

NAG Library

How to Use the NAG Library and its Documentation

Contents

1	Library Identification	3
2	How to Find a NAG Library Function	3
3	How to Use the NAG Library	3
3.1	Structure of the Library	3
3.1.1	Experimental Routines	4
3.2	General Advice	4
3.3	Programming Advice	4
3.3.1	The NAG C environment	4
3.3.1.1	NAG data types	5
3.3.1.2	Memory management in the Library	5
3.3.1.3	The Nag_Order Argument	6
3.3.1.4	Array references	6
3.3.1.5	Internal data structures in the NAG C Library	11
3.3.1.6	Chapter header files	11
3.3.2	Direct and Reverse Communication Functions	12
3.4	Use of NAG Long Names	12
3.5	Input/Output	13
3.6	Auxiliary Functions	13
3.7	NAG Error Handling and the fail Argument	13
3.7.1	Use of NAGERR_DEFAULT	13
3.7.2	Use of the fail Argument	13
3.7.3	The NagError structure	14
3.7.4	Structure of the NAG error messages	15
3.7.5	License Management	15
3.7.6	Unexpected Errors	16
3.8	Calling the Library from Other Languages	16
3.9	Arithmetic Considerations and Reproducibility of Results	16
3.9.1	Bit-wise Reproducibility (BWR)	18
3.9.1.1	Vendor Math Libraries and Conditional Bitwise Reproducibility (CBWR)	18
3.10	Multithreading	19
3.10.1	Thread Safety	19
3.10.1.1	Functions with Function Arguments	19
3.10.1.2	Input/Output	19
3.10.1.3	Implementation Issues	19
3.10.2	Parallelism	20
3.10.2.1	Introduction	20
3.10.2.2	How is Parallelism Used in the NAG Library?	21
3.10.3	Multithreaded Functions	22
4	How to Use NAG Documentation	22
4.1	Using the Manual	22
4.2	Structure of the Documentation	23

4.3	Specification of Arguments	24
4.3.1	Classification of Arguments	24
4.3.2	Constraints and suggested values	24
4.4	Example Programs and Results	24
4.5	Online Documentation	25
4.5.1	HTML Format	25
4.5.1.1	Viewing HTML5 Files	25
4.5.1.2	Firefox (and other Mozilla based browsers)	25
4.5.1.3	Other Browsers	25
4.5.1.4	Navigating HTML5 Files	26
4.5.1.5	Printing HTML5 Files	26
4.5.1.6	Windows HTML Help	26
4.5.2	PDF Format	26
4.5.2.1	Viewing and Printing PDF Files	26
4.5.2.2	Navigating the PDF Files	26
5	Background to the Problems	27
6	NAG Library Standards	27
7	References	27

1 Library Identification

Periodically a new **Mark** of the NAG Library is released: new functions are added, corrections and/or improvements are made to existing functions; and, occasionally, functions are withdrawn if they have been superseded by improved functions. A point release such as 26.1, only differs from a new Mark (e.g., 27) in that no functions are withdrawn.

You must know **which implementation** and **which mark and revision** of the Library you are using or intend to use. To find out these Library details, you can run a program which calls the NAG C Library function **nag_implementation_details (a00aac)**.

This function has no arguments; it simply outputs text to the standard output, if available. Details can also be found by a call to **nag_implementation_separated_details (a00adc)**.

2 How to Find a NAG Library Function

All users both familiar or unfamiliar with this Library who are thinking of using a function from it, are asked to please follow these instructions:

- (a) read **How to Use the NAG Library and its Documentation** as it provides valuable background information;
- (b) select an appropriate chapter or function by searching through the **Keyword and GAMS Search**;
- (c) read the relevant **Chapter Introduction**;
- (d) choose a function, and read the **function document**. If the function does not after all meet your needs, return to step (b);
- (e) read the **Users' Note** for your implementation (this contains instructions on how to compile and run a program);
- (f) consult local documentation, which should be provided by your local support staff, about access to the Library on your computing system;
- (g) obtain a copy of the example program (see Section 4.4) for the particular function of interest and experiment with it.

You should now be in a position to include a call to the function in a program, and to attempt to compile and run it. You may of course need to refer back to the relevant documentation in the case of difficulties, for advice on assessment of results, and so on.

As you become familiar with the Library, some of steps (a) to (g) can be omitted, but it is useful to keep up to date with the following documents as they are subject to change:

- How to Use the NAG Library and its Documentation**;
- the **Chapter Introduction**;
- the **function document**;
- the **Users' Note** for your implementation.

3 How to Use the NAG Library

3.1 Structure of the Library

The NAG Library is a comprehensive collection of **functions** for the solution of numerical and statistical problems.

The Library is divided into **chapters**, each devoted to a branch of numerical analysis or statistics. Each chapter has a three-character name and a title, e.g.,

Chapter d01 – Quadrature

Exceptionally, Chapters h and s have one-character names. The chapters and their names are based on the ACM modified SHARE classification index (see ACM (1960–1976)).

All documented functions have two names. One is based on the SHARE index classification and consists of a six-character name which begins with the characters of the chapter/subchapter name, for example

c06pcc

The letters of this type of function name are always lower case, the second and third characters being digits and the last letter being c. This function name is referred to as the short name. Each function also has a more meaningful and longer name, for example

nag_sum_fft_complex_1d

which we refer to as the long name. The long name may be used as an alternative to the short name when calling the function. See Section 3.4 for further details.

3.1.1 Experimental Routines

Some functions in the Library may be classified as experimental. These functions will be flagged by a note at the top of the documentation.

Functions may be classified as experimental if:

- (a) The interface and / or functionality of the function may change between Marks of the Library. The function will have been designed in such a way as to minimize the number of such changes.
- (b) The complexity of the function, or suite of functions it is part of, are such that comprehensive testing is not practical. Functions classified as experimental have gone through the same testing and review processes as a normal Library function, however it is recommended that additional care is taken when using the function.

A function classified as experimental may be reclassified as no longer being experimental, at which point the interface will become fixed as per a normal Library function.

3.2 General Advice

A NAG Library function **cannot** be guaranteed to return meaningful results irrespective of the data supplied to it. Care and thought **must** be exercised in:

- (a) formulating the problem;
- (b) programming the use of Library functions;
- (c) assessing the significance of the results.

The remainder of this document is concerned with (b) and (c).

3.3 Programming Advice

The Library and its documentation are designed with the assumption that you will write a calling program in C (although it may be called from other languages – see Section 3.8).

When a suitable NAG function has been selected, (see Section 4.1) the function must be called from the C Library via a suitable user-written program, the calling program. This manual assumes that you have sufficient knowledge of the C programming language to be able to write such a program. Each C Library function document contains an example of a suitable calling program (see Section 4.4).

When writing a calling program, a number of environmental features common to all such NAG programs must be observed in addition to specific features which are relevant to the particular NAG function being called. These features are discussed below; you should also refer to the example program.

3.3.1 The NAG C environment

The environment for the NAG Library is defined in a number of include files; a list is given in Section 3.3.1.6. The most important of the header files is `nag.h`, which must be included in any program that calls a NAG Library function and must precede any other NAG header file.

These include files are placed in <product_folder>/include folder by the installer. For its exact location please see the **Users' Note** or other local documentation.

The file `nag.h` defines data types and error codes used in the NAG Library together with a number of macros used in example programs.

You will also need to include the header file `nag_stdlib.h` in the calling program, if any memory management is required; see Section 3.3.1.2.

3.3.1.1 NAG data types

Integer

This data type is used for integer arguments to NAG Library functions. It is normally defined to be `long`. Please refer to the Users' Note for its exact type.

Nag_Boolean

This is an enumeration type defined as

```
typedef enum{Nag_FALSE=0; Nag_TRUE=1} Nag_Boolean;
```

Complex

This data type is a structure defined for use with complex numbers:

```
typedef struct {double re, im;} Complex;
```

Pointer

This data type represents the generic pointer `void *`.

Nag_Comm

This data type is a structure with a number of fields among which are a generic pointer, an Integer pointer and a double pointer which may be used for communicating information between a user-defined function and your calling program, where the user-defined function is supplied as an argument to the NAG function. This avoids the necessity of using global variables for such communication. The use of Integer and double pointers allows double and Integer arrays to be communicated with no casting. For an example of use, please see Section 10 in **nag_opt_simplex_easy (e04cbc)**.

Nag_User

This is also a communication structure with a generic pointer with usage similar to `Nag_Comm`.

Nag_FileID

This data type is an integer type used by certain NAG Library functions which deal with input or output files. Functions in Chapter x04 must be used to associate a file name with the file ID.

Enumeration Types

A number of other enumeration types are defined in `nag_types.h` (included by `nag.h`) for use in calls to various NAG Library functions. You should use these enumeration types in your C/C++ calling programs. If the NAG Library is being called from languages other than C/C++ we have provided a function to map the enumeration members to integer values (see **nag_enum_name_to_value (x04nac)**).

Other structure types

A number of structures have been defined to facilitate calls to NAG functions. Please note that only those components of a structure that are relevant to a call to a specific NAG function are described in the corresponding function document. A complete specification of a NAG defined structure can be found in the NAG Library header file `nag_types.h` included in the distribution materials.

3.3.1.2 Memory management in the Library

Memory is frequently dynamically allocated within NAG Library functions. All requests for memory are checked for success or failure. In the unlikely event of failure occurring the Library function returns or terminates with the error state **NE_ALLOC_FAIL** (details of error handling in the Library are given in Section 3.7).

The macros `NAG_ALLOC`, `NAG_REALLOC` and `NAG_FREE` are defined to select suitable memory management functions for the NAG Library. `NAG_ALLOC` has two arguments; the first specifies the number of elements to be allocated while the second specifies the type of element. The statement

```
p = NAG_ALLOC(n, double);
```

allocates `n` elements of memory of type `double` to `p`, a pointer to `double`.

`NAG_REALLOC` has three arguments; the first specifies the name of the pointer whose memory is to be extended, the second specifies the number of elements and the third specifies the type of element.

The statement

```
p = NAG_REALLOC(p, n, double);
```

allocates `n` elements of memory of type `double` to `p`, a pointer to `double`.

`NAG_FREE` frees memory allocated by `NAG_ALLOC` or `NAG_REALLOC`; its single argument is the pointer which specifies the memory to be deallocated. The statement

```
NAG_FREE(p);
```

deallocates memory pointed to by `p` and sets its value to `NULL`.

These macros are defined in the header file `nag_stdlib.h` which must be included if these macros are used in the calling program. `NAG_FREE` must be used to free memory allocated and returned from a NAG function. If memory is allocated using `NAG_ALLOC` for whatever reason, it must be freed using `NAG_FREE`. For an illustration of their use, see Section 10.1 in **nag_1d_cheb_fit_constr (e02agc)**. The use of `NAG_ALLOC`, `NAG_REALLOC` and `NAG_FREE` is strongly recommended.

3.3.1.3 The Nag_Order Argument

Different programming languages lay out two-dimensional data in memory in different ways. The C language treats a two-dimensional array as a single block of memory arranged in rows, the so called 'row-major' ordering. Some other languages, notably Fortran, arrange two-dimensional arrays by column ('column-major' ordering). Those functions in the NAG C Library that deal with two-dimensional arrays and where deemed appropriate have an extra argument, called **order**, which allows you to specify that your data is arranged in rows (by setting order to `Nag_RowMajor`) or in columns (by setting order to `Nag_ColMajor`). This is particularly useful if the NAG C Library is being called from a language which uses the column-major ordering or if you wish to call a function from a language, such as Visual Basic 7 onwards, which supports column-major ordering.

3.3.1.4 Array references

In C it is possible to declare a two-dimensional variable using notation of the form:

```
double a[dim1][dim2];
```

When this variable is an argument to a function, it is effectively treated by the compiler as a pointer, `*a`, of type `double` with an allocated memory of `dim1*dim2` on the stack. The address of an element of this array, say `a[3][5]` is then an explicit address computed to be `*(a+3*dim2+5)`, since C stores data in row-major order.

Alternatively it is possible for you to allocate memory explicitly (on the heap) to a pointer of type `double *`, using the form:

```
a=(double *)malloc((size_t)(dim1*dim2*sizeof(double)));
```

In this case the C preprocessor allows a succinct notation for computing this explicit address by using a macro definition:

```
#define A(I,J) a[I*dim2+J]
```

The element of this array if indexed `ij` is then indexed using the pointer notation `*(a+i*dim2+j)` or by using the array notation `a[i*dim2+j]`; or by using `A(I,J)` assuming the macro [2] is already defined.

We often wish to refer to the storage of elements representing a submatrix of the matrix `A`, for example, the submatrix comprising rows `k` to `l` and columns `m` to `n`. For convenience we use the notation `A(k:l,`

$m:n$) to refer to those elements of a storing the given submatrix with elements $A_{ij} = A(i-1, j-1)$ (see [2] and [4]), for $i = k, \dots, l$ and $j = m, \dots, n$. That is,

$$A(k:l, m:n) = \{a[p], p = (i-1)*pda+j-1, i=k, \dots, l \text{ and } j=m, \dots, n\},$$

for row major ordering.

If the data is to be stored using column-major ordering and we have declared an array variable as

```
double a[dim1][dim2];
```

then the element indexed ij is effectively transposed. That is, the element $a[i][j]$ under row-major ordering is the element $a[j][i]$ under column-major ordering.

As another alternative you may choose to malloc the required memory as in [1] above. In this case, the element indexed ij is using the pointer notation $*(a+j*\text{dim1}+i)$; or by using the array notation $a[j*\text{dim1}+i]$; or by using $A(I, J)$ if the macro is defined as

```
#define A(I, J) A[J*dim1+I]
```

Note the difference in definition between [2] and [4] above.

In order to simplify the documentation, we refer to array elements using the macro definitions above. Further, we note that in the preprocessor directive [2] and [4], the critical dimension is either dim2 if the row-major ordering is used or dim1 if column-major ordering is used. We designate either dim2 or dim1 as the principal dimension depending on the storage ordering scheme. The principal dimension can be thought of as a stride which must be taken to traverse either a row or a column depending on the **order** argument.

Typically in the NAG Library we use the convention '**pda**' if the function has the **order** argument and **pda** can have different values depending on whether the array is stored in row major or in column major order. If the specification of an array is such that its elements must be stored in row major order we use the term '**tda**' (meaning trailing dimension). For arrays whose elements must be stored in column major order we use the term '**lda**' (meaning leading dimension). As of Mark 24, new functions in the library generally store two-dimensional data in the column order. We are standardizing on the phrase 'stride separating row elements' to mean column order storage and the phrase 'stride separating column elements' to mean row order storage.

In order to facilitate calling functions in which data has to be stored in a mutually exclusive manner, such as for example, function A requires data to be in row order and while function B requires data to be stored in column order and the 'order' parameter has not been provided, then functions provided in the f16 chapter will have to be used. For example, **nag_dge_copy (f16qfc)** can be used to change from row order to column order by performing a transposed copy. Functions are available in this chapter for performing a variety of copying tasks such as triangular copy, etc..

We illustrate these concepts using two examples. In the first example, memory is allocated while in the second example memory is declared. In both examples, row or column modes are demarcated using the preprocessor macro `NAG_ROW_MAJOR`. Memory is allocated using `NAG_ALLOC`, which is a macro defined in `nag_stdlib.h`, in preference to explicit malloc calls. This macro maps to calls to internal Library functions to allocate memory. Also note the use of `NAG_FREE` to free the memory.

Example 1

```
/* Example Program, with memory allocated, based on:
 *
 * nag_dorgqr (f08afc) Example Program.
 *
 * Copyright Numerical Algorithms Group.
 *
 */

#include <stdio.h>
#include <string.h>
#include <nag.h>
#include <nag_stdlib.h>
#include <nagf08.h>
#include <nagx04.h>
```

```

int main(void)
{
    /* Scalars */
    Integer i, j, m, n, pda_row, pda_column, tau_len;
    Integer exit_status=0;
    NagError fail;

    /* Arrays */
    char *title=0;
    double *a_row=0, *a_column=0, *tau=0;
    char matrix_data [] = {
        " -0.57 -1.28 -0.39  0.25 "
        " -1.93  1.08 -0.31 -2.14 "
        "  2.30  0.24  0.40 -0.35 "
        " -1.93  0.64 -0.66  0.08 "
        "  0.15  0.30  0.15 -2.13 "
        " -0.02  1.03 -1.43  0.50 "
    }, *matrix_data_ptr = matrix_data;

    /* Initialize strtok */
    matrix_data_ptr = strtok(matrix_data_ptr, " \t\n");

#define A_COLUMN(I,J) a_column[(J-1)*pda_column + I - 1]
#define A_ROW(I,J) a_row[(I-1)*pda_row + J - 1]

    INIT_FAIL(fail);

    m = 6;
    n = 4;;
    pda_column = m;
    pda_row = n;
    tau_len = MIN(m, n);

    /* Allocate memory */
    if ( !(title = NAG_ALLOC(31, char)) ||
        !(a_row = NAG_ALLOC(m * n, double)) ||
        !(a_column = NAG_ALLOC(m * n, double)) ||
        !(tau = NAG_ALLOC(tau_len, double)) )
    {
        printf("Allocation failure\n");
        exit_status = -1;
        goto End;
    }

#ifdef NAG_ROW_MAJOR
    printf("Using row major storage, allocated memory\n");
    /* Read A from data above */
    for (i = 1; i <= m; ++i)
    {
        for (j = 1; j <= n; j++)
        {
            sscanf(matrix_data_ptr, "%lf", &A_ROW(i,j));
            matrix_data_ptr = strtok(0, " \t\n");
        }
    }

    /* Compute the QR factorization of A */
    f08aec(Nag_RowMajor, m, n, a_row, pda_row, tau, &fail);
    if (fail.code != NE_NOERROR)
    {
        printf("Error from f08aec.\n%s\n", fail.message);
        exit_status = 1;
        goto End;
    }

    /* Form the leading N columns of Q explicitly */
    f08afc(Nag_RowMajor, m, n, n, a_row, pda_row, tau, &fail);
    if (fail.code != NE_NOERROR)
    {
        printf("Error from f08afc.\n%s\n", fail.message);
        exit_status = 2;
    }

```



```

        goto End;
    }
    /* Print the leading N columns of Q only */
    sprintf(title, "The leading %2ld columns of Q\n", n);
    x04cac(Nag_RowMajor, Nag_GeneralMatrix, Nag_NonUnitDiag, m, n,
a_row, pda_row, title, 0, &fail);
    if (fail.code != NE_NOERROR)
    {
        printf("Error from x04cac.\n%s\n", fail.message);
        exit_status = 3;
        goto End;
    }
#else
    printf("Using column major storage, allocated memory\n");
    /* Read A from data above */
    for (i = 1; i <= m; ++i)
    {
        for (j = 1; j <= n; j++)
        {
            sscanf(matrix_data_ptr, "%lf", &A_COLUMN(i,j));
            matrix_data_ptr = strtok(0, " \t\n");
        }
    }

    f08aec(Nag_ColMajor, m, n, a_column, pda_column, tau, &fail);
    if (fail.code != NE_NOERROR)
    {
        printf("Error from f08aec.\n%s\n", fail.message);
        exit_status = 1;
        goto End;
    }
    /* Form the leading N columns of Q explicitly */
    f08afc(Nag_ColMajor, m, n, n, a_column, pda_column, tau, &fail);
    if (fail.code != NE_NOERROR)
    {
        printf("Error from f08afc.\n%s\n", fail.message);
        exit_status = 2;
        goto End;
    }
    /* Print the leading N columns of Q only */
    sprintf(title, "The leading %2ld columns of Q\n", n);
    x04cac(Nag_ColMajor, Nag_GeneralMatrix, Nag_NonUnitDiag, m, n,
        a_column, pda_column, title, 0, &fail);
    if (fail.code != NE_NOERROR)
    {
        printf("Error from x04cac.\n%s\n", fail.message);
        exit_status = 3;
        goto End;
    }
}
#endif
End:
    if (title) NAG_FREE(title);
    if (a_row) NAG_FREE(a_row);
    if (a_column) NAG_FREE(a_column);
    if (tau) NAG_FREE(tau);

    return exit_status;
}

```

Example 2

```

/* Example Program, with memory declared, based on:
 *
 * nag_dorgqr (f08afc) Example Program.
 *
 * Copyright Numerical Algorithms Group.
 *
 */

#include <stdio.h>
#include <string.h>

```

```

#include <nag.h>
#include <nag_stdlib.h>
#include <nagf08.h>
#include <nagx04.h>

#define MMAX 10
#define NMAX 8

int main(void)
{
    /* Scalars */
    Integer i, j, m, n, pda, tau_len;
    Integer exit_status=0;
    NagError fail;

    /* Arrays */
    char title[30];
    double a_row[MMAX][NMAX], a_column[NMAX][MMAX], tau[NMAX];
    char matr_data [] = {
        " -0.57 -1.28 -0.39  0.25 "
        " -1.93  1.08 -0.31 -2.14 "
        "  2.30  0.24  0.40 -0.35 "
        " -1.93  0.64 -0.66  0.08 "
        "  0.15  0.30  0.15 -2.13 "
        " -0.02  1.03 -1.43  0.50 "
    }, *matrix_data_ptr = matr_data;

    /* Initialize strtok */
    matrix_data_ptr = strtok(matrix_data_ptr, " \t\n");

    INIT_FAIL(fail);

    m = 6;
    n = 4;
    pda = NMAX;
    tau_len = MIN(m, n);

    /* Read A from data above */
#ifdef NAG_ROW_MAJOR
    for (i = 0; i < m; ++i)
    {
        for (j = 0; j < n; j++)
        {
            sscanf(matrix_data_ptr, "%lf", &a_row[i][j]);
            matrix_data_ptr = strtok(0, " \t\n");
        }
    }

    /* Compute the QR factorization of A */
    printf("Using row major storage, declared memory\n");
    f08aec(Nag_RowMajor, m, n, &a_row[0][0], pda, tau, &fail);
    if (fail.code != NE_NOERROR)
    {
        printf("Error from f08aec.\n%s\n", fail.message);
        exit_status = 1;
        goto End;
    }
    /* Form the leading N columns of Q explicitly */
    f08afc(Nag_RowMajor, m, n, n, &a_row[0][0], pda, tau, &fail);
    if (fail.code != NE_NOERROR)
    {
        printf("Error from f08afc.\n%s\n", fail.message);
        exit_status = 2;
        goto End;
    }
    /* Print the leading N columns of Q only */
    sprintf(title, "The leading %2ld columns of Q\n", n);
    x04cac(Nag_RowMajor, Nag_GeneralMatrix, Nag_NonUnitDiag, m, n,
    &a_row[0][0], pda, title, 0, &fail);
    if (fail.code != NE_NOERROR)

```

```

    {
        printf("Error from x04cac.\n%s\n", fail.message);
        exit_status = 3;
        goto End;
    }
#else
    printf("Using column major storage, declared memory\n");
    for (i = 0; i < m; ++i)
        for (j = 0; j < n; j++)
            {
/* Note column data is transposed */
                sscanf(matrix_data_ptr, "%lf", &a_column[j][i]);
                matrix_data_ptr = strtok(0, " \t\n");
            }

    f08aec(Nag_ColMajor, m, n, &a_column[0][0], pda, tau, &fail);
    if (fail.code != NE_NOERROR)
        {
            printf("Error from f08aec.\n%s\n", fail.message);
            exit_status = 1;
            goto End;
        }
/* Form the leading N columns of Q explicitly */
    f08afc(Nag_ColMajor, m, n, &a_column[0][0], pda, tau, &fail);
    if (fail.code != NE_NOERROR)
        {
            printf("Error from f08afc.\n%s\n", fail.message);
            exit_status = 2;
            goto End;
        }
/* Print the leading N columns of Q only */
    sprintf(title, "The leading %2ld columns of Q\n", n);
    x04cac(Nag_ColMajor, Nag_GeneralMatrix, Nag_NonUnitDiag, m, n,
&a_column[0][0], pda, title, 0, &fail);
    if (fail.code != NE_NOERROR)
        {
            printf("Error from x04cac.\n%s\n", fail.message);
            exit_status = 3;
            goto End;
        }
#endif
End:
    return exit_status;
}

```

3.3.1.5 Internal data structures in the NAG C Library

For efficiency, and wherever possible, the NAG C Library is designed to make use of the Basic Linear Algebra Subprograms (BLAS), a suite of low-level functions tuned by many computer manufacturers for their particular hardware. Since the BLAS are specified in Fortran, and therefore use the column-major storage order, the NAG C Library also uses this scheme, internally, for new functions wherever it is practical. Thus any two-dimensional arrays you have provided may be re-ordered on entry and/or on exit from the NAG functions, as appropriate. It is therefore slightly more efficient to use the column-major ordering; however, except for very large data sets, the effect is negligible in practice.

3.3.1.6 Chapter header files

Chapter header files contain the function declarations for the NAG Library with ANSI function prototyping. The appropriate chapter header file must be included for each NAG function called by your program. For example, to call the function **nag_sum_fft_complex_1d (c06pcc)** the chapter header file `nagc06.h` must be included as

```
#include <nagc06.h>
```

The naming convention is to prefix the first three characters of the function name in lower case by `nag` and use `.h` as the postfix as in normal C practice, except that all functions in Chapter *s* use the header file `nags.h`.

(a) Header files intended for your inclusion within calling programs to the NAG Library

`nag.h` defines the basic environment for use of the NAG Library. This header file must be included in each calling program to the NAG Library and must precede all other header files that are included. This must be followed by one or more of the following chapter header files.

```
naga00.h  naga02.h  nagc02.h  nagc05.h  nagc06.h  nagc09.h
nagd01.h  nagd02.h  nagd03.h  nagd04.h  nagd05.h  nagd06.h
nage01.h  nage02.h  nage04.h  nage05.h  nagf01.h  nagf02.h
nagf03.h  nagf04.h  nagf06.h  nagf07.h  nagf08.h  nagf11.h
nagf12.h  nagf16.h  nagg01.h  nagg02.h  nagg03.h  nagg04.h
nagg05.h  nagg07.h  nagg08.h  nagg10.h  nagg11.h  nagg12.h
nagg13.h  nagh02.h  nagh03.h  nagm01.h  nags.h    nagx01.h
nagx02.h  nagx04.h  nagx06.h  nagx07.h
```

`nag_stdlib.h` defines the memory allocation macro `NAG_ALLOC` and `NAG_FREE`. You must include this header file if the NAG definitions of `NAG_ALLOC` and `NAG_FREE`, as used in the example programs, are required.

(b) The following three header files are included by `nag.h` (you do not need to supply a specific statement to include them)

`nag_types.h` defines the NAG types used in the Library.

`nag_errlist.h` defines the NAG error codes and messages used in the Library.

`nag_names.h` maps the NAG long names to short names.

3.3.2 Direct and Reverse Communication Functions

Functions in the Library that require a user-supplied function may be classified as either direct communication or reverse communication.

Direct communication functions require a user-supplied function to be provided as an actual argument to the NAG Library function. You must write this function using a very rigid interface as specified in the relevant function document. For the majority of applications this is the simplest and most convenient usage. Sometimes however this approach can be restrictive:

- (i) when the required format of the function does not allow useful information to be passed conveniently to and from your calling program;
- (ii) when the direct communication function is being called from another computer language which does not fully support procedure arguments in a way that is compatible with the Library.

These restrictions can be removed by using a reverse communication function. Instead of obtaining the solution in one call, reverse communication functions perform one step of the solution process before returning to the calling program with an appropriate flag (`irevcm`) set. The value of `irevcm` determines whether the process has finished or whether fresh information is required. In the latter case the required information must be calculated before re-entering the reverse communication function. Thus you have the responsibility for providing an iterative loop. Although reverse communication functions will typically be more complicated to use than direct communication equivalents they do provide greater flexibility for the evaluation of the function.

3.4 Use of NAG Long Names

The long names defined in the header file `nag_names.h` are `#defines`. As the header file `nag_names.h` is already included via `nag.h`, you need not include `nag_names.h` in their calling programs.

3.5 Input/Output

NAG Library functions output all error and warning messages to the C standard error stream `stderr`. Chapters e04, e05, g02 and g13 will optionally output results to the C standard output stream `stdout` or to an alternative user-specified file. A number of functions in Minimizing or Maximizing a Function (Chapter e04) and Operations Research (Chapter h) read input from external files.

3.6 Auxiliary Functions

In addition to the documented functions, the NAG Library contains a much larger number of auxiliary functions. You do not normally need to concern yourself with these functions, as they will automatically be called as required by the user-callable function you have selected.

3.7 NAG Error Handling and the `fail` Argument

All functions that have error exits have an argument that allows you control over the printing of error messages when an error is detected. There is a further option which allows you to either continue running your program, having returned from the NAG function, or to stop with either an exit statement or an abort within the NAG function. The different ways of using these error handling facilities are described below.

Note that in some implementations, the Library is linked with the vendor library containing LAPACK functions and the Chapters f07 and f08 function interfaces, where appropriate, act as wrappers to the corresponding vendor LAPACK functions. In this case, the **fail** argument passed through the Chapter f07 and Chapter f08 interfaces does not have full control over the printing of error messages; nor does it determine whether or not control is returned to the calling program when an error is detected.

3.7.1 Use of `NAGERR_DEFAULT`

The simplest method of using the error handling facility is to put `NAGERR_DEFAULT` in place of the **fail** argument in calls to the NAG C functions. If an error is detected the appropriate NAG error message is output on `stderr` and the program is stopped by the use of `exit`.

3.7.2 Use of the `fail` Argument

The two remaining ways of using the NAG error handling facility both involve defining the **fail** argument in the calling program. The **fail** argument is of type `NagError` which is a structure fully defined in `nag_types.h`. The fields of this structure of relevance to you as a user of the NAG C Library are:

```
int code;
Nag_Boolean print;
char message[NAG_ERROR_BUF_LEN];
Integer errnum;
void (*handler)(char*,int,char*);
```

where the symbol `NAG_ERROR_BUF_LEN` is normally defined to be 512.

This structure will contain the NAG error code and message on return from a call to a NAG Library function. The NAG error codes and some associated NAG error messages are defined in `nag_err-list.h`. A detailed description of the individual members of this structure is given below (see Section 3.7.3).

The NAG error argument **fail** is declared in the calling program as:

```
NagError fail;
```

The address of the argument is then passed to the NAG C function being called. Relevant members of the structure must be initialized before passing the argument to the called function, even though you may not actually require all members. It is recommended that the NAG defined macro `INIT_FAIL` be used for this purpose.

The `INIT_FAIL` macro sets:

```

fail.code = NE_NOERROR
fail.print = Nag_FALSE
fail.errnum = 0
fail.handler = 0

```

The `SET_FAIL` macro is also available. It sets the contents of `fail` in the same way as `INIT_FAIL` except that `fail.print` is set to `Nag_TRUE`.

(a) Use of the `fail` argument with the `print` member set to `Nag_TRUE`

If you require that the NAG error message be printed when an error is found, but that the called function should return control to the calling program, then the `fail` argument must be declared with all members initialized and the `print` member set to `Nag_TRUE`. Use of the NAG-defined macro `SET_FAIL` with the statement `SET_FAIL(fail);` performs the appropriate assignments. Alternatively the initialization could be done by declaring the `fail` argument with `static` and then setting `fail.print` to `Nag_TRUE`.

If no error occurs, `fail.code` will contain the error code `NE_NOERROR` on return from the called function. However, if an error is found, the appropriate NAG error message will be output on `stderr` before returning control to the calling program; `fail.code` will contain the relevant NAG error code. You must ensure that the calling program tests the `code` member of the `fail` argument on return from the NAG C function; you may then choose whether to exit the calling program or continue. See the example program for `nag_zero_nonlin_eqns_easy (c05qbc)` for such a case. The option of continuing may be advantageous if the results being returned are of some value even when an error has been detected. In the case of `nag_zero_nonlin_eqns_easy (c05qbc)` the code could be altered to allow the program to continue if the specific error codes `NE_TOO_MANY_FEVALS`, `NE_TOO_SMALL` and `NE_NO_IMPROVEMENT` occur, as in such a case useful partial results are returned (see the function document for `nag_zero_nonlin_eqns_easy (c05qbc)`).

(b) Use of the `fail` argument with the `print` member set to `Nag_FALSE`

If you do not wish the NAG error messages to be printed automatically when an error is found then the `fail` argument must be declared with all members initialized and the `print` member set to `Nag_FALSE`. Use of `static` in the declaration of `fail` will automatically leave the `print` member as `Nag_FALSE` as will the use of `INIT_FAIL(fail)`.

This method is suitable for those of you who wish to produce your own error messages rather than use the NAG Library versions. Alternative error messages may be coded directly into the calling program or be produced via a user-written error-handling function which is assigned to the `handler` member of the `fail` argument (see the description of the `handler` member below).

3.7.3 The `NagError` structure

The members of the `NagError` structure, of relevance to you as the user of the NAG C Library, are described in full below.

`code`

On successful exit, `code` contains the NAG error code `NE_NOERROR`; if an error or warning has been detected, then `code` contains the specific error or warning code. Error codes are prefixed with `NE_` whereas warning codes have the prefix `NW_`.

`print`

`print` must be set before calling any NAG Library function with a `fail` argument. It should be set to `Nag_TRUE` if the NAG error message is to be printed, otherwise `Nag_FALSE`. It is not changed by the NAG Library function.

`message`

On successful exit the array `message` contains the character string `"NE_NOERROR:\n No error"`. If an error has been detected, then `message` contains the error message text, whether or not this is printed.

`errnum`

On successful exit, `errno` is unchanged. For certain error or warning exits `errno` will contain a value specifying additional information concerning the error. For example if a vector is supplied incorrectly, then `errno` may specify which component of the vector is wrong. Cases where `errno` returns information are described in the relevant function documents.

`handler`

`handler` must be set to 0 if control is to be returned to the calling function after an error has been detected. Otherwise it must point to a user-supplied error-handling function. An example of the ANSI C declaration of a user-supplied error function (here called `errhan`) is:

```
void errhan(const char *string, int code, const char *name)
```

where `string` contains the NAG error message on input, `code` is the NAG error code and `name` is the short name of the NAG Library function which detected the error. If `print` (see above) is `Nag_TRUE`, then the NAG error message is printed before the user-supplied error handler is called. If the user-supplied error handler returns control, then the NAG error handler will return control to the calling program; otherwise the user-supplied error handler may exit.

An elementary example of where this feature might be used is if it is preferred to print error messages on `stdout` rather than the default `stderr`. In this case `errhan` could be defined as:

```
void errhan(const char *string, int code, const char *name)
{
    if (code != NE_NOERROR)
    {
        printf("\nError or warning from %s.\n", name);
        printf("%s\n", string);
    }
}
```

3.7.4 Structure of the NAG error messages

For illustrative purposes, let us consider two examples of the format of the NAG Library error messages in the NAG Library documentation:

NE_INT

On entry, `n = <value>`.
Constraint: `n > 1`.

NE_BAD_PARAM

On entry, argument `<value>` had an illegal value.

If the NAG function in question detects an error and error messages are being displayed, either by using the default error handler `NAGERR_DEFAULT` or by setting `fail.print = Nag_TRUE`, then text of the following form would be displayed:

```
NE_INT:
On entry, n = 1
Constraint: n > 1.
```

or

```
NE_BAD_PARAM:
On entry, argument order had an illegal value.
```

i.e., the notation `<value>` appearing in the documented error message is a place holder that will be populated by the value of a variable, argument name or some other piece of information when that error message is displayed.

3.7.5 License Management

If your implementation is license managed then your local site will have details on how the license management is implemented; please contact your site installer. To determine whether a valid license is available on your machine run the example program for `nag_licence_query (a00acc)`.

Should a valid license not be found when calling license managed functions from the Library then the function returns or terminates with the error state `NE_NO_LICENCE`. On Unix based systems, the appropriate environment variables should then be checked (e.g., `NAG_KUSARI_FILE`) to make sure this points to the licence file containing a valid licence and the licence file should be checked for any obvious errors (e.g., the licence refers to a different implementation). If everything appears to be correct then please contact NAG (see Support from NAG for details).

3.7.6 Unexpected Errors

Internal calls to Library functions are checked for error exits even when these exits are not to be expected. Should an unexpected error exit occur the function returns or terminates with the error state `NE_INTERNAL_ERROR`.

3.8 Calling the Library from Other Languages

In general the NAG Library can be called from other computer languages (such as C++, C#, Java, Visual Basic and Python) provided that appropriate mappings exist between the NAG Library data types and their data types (see NAG C Library Associated Information (<http://www.nag.co.uk/numeric/CL/classocinfo.asp>)).

3.9 Arithmetic Considerations and Reproducibility of Results

The results obtained when calling a NAG Library function depend not only on the algorithm used to solve the problem, but also on the compiler used to build the library, compiler run-time libraries and also the arithmetic properties of the machine on which the code is run.

Historically, different kinds of computer hardware tended to have different kinds of arithmetic. Some machines would store floating-point numbers using a base 16 significand and exponent system, others would use base 2, and some even used base 8 or 10. Such differences caused major headaches for software library providers because code that worked well on one arithmetic system might not behave in exactly the same way on another. This meant that great care had to be taken to make the library code **portable**.

In addition, it was not unheard of for machine arithmetic to have flaws or errors where basic operations such as multiplication or division could sometimes give incorrect results, especially on numbers that were in some way 'extreme', such as being very large or small.

After the first of the IEEE standards for floating-point arithmetic (ANSI/IEEE (1985)) was introduced in the 1980s, the situation improved greatly. Nowadays most significant hardware, and certainly most hardware that NAG libraries run on, will use IEEE-style base 2 arithmetic. This makes production of portable code easier, but there are still problems, partly due to the latitude allowed by the IEEE standards. For example, hardware which uses extra-precise 80-bit internal registers for arithmetic, as originally introduced in the Intel 8087 coprocessor in the 1980s, behaves slightly differently from hardware that uses 64-bit registers, particularly if a compiler generates optimized code which holds arithmetic subexpressions in the extra-precise registers.

Since, for performance reasons, computer arithmetic is generally finite precision (as is certainly the case for IEEE standard arithmetic) most of the numerical methods implemented by NAG Library functions can only return an approximation to the true solution, simply due to the accumulation of rounding errors.

It should therefore be clear that running a program which calls a NAG Library function with the same data on two different machines can give different results, due to compiler, hardware and run-time library considerations. Usually these differences are small – it may be that a result computed on one machine differs only in the last few significant bits from the same result computed on another machine – for example, when solving a well-conditioned set of linear equations on two different machines. Occasionally small differences may be magnified, for example if a conditional test depends on an imprecise result. A function that searches for a minimum of an optimization problem may converge to a different local minimum, but in general, so long as the function's documentation doesn't claim that the **same** local minimum will always be obtained, this should be acceptable. Even if an algorithm converges to the same local minimum, arithmetic differences may mean that a different number of iterations is taken to get there.

Modern hardware and optimizing compilers have introduced further scope for arithmetic quirks. An example is in the use of **Streaming SIMD Extension (SSE)** instructions. These low-level machine instructions allow hardware to operate on more than one number in parallel, if your compiler is smart enough to generate and use them correctly, or if you hand-code your own assembly language functions.

SSE instructions enable low-level parallelism of floating-point arithmetic operations. For example, a 128-bit SSE register can hold two 64-bit double precision (or four 32-bit single precision) numbers at the same time, and operate on them all simultaneously. This can lead to big time savings when working on large amounts of data.

But this may come at a price. Efficient use of SSE instructions can sometimes depend on exactly how the memory used to store data is aligned. Some SSE instructions for moving data to and from memory need memory to be aligned on a 16-byte boundary. If it happens that the memory (for example, a pointer to an array of numbers) that a NAG function uses is **not** aligned nicely, then it may not be possible to use those SSE instructions. An optimizing compiler might well generate two instruction streams, one for when it detects that memory is aligned and one for when it is not.

An example should serve to make things clearer. Suppose we wish to compute the inner product of two vectors, \mathbf{x} and \mathbf{y} , each of length \mathbf{n} . The inner product (or dot product) of two vectors is computed by multiplying together corresponding elements of the two vectors, and summing the individual products to get the result. A function compiled by a good optimizing compiler would load numbers two or four at a time, multiply them together two or four at a time, and accumulate the results into the final result.

But if the memory is not nicely aligned – and it may well not be – the compiler needs to generate a different code path to deal with the situation. Here the result will take longer to get because the products must be computed and accumulated one at a time. At run-time, the code checks whether it can take the fast path or not, and works appropriately.

The problem is that by altering the order of the accumulations, we are quite possibly changing the final result, simply due to rounding differences when working with finite precision computer arithmetic. Instead of getting the inner product

$$s = x_1 \times y_1 + x_2 \times y_2 + x_3 \times y_3 + \cdots + x_n \times y_n$$

we may get

$$s = (x_1 \times y_1 + x_3 \times y_3) + (x_2 \times y_2 + x_4 \times y_4) + \cdots$$

It is likely that the result will be just as accurate either way – neither result will be precise due to finite arithmetic – but they may differ by a tiny amount. And if that tiny difference leads to a different decision being made by the code that called the inner product function, the difference may be magnified.

Furthermore, it is possible that the same program running with bitwise identical data on the same machine may give different results when run twice in a row simply because, when the program is loaded, by chance some piece of memory may or may not be aligned on a particular boundary. Such non-deterministic results can be frustrating if you depend on always getting identical results for the same data.

On even newer hardware, **AVX** instructions use 256-bit and 512-bit registers, and can therefore operate on more numbers at a time. For AVX instructions, memory may need to be 32-byte aligned.

Some memory used by NAG Library functions is allocated inside the NAG Library. In order to minimize differences due to effects like that described above, we can try to make sure the memory is always aligned nicely – for example, by use of more controllable memory allocation functions where available – but that is not always possible since it partly depends on the support of the compiler.

Of course, no Library function has control over memory you have allocated before being passed to the function. If you do observe non-deterministic results which you suspect are due to memory considerations, and you are unable to accept this variation, then you are advised to make sure that any memory you allocate is aligned nicely; unfortunately, precisely how you do this is dependent on your system, but you may be able to get advice through NAG's usual support channels (see Support from NAG).

Parallelism, coming from a multithreaded implementation of the NAG Library and/or a multithreaded vendor library is another potential source of non-determinism in numerical results. Some functions may give different results when run on different numbers of cores, or even different results when a calculation is repeated on the same number of cores. Where reproducibility of results is vital, a purely serial NAG Library, without parallelism in either NAG functions or calls to parallel vendor library functions will generally be available in an appropriate implementation, and may be the best choice. You are advised to contact NAG (see Support from NAG) for advice.

3.9.1 Bit-wise Reproducibility (BWR)

Mathematical operations on fixed-length floating point numbers (e.g., 32-bit floats or 64-bit doubles) are not associative. This means that a computer may produce different results for $a + (b + c)$ and $(a + b) + c$. For example, an IEEE 754 32-bit floating point number has a mantissa of 23 bits. Therefore in this number format $2^{24} + 1 = 2^{24}$, which means that for instance $(2^{24} + 1) - 2^{24} = 0$ while $2^{24} + (1 - 2^{24}) = 1$. BWR is a term which refers to the case in which a given computer program (e.g., a set of source codes) produces bit-for-bit the exact same answer in different computing environments such as

1. Different operating systems (e.g., answers produced on Windows vs answers produced on Linux).
2. Different CPU architectures (e.g., Intel vs AMD or Intel Sandy Bridge vs Intel Ivy Bridge etc.).
3. Different compiler versions.
4. Different numbers of threads.

Users often desire BWR however it is extremely difficult to achieve. Typically you should ensure that:

- (a) Instructions are always executed in exactly the same order.
- (b) No advanced CPU features are used which may not be available on other processors (e.g., SSE3, SSE4, AVX).
- (c) A fixed number of threads is always used.

Often condition (a) is equivalent to compiling with no (or very limited) compiler optimizations, since newer versions of compilers typically improve their code optimization algorithms, which means one version of a compiler may optimize a set of operations one way while the next version may optimize it a different way. Condition (b) typically means that only basic SSE instructions are allowed, such as are supported across the widest range of processors and the enhanced SIMD instructions present in newer processors are not exploited.

The result is that to achieve BWR across a wide range of computing environments one often has to sacrifice a lot of performance.

3.9.1.1 Vendor Math Libraries and Conditional Bitwise Reproducibility (CBWR)

An implementation of the NAG Library that is not self-contained will make calls to an appropriate vendor library containing, in particular, high performance linear algebra functions. The NAG Library has no direct control over BWR with respect to results obtained from calls to the vendor library. However, for at least one such vendor library, CBWR has been introduced such that if an environment variable is set and a set of conditions adhered to in the code calling the vendor library then BWR can be forced. Where CBWR is available for a vendor library used by an implementation of the NAG Library, details will be given in the Users' Note for that implementation.

It should be noted that many NAG functions do not adhere to the conditions set out by vendor library CBWR and so it may not be possible to ensure BWR for all NAG Library functions across different CPU architectures for implementations that are not self-contained.

3.10 Multithreading

3.10.1 Thread Safety

In multithreaded applications, each thread in a team processes instructions independently while sharing the same memory address space. For these applications to operate correctly any functions called from them must be **thread safe**. That is, any global variables they contain are guaranteed not to be accessed simultaneously by different threads, as this can compromise results. This can be ensured through appropriate synchronization, such as that found in OpenMP.

When a function is described as thread safe we are considering its behaviour when it is *called* by multiple threads. It is worth noting that a thread unsafe function can still, itself, be multithreaded. A team of threads can be created inside the routine to share the workload as described in Section 3.10.2.

The NAG C Library is thread safe by design: the functions do not use global variables and all communication between them is via argument lists, and thus can be safely called simultaneously by multiple threads in your program.

3.10.1.1 Functions with Function Arguments

Some Library functions require you to supply a function and to pass the name of the function as an actual argument in the call to the Library function. For many of these Library functions, the supplied function interface includes an array parameter (called **comm**) specifically for you to pass information to the supplied function without the need for global variables.

If you need to provide your supplied function with more information than can be given via the interface argument list, then you are advised to check, in the relevant Chapter Introduction, whether the Library function you intend to call has an equivalent reverse communication interface. These have been designed specifically for problems where user-supplied function interfaces are not flexible enough for a given problem, and their use should eliminate the need to provide data through global variables. Where reverse communication interfaces are not available, it is usual to use global variables containing the required data that is accessible from both the supplied function and from the calling program. It is thread safe to do this only if any global data referenced is made threadprivate by OpenMP or is updated using appropriate synchronisation, thus avoiding the possibility of simultaneous modification by different threads.

Thread safety of user-supplied functions is also an issue with a number of functions in multi-threaded implementations of the NAG Library, which may internally parallelize around the calls to the user-supplied functions. This issue affects not just global variables but also how the **comm** array may be used. In these cases, synchronisation may be needed to ensure thread safety. Chapter x06 provides functions which can be used in your supplied function to determine whether it is being called from within an OpenMP parallel region. If you are in doubt over the thread safety of your program you are advised to contact NAG for assistance.

3.10.1.2 Input/Output

When using the NAG C Library in multi-threaded applications we recommend that when using the C Library error mechanism, the output is switched off (by setting `fail:print=Nag_FALSE`).

3.10.1.3 Implementation Issues

In very rare cases we are unable to guarantee the thread safety of a particular specific implementation. Note also that in some implementations, the Library is linked with one or more vendor libraries to provide, for example, efficient BLAS functions. NAG cannot guarantee that any such vendor library is thread safe. Please consult the Users' Note for your implementation for any additional implementation-specific information.

3.10.2 Parallelism

3.10.2.1 Introduction

The time taken to execute a function from the NAG Library has traditionally depended, to a large degree, on the serial performance capabilities of the processor being used. In an effort to go beyond the performance limitations of a single core processor, multithreaded implementations of the NAG Library are available. These implementations divide the computational workload of some functions between multiple cores and executes these tasks in parallel. Traditionally, such systems consisted of a small number of processors each with a single core. Improvements in the performance capabilities of these processors happened in line with increases in clock frequencies. However, this increase reached a limit which meant that processor designers had to find another way in which to improve performance; this led to the development of **multicore** processors, which are now ubiquitous. Instead of consisting of a single compute core, multicore processors consist of two or more, which typically comprise at least a Central Processing Unit and a small cache. Thus making effective use of parallelism, wherever possible, has become imperative in order to maximize the performance potential of modern hardware resources, and the multithreaded implementations.

The effectiveness of parallelism can be measured by how much faster a parallel program is compared to an equivalent serial program. This is called the parallel **speedup**. If a serial program has been parallelized then the speedup of the parallel implementation of the program is defined by dividing the time taken by the original serial program on a given problem by the time taken by the parallel program using n cores to compute the same problem. Ideal speedup is obtained when this value is n (i.e., when the parallel program takes $\frac{1}{n}$ th the time of the original serial program). If speedup of the parallel program is close to ideal for increasing values of n then we say the program has good **scalability**.

The scalability of a parallel program may be less than the ideal value because of two factors:

- (a) the overheads introduced as part of the parallel implementation, and
- (b) inherently serial parts of the program.

Overheads include communication and synchronisation as well as any extra setup required to allow parallelism. Such overheads depend on the efficiency of the compiler and operating system libraries and the underlying hardware. The impact on performance of inherently serial fractions of a program is explained theoretically (i.e., assuming an idealised system in which overheads are zero) by **Amdahl's law**. Amdahl's law places an upper bound on the speedup of a parallel program with a given inherently serial fraction. If r is the parallelizable fraction of a program and $s = 1 - r$ is the inherently serial fraction then the speedup using n sub-tasks, S_n , satisfies the following:

$$S_n \leq \frac{1}{(s + \frac{r}{n})}$$

Thus, for example, this says that a program with a serial fraction of one quarter can only ever achieve a speedup of 4 since as $n \rightarrow \infty$, $S_n \leq 4$.

Parallelism may be utilised on two classes of systems: shared memory and distributed memory machines, which require different programming techniques. Distributed memory machines are composed of processors located in multiple components which each have their own memory space and are connected by a network. Communication and synchronisation between these components is explicit. Shared memory machines have multiple processors (or a single multicore processor) which can all access the same memory space, and this shared memory is used for communication and synchronisation. The NAG Library makes use of shared memory parallelism using OpenMP as described in Section 3.10.2.2.

Parallel programs which use OpenMP create (or "fork") a number of **threads** from a single process when required at run-time. (Programs which make use of shared memory parallelism are also called **multithreaded** programs.) The threads form a **team** comprising of a single **master** thread and a number of **slave** threads. These threads are capable of executing program instructions independently of one another in parallel. Once the parallel work has been completed the slave threads return control to the master thread and become inactive (or "join") until the next **parallel region** of work. The threads share the same memory address space, i.e., that of the parent process, and this shared memory is used for communication and synchronisation. OpenMP provides some mechanisms for access control so that, as well as allowing all threads to access shared variables, it is possible for each thread to have private

copies of other variables that only it can access. Threads in a team can create their own parallel regions within the current parallel region. At this next **level** of parallelism, the thread creating the new team becomes the master thread of that team. We call this **nested** parallelism.

Something to be aware of for multithreaded programs, compared to serial ones, is that identical results cannot be guaranteed, nor should be expected. Identical results are often impossible in a parallel program since using different numbers of threads may cause floating-point arithmetic to be evaluated in a different (but equally valid) order, thus changing the accumulation of rounding errors. For a more in-depth discussion of reproducibility of results see Section 3.9.

3.10.2.2 How is Parallelism Used in the NAG Library?

The multithreaded implementations differ from the serial implementations of the NAG Library in that it makes use of multithreading through use of OpenMP, which is a portable specification for shared memory programming that is available in many different compilers on a wide range of different hardware platforms (see OpenMP).

Note that not all functions are parallelized; you should check Section 8 of the function documents to find details about parallelism and performance of functions of interest.

There are two situations in which a call to a function in the NAG Library makes use of multithreading:

1. The function being called is a NAG-specific function that has been threaded using OpenMP, or that internally calls another NAG-specific function that is threaded. This applies to multithreaded implementations of the NAG Library only.
2. The function being called calls through to BLAS or LAPACK functions. The vendor library recommended for use with your implementation of the NAG Library (whether the NAG Library is threaded or not) may be threaded. Please consult the Users' Note for further information.

A complete list of all the functions in the NAG Library, and their threaded status is given in Section 3.10.3.

It is useful to understand how OpenMP is used within the Library in order to avoid the potential pitfalls which lead to making inefficient use of the Library.

A call to a threaded NAG-specific function may, depending on input and at one or more points during execution, use OpenMP to create a team of threads for a parallel region of work. The team of threads will fork at the start of the parallel region before joining at the end of the parallel region. Both the fork and the join will happen internally within the function call. However, there are situations in which the teams of threads may be made available to OpenMP directives in your code via user-supplied subprograms, we refer to directives not contained within a parallel region as **orphaned** directives. (See Section 8 of the function documents for further information.) Furthermore, OpenMP constructs within NAG functions are executed by teams of threads created within the NAG code, that is, there are no orphaned directives in the Library itself. Throughout this documentation we assume the use of the recommended compiler as given in the Users' Note, and in particular the use of a single OpenMP runtime library. Thus all OpenMP environment variables will apply to your own code and to NAG functions. However, they may not be respected by vendor libraries that have a mechanism for overriding them. NAG provides functions in Chapter x06 to control threads for your **whole** program, including any specific to a vendor library being called by NAG. You should take care when calling these NAG functions from within your own parallel regions, since if nested parallelism is enabled (it is disabled by default) the NAG function will fork-and-join a team of threads for each calling thread, which may lead to contention on system resources and very poor performance. Poor performance due to contention can also occur if the number of threads requested exceeds the number of physical cores in your machine, or if some hardware resources are busy executing other processes (which may belong to other users in a shared system). For these reasons you should be aware of the number of physical cores available to your program on your machine, and use this information in selecting a number of threads which minimizes contention on resources. Please read the Users' Note for advice about setting the number of threads to use, or contact NAG (see Support from NAG) for advice.

If you are calling multithreaded NAG functions from within another threading mechanism you need to be aware of whether or not this threading mechanism is compatible with the OpenMP compiler runtime used to build the multithreaded implementation of the NAG Library on your platform(s) of choice. The

Users' Note document for each of the implementations in question will include some guidance on this, and you should contact NAG for further advice if required.

Parallelism is used in many places throughout the NAG Library since, although many functions have not been the focus of parallel development by NAG, they may benefit by calling functions that have, and/or by calling parallel vendor functions (e.g., BLAS, LAPACK). Thus, the performance improvement due to multithreading, if any, will vary depending upon which function is called, problem sizes and other parameters, system design and operating system configuration. If you frequently call a function with similar data sizes and other parameters, it may be worthwhile to experiment with different numbers of threads to determine the choice that gives optimal performance. Please contact NAG for further advice if required.

As a general guide, many key functions in the following areas are known to benefit from shared memory parallelism:

- Dense and Sparse Linear Algebra
- FFTs
- Random Number Generators
- Quadrature
- Partial Differential Equations
- Interpolation
- Curve and Surface Fitting
- Correlation and Regression Analysis
- Multivariate Methods
- Time Series Analysis
- Financial Option Pricing
- Global Optimization
- Wavelets

3.10.3 Multithreaded Functions

Many functions are threaded using OpenMP in multithreaded implementations of the NAG Library. These implementations are denoted by having a product code of the form 'CS_____'; rather than 'CL_____'. Please consult Section 8 of each routine document for further information. A table available in the HTML documentation lists which routines have been threaded by NAG. The list also includes functions which internally call BLAS or LAPACK routines, which may be threaded within the vendor library used by both serial and multithreaded NAG Library implementations. You are advised to consult the documentation for the vendor library for further information. Please consult the Users' Note for your implementation for any additional implementation-specific information.

4 How to Use NAG Documentation

4.1 Using the Manual

The Manual is designed to serve the following functions for the NAG Library:

- to give background information about different areas of numerical and statistical computation;
- to advise on the choice of the most suitable NAG Library function or functions to solve a particular problem;
- to give all the information needed to call a NAG Library function correctly from a C program, and to assess the results.

At the beginning of the Manual are some general introductory documents which provide some background and additional information.

The document entitled 'NAG C Library News, Mark 26' provides details of new functions added, details of functions scheduled for withdrawal and details of functions withdrawn at this mark. This document also provides details of internal changes affecting the user at this Mark.

The document entitled 'Advice on Replacement Calls for Withdrawn/Superseded Functions' provides advice on how to modify your program.

The online documentation includes a Keyword and GAMS Search which is available as a keyword search box at the top of every page and, additionally, a separate page containing a form and search guidelines.

Having found a likely chapter or function, you should read the corresponding **Chapter Introduction**, which gives background information about that area of numerical computation, and recommendations on the choice of a function, including indexes, tables and decision trees.

When you have chosen a function, you must consult the **function document**. Each function document is essentially self-contained (it may, however, contain references to related documents). It includes a description of the method, detailed specifications of each argument, explanations of each error exit, remarks on accuracy, and (in most cases) an example program to illustrate the use of the function. In some cases a plot accompanies an example program to illustrate the results from running the example program (possibly amended from the original to output more data points).

4.2 Structure of the Documentation

The NAG Library Manual is the principal documentation for the NAG C Library. It has the same chapter structure as the Library: each chapter of functions in the Library has a corresponding chapter (of the same name) in the Manual. The chapters occur in alphanumeric order. General introductory documents appear at the beginning of the Manual.

Each chapter consists of the following documents:

Chapter Contents, e.g., d01 Chapter Contents;

Chapter Introduction, e.g., d01 Chapter Introduction;

Function Documents, one for each documented function in the chapter.

A function document has the same short name as the function which it describes. Within each chapter, function documents occur in alphanumeric order of short names. It should be noted that all the computational functions from LAPACK, Release 3 are included in the NAG Library and can be called by the NAG Library provided C Library interfaces to LAPACK. The NAG Library names follow the naming convention of LAPACK except that the names are in lower case and 'nag_' is prepended.

All function documents have the same structure consisting of ten numbered sections:

1. **Purpose**
2. **Specification**
3. **Description**
4. **References**
5. **Arguments** (see Section 4.3)
6. **Error Indicators and Warnings**
7. **Accuracy**
8. **Parallelism and Performance**
9. **Further Comments**
10. **Example** (see Section 4.4)

In some documents (notably Chapters e04, e05 and h) there are a further three sections:

11. **Algorithmic Details**
12. **Optional Parameters**

13. Description of Monitoring Information

The sections numbered 11. and 13. above are optional; thus, the section titled **Optional Parameters** may appear as (the possibly final) Section 11.

4.3 Specification of Arguments

Section 5 of each function document contains the specification of the arguments, in the order of their appearance in the argument list.

4.3.1 Classification of Arguments

Arguments are classified as follows.

Input: you must assign values to these arguments on or before entry to the function, and these values are unchanged on exit from the function.

Output: you need not assign values to these arguments before entry to the function; the function may assign values to them.

Input/Output: you must assign values to these arguments before entry to the function, and the function may then change these values.

Communication Structure and Arrays: arguments which are used to communicate data from one function call to another.

External Function: a function which must be supplied (e.g., to evaluate an integrand or to print intermediate output). Usually it must be supplied as part of your calling program, in which case its specification includes full details of its argument list and specifications of its arguments (all enclosed in a box). Its arguments are classified in the same way as those of the Library function, but because you must write the procedure rather than call it, the significance of the classification is different.

Input: values may be supplied on entry.

Output: you may or must assign values to these arguments before exit from your procedure.

Input/Output: values may be supplied on entry, and you may or must assign values to them before exit from your procedure.

4.3.2 Constraints and suggested values

The word '*Constraint:*' or '*Constraints:*' in the specification of an *Input* argument introduces a statement of the range of valid values for that argument, e.g.,

Constraint: $n > 0$.

If the function is called with an invalid value for the argument (e.g., $n = 0$), the function will usually take an error exit.

Occasionally, an enhancement of an existing function at a given Mark may weaken some constraints on some arguments, this will not change the behaviour of existing code that calls the function, but will allow new code to take advantage of enhanced functionality.

The phrase '*Suggested value:*' introduces a suggestion for a reasonable initial setting for an *Input* argument (e.g., accuracy or maximum number of iterations) in case you are unsure what value to use; you should be prepared to use a different setting if the suggested value turns out to be unsuitable for your problem.

4.4 Example Programs and Results

The **example program** in Section 10 of most function documents illustrates a simple call of the function. The programs are designed so that they can be fairly easily modified, and so serve as the basis for a simple program to solve your problem.

For each implementation of the Library, NAG distributes the example programs in machine-readable form, with all necessary modifications already applied. Many sites make the programs accessible to you

in this form. These programs can also be obtained by using the `nagc_example` scripts, provided with the product. Generic forms of the programs, without implementation-specific modifications, may be obtained directly from the NAG web site. The Users' Note for your implementation will mention any special changes which need to be made to the example programs from the generic form.

These example programs may contain preprocessor identifiers such as `NAG_CALL` to enable cross platform portability. Such identifiers are subsequently replaced by appropriate implementation-specific tokens via the header files `nag.h` or `nag_types.h`, using the C preprocessor `#defines`.

Note that the results obtained from running the example programs may not be identical in all implementations and may not agree exactly with the results in the Manual.

For many function documents, a plot of the example program results is also provided. In some cases the example program has been modified slightly to produce larger sets of results to give a more representative plot of the solution profile produced.

4.5 Online Documentation

The complete NAG C Library Manual, Mark 26 can be viewed online in the following formats:

HTML, a fully linked version of the manual using HTML, SVG and MathML (recommended for browsing) and providing links to the PDF version of each document (recommended for printing);

PDF, a full PDF manual browsed using the PDF bookmarks, or via HTML index files;

Single file PDF, the manual as a single PDF file;

Windows HTML help, Windows HTML help version as a single file.

The two single file formats are more compact than the formats that use one file per function and, for example, allow text searches across the entire manual, but of course the larger files may not be so convenient if you only need to view the documentation for a few functions.

The following sections describe how to obtain the software required to view the documentation and advises you how best to navigate the files with or without a browser.

4.5.1 HTML Format

4.5.1.1 Viewing HTML5 Files

These files do not use any proprietary browser specific features, and conform to relevant W3C Recommendations (HTML, MathML, SVG, CSS).

Support for these languages may require that your browser be updated and/or the installation of additional fonts. This information is restricted to the more widely used browsers. If you require information for additional browsers please contact NAG.

4.5.1.2 Firefox (and other Mozilla based browsers)

Versions of Firefox from Firefox 4 onwards should display MathML in HTML files by default.

Rendering of the mathematics is improved if you install the STIX or other OpenType math fonts if they are not already included on your system (as is the case with OS X and some Linux distributions). Full details of the installers available for these fonts on all the major platforms are included in the Firefox MathML fonts page:

<http://www.mozilla.org/projects/mathml/fonts/>

4.5.1.3 Other Browsers

If Firefox is not being used, then the javascript on the page loads the MathJax javascript library (<http://www.mathjax.org>) to enable MathML rendering. By default this is loaded from the web using the MathJax Content Distribution Network. If you require the documentation to work without an Internet connection then you may either use Firefox as described in the previous section or you may download a local copy of MathJax (<http://docs.mathjax.org/en/latest/installation.html>) which needs to be unpacked

on to your local fileserver or file system, and then edit the file `../styles/nagmathml.js` changing the line `http://cdn.mathjax.org/mathjax/latest/` to refer to your local installation.

4.5.1.4 Navigating HTML5 Files

A main index file has been provided (`html/frontmatter/manconts.html`) which links to individual Chapter Contents documents, which in turn link to a complete set of HTML files. Use your browser to navigate from this main index file. For each function document in HTML format you are provided with a link to its equivalent PDF file, this file has been provided primarily for printing purposes.

Each library document contains a number of hyperlinks to particular elements, e.g., arguments, sections, chapter contents, etc. The following key identifies the colour used for each element:

CSS colour	CSS name
black	NAG type
green	appendix, chap, chapter introduction, decision tree, general introduction, section
grey	withdrawn document
pale blue	equation, figure, item in a list, note, bibliographic reference, table, url, verbatim item, website
navy blue	ifail value
red	parameter name
pink	member
purple	optional parameter
royal blue	html table of contents, example plot, routine document, link to a routine example from a table of contents

4.5.1.5 Printing HTML5 Files

It is possible to print your HTML5 files from the browser, however support for printing from browsers, especially support for printing mathematics, varies considerably between versions of browsers and platforms and printer drivers in use.

4.5.1.6 Windows HTML Help

The Windows HTML Help version of the manual is essentially a compressed version of the HTML5 help, customised for the Windows HTML Help viewer (with a bundled copy of MathJax). This format can be very convenient as it is a small compressed single file version allowing full text search over the entire library. You may find this useful if you have a Microsoft Windows desktop, even if you have the NAG Library installed on a different platform.

4.5.2 PDF Format

4.5.2.1 Viewing and Printing PDF Files

If you do not already have a copy of Adobe Acrobat Reader, a free copy can be downloaded from <http://www.adobe.com/reader>. Please check this site for availability of a reader for your platform. While we recommend the use of Acrobat Reader, there are alternative PDF viewers available which can also be used, such as `xpdf` or `ghostview`.

If Acrobat is not running as a plug-in then the bookmark links will not work correctly if you are browsing the PDF files via `http` rather than the local filesystem. You are advised to reinstall Adobe Acrobat which should rectify the problem.

4.5.2.2 Navigating the PDF Files

The manual is supplied as a set of individual PDF files, one for each function document, chapter introduction, etc.. Each PDF file contains bookmarks that can be used to navigate between the files. Alternatively, and often more conveniently, HTML tables of contents are supplied which allow you to navigate to the desired file using a browser, and then using Acrobat as a browser plugin to read or print the document.

Alternatively the single file version of the PDF may be used. In this case the bookmarks will provide links to every function in the Library, and text search may be used to search the entire Library contents.

5 NAG Library Design and Development

Various aspects of the design and development of the NAG Library, and NAG's technical policies and organization are given in Ford (1982), Ford *et al.* (1979), Ford and Pool (1984), and Hague *et al.* (1982).

6 NAG Library Standards

NAG Library development adheres to a number of international standards, see , ISO/IEC (1990), Kernighan and Ritchie (1988) ANSI/IEEE POSIX (1995), Basic Linear Algebra Subprograms Technical (BLAST) Forum (2001).

7 References

ACM (1960–1976) Collected algorithms from ACM index by subject to algorithms

Anderson E, Bai Z, Bischof C, Blackford S, Demmel J, Dongarra J J, Du Croz J J, Greenbaum A, Hammarling S, McKenney A and Sorensen D (1999) *LAPACK Users' Guide* (3rd Edition) SIAM, Philadelphia <http://www.netlib.org/lapack/lug>

ANSI/IEEE (1985) IEEE standard for binary floating-point arithmetic *Std 754-1985* IEEE, New York

ANSI/IEEE POSIX (1995) *POSIX Standard Thread Library* ANSI/IEEE POSIX 1003.1c:1995

Basic Linear Algebra Subprograms Technical (BLAST) Forum (2001) *Basic Linear Algebra Subprograms Technical (BLAST) Forum Standard* University of Tennessee, Knoxville, Tennessee <http://www.netlib.org/blas/blast-forum/blas-report.pdf>

Blackford L S, Demmel J, Dongarra J J, Duff I S, Hammarling S, Henry G, Heroux M, Kaufman L, Lumsdaine A, Petitet A, Pozo R, Remington K and Whaley R C (2002) An updated set of *Basic Linear Algebra Subprograms (BLAS)* *ACM Trans. Math. Software* **28** 135–151

Dongarra J J, Du Croz J J, Duff I S and Hammarling S (1990) A set of Level 3 basic linear algebra subprograms *ACM Trans. Math. Software* **16** 1–28

Dongarra J J, Du Croz J J, Hammarling S and Hanson R J (1988) An extended set of FORTRAN basic linear algebra subprograms *ACM Trans. Math. Software* **14** 1–32

Ford B (1982) Transportable numerical software *Lecture Notes in Computer Science* **142** 128–140 Springer–Verlag

Ford B, Bentley J, Du Croz J J and Hague S J (1979) The NAG Library ‘machine’ *Softw. Pract. Exper.* **9(1)** 65–72

Ford B and Pool J C T (1984) The evolving NAG Library service *Sources and Development of Mathematical Software* (ed W Cowell) 375–397 Prentice–Hall

Hague S J, Nugent S M and Ford B (1982) Computer-based documentation for the NAG Library *Lecture Notes in Computer Science* **142** 91–127 Springer–Verlag

ISO/IEC (1990) Information technology – programming language C *Current C Language Standard* ISO/IEC 9899:1990

Kernighan B W and Ritchie D M (1988) *The C Programming Language* (2nd Edition) Prentice–Hall
The BLAS Technical Forum Standard (2001) <http://www.netlib.org/blas/blast-forum>
