}^{-1}\mathbf{v}$ .

#### 2.6 Reverse Communication Information

A positive value of `inform%status` on exit from `SLLS_solve` indicates that `GALAHAD_SLLS_solve` is seeking further information—this will happen if the user has chosen to evaluate matrix-vector products by reverse communication. The user should compute the required information and re-enter `GALAHAD_SLLS_solve` with `inform%status` and all other arguments (except those specifically mentioned below) unchanged.

Possible values of `inform%status` and the information required are

2. The user should compute the product  $\mathbf{A}_o\mathbf{v}$  involving the product of the residual Jacobian,  $\mathbf{A}_o$ , with a given vector  $\mathbf{v}$ . The vector  $\mathbf{v}$  is given in `reverse%V`, and the product  $\mathbf{p} = \mathbf{A}_o\mathbf{v}$  should be written to `reverse%P`, and `reverse%eval_status` should be set to 0. If the user is unable to evaluate the product, the user need not set `reverse%P`, but should then set `reverse%eval_status` to a non-zero value.
3. The user should compute the product  $\mathbf{A}_o^T\mathbf{v}$  involving the product of the transpose of the residual Jacobian,  $\mathbf{A}_o$ , with a given vector  $\mathbf{v}$ . The vector  $\mathbf{v}$  is given in `reverse%V`, and the product  $\mathbf{p} = \mathbf{A}_o^T\mathbf{v}$  should be written to `reverse%P`, and `reverse%eval_status` should be set to 0. If the user is unable to evaluate the product, the user need not set `reverse%P`, but should then set `reverse%eval_status` to a non-zero value.
4. The user should compute the nonzero components of the  $j$ -th column of the matrix  $\mathbf{A}_o$ , where `reverse%index` holds the index  $j$ . The values and corresponding row indices of the selected column should be set in `reverse%P(1:reverse%lp)` and `reverse%IP(1:reverse%lp)` respectively, and `reverse%lp` set accordingly. The component `reverse%eval_status` should set to 0 unless the user is unable to provide the selected column, in which case `reverse%eval_status` should be set to a non-zero value, and the remaining components left unaltered.
5. The user should compute the matrix-vector product  $\mathbf{p} = \mathbf{A}_o\mathbf{v}$  using the vector  $\mathbf{v}$  whose nonzero components are stored in positions `reverse%IV(reverse%lvl:reverse%lvu)` of `reverse%V`. The remaining components of `reverse%V` should be ignored. The resulting vector  $\mathbf{p}$  should be assigned to `reverse%P`. The component `reverse%eval_status` should set to 0 unless the user is unable to evaluate the product, in which case `reverse%eval_status` should be set to a non-zero value, and the remaining components left unaltered.

---

**All use is subject to the conditions of a BSD-3-Clause License.**  
See <https://www.galahad.rl.ac.uk/download/> for full details.

6. The user should compute the matrix-vector product  $\mathbf{A}_o^T \mathbf{v}$  using the vector  $\mathbf{v}$  given in `reverse%V`. The components  $(\mathbf{A}_o^T \mathbf{v})_j$  of the product  $\mathbf{A}_o^T \mathbf{v}$ , for  $j = \text{reverse\%IV}(\text{reverse\%lv1}:\text{reverse\%lvu})$ , should be written to `reverse\%P(j)`. The component `reverse\%eval_status` should be set to 0 unless the user is unable to evaluate the product, in which case `reverse\%eval_status` should be set to a non-zero value, and the remaining components left unaltered.
7. The user should compute the product  $\mathbf{P}^{-1} \mathbf{v}$  involving a symmetric, positive definite preconditioner  $\mathbf{P}$  and a given vector  $\mathbf{v}$ . The vector  $\mathbf{v}$  is given in `reverse%V`, and the product  $\mathbf{p} = \mathbf{P}^{-1} \mathbf{v}$  should be written to `reverse\%P`, and `reverse\%eval_status` should be set to 0. If the user is unable to evaluate the product, the user need not set `reverse\%P`, but should then set `reverse\%eval_status` to a non-zero value. This value of `inform\%status` can only occur if the user has set `control\%preconditioner = 2`.

## 2.7 Warning and error messages

A negative value of `inform\%status` on exit from `SLLS_solve` or `SLLS_terminate` indicates that an error has 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. One of the restrictions `prob%n > 0`, `prob%o > 0` or the requirement that `prob\%Ao_type` contain its relevant string 'DENSE\_BY\_ROWS', 'DENSE\_BY\_COLUMNS', 'COORDINATE', 'SPARSE\_BY\_ROWS' or 'SPARSE\_BY\_COLUMNS', when  $\mathbf{A}_o$  is available, has been violated.
- 4. The bound constraints are inconsistent.
- 9. The analysis phase of the factorization failed; the return status from the factorization package is given in the component `inform\%factor_status`.
- 10. The factorization failed; the return status from the factorization package is given in the component `inform\%factor_status`.
- 18. Too many iterations have been performed. This may happen if `control\%maxit` is too small, but may also be symptomatic of a badly scaled problem.
- 19. The CPU time limit has been reached. This may happen if `control\%cpu_time_limit` is too small, but may also be symptomatic of a badly scaled problem.

## 2.8 Further features

In this section, we describe an alternative means of setting control parameters, that is components of the variable `control` of type `SLLS_control_type` (see Section 2.3.3), by reading an appropriate data specification file using the subroutine `SLLS_read_specfile`. This facility is useful as it allows a user to change SLLS control parameters without editing and recompiling programs that call SLLS.

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

---

**All use is subject to the conditions of a BSD-3-Clause License.**  
**See <https://www.galahad.rl.ac.uk/download/> for full details.**

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 `SLLS_read_specfile` must start with a "BEGIN SLLS" command and end with an "END" command. The syntax of the specfile is thus defined as follows:

```
( .. lines ignored by SLLS_read_specfile .. )
BEGIN SLLS
  keyword      value
  .....
  keyword      value
END
( .. lines ignored by SLLS_read_specfile .. )
```

where keyword and value are two strings separated by (at least) one blank. The "BEGIN SLLS" 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 SLLS SPECIFICATION
```

and

```
END SLLS SPECIFICATION
```

are acceptable. Furthermore, between the "BEGIN SLLS" 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 `SLLS_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 `SLLS_read_specfile`.

### 2.8.1 To read control parameters from a specification file

Control parameters may be read from a file as follows:

```
CALL SLLS_read_specfile( control, device )
```

`control` is a scalar `INTENT(INOUT)` argument of type `SLLS_control_type` (see Section 2.3.3). Default values should have already been set, perhaps by calling `SLLS_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 2.3.3) of `control` that each affects are given in Table 2.1.

`device` is a scalar `INTENT(IN)` argument of type `INTEGER(ip_)`, 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`.

---

**All use is subject to the conditions of a BSD-3-Clause License.**  
**See <https://www.galahad.rl.ac.uk/download/> for full details.**

| command                                         | component of control      | value type |
|-------------------------------------------------|---------------------------|------------|
| error-printout-device                           | %error                    | integer    |
| printout-device                                 | %out                      | integer    |
| print-level                                     | %print_level              | integer    |
| start-print                                     | %start_print              | integer    |
| stop-print                                      | %stop_print               | integer    |
| iterations-between-printing                     | %print_gap                | integer    |
| maximum-number-of-iterations                    | %maxit                    | integer    |
| cold-start                                      | %cold_start               | integer    |
| preconditioner                                  | %preconditioner           | integer    |
| max-change-to-working-set-for-subspace-solution | %change_max               | integer    |
| maximum-number-of-cg-iterations-per-iteration   | %cg_maxit                 | integer    |
| maximum-number-of-arcsearch-steps               | %arcsearch_max_steps      | integer    |
| primal-accuracy-required                        | %stop_p                   | real       |
| dual-accuracy-required                          | %stop_d                   | real       |
| complementary-slackness-accuracy-required       | %stop_c                   | real       |
| cg-relative-accuracy-required                   | %stop_cg_relative         | real       |
| cg-absolute-accuracy-required                   | %stop_cg_absolute         | real       |
| maximum-arcsearch-stepsize                      | %alpha_max                | real       |
| initial-arcsearch-stepsize                      | %alpha_initial            | real       |
| arcsearch-reduction-factor                      | %alpha_reduction          | real       |
| arcsearch-acceptance-tolerance                  | %arcsearch_acceptance_tol | real       |
| regularization-weight                           | %regularization_weight    | real       |
| maximum-cpu-time-limit                          | %cpu_time_limit           | real       |
| direct-subproblem-solve                         | %direct_subproblem_solve  | logical    |
| exact-arc-search-used                           | %exact_arc_search         | logical    |
| space-critical                                  | %space_critical           | logical    |
| deallocate-error-fatal                          | %deallocate_error_fatal   | logical    |
| output-line-prefix                              | %prefix                   | character  |

Table 2.1: Specfile commands and associated components of control.

## 2.9 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. This will give the total number of CG iterations performed (if any), the value of the objective function and the norm of the projected gradient, the stepsize taken, the number of free variables (i.e., those that have a positive value), the change in the number of free variables since the last iteration, and the elapsed CPU time in seconds. If `control%print_level ≥ 2` this output will be increased to provide significant detail of each iteration. This extra output includes residuals of the linear systems solved, and, for larger values of `control%print_level`, values of the primal and dual variables and Lagrange multiplier.

## 3 GENERAL INFORMATION

**Use of common:** None.

**Workspace:** Provided automatically by the module.

**Other routines called directly:** None.

---

**All use is subject to the conditions of a BSD-3-Clause License.**  
**See <https://www.galahad.rl.ac.uk/download/> for full details.**

**Other modules used directly:** SLLS\_solve calls the GALAHAD packages GALAHAD\_CPU\_time, GALAHAD\_SYMBOLS, GALAHAD\_SPACE, GALAHAD\_STRING, GALAHAD\_SORT, GALAHAD\_NORMS, GALAHAD\_CONVERT, GALAHAD\_SBLS, GALAHAD\_QPT, GALAHAD\_QPD, GALAHAD\_USERDATA, GALAHAD\_REVERSE and GALAHAD\_SPECFILE.

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

**Restrictions:** `prob%n > 0`, `prob%o > 0`, `prob%A_type`  $\in \{ 'DENSE\_BY\_ROWS', 'DENSE\_BY\_COLUMNS', 'COORDINATE', 'SPARSE\_BY\_ROWS', 'SPARSE\_BY\_COLUMNS' \}$  (if  $\mathbf{A}_o$  is explicit).

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

## 4 METHOD

The required solution  $\mathbf{x}$  necessarily satisfies the primal optimality conditions

$$\mathbf{e}_{G_i}^T \mathbf{x}_{G_i} = 1 \text{ and } \mathbf{x}_{G_i} \geq 0 \text{ for } i = 1, \dots, m, \quad (4.1)$$

the dual optimality conditions

$$\mathbf{A}_o^T \mathbf{W}(\mathbf{A}_o \mathbf{x} - \mathbf{b}) + \boldsymbol{\sigma} \mathbf{x} = \sum_{i=1}^m \mathbf{e}_{G_i} y_i + \mathbf{z}, \quad (4.2)$$

for some Lagrange multipliers  $\mathbf{y}$  and dual variables

$$\mathbf{z} \geq 0, \quad (4.3)$$

and the complementary slackness conditions

$$\mathbf{x}^T \mathbf{z} = 0, \quad (4.4)$$

where the vector inequalities hold componentwise. Projected-gradient methods iterate towards a point that satisfies these conditions by ultimately aiming to satisfy (4.2), while ensuring that (4.1), and (4.3) and (4.4) are satisfied at each stage. Appropriate norms of the amounts by which (4.1), (4.2) and (4.4) fail to be satisfied are known as the primal and dual infeasibility, and the violation of complementary slackness, respectively.

The method is iterative. Each iteration proceeds in two stages. Firstly, a search direction  $\mathbf{s}$  from the current estimate of the solution  $\mathbf{x}$  is computed. This may be in a scaled steepest-descent direction, or, if the working set of variables on bounds has not changed dramatically, in a direction that provides an approximate minimizer of the objective over a subspace comprising the currently free-variables. The latter is computed either using an appropriate sparse factorization by the package GALAHAD\_SBLS, or by the conjugate-gradient least-squares (CGLS) method; it may be necessary to regularize the subproblem very slightly to avoid a ill-posedness. Thereafter, a piecewise linesearch (arc search) is carried out along the arc  $\mathbf{x}(\alpha) = P(\mathbf{x} + \alpha \mathbf{s})$  for  $\alpha > 0$ , where the projection operator  $P(\mathbf{v})$  gives the nearest point to  $\mathbf{v}$  within the intersection of specified simplices; thus this arc bends the search direction into the feasible region. The arc search is performed either exactly, by passing through a set of increasing breakpoints at which it changes direction, or inexactly, by evaluating a sequence of different  $\alpha$  on the arc. All computation is designed to exploit sparsity in  $\mathbf{A}_o$ .

## References:

Full details are provided in

H. Al Daas, J. M. Fowkes, N. I. M. Gould and J. Huntley (2026). Linear least-squares over unit simplices, RAL-NA-2026-1.

---

**All use is subject to the conditions of a BSD-3-Clause License.**

See <https://www.galahad.rl.ac.uk/download/> for full details.

## 5 EXAMPLE OF USE

Suppose we wish to minimize

$$\frac{1}{2} \left\| \begin{pmatrix} x_1 \\ x_1 + x_2 - 2 \\ x_3 - 1 \\ x_3 - 2 \end{pmatrix} \right\|_2^2$$

within the simplex

$$x_1 + x_2 + x_3 = 1, \quad x_1, x_2, x_3 \geq 0.$$

Then, on writing the data for this problem as

$$\mathbf{A}_o = \begin{pmatrix} 1 & & & \\ 1 & 1 & & \\ & & 1 & \\ & & & 1 \end{pmatrix}, \quad \text{and } \mathbf{b} = \begin{pmatrix} 0 \\ 2 \\ 1 \\ 2 \end{pmatrix}$$

in sparse coordinate format, we may use the following code:

```
! THIS VERSION: GALAHAD 5.5 - 2026-02-01 AT 08:40 GMT
PROGRAM GALAHAD_SLLS_EXAMPLE
USE GALAHAD_SLLS_double           ! double precision version
IMPLICIT NONE
INTEGER, PARAMETER :: wp = KIND( 1.0D+0 ) ! set precision
TYPE ( QPT_problem_type ) :: p
TYPE ( SLLS_data_type ) :: data
TYPE ( SLLS_control_type ) :: control
TYPE ( SLLS_inform_type ) :: inform
TYPE ( USERDATA_type ) :: userdata
INTEGER :: s
INTEGER, PARAMETER :: n = 3, o = 4, m = 1, ao_ne = 5
! start problem data
ALLOCATE( p%B( o ), p%X( n ), p%X_status( n ) )
p%n = n ; p%o = o ; p%m = m           ! dimensions
p%B = (/ 0.0_wp, 2.0_wp, 1.0_wp, 2.0_wp /) ! right-hand side
p%X = 0.0_wp ! start from zero
! sparse co-ordinate storage format
CALL SMT_put( p%Ao%type, 'COORDINATE', s ) ! Co-ordinate storage for Ao
ALLOCATE( p%Ao%val( ao_ne ), p%Ao%row( ao_ne ), p%Ao%col( ao_ne ) )
p%Ao%m = o ; p%Ao%n = n ; p%Ao%ne = ao_ne ! design matrix Ao
p%Ao%row = (/ 1, 2, 2, 3, 4 /) ; p%Ao%col = (/ 1, 1, 2, 3, 3 /)
p%Ao%val = (/ 1.0_wp, 1.0_wp, 1.0_wp, 1.0_wp, 1.0_wp /)
! problem data complete
CALL SLLS_initialize( data, control, inform ) ! Initialize control parameters
control%SLLS_control%symmetric_linear_solver = 'sytr' ! non-default solver
control%SLLS_control%definite_linear_solver = 'potr' ! non-default solver
! control%print_level = 1           ! print one line/iteration
inform%status = 1
CALL SLLS_solve( p, data, control, inform, userdata )
IF ( inform%status == 0 ) THEN           ! Successful return
  WRITE( 6, "( /, ' SLLS: ', I0, ' iterations ', /,
    &      ' Optimal objective value = ',
    &      ES12.4, /, ' Optimal solution = ', ( 5ES12.4 ) )" )
    &      inform%iter, inform%obj, p%X
  WRITE( 6, "( ' Lagrange multiplier estimate = ', ES12.4 )" ) p%Y
```

---

**All use is subject to the conditions of a BSD-3-Clause License.**  
**See <https://www.galahad.rl.ac.uk/download/> for full details.**

```

ELSE                                     ! Error returns
  WRITE( 6, "( /, ' SLLS_solve exit status = ', I0 ) " ) inform%status
  WRITE( 6, * ) inform%alloc_status, inform%bad_alloc
END IF
CALL SLLS_terminate( data, control, inform ) ! delete workspace
DEALLOCATE( p%B, p%X, p%Y, p%Z, p%R, p%G, p%X_status )
DEALLOCATE( p%Ao%val, p%Ao%row, p%Ao%col, p%Ao%type )
END PROGRAM GALAHAD_SLLS_EXAMPLE

```

This produces the following output:

```

SLLS: 3 iterations
Optimal objective value = 2.3333E+00
Optimal solution = 0.0000E+00 3.3333E-01 6.6667E-01
Lagrange multiplier estimate = -1.6667E+00

```

The same problem may be solved holding the data in a sparse row-wise storage format by replacing the lines

```

! sparse coordinate storage format
...
! problem data complete

```

by

```

! sparse row-wise storage format
CALL SMT_put( p%A%type, 'SPARSE_BY_ROWS', s ) ! Specify sparse-by-rows
ALLOCATE( p%A%val( a_ne ), p%A%col( a_ne ), p%A%ptr( m + 1 ) )
p%A%val = (/ 1.0_wp, 1.0_wp, 1.0_wp, 1.0_wp, 1.0_wp /) ! Jacobian A
p%A%col = (/ 1, 1, 2, 3, 3 /) ! Column indices
p%A%ptr = (/ 1, 2, 4, 5, 6 /) ! Set row pointers
! problem data complete

```

a sparse column-wise storage format by replacing the lines

```

! sparse column-wise storage format
CALL SMT_put( p%A%type, 'SPARSE_BY_COLUMNS', s ) ! Specify sparse-by-columns
ALLOCATE( p%A%val( a_ne ), p%A%row( a_ne ), p%A%ptr( n + 1 ) )
p%A%val = (/ 1.0_wp, 1.0_wp, 1.0_wp, 1.0_wp, 1.0_wp /) ! Jacobian A
p%A%row = (/ 1, 2, 2, 3, 4 /) ! Row indices
p%A%ptr = (/ 1, 3, 4, 6 /) ! Set column pointers
! problem data complete

```

a dense-by-rows storage format with the replacement lines

```

! dense storage format
CALL SMT_put( p%A%type, 'DENSE_BY_ROWS', s ) ! Specify dense-by-rows
ALLOCATE( p%A%val( m * n ) )
p%A%val = (/ 1.0_wp, 0.0_wp, 0.0_wp, 1.0_wp, 1.0_wp, 0.0_wp, 0.0_wp, &
            0.0_wp, 1.0_wp, 0.0_wp, 0.1_wp, 1.0_wp /)
! problem data complete

```

or a dense-by-columns storage format using the replacement

```

! dense storage format
CALL SMT_put( p%A%type, 'DENSE_BY_COLUMNS', s ) ! Specify dense-by-columns
ALLOCATE( p%A%val( m * n ) )
p%A%val = (/ 1.0_wp, 1.0_wp, 0.0_wp, 0.0_wp, 0.0_wp, 1.0_wp, 0.0_wp, &
            0.0_wp, 0.0_wp, 0.0_wp, 1.0_wp, 1.0_wp /)
! problem data complete

```

---

**All use is subject to the conditions of a BSD-3-Clause License.**  
**See <https://www.galahad.rl.ac.uk/download/> for full details.**

respectively.

The same problem may be solved using reverse communication with the following code:

```
! THIS VERSION: GALAHAD 5.5 - 2026-02-19 AT 10:10 GMT.
PROGRAM GALAHAD_SLLS_SECOND_EXAMPLE ! reverse communication interface
USE GALAHAD_SLLS_double             ! double precision version
IMPLICIT NONE
INTEGER, PARAMETER :: wp = KIND( 1.0D+0 ) ! set precision
TYPE ( QPT_problem_type ) :: p
TYPE ( SLLS_data_type ) :: data
TYPE ( SLLS_control_type ) :: control
TYPE ( SLLS_inform_type ) :: inform
TYPE ( REVERSE_type ) :: reverse
TYPE ( USERDATA_type ) :: userdata
INTEGER :: i, j, k, l
REAL ( KIND = wp ) :: val
INTEGER, PARAMETER :: n = 3, o = 4, m = 1, Ao_ne = 5
INTEGER, ALLOCATABLE, DIMENSION( : ) :: Ao_row, Ao_ptr
REAL ( KIND = wp ), ALLOCATABLE, DIMENSION( : ) :: Ao_val
! start problem data
ALLOCATE( p%B( o ), p%X( n ), p%X_status( n ) )
p%n = n ; p%o = o ; p%m = m ! dimensions
p%B = (/ 0.0_wp, 2.0_wp, 1.0_wp, 2.0_wp /) ! right-hand side
p%X = 0.0_wp ! start from zero
! sparse column storage format
ALLOCATE( Ao_val( Ao_ne ), Ao_row( Ao_ne ), Ao_ptr( n + 1 ) )
Ao_row = (/ 1, 2, 2, 3, 4 /) !! design matrix Ao by columns: row indices
Ao_ptr = (/ 1, 3, 4, 6 /) ! pointers to column starts
Ao_val = (/ 1.0_wp, 1.0_wp, 1.0_wp, 1.0_wp, 1.0_wp /)
! problem data complete
CALL SLLS_initialize( data, control, inform ) ! Initialize control parameters
! control%print_level = 1 ! print one line/iteration
control%exact_arc_search = .FALSE.
inform%status = 1
10 CONTINUE ! Solve problem - reverse communication loop
CALL SLLS_solve( p, data, control, inform, userdata, reverse = reverse )
SELECT CASE ( inform%status )
CASE ( 0 ) ! successful return
WRITE( 6, "( /, ' SLLS: ', I0, ' iterations ', /,
& ' Optimal objective value = ',
& ES12.4, /, ' Optimal solution = ', ( 5ES12.4 ) )" )
inform%iter, inform%obj, p%X
WRITE( 6, "( ' Lagrange multiplier estimate = ', ES12.4 )" ) p%Y
CASE ( 2 ) ! compute Ao * v
reverse%P( : o ) = 0.0_wp
DO j = 1, n
val = reverse%V( j )
DO k = Ao_ptr( j ), Ao_ptr( j + 1 ) - 1
i = Ao_row( k )
reverse%P( i ) = reverse%P( i ) + Ao_val( k ) * val
END DO
END DO
GO TO 10
CASE ( 3 ) ! compute Ao^T * v
reverse%P( : n ) = 0.0_wp
```

---

**All use is subject to the conditions of a BSD-3-Clause License.**  
**See <https://www.galahad.rl.ac.uk/download/> for full details.**

```
DO j = 1, n
  val = 0.0_wp
  DO k = Ao_ptr( j ), Ao_ptr( j + 1 ) - 1
    val = val + Ao_val( k ) * reverse%V( Ao_row( k ) )
  END DO
  reverse%P( j ) = val
END DO
GO TO 10
CASE ( 4 ) ! compute a column of Ao
  reverse%lp = 0
  DO k = Ao_ptr( reverse%index ), Ao_ptr( reverse%index + 1 ) - 1
    reverse%lp = reverse%lp + 1
    reverse%P( reverse%lp ) = Ao_val( k )
    reverse%IP( reverse%lp ) = Ao_row( k )
  END DO
  GO TO 10
CASE ( 5 ) ! compute Ao * sparse v
  reverse%P( : o ) = 0.0_wp
  DO l = reverse%lv1, reverse%lvu
    j = reverse%IV( l )
    val = reverse%V( j )
    DO k = Ao_ptr( j ), Ao_ptr( j + 1 ) - 1
      i = Ao_row( k )
      reverse%P( i ) = reverse%P( i ) + Ao_val( k ) * val
    END DO
  END DO
  GO TO 10
CASE ( 6 ) ! compute sparse( Ao^T * v )
  reverse%P( : n ) = 0.0_wp
  DO l = reverse%lv1, reverse%lvu
    j = reverse%IV( l )
    val = 0.0_wp
    DO k = Ao_ptr( j ), Ao_ptr( j + 1 ) - 1
      val = val + Ao_val( k ) * reverse%V( Ao_row( k ) )
    END DO
    reverse%P( j ) = val
  END DO
  GO TO 10
CASE DEFAULT ! error returns
  WRITE( 6, "( /, ' SLLS_solve exit status = ', I0 ) " ) inform%status
END SELECT
! delete workspace
CALL SLLS_terminate( data, control, inform, reverse = reverse )
DEALLOCATE( p%B, p%X, p%Y, p%Z, p%R, p%X_status )
DEALLOCATE( Ao_val, Ao_row, Ao_ptr )
END PROGRAM GALAHAD_SLLS_SECOND_EXAMPLE
```

This produces the following output:

```
SLLS: 5 iterations
Optimal objective value = 2.3333E+00
Optimal solution = 4.7184E-16 3.3333E-01 6.6667E-01
```

The same problem may also be solved by user-provided matrix-vector products as follows:

```
! THIS VERSION: GALAHAD 5.5 - 2026-02-19 AT 10:20 GMT.
PROGRAM GALAHAD_SLLS_THIRD_EXAMPLE ! subroutine evaluation interface
```

---

**All use is subject to the conditions of a BSD-3-Clause License.**  
**See <https://www.galahad.rl.ac.uk/download/> for full details.**

```

USE GALAHAD_SLLS_double          ! double precision version
IMPLICIT NONE
INTEGER, PARAMETER :: wp = KIND( 1.0D+0 ) ! set precision
TYPE ( QPT_problem_type ) :: p
TYPE ( SLLS_data_type ) :: data
TYPE ( SLLS_control_type ) :: control
TYPE ( SLLS_inform_type ) :: inform
TYPE ( USERDATA_type ) :: userdata
INTEGER, PARAMETER :: n = 3, o = 4, m = 1, Ao_ne = 5
! partition userdata%integer so that it holds
!   o n Ao_ptr      Ao_row
!   |1|2|3 to n+3|n+4 to n+3+a_ne|
! partition userdata%real so that it holds
!   Ao_val
!   |1 to Ao_ne|
INTEGER, PARAMETER :: on = MAX( o, n )
INTEGER, PARAMETER :: st_ptr = 3
INTEGER, PARAMETER :: st_row = st_ptr + n + 1, st_val = 0
INTEGER, PARAMETER :: len_integer = st_row + Ao_ne + 1, len_real = Ao_ne
EXTERNAL :: APROD, ASCOL, AFPROD
! start problem data
ALLOCATE( p%B( o ), p%X( n ), p%X_status( n ) )
p%n = n ; p%o = o ; p%m = m ! dimensions
p%B = (/ 0.0_wp, 2.0_wp, 1.0_wp, 2.0_wp /) ! right-hand side
p%X = 0.0_wp ! start from zero
! sparse co-ordinate storage format
ALLOCATE( userdata%integer( len_integer ), userdata%real( len_real ) )
userdata%integer( 1 ) = o ! load design matrix Ao data into userdata
userdata%integer( 2 ) = n
userdata%integer( st_ptr + 1 : st_ptr + n + 1 ) = (/ 1, 3, 4, 6 /)
userdata%integer( st_row + 1 : st_row + Ao_ne ) = (/ 1, 2, 2, 3, 4 /)
userdata%real( st_val + 1 : st_val + Ao_ne )
= (/ 1.0_wp, 1.0_wp, 1.0_wp, 1.0_wp, 1.0_wp /)
! problem data complete
CALL SLLS_initialize( data, control, inform ) ! Initialize control parameters
! control%print_level = 1 ! print one line/iteration
control%exact_arc_search = .FALSE.
! load workspace into userdata
inform%status = 1
CALL SLLS_solve( p, data, control, inform, userdata, eval_APROD = APROD, &
eval_ASCOL = ASCOL, eval_AFPROD = AFPROD )
IF ( inform%status == 0 ) THEN ! Successful return
WRITE( 6, "( /, ' SLLS: ', I0, ' iterations ', /, &
& ' Optimal objective value =', &
& ES12.4, /, ' Optimal solution = ', ( 5ES12.4 ) )" ) &
inform%iter, inform%obj, p%X
WRITE( 6, "( ' Lagrange multiplier estimate =', ES12.4 )" ) p%Y
ELSE ! Error returns
WRITE( 6, "( /, ' SLLS_solve exit status = ', I0 )" ) inform%status
END IF
CALL SLLS_terminate( data, control, inform ) ! delete workspace
DEALLOCATE( p%B, p%X, p%Y, p%Z, p%R, p%X_status )
DEALLOCATE( userdata%integer, userdata%real )
END PROGRAM GALAHAD_SLLS_THIRD_EXAMPLE

```

**All use is subject to the conditions of a BSD-3-Clause License.**  
**See <https://www.galahad.rl.ac.uk/download/> for full details.**

```
SUBROUTINE APROD( status, userdata, transpose, V, P )
USE GALAHAD_USERDATA_double
INTEGER, PARAMETER :: wp = KIND( 1.0D+0 )
INTEGER, INTENT( OUT ) :: status
TYPE ( USERDATA_type ), INTENT( INOUT ) :: userdata
LOGICAL, INTENT( IN ) :: transpose
REAL ( KIND = wp ), DIMENSION( : ), INTENT( IN ) :: V
REAL ( KIND = wp ), DIMENSION( : ), INTENT( INOUT ) :: P
INTEGER :: i, j, k
REAL ( KIND = wp ) :: val
! recover problem data from userdata
INTEGER :: o, n, st_ptr, st_row, st_val
o = userdata%integer( 1 )
n = userdata%integer( 2 )
st_ptr = 3
st_row = st_ptr + n + 1
st_val = 0
IF ( transpose ) THEN
  DO j = 1, n
    DO k = userdata%integer( st_ptr + j ),
      userdata%integer( st_ptr + j + 1 ) - 1
      P( j ) = P( j ) + userdata%real( st_val + k ) *
        V( userdata%integer( st_row + k ) )
    END DO
  END DO
ELSE
  DO j = 1, n
    val = V( j )
    DO k = userdata%integer( st_ptr + j ),
      userdata%integer( st_ptr + j + 1 ) - 1
      i = userdata%integer( st_row + k )
      P( i ) = P( i ) + userdata%real( st_val + k ) * val
    END DO
  END DO
END IF
status = 0
RETURN
END SUBROUTINE APROD

SUBROUTINE ASCOL( status, userdata, index, P, IP, lp )
USE GALAHAD_USERDATA_double
INTEGER ( KIND = ip_ ), INTENT( OUT ) :: status
TYPE ( USERDATA_type ), INTENT( INOUT ) :: userdata
INTEGER ( KIND = ip_ ), INTENT( IN ) :: index
REAL ( KIND = rp_ ), DIMENSION( : ), INTENT( OUT ) :: P
INTEGER ( KIND = ip_ ), DIMENSION( : ), INTENT( INOUT ) :: IP
INTEGER ( KIND = ip_ ), INTENT( INOUT ) :: lp
INTEGER ( KIND = ip_ ) :: k
! recover problem data from userdata
INTEGER ( KIND = ip_ ) :: o, n, st_ptr, st_row, st_val
o = userdata%integer( 1 )
n = userdata%integer( 2 )
st_ptr = 3
st_row = st_ptr + n + 1
st_val = 0
```

---

**All use is subject to the conditions of a BSD-3-Clause License.**  
**See <https://www.galahad.rl.ac.uk/download/> for full details.**

```

lp = 0
DO k = userdata%integer( st_ptr + index ),
    userdata%integer( st_ptr + index + 1 ) - 1
    lp = lp + 1
    P( lp ) = userdata%real( st_val + k )
    IP( lp ) = userdata%integer( st_row + k )
END DO
status = 0
RETURN
END SUBROUTINE ASCOL

SUBROUTINE AFFPROD( status, userdata, transpose, V, P, FREE, n_free )
USE GALAHAD_USERDATA_double
INTEGER, PARAMETER :: wp = KIND( 1.0D+0 )
INTEGER, INTENT( OUT ) :: status
TYPE ( USERDATA_type ), INTENT( INOUT ) :: userdata
LOGICAL, INTENT( IN ) :: transpose
INTEGER, INTENT( IN ) :: n_free
INTEGER, INTENT( IN ), DIMENSION( : ) :: FREE
REAL ( KIND = wp ), DIMENSION( : ), INTENT( IN ) :: V
REAL ( KIND = wp ), DIMENSION( : ), INTENT( OUT ) :: P
INTEGER :: i, j, k, l
REAL ( KIND = wp ) :: val
! recover problem data from userdata
INTEGER :: o, n, st_ptr, st_row, st_val
o = userdata%integer( 1 )
n = userdata%integer( 2 )
st_ptr = 3
st_row = st_ptr + n + 1
st_val = 0
IF ( transpose ) THEN
    DO l = 1, n_free
        j = FREE( l )
        val = 0.0_wp
        DO k = userdata%integer( st_ptr + j ),
            userdata%integer( st_ptr + j + 1 ) - 1
            val = val + userdata%real( st_val + k ) *
                V( userdata%integer( st_row + k ) )
        END DO
        P( j ) = val
    END DO
ELSE
    P( : o ) = 0.0_wp
    DO l = 1, n_free
        j = FREE( l )
        val = V( j )
        DO k = userdata%integer( st_ptr + j ),
            userdata%integer( st_ptr + j + 1 ) - 1
            i = userdata%integer( st_row + k )
            P( i ) = P( i ) + userdata%real( st_val + k ) * val
        END DO
    END DO
END IF
status = 0
RETURN

```

---

**All use is subject to the conditions of a BSD-3-Clause License.**  
**See <https://www.galahad.rl.ac.uk/download/> for full details.**

END SUBROUTINE AFPROD

This produces the same output. Now notice how the matrix  $\mathbf{A}_o$  is passed to the matrix-vector product evaluation routines via the integer and real components of the derived type `USERDATA_type`.

Now, suppose instead that we wish to minimize

$$\frac{1}{2} \left\| \begin{pmatrix} x_1 \\ x_1 + x_2 - 2 \\ x_3 + x_2 - 1 \\ x_4 + x_3 - 2 \\ x_5 + x_4 - 1 \\ x_5 - 2 \end{pmatrix} \right\|_2^2$$

within the two non-overlapping simplices

$$x_1 + x_4 = 1, \quad x_2 + x_5 = 1, \quad x_1, x_2, x_4, x_5 \geq 0.$$

Then, on writing the data for this problem as

$$\mathbf{A}_o = \begin{pmatrix} 1 & & & & & \\ & 1 & & & & \\ & & 1 & & & \\ & & & 1 & & \\ & & & & 1 & \\ & & & & & 1 \end{pmatrix}, \quad \text{and } \mathbf{b} = \begin{pmatrix} 0 \\ 2 \\ 1 \\ 2 \\ 1 \\ 2 \end{pmatrix}$$

in sparse coordinate format, and noticing that the constraints involve two cohorts  $\{1,4\}$  and  $\{2,5\}$  with no constraint on variable  $x_3$ , we may use the following code:

```
! THIS VERSION: GALAHAD 5.5 - 2026-02-01 AT 08:40 GMT
PROGRAM GALAHAD_SLLS_MULTIPLE_EXAMPLE
USE GALAHAD_SLLS_double ! double precision version
IMPLICIT NONE
INTEGER, PARAMETER :: wp = KIND( 1.0D+0 ) ! set precision
TYPE ( QPT_problem_type ) :: p
TYPE ( SLLS_data_type ) :: data
TYPE ( SLLS_control_type ) :: control
TYPE ( SLLS_inform_type ) :: inform
TYPE ( USERDATA_type ) :: userdata
INTEGER :: i, j, m, s
INTEGER, PARAMETER :: n = 12, o = n + 1, a_ne = 2 * n + 2
! set problem data
ALLOCATE( p%B( o ), p%X( n ), p%X_status( n ), p%COHORT( n ) )
DO j = 1, n ! cohorts
  SELECT CASE( MOD( j, 3 ) )
    CASE ( 1 )
      p%COHORT( j ) = 1
    CASE ( 2 )
      p%COHORT( j ) = 2
    CASE DEFAULT
      p%COHORT( j ) = 0
  END SELECT
END DO
m = MAXVAL( p%COHORT ) ! number of cohorts
p%B( 1 ) = 0.0_wp ! observations
```

---

**All use is subject to the conditions of a BSD-3-Clause License.**  
**See <https://www.galahad.rl.ac.uk/download/> for full details.**

```

DO i = 2, o
  SELECT CASE( MOD( i, 3 ) )
  CASE ( 1 )
    p%B( i ) = 2.0_wp
  CASE DEFAULT
    p%B( i ) = 1.0_wp
  END SELECT
END DO
p%n = n ; p%o = o ; p%m = m ! dimensions
p%X = 0.0_wp ! start from zero
! sparse co-ordinate storage format for the design matrix Ao
CALL SMT_put( p%Ao%type, 'COORDINATE', s ) ! Co-ordinate storage for Ao
ALLOCATE( p%Ao%val( a_ne ), p%Ao%row( a_ne ), p%Ao%col( A_ne ) )
p%Ao%m = o ; p%Ao%n = n ; p%Ao%ne = 0
DO j = 1, n
  p%Ao%ne = p%Ao%ne + 1
  p%Ao%row( p%Ao%ne ) = j
  p%Ao%col( p%Ao%ne ) = j
  p%Ao%val( p%Ao%ne ) = 1.0_wp
  p%Ao%ne = p%Ao%ne + 1
  p%Ao%row( p%Ao%ne ) = j + 1
  p%Ao%col( p%Ao%ne ) = j
  p%Ao%val( p%Ao%ne ) = - 1.0_wp
END DO
p%regularization_weight = 0.1_wp
ALLOCATE( p%X_s( n ) )
p%X_s( n ) = 1.0_wp
! problem data complete
CALL SLLS_initialize( data, control, inform ) ! Initialize control parameters
control%print_level = 3 ! print one line/iteration
control%maxit = 20 ! maximum of 20 iterations
! control%maxit = 2 ! maximum of 2 iterations
control%SBLS_control%symmetric_linear_solver = 'sytr '
control%SBLS_control%definite_linear_solver = 'potr '
! control%direct_subproblem_solve = .FALSE.
inform%status = 1
CALL SLLS_solve( p, data, control, inform, userdata )
IF ( inform%status == 0 ) THEN ! Successful return
  WRITE( 6, "( /, ' SLLS: ', I0, ' iterations ', /,
    & ' Optimal objective value =',
    & ES12.4, /, ' Optimal solution =', /, ( 5ES12.4 ) )" )
    inform%iter, inform%obj, p%X
  WRITE( 6, "( ' Lagrange multiplier estimates =', ( 5ES12.4 ) )" ) p%Y
ELSE ! Error returns
  WRITE( 6, "( /, ' SLLS_solve exit status = ', I0 ) )" inform%status
  WRITE( 6, "( /, ' objective value =',
    & ES12.4, /, ' Current solution =', ( 5ES12.4 ) )" ) inform%obj, p%X
END IF
CALL SLLS_terminate( data, control, inform ) ! delete workspace
DEALLOCATE( p%B, p%X, p%Y, p%Z, p%R, p%G, p%X_status, p%COHORT )
DEALLOCATE( p%Ao%val, p%Ao%row, p%Ao%col, p%Ao%type )
END PROGRAM GALAHAD_SLLS_MULTIPLE_EXAMPLE

```

This produces the following output:

```
constraints are the intersection of 2 non-overlapping simplices
```

---

**All use is subject to the conditions of a BSD-3-Clause License.**  
**See <https://www.galahad.rl.ac.uk/download/> for full details.**

diagonal preconditioner, min, max = 2.1000E+00 2.1000E+00

S=steepest descent, F=factorization used

| #its | #cg | f                    | proj gr   | step | #free | change | time |
|------|-----|----------------------|-----------|------|-------|--------|------|
| 0    | -   | 1.18250000000000E+01 | 3.005E-01 | -    | -     | -      | 0.0  |

steepest descent search direction

| segment | phi_0      | phi_1       | phi_2      | t_break    | t_opt      | fixed |
|---------|------------|-------------|------------|------------|------------|-------|
| 1       | 1.1825E+01 | -1.1300E+00 | 2.2818E+00 | 2.3452E+00 | 4.9520E-01 | 1     |

phi\_opt = 1.15452191235060E+01 at t = 4.95203257511239E-01 in segment 1

| #its | #cg | f                    | proj gr   | step      | #free | change | time |
|------|-----|----------------------|-----------|-----------|-------|--------|------|
| 1    | S   | 1.15452191235060E+01 | 2.209E-01 | 2.112E-01 | 12    | -      | 0.0  |

steepest descent search direction

| segment | phi_0      | phi_1       | phi_2      | t_break    | t_opt      | fixed |
|---------|------------|-------------|------------|------------|------------|-------|
| 1       | 1.1545E+01 | -7.0167E-01 | 2.8606E+00 | 2.4038E+00 | 2.4529E-01 | 1     |

phi\_opt = 1.14591643421733E+01 at t = 2.45286442270197E-01 in segment 1

| #its | #cg | f                    | proj gr   | step      | #free | change | time |
|------|-----|----------------------|-----------|-----------|-------|--------|------|
| 2    | S   | 1.14591643421732E+01 | 1.771E-01 | 1.020E-01 | 12    | 0      | 0.0  |

augmented system search direction

| segment | phi_0      | phi_1       | phi_2      | t_break    | t_opt      | fixed |
|---------|------------|-------------|------------|------------|------------|-------|
| 1       | 1.1459E+01 | -9.7818E-01 | 6.2345E-01 | 3.0797E-01 | 1.5690E+00 | 1     |
| 2       | 1.1187E+01 | -1.4145E-01 | 5.2667E-01 | 7.7390E-01 | 5.7655E-01 | 2     |

phi\_opt = 1.11684827450622E+01 at t = 5.76548744080473E-01 in segment 2

| #its | #cg | f                    | proj gr   | step      | #free | change | time |
|------|-----|----------------------|-----------|-----------|-------|--------|------|
| 3    | F   | 1.11684827450622E+01 | 5.585E-02 | 2.732E-01 | 11    | 1      | 0.0  |

augmented system search direction

| segment | phi_0      | phi_1       | phi_2      | t_break    | t_opt      | fixed |
|---------|------------|-------------|------------|------------|------------|-------|
| 1       | 1.1168E+01 | -1.5215E-01 | 2.1161E+00 | 1.0895E+00 | 7.1899E-02 | 4     |

phi\_opt = 1.11630131355241E+01 at t = 7.18993656763907E-02 in segment 1

| #its | #cg | f                    | proj gr   | step      | #free | change | time |
|------|-----|----------------------|-----------|-----------|-------|--------|------|
| 4    | F   | 1.11630131355241E+01 | 8.049E-16 | 4.179E-02 | 11    | 0      | 0.0  |

SLLS: 4 iterations

Optimal objective value = 1.1163E+01

Optimal solution =

5.8870E-17 1.0175E-01 -2.8752E-01 2.9445E-01 3.0738E-01  
-1.5014E-01 3.7732E-01 3.4402E-01 -1.5607E-01 3.2823E-01  
2.4686E-01 -3.1102E-01

Lagrange multiplier estimates = -4.0150E-01 5.0119E-01

---

**All use is subject to the conditions of a BSD-3-Clause License.**  
**See <https://www.galahad.rl.ac.uk/download/> for full details.**