NAG Library Function Document

nag_specfun_2f1_real_scaled (s22bfc)

1   Purpose

nag_specfun_2f1_real_scaled (s22bfc) returns a value for the Gauss hypergeometric function $_2F_1(a, b; c; x)$ for real parameters a, b and c, and real argument x. The result is returned in the scaled form $_2F_1(a, b; c; x) = f_{\mathrm{fr}} \times 2^{f_{\mathrm{sc}}}$.

2   Specification

```
#include <nag.h>
#include <nags.h>
void nag_specfun_2f1_real_scaled (double ani, double adr, double bni,
     double bdr, double cni, double cdr, double x, double *frf, Integer *scf,
     NagError *fail)
```

3   Description

nag_specfun_2f1_real_scaled (s22bfc) returns a value for the Gauss hypergeometric function $_2F_1(a, b; c; x)$ for real parameters a, b and c, and for real argument x.

The Gauss hypergeometric function is a solution to the hypergeometric differential equation,

$$x(1-x)\frac{d^2 f}{dx^2} + (c - (a + b + 1)x)\frac{df}{dx} - abf = 0. \tag{1}$$

For $|x| < 1$, it may be defined by the Gauss series,

$$_2F_1(a, b; c; x) = \sum_{s=0}^{\infty} \frac{(a)_s (b)_s}{(c)_s s!} x^s = 1 + \frac{ab}{c} x + \frac{a(a+1)b(b+1)}{c(c+1)2!} x^2 + \cdots, \tag{2}$$

where $(a)_s = 1(a)(a+1)(a+2)\ldots(a+s-1)$ is the rising factorial of a. $_2F_1(a, b; c; x)$ is undefined for $c = 0$ or c a negative integer.

For $|x| < 1$, the series is absolutely convergent and $_2F_1(a, b; c; x)$ is finite.

For $x < 1$, linear transformations of the form,

$$_2F_1(a, b; c; x) = C_1(a_1, b_1, c_1, x_1)_2F_1(a_1, b_1; c_1; x_1) + C_2(a_2, b_2, c_2, x_2)_2F_1(a_2, b_2; c_2; x_2) \tag{3}$$

exist, where $x_1, x_2 \in (0, 1]$. C_1 and C_2 are real valued functions of the parameters and argument, typically involving products of gamma functions. When these are degenerate, finite limiting cases exist. Hence for $x < 0$, $_2F_1(a, b; c; x)$ is defined by analytic continuation, and for $x < 1$, $_2F_1(a, b; c; x)$ is real and finite.

For $x = 1$, the following apply:

If $c > a + b$, $_2F_1(a, b; c; 1) = \dfrac{\Gamma(c)\Gamma(c - a - b)}{\Gamma(c - a)\Gamma(c - b)}$, and hence is finite. Solutions also exist for the degenerate cases where $c - a$ or $c - b$ are negative integers or zero.

If $c \le a + b$, $_2F_1(a, b; c; 1)$ is infinite, and the sign of $_2F_1(a, b; c; 1)$ is determinable as x approaches 1 from below.

In the complex plane, the principal branch of $_2F_1(a, b; c; z)$ is taken along the real axis from $x = 1.0$ increasing. $_2F_1(a, b; c; z)$ is multivalued along this branch, and for real parameters a, b and c is typically not real valued. As such, this function will not compute a solution when $x > 1$.

The solution strategy used by this function is primarily dependent upon the value of the argument x. Once trivial cases and the case $x = 1.0$ are eliminated, this proceeds as follows.

For $0 < x \leq 0.5$, sets of safe parameters $\{\alpha_{i,j}, \beta_{i,j}, \zeta_{i,j}, \chi_j | 1 \leq j \leq 2|, 1 \leq i \leq 4\}$ are determined, such that the values of $_2F_1(a_j, b_j; c_j; x_j)$ required for an appropriate transformation of the type (3) may be calculated either directly or using recurrence relations from the solutions of $_2F_1(\alpha_{i,j}, \beta_{i,j}; \zeta_{i,j}; \chi_j)$. If c is positive, then only transformations with $C_2 = 0.0$ will be used, implying only $_2F_1(a_1, b_1; c_1; x_1)$ will be required, with the transformed argument $x_1 = x$. If c is negative, in some cases a transformation with $C_2 \neq 0.0$ will be used, with the argument $x_2 = 1.0 - x$. The function then cycles through these sets until acceptable solutions are generated. If no computation produces an accurate answer, the least inaccurate answer is selected to complete the computation. See Section 7.

For $0.5 < x < 1.0$, an identical approach is first used with the argument x. Should this fail, a linear transformation resulting in both transformed arguments satisfying $x_j = 1.0 - x$ is employed, and the above strategy for $0 < x \leq 0.5$ is utilized on both components. Further transformations in these sub-computations are however limited to single terms with no argument transformation.

For $x < 0$, a linear transformation mapping the argument x to the interval $(0, 0.5]$ is first employed. The strategy for $0 < x \leq 0.5$ is then used on each component, including possible further two term transforms. To avoid some degenerate cases, a transform mapping the argument x to $[0.5, 1)$ may also be used.

For improved precision in the final result, this function accepts a, b and c split into an integral and a decimal fractional component. Specifically, $a = a_i + a_r$, where $|a_r| \leq 0.5$ and $a_i = a - a_r$ is integral. The other parameters b and c are similarly deconstructed.

In addition to the above restrictions on c and x, an artificial bound, $arbnd$, is placed on the magnitudes of a, b, c and x to minimize the occurrence of overflow in internal calculations, particularly those involving real to integer conversions. $arbnd = 0.0001 \times I_{max}$, where I_{max} is the largest machine integer (see nag_max_integer (X02BBC)). It should however not be assumed that this function will produce accurate answers for all values of a, b, c and x satisfying this criterion.

This function also tests for non-finite values of the parameters and argument on entry, and assigns non-finite values upon completion if appropriate. See Section 9 and Chapter x07.

Please consult the NIST Digital Library of Mathematical Functions or the companion (2010) for a detailed discussion of the Gauss hypergeometric function including special cases, transformations, relations and asymptotic approximations.

4 References

NIST Handbook of Mathematical Functions (2010) (eds F W J Olver, D W Lozier, R F Boisvert, C W Clark) Cambridge University Press

Pearson J (2009) Computation of hypergeometric functions *MSc Dissertation, Mathematical Institute, University of Oxford*

5 Arguments

1: **ani** – double *Input*

On entry: a_i, the nearest integer to a, satisfying $a_i = a - a_r$.

Constraints:

$\quad$ **ani** $= \lfloor$**ani**$\rfloor$;
$\quad$ $|$**ani**$| \leq arbnd$.

2: **adr** – double *Input*

On entry: a_r, the signed decimal remainder satisfying $a_r = a - a_i$ and $|a_r| \leq 0.5$.

Constraint: $|$**adr**$| \leq 0.5$.

3:     **bni** – double   *Input*

On entry: b_i, the nearest integer to b, satisfying $b_i = b - b_r$.

Constraints:

$\mathbf{bni} = \lfloor \mathbf{bni} \rfloor$;
$|\mathbf{bni}| \leq arbnd$.

4:     **bdr** – double   *Input*

On entry: b_r, the signed decimal remainder satisfying $b_r = b - b_i$ and $|b_r| \leq 0.5$.

Constraint: $|\mathbf{bdr}| \leq 0.5$.

5:     **cni** – double   *Input*

On entry: c_i, the nearest integer to c, satisfying $c_i = c - c_r$.

Constraints:

$\mathbf{cni} = \lfloor \mathbf{cni} \rfloor$;
$|\mathbf{cni}| \leq arbnd$;
if $|\mathbf{cdr}| < 16.0\epsilon$, $\mathbf{cni} \geq 1.0$.

6:     **cdr** – double   *Input*

On entry: c_r, the signed decimal remainder satisfying $c_r = c - c_i$ and $|c_r| \leq 0.5$.

Constraint: $|\mathbf{cdr}| \leq 0.5$.

7:     **x** – double   *Input*

On entry: the argument x.

Constraint: $-arbnd < \mathbf{x} \leq 1$.

8:     **frf** – double *   *Output*

On exit: f_{fr}, the scaled real component of the solution satisfying $f_{\mathrm{fr}} = {}_2F_1(a, b; c; x) \times 2^{-f_{\mathrm{sc}}}$, i.e., ${}_2F_1(a, b; c; x) = f_{\mathrm{fr}} \times 2^{f_{\mathrm{sc}}}$. See Section 9 for the behaviour of f_{fr} when a finite or non-finite answer is returned.

9:     **scf** – Integer *   *Output*

On exit: f_{sc}, the scaling power of two, satisfying $f_{\mathrm{sc}} = \log_2 \left(\dfrac{{}_2F_1(a, b; c; x)}{f_{\mathrm{fr}}} \right)$, i.e., ${}_2F_1(a, b; c; x) = f_{\mathrm{fr}} \times 2^{f_{\mathrm{sc}}}$. See Section 9 for the behaviour of f_{sc} when a non-finite answer is returned.

10:    **fail** – NagError *   *Input/Output*

The NAG error argument (see Section 2.7 in How to Use the NAG Library and its Documentation).

6   Error Indicators and Warnings

NE_ALLOC_FAIL

Dynamic memory allocation failed.
See Section 3.2.1.2 in How to Use the NAG Library and its Documentation for further information.

NE_BAD_PARAM

On entry, argument $\langle value \rangle$ had an illegal value.

NE_CANNOT_CALCULATE

An internal calculation has resulted in an undefined result.

NE_COMPLEX

On entry, $\mathbf{x} = \langle value \rangle$.
In general, $_2F_1(a, b; c; x)$ is not real valued when $x > 1$.

NE_INFINITE

On entry, $\mathbf{x} = \langle value \rangle$, $c = \langle value \rangle$, $a + b = \langle value \rangle$.
$_2F_1(a, b; c; 1)$ is infinite in the case $c \leq a + b$.

NE_INTERNAL_ERROR

An internal error has occurred in this function. Check the function call and any array sizes. If the call is correct then please contact NAG for assistance.

An unexpected error has been triggered by this function. Please contact NAG.
See Section 3.6.6 in How to Use the NAG Library and its Documentation for further information.

NE_NO_LICENCE

Your licence key may have expired or may not have been installed correctly.
See Section 3.6.5 in How to Use the NAG Library and its Documentation for further information.

NE_OVERFLOW

Overflow occurred in a subcalculation of $_2F_1(a, b; c; x)$. The answer may be completely incorrect.

NE_REAL

On entry, **adr** does not satisfy $|\mathbf{adr}| \leq 0.5$.

On entry, **bdr** does not satisfy $|\mathbf{bdr}| \leq 0.5$.

On entry, **cdr** does not satisfy $|\mathbf{cdr}| \leq 0.5$.

NE_REAL_2

On entry, $c = \mathbf{cni} + \mathbf{cdr} = \langle value \rangle$.
$_2F_1(a, b; c; x)$ is undefined when c is zero or a negative integer.

NE_REAL_ARG_NON_INTEGRAL

ANI is non-integral.
On entry, $\mathbf{ani} = \langle value \rangle$.
Constraint: $\mathbf{ani} = \lfloor \mathbf{ani} \rfloor$.

bni is non-integral.
On entry, $\mathbf{bni} = \langle value \rangle$.
Constraint: $\mathbf{bni} = \lfloor \mathbf{bni} \rfloor$.

cni is non-integral.
On entry, $\mathbf{cni} = \langle value \rangle$.
Constraint: $\mathbf{cni} = \lfloor \mathbf{cni} \rfloor$.

NE_REAL_RANGE_CONS

On entry, **ani** does not satisfy $|\mathbf{ani}| \leq arbnd = \langle value \rangle$.

On entry, **bni** does not satisfy $|\mathbf{bni}| \leq arbnd = \langle value \rangle$.

On entry, **cni** does not satisfy $|\mathbf{cni}| \leq arbnd = \langle value \rangle$.

On entry, **x** does not satisfy $|\mathbf{x}| \leq arbnd = \langle value \rangle$.

NE_TOTAL_PRECISION_LOSS

All approximations have completed, and the final residual estimate indicates no accuracy can be guaranteed.
Relative residual $= \langle value \rangle$.

NW_OVERFLOW_WARN

On completion, overflow occurred in the evaluation of $_2F_1(a, b; c; x)$.

NW_SOME_PRECISION_LOSS

All approximations have completed, and the final residual estimate indicates some precision may have been lost.
Relative residual $= \langle value \rangle$.

NW_UNDERFLOW_WARN

Underflow occurred during the evaluation of $_2F_1(a, b; c; x)$. The returned value may be inaccurate.

7 Accuracy

In general, if **fail.code** = NE_NOERROR, the value of $_2F_1(a, b; c; x)$ may be assumed accurate, with the possible loss of one or two decimal places. Assuming the result does not overflow, an error estimate res is made internally using equation (1). If the magnitude of this residual res is sufficiently large, a different **fail.code** will be returned. Specifically,

fail.code = NE_NOERROR or NW_UNDERFLOW_WARN	$res \leq 1000\epsilon$
fail.code = NW_SOME_PRECISION_LOSS	$1000\epsilon < res \leq 0.1$
fail.code = NE_TOTAL_PRECISION_LOSS	$res > 0.1$

where ϵ is the ***machine precision*** as returned by nag_machine_precision (X02AJC). Note that underflow may also have occurred if **fail.code** = NE_TOTAL_PRECISION_LOSS or NW_SOME_PRECISION_LOSS.

A further estimate of the residual can be constructed using equation (1), and the differential identity,

$$
\begin{aligned}
\frac{d\left(_2F_1(a, b; c; x)\right)}{dx} &= \frac{ab}{c}{_2}F_1(a + 1, b + 1; c + 1; x) \\
\frac{d^2\left(_2F_1(a, b; c; x)\right)}{dx^2} &= \frac{a(a + 1)b(b + 1)}{c(c + 1)}{_2}F_1(a + 2, b + 2; c + 2; x)
\end{aligned}
\tag{4}
$$

This estimate is however dependent upon the error involved in approximating $_2F_1(a + 1, b + 1; c + 1; x)$ and $_2F_1(a + 2, b + 2; c + 2; x)$.

8 Parallelism and Performance

nag_specfun_2f1_real_scaled (s22bfc) is not threaded in any implementation.

9 Further Comments

nag_specfun_2f1_real_scaled (s22bfc) returns non-finite values when appropriate. See Chapter x07 for more information on the definitions of non-finite values.

Should a non-finite value be returned, this will be indicated in the value of **fail**, as detailed in the following cases.

If **fail.code** = NE_NOERROR or **fail.code** = NE_TOTAL_PRECISION_LOSS, NW_SOME_PRECISION_LOSS or NW_UNDERFLOW_WARN, a finite value will have been returned with approximate accuracy as detailed in Section 7.

The values of f_{fr} and f_{sc} are implementation dependent. In most cases, if $_2F_1(a,b;c;x) = 0$, $f_{fr} = 0$ and $f_{sc} = 0$ will be returned, and if $_2F_1(a,b;c;x)$ is finite, the fractional component will be bound by $0.5 \le |f_{fr}| < 1$, with f_{sc} chosen accordingly.

The values returned in **frf** (f_{fr}) and **scf** (f_{sc}) may be used to explicitly evaluate $_2F_1(a,b;c;x)$, and may also be used to evaluate products and ratios of multiple values of $_2F_1$ as follows,

$$_2F_1(a,b;c;x) = f_{fr} \times 2^{f_{sc}}$$
$$_2F_1(a_1,b_1;c_1;x_1) \times {}_2F_1(a_2,b_2;c_2;x_2) = (f_{fr1} \times f_{fr2}) \times 2^{(f_{sc1}+f_{sc2})}$$
$$\frac{_2F_1(a_1,b_1;c_1;x_1)}{_2F_1(a_2,b_2;c_2;x_2)} = \frac{f_{fr1}}{f_{fr2}} \times 2^{(f_{sc1}-f_{sc2})}$$
$$\ln\left|_2F_1(a,b;c;x)\right| = \ln|f_{fr}| + f_{sc} \times \ln(2).$$

If **fail.code** = NE_INFINITE then $_2F_1(a,b;c;x)$ is infinite. A signed infinity will have been returned for **frf**, and **scf** = 0. The sign of **frf** should be correct when taking the limit as x approaches 1 from below.

If **fail.code** = NW_OVERFLOW_WARN then upon completion, $\left|_2F_1(a,b;c;x)\right| > 2^{I_{\max}}$, where $I_{\max}$ is given by nag_max_integer (X02BBC), and hence is too large to be representable even in the scaled form. The scaled real component returned in **frf** may still be correct, whilst **scf** = $I_{\max}$ will have been returned.

If **fail.code** = NE_OVERFLOW then overflow occurred during a subcalculation of $_2F_1(a,b;c;x)$. The same result as for **fail.code** = NW_OVERFLOW_WARN will have been returned, however there is no guarantee that this is representative of either the magnitude of the scaling power f_{sc}, or the scaled component f_{fr} of $_2F_1(a,b;c;x)$.

If **fail.code** = NE_NOERROR, **frf** and **scf** were inaccessible to nag_specfun_2f1_real_scaled (s22bfc), and as such it is not possible to determine what their values may be following the call to nag_specfun_2f1_real_scaled (s22bfc).

For all other error exits, **scf** = 0 will be returned and **frf** will be returned as a signalling NaN (see nag_create_nan (x07bbc)).

If **fail.code** = NE_CANNOT_CALCULATE an internal computation produced an undefined result. This may occur when two terms overflow with opposite signs, and the result is dependent upon their summation for example.

If **fail.code** = NE_REAL_2 then c is too close to a negative integer or zero on entry, and $_2F_1(a,b;c;x)$ is undefined. Note, this will also be the case when c is a negative integer, and a (possibly trivial) linear transformation of the form (3) would result in either:

(i)   all c_j not being negative integers,

(ii) for any c_j which remain as negative integers, one of the corresponding parameters a_j or b_j is a negative integer of magnitude less than c_j.

In the first case, the transformation coefficients $C_j(a_j, b_j, c_j, x_j)$ are typically either infinite or undefined, preventing a solution being constructed. In the second case, the series (2) will terminate before the degenerate term, resulting in a polynomial of fixed degree, and hence potentially a finite solution.

If **fail.code** = NE_REAL_RANGE_CONS then no computation will have been performed due to the risk of integer overflow. The actual solution may however be finite.

fail.code = NE_COMPLEX indicates $x > 1$, and hence the requested solution is on the boundary of the principal branch of $_2F_1(a,b;c;x)$. Hence it is multivalued, typically with a nonzero imaginary component. It is however strictly finite.

10   Example

This example evaluates the Gauss hypergeometric function at two points in scaled form using nag_specfun_2f1_real_scaled (s22bfc), and subsequently calculates their product and ratio implicitly.

10.1 Program Text

```
/* nag_specfun_2f1_real_scaled (s22bfc) Example Program.
 *
 * NAGPRODCODE Version.
 *
 * Copyright 2016 Numerical Algorithms Group.
 *
 * Mark 26, 2016.
 */

#include <stdio.h>
#include <math.h>
#include <nag.h>
#include <nag_stdlib.h>
#include <nags.h>
#include <nagx02.h>

int main(void)
{
  /* Scalars */
  Integer exit_status = 0;
  Integer k, imax, scf;
  double ani, adr, bni, bdr, cni, cdr, delta, frf, x;
  /* Arrays */
  double frfv[2];
  Integer scfv[2];
  /* Nag Types */
  Nag_Boolean finite_solutions;
  NagError fail;

  imax = X02BLC;
  printf("nag_specfun_2f1_real_scaled (s22bfc) Example Program Results\n\n");

  ani = -10.0;
  bni = 2.0;
  cni = -5.0;
  delta = 1.0E-4;
  adr = delta;
  bdr = -delta;
  cdr = delta;
  x = 0.45;
  finite_solutions = Nag_TRUE;
  printf("%11s%11s%11s%11s%14s%7s%14s\n",
         "a", "b", "c", "x", "frf", "scf", "2F1(a,b;c;x)");
  for (k = 0; k < 2; k++) {
    INIT_FAIL(fail);
    /* Compute the real Gauss hypergeometric function 2F1(a,b;c;x) in scaled
     * form using nag_specfun_2f1_real_scaled (s22bfc).
     */
    nag_specfun_2f1_real_scaled(ani, adr, bni, bdr, cni, cdr, x,
                                &frf, &scf, &fail);
    switch (fail.code) {
    case NE_NOERROR:
    case NW_UNDERFLOW_WARN:
    case NW_SOME_PRECISION_LOSS:
      /* A finite result has been returned. */
      if (scf < imax)
        printf(" %10.4f %10.4f %10.4f %10.4f %13.5e %6" NAG_IFMT " %13.5e\n",
               ani + adr, bni + bdr, cni + cdr, x, frf, scf, frf * pow(2.0,
                                                                       scf));
      else
        printf(" %10.4f %10.4f %10.4f %10.4f %13.5e %6" NAG_IFMT " %17s\n",
               ani + adr, bni + bdr, cni + cdr, x, frf, scf,
               "Not Representable");
      frfv[k] = frf;
      scfv[k] = scf;
      break;
    case NE_INFINITE:
      /* The result is analytically infinite. */
      finite_solutions = Nag_FALSE;
```

```
        if (frf >= 0.0)
          printf(" %10.4f %10.4f %10.4f %10.4f %13s %6" NAG_IFMT " %13s\n",
                 ani + adr, bni + bdr, cni + cdr, x, "Inf", scf, "Inf");
        else
          printf(" %10.4f %10.4f %10.4f %10.4f %13s %6" NAG_IFMT " %13s\n",
                 ani + adr, bni + bdr, cni + cdr, x, "-Inf", scf, "-Inf");
        break;
      case NW_OVERFLOW_WARN:
      case NE_OVERFLOW:
        /* The final result has overflowed. */
        finite_solutions = Nag_FALSE;
        if (frf >= 0.0)
          printf(" %10.4f %10.4f %10.4f %10.4f %13.5e %6s %13s\n",
                 ani + adr, bni + bdr, cni + cdr, x, frf, "imax",
                 ">pow(2,imax)");
        else
          printf(" %10.4f %10.4f %10.4f %10.4f %13.5e %6s %13s\n",
                 ani + adr, bni + bdr, cni + cdr, x, frf, "imax",
                 "<-pow(2,imax)");
        break;
      case NE_CANNOT_CALCULATE:
        /* An internal calculation resulted in an undefined result. */
        finite_solutions = Nag_FALSE;
        printf(" %10.4f %10.4f %10.4f %10.4f %13s %6" NAG_IFMT " %13s\n",
               ani + adr, bni + bdr, cni + cdr, x, "NaN", scf, "NaN");
        break;
      default:
        /* An input error has been detected. */
        printf(" %10.4f %10.4f %10.4f %10.4f %17s\n",
               ani + adr, bni + bdr, cni + cdr, x, "FAILED");
        exit_status = 1;
        goto END;
        break;
      }
      adr = -adr;
      bdr = -bdr;
      cdr = -cdr;
    }
    if (finite_solutions) {
      /* Calculate the product M1*M2. */
      frf = frfv[0] * frfv[1];
      scf = scfv[0] + scfv[1];
      printf("\n");
      if (scf < imax)
        printf("%-34s%13.5e %6" NAG_IFMT " %13.5e\n",
               " Solution product", frf, scf, frf * pow(2.0, scf));
      else
        printf("%-34s%13.5e %6" NAG_IFMT "%17s\n",
               " Solution product", frf, scf, "Not Representable");

      /* Calculate the ratio M1/M2. */
      if (frfv[1] != 0.0) {
        frf = frfv[0] / frfv[1];
        scf = scfv[0] - scfv[1];
        printf("\n");
        if (scf < imax)
          printf("%-34s%13.5e %6" NAG_IFMT " %13.5e\n",
                 " Solution ratio", frf, scf, frf * pow(2.0, scf));
        else
          printf("%-34s%13.5e %6" NAG_IFMT "%17s\n",
                 " Solution ratio", frf, scf, "Not Representable");
      }
    }
END:
  return exit_status;
}
```

10.2 Program Data

None.

10.3 Program Results

```
nag_specfun_2f1_real_scaled (s22bfc) Example Program Results

          a          b          c          x          frf    scf  2F1(a,b;c;x)
    -9.9999     1.9999    -4.9999     0.4500  -5.44477e-01     16  -3.56828e+04
   -10.0001     2.0001    -5.0001     0.4500   5.44547e-01     16   3.56875e+04


 Solution product                          -2.96494e-01     32  -1.27343e+09


 Solution ratio                            -9.99871e-01      0  -9.99871e-01
```