

## The Cumulative Normal Distribution for Dimensions up to 3 using the qfloat floating-point Library from LCC-WIN32: Some Tests for Dimension = 3

This implementation gives exactness over almost the 104 digits which the library provides.

The system library `qfloat.dll` should be in Window's system directory and the concurrent `cdfn123.dll` should be placed in the directory of this worksheet.

AVt, Jan 2006

```
> restart;
kernelopts(version);
                               Maple 10.02, IBM INTEL NT, Nov 8 2005 Build ID 208934
> Digits_lcc:=105;
                               Digits_lcc := 105
> Digits:=2*Digits_lcc; # greater precision to check results
                               Digits := 210
```

For using the DLL locate its directory to call external functions from there:

```
> currentdir(): myDLL:=cat(%,`\\`cdfs123.dll`);
myDLL := "C:\_Work\Maple_Work_allTheStuff\Statistik\Verteilungen\TrivariateNormal\cdfn123.dll"
```

Accessing the DLL functions is through strings:

```
> lccstr:=proc(var)
    convert( evalf(parse(convert(var,string)),Digits_lcc + 10), string);
end proc;
```

### The cumulative normal distribution

Define the cumulative normal distribution within Maple

```
> cdfN := x -> 1/2+1/2*erf(1/2*x*2^(1/2));
pdfN := x -> 1/2*1/Pi^(1/2)*exp(-1/2*x^2)*2^(1/2);
```

$$\text{cdfN} := x \rightarrow \frac{1}{2} + \frac{1}{2} \operatorname{erf}\left(\frac{1}{2} x \sqrt{2}\right)$$

$$\text{pdfN} := x \rightarrow \frac{1}{2} \frac{e^{(-1/2x^2)} \sqrt{2}}{\sqrt{\pi}}$$

```
> fct_cdfN := define_external(
    'cdfN_string_Maple',
    'C',
    'x_str'::string[],
    'ret_str'::string[],
    RETURN::integer[4],
    LIB=myDLL);

cdfN_lcc:=proc(x)
    local X::string, result::string;
    result:=StringTools:-Fill( `0` , Digits_lcc+10);
    X:=lccstr(x);
    if type(parse(X),numeric) then
        fct_cdfN(X,result);
    return parse(result);
```

```

else
  return 'cdfN_lcc(x)';
end if;
end proc: #maplemint(%);

```

This will provide the DLL with memory space (given as a string `y_str`) to store the results in the DLL. Since update is 'inplace' this will modify the string and its length. As Maple is a symbolic system one should never call this result directly, since this inconsistency for the same object will crash it (just try it and restart ...). So a simple procedure is used as interface.

## The Bivariate Case

The bivariate normal distribution can be written as:

```

> pdfN2 := (x,y,rho) ->
  1/sqrt(1-rho^2)/(2*Pi)*exp(-(x^2-2*rho*x*y+y^2)/(2*(1-rho^2)));
``;
cdfN2 := (x,y,rho) ->
  Int(Int(pdfN2(xi,eta,rho), eta=-infinity..y),xi=-infinity..x);

```

$$\text{pdfN2} := (x, y, \rho) \rightarrow \frac{1}{2} \frac{e^{-\frac{x^2 - 2\rho xy + y^2}{2(1-\rho^2)}}}{\sqrt{1-\rho^2} \pi}$$

$$\text{cdfN2} := (x, y, \rho) \rightarrow \int_{-\infty}^x \int_{-\infty}^y \text{pdfN2}(\xi, \eta, \rho) \, d\eta \, d\xi$$

It is coded within the DLL giving 104 decimal points of precision and can be accessed as follows:

```

> fct_cdfN2 := define_external(
  'cdfN2_string_Maple',
  'C',
  'x_str'::string[],
  'y_str'::string[],
  'r_str'::string[],
  'ret_str'::string[],
  RETURN::integer[4],
  LIB=myDLL):

cdfN2_lcc:=proc(x,y,r)
  local X::string, Y::string, R::string, result;

  result:=StringTools:-Fill( `0` , Digits_lcc+10);
  X:=lccstr(x);
  Y:=lccstr(y);
  R:=lccstr(r);

  if type(parse(X),numeric) and type(parse(Y),numeric) and
  type(parse(R),numeric) then
    fct_cdfN2(X,Y,R,result);
    return parse(result);
  else
    return 'cdfN2_lcc(x,y,r)';
  end if;
end proc:

```

This is fast, about 10 milli seconds.

But difficult to test. For that use a re-formulation:

```
> 'cdfN2(x,y,rho)'= 'Int(pdfN(tau)*cdfN((y-rho*tau)/sqrt(1-rho^2)),tau =  
-infinity .. x)';
```

$$\text{cdfN2}(x, y, \rho) = \int_{-\infty}^x \text{pdfN}(\tau) \text{cdfN}\left(\frac{y - \rho \tau}{\sqrt{1 - \rho^2}}\right) d\tau$$

For dimension 3 a similar reduction is available and attributed to D B Owen.

>

### Dimension = 3

The trivariate case can be handled through semi-definite integrals using bivariate normal distributions, cf Alan Genz, *Numerical Computation of Rectangular Bivariate and Trivariate ...*  
<http://www.sci.wsu.edu/math/faculty/genz/homepage>

To see that the used integration routine is highly exact as an example just compare what happens for integrating over the usual pdfN (that is coded in the DLL only for a test):

For directly inputting constants it not clear that they belong to a correlation matrix, so provide a check:

```
> check_coefficients:=proc(_r12,_r13,_r23)  
  local R, det, ev, remDigits;  
  
  remDigits:=Digits;  
  Digits:=18;  
  R := Matrix(  
    [[1,_r12,_r13],  
     [1,_r23],  
     [1]],  
    shape=symmetric,  
    scan=triangular[upper]);  
  
  if (LinearAlgebra:-IsDefinite(R, query=positive_definite)) then  
    det:=evalf(LinearAlgebra:-Determinant(R), 2*Digits);  
    det:=evalf(det,6);  
    print(`valid correlation matrix`, `determinant` =det);  
    Digits:=remDigits;  
    return 0;  
  else  
    print(`not a correlation matrix!`);  
    Digits:=remDigits;  
    return -1;  
  end if;  
  Digits:=remDigits;  
  return;  
end proc;
```

Provide an interface to the DLL ...

```
> fct3 := define_external(  
  'cdfN3_string_Maple',  
  'C',  
  'x1_str'::string[], 'x2_str'::string[], 'x3_str'::string[],
```

```

'r12_str'::string[], 'r13_str'::string[], 'r23_str'::string[],
'ret_str'::string[],
RETURN::integer[4],
LIB=myDLL):

cdfN3_lcc:=proc(x1,x2,x3,r12,r13,r23)
#local X::string, Y::string, R::string, result;
local X1, X2, X3, R12, R13, R23, result, lccstr;

lccstr:=proc(var)
  convert( evalf(parse(convert(var,string)),105), string);
end proc:

result:=StringTools:-Fill( `0` , Digits_lcc+10);
#X1:=convert(x1,string); X2:=convert(x2,string); X3:=convert(x3,string);
#R12:=convert(r12,string); R13:=convert(r13,string);
R23:=convert(r23,string);
X1:=lccstr(x1): X2:=lccstr(x2): X3:=lccstr(x3):
R12:=lccstr(r12): R13:=lccstr(r13): R23:=lccstr(r23):

if type(parse(X1),numeric) then
  fct3(X1, X2, X3, R12, R13, R23,result);
  return parse(result);
else
  return 'x';
end if;
end proc:

```

To cross check results one can use a formula, which is attributed to Owen:

```

> 'cdfN3'='Int(pdfN(xi)*F_Owen(xi), xi= -infinity .. x1)';
F_Owen:= (xi,x2,x3, rho12,rho13,rho23) ->
  cdfN2(
    (x2 - rho12 * xi) / sqrt(1 - rho12*rho12),
    (x3 - rho13 * xi) / sqrt(1 - rho13*rho13),
    (rho23 - rho13 * rho12) / sqrt( (1 - rho12*rho12)*(1 - rho13*rho13) )
  );

```

$$\text{cdfN3} = \int_{-\infty}^{x1} \text{pdfN}(\xi) F_{\text{Owen}}(\xi) d\xi$$

F\_Owen :=

$$(\xi, x2, x3, \rho12, \rho13, \rho23) \rightarrow \text{cdfN2}\left(\frac{x2 - \rho12 \xi}{\sqrt{1 - \rho12 \rho12}}, \frac{x3 - \rho13 \xi}{\sqrt{1 - \rho13 \rho13}}, \frac{\rho23 - \rho13 \rho12}{\sqrt{(1 - \rho12 \rho12)(1 - \rho13 \rho13)}}\right)$$

That reduces the problem to dimension 2. But it will be very slow for 100 Digits and for improving speed I would use an external implementation for cdfN2 - which would make the test depend on a repeated use my implementation in dim 2.

But one can reduce to dimension one through partial integration:

```

> 'Int(pdfN(x)*F(x),x=-infinity..x1)' =
'F(x1)*cdfN(x1)-int(D(F)(x)*cdfN(x),x = -infinity .. x1)';
``;
'diff(cdfN2(a1+b1*x, a2+b2*x,r),x)' =
'b2*pdfN(b2*x+a2)*cdfN((a1-r*a2+(b1-r*b2)*x)/(1-r^2)^(1/2)) +
b1*pdfN(b1*x+a1)*cdfN((a2-r*a1+(b2-r*b1)*x)/(1-r^2)^(1/2))';

```

$$\int_{-\infty}^{x_1} \text{pdfN}(x) F(x) dx = F(x_1) \text{cdfN}(x_1) - \int_{-\infty}^{x_1} D(F)(x) \text{cdfN}(x) dx$$

$$\frac{\partial}{\partial x} \text{cdfN2}(a_1 + b_1 x, a_2 + b_2 x, r) = b_2 \text{pdfN}(a_2 + b_2 x) \text{cdfN}\left(\frac{a_1 - r a_2 + (b_1 - r b_2) x}{\sqrt{1 - r^2}}\right) + b_1 \text{pdfN}(a_1 + b_1 x) \text{cdfN}\left(\frac{a_2 - r a_1 + (b_2 - r b_1) x}{\sqrt{1 - r^2}}\right)$$

where the constant have to be evaluated as

```
> L:={
  b1 = -r12/(1-r12^2)^(1/2), a1 = x2/(1-r12^2)^(1/2),
  b2 = -r13/(1-r13^2)^(1/2), a2 = x3/(1-r13^2)^(1/2),
  r = (r23 - r13 * r12) / sqrt( (1 - r12*r12)*(1 - r13*r13) )};
```

L :=

$$\left\{ b_2 = -\frac{r_{13}}{\sqrt{1 - r_{13}^2}}, a_2 = \frac{x_3}{\sqrt{1 - r_{13}^2}}, b_1 = -\frac{r_{12}}{\sqrt{1 - r_{12}^2}}, a_1 = \frac{x_2}{\sqrt{1 - r_{12}^2}}, r = \frac{r_{23} - r_{13} r_{12}}{\sqrt{(1 - r_{12}^2)(1 - r_{13}^2)}} \right\}$$

For short let us write it as

```
> theIntegrand:= 'cdfN(xi)' *
  '(b2*pdfN(b2*xi+a2)*cdfN((a1-r*a2+(b1-r*b2)*xi)/(1-r^2)^(1/2)) +
  b1*pdfN(b1*xi+a1)*cdfN((a2-r*a1+(b2-r*b1)*xi)/(1-r^2)^(1/2)))';
``;
'cdfN3' =
'F(x1)*cdfN(x1)-Int(theIntegrand,x = -infinity .. x1)';
`F(x)` = 'cdfN2(a1+b1*x, a2+b2*x,r)';
```

$$\text{theIntegrand} := \text{cdfN}(\xi) \left( b_2 \text{pdfN}(b_2 \xi + a_2) \text{cdfN}\left(\frac{a_1 - r a_2 + (b_1 - r b_2) \xi}{\sqrt{1 - r^2}}\right) + b_1 \text{pdfN}(b_1 \xi + a_1) \text{cdfN}\left(\frac{a_2 - r a_1 + (b_2 - r b_1) \xi}{\sqrt{1 - r^2}}\right) \right)$$

$$\text{cdfN3} = F(x_1) \text{cdfN}(x_1) - \int_{-\infty}^{x_1} \text{theIntegrand} dx$$

$$F(x) = \text{cdfN2}(a_1 + b_1 x, a_2 + b_2 x, r)$$

## Test 1

```
> x1 := 1.00000000;
  x2 := 0.33000000;
  x3 := -0.50000000;

  r12 := 0.9;
  r13 := 0.7;
  r23 := 0.8;

  st:=time():
  if 0 <= check_coefficients(r12,r13,r23) then
    cdfN3_lcc(x1,x2,x3,r12,r13,r23);
    #evalf(%,16);
  end if;
  `seconds`=time()-st;
```

```

x1 := 1.000000000
x2 := 0.330000000
x3 := -0.500000000
r12 := 0.9
r13 := 0.7
r23 := 0.8
valid correlation matrix, determinant = 0.068
0.296114833713582696020468935168923545697048449068019828037433866742829296013568740952\
543110441612268973184
seconds = 5.419

```

Test it with the modification of Owen's formula

```

> TheIntegrand:= 'eval(theIntegrand, L)';
TheIntegral:= 'Int(theIntegrand,xi = -infinity .. x1)';
TheIntegrand := eval(theIntegrand, L)
TheIntegral :=  $\int_{-\infty}^{x1} \text{theIntegrand } d\xi$ 

```

Since the integrand contains  $\text{cdfN}(\xi)$  as factor I will cut off against  $-\infty$  (as it will be below 1E-105):

```

> TheIntegral:= 'Int(TheIntegrand,xi = -23.0 .. x1, method = _Gquad)';
TheIntegral := Int(TheIntegrand,  $\xi = -23.0 .. x1$ , method = _Gquad)

```

Then Maple gives the same result as the DLL implementation:

```

> 'TheIntegral'=evalf[105](TheIntegral):
`eval(cdfN2_lcc(a1+b1*x1, a2+b2*x1,r)*cdfN(x1), L) - TheIntegral =`;
evalf(eval(cdfN2_lcc(a1+b1*x1, a2+b2*x1,r)*cdfN(x1), L) -TheIntegral,105);
eval(cdfN2_lcc(a1+b1*x1, a2+b2*x1,r)*cdfN(x1), L) - TheIntegral =
0.296114833713582696020468935168923545697048449068019828037433866742829296013568740952\
543110441612268973188

```

## Test 2

If the determinant is close to zero the correlation coefficients are chosen from a valid matrix R (cf the Genz paper):

```

> theta1 := 0.01;
theta2 := 0.01; #1 - 1e-6;
theta3 := -0.02;

R:='R': A:='A': `R`= A*A^t;
`A`=Matrix('[[1,0,0],
[cos(theta1*Pi),sin(theta1*Pi),0],
[cos(theta2*Pi)*cos(theta3*Pi),cos(theta2*Pi)*sin(theta3*Pi),sin(theta2*Pi)]]');
rhs(%):
A:=evalf(%): #evalf(evalhf(%),15);
LinearAlgebra:-Multiply(A, LinearAlgebra:-Transpose(A)):
R:=evalf(evalhf(%),15); A:='A':
``;
r12:=R[1,2];

```

```
r13:=R[1,3];
r23:=R[2,3];
```

```
check_coefficients(r12,r13,r23):
```

```
θ1 := 0.01
```

```
θ2 := 0.01
```

```
θ3 := -0.02
```

```
R = A At
```

$$A = \begin{bmatrix} 1 & 0 & 0 \\ \cos(\theta_1 \pi) & \sin(\theta_1 \pi) & 0 \\ \cos(\theta_2 \pi) \cos(\theta_3 \pi) & \cos(\theta_2 \pi) \sin(\theta_3 \pi) & \sin(\theta_2 \pi) \end{bmatrix}$$

$$R := \begin{bmatrix} 1. & 0.999506560365732 & 0.997534262484406 \\ 0.999506560365732 & 1. & 0.995070714871375 \\ 0.997534262484406 & 0.995070714871375 & 1. \end{bmatrix}$$

```
r12 := 0.999506560365732
```

```
r13 := 0.997534262484406
```

```
r23 := 0.995070714871375
```

```
valid correlation matrix, determinant = 0.973450 10-6
```

```
> x1 := 1.000000000;
x2 := 0.330000000;
x3 := -0.400000000;
```

```
x1 := 1.000000000
```

```
x2 := 0.330000000
```

```
x3 := -0.400000000
```

Here the computation needs not much more time to achieve the desired accuracy:

```
> st:=time():
```

```
if 0 <= check_coefficients(r12,r13,r23) then
  print(`cdfN3 =`);
  cdfN3_lcc(x1,x2,x3,r12,r13,r23);
  #evalf(%,16);
end if;
`seconds`=time()-st;
```

```
valid correlation matrix, determinant = 0.973450 10-6
```

```
cdfN3 =
```

```
0.344578258389675325504053554349651270204065515985560001416107691004781078931187426451\
380957022207072550523
```

```
seconds = 6.795
```

```
> TheIntegrand:='eval(theIntegrand, L)';
TheIntegral:= 'Int(TheIntegrand,xi = -23.0 .. x1, method = _Gquad)';
```

```
TheIntegrand := eval(theIntegrand, L)
```

```
TheIntegral := Int(TheIntegrand, ξ = -23.0 .. x1, method = _Gquad)
```

```
> 'TheIntegral'=evalf[105](TheIntegral):
```

```
`eval(cdfN2_lcc(a1+b1*x1, a2+b2*x1,r)*cdfN(x1), L) - TheIntegral =`;
evalf(eval(cdfN2_lcc(a1+b1*x1, a2+b2*x1,r)*cdfN(x1), L) -TheIntegral,105);
```

```
eval(cdfN2_lcc(a1+b1*x1, a2+b2*x1,r)*cdfN(x1), L) - TheIntegral =
```

```
0.344578258389675325504053554349651270204065515985560001416107691004781078931187426451\
380957022207072550524
```

Again up to the last decimal place the results coincide.

### Test 3

```
> theta1 := 0.01;
theta2 := 1 - 1e-5; #1 - 1e-6;
theta3 := -0.02;

R:='R': A:='A': `R`= A*A^t;
`A`=Matrix('[[1,0,0],
  [cos(theta1*Pi),sin(theta1*Pi),0],

[cos(theta2*Pi)*cos(theta3*Pi),cos(theta2*Pi)*sin(theta3*Pi),sin(theta2*Pi
)]]'):
rhs(%):
A:=evalf(%): #evalf(evalhf(%),15);
LinearAlgebra:-Multiply(A, LinearAlgebra:-Transpose(A)):
R:=evalf(evalhf(%),15); A:='A':
``;
r12:=R[1,2];
r13:=R[1,3];
r23:=R[2,3];

check_coefficients(r12,r13,r23):

          theta1 := 0.01
          theta2 := 0.99999
          theta3 := -0.02
          R = A^t
R := 
$$\begin{bmatrix} 1. & 0.999506560365732 & -0.998026727935765 \\ 0.999506560365732 & 1. & -0.995561964111790 \\ -0.998026727935765 & -0.995561964111790 & 1. \end{bmatrix}$$

          r12 := 0.999506560365732
          r13 := -0.998026727935765
          r23 := -0.995561964111790
          valid correlation matrix, determinant = 0.973766 10-12

> x1 := 1.000000000;
x2 := 0.330000000;
x3 := -0.400000000;

          x1 := 1.000000000
          x2 := 0.330000000
          x3 := -0.400000000

> st:=time():
if 0 <= check_coefficients(r12,r13,r23) then
  print(`cdfN3 =`);
  cdfN3_lcc(x1,x2,x3,r12,r13,r23);
  #evalf(% ,16);
end if;
`seconds`=time()-st;

          valid correlation matrix, determinant = 0.973766 10-12
          cdfN3 =
0.004670924202417529607877661369138761508077131238829783499498108536734169803673820850\
69678355927915376347992

          seconds = 4.786
```



For that case Owen's function is already difficult to use:

```
> theF := proc(xi,x2,x3, rho12,rho13,rho23)
  local a1,b1,a2,b2,corr, result,r;

  b1 := -rho12/(1-rho12^2)^(1/2); a1 := x2/(1-rho12^2)^(1/2);
  b2 := -rho13/(1-rho13^2)^(1/2); a2 := x3/(1-rho13^2)^(1/2);
  corr := (rho23 - rho13 * rho12) / sqrt( (1 - rho12*rho12)*(1 -
rho13*rho13) );

  r:=corr;
  #[(a1-r*a2+(b1-r*b2)*x)/(1-r^2)^(1/2),(a2-r*a1+(b2-r*b1)*x)/(1-r^2)^(1/2)]
  ;

  return cdfN2_lcc(a1+b1*xi, a2+b2*xi,corr); #cdfN2(a1+b1*x, a2+b2*x,corr);
end proc;
theF := proc(ξ, x2, x3, ρ12, ρ13, ρ23)
local a1, b1, a2, b2, corr, result, r;
  b1 := -ρ12 / (1 - ρ12^2)^(1 / 2);
  a1 := x2 / (1 - ρ12^2)^(1 / 2);
  b2 := -ρ13 / (1 - ρ13^2)^(1 / 2);
  a2 := x3 / (1 - ρ13^2)^(1 / 2);
  corr := (ρ23 - ρ13*ρ12) / sqrt((1 - ρ12*ρ12)*(1 - ρ13*ρ13));
  r := corr;
  return cdfN2_lcc(a1 + b1*ξ, a2 + b2*ξ, corr)
end proc
> theF(xi,x2,x3, r12,r13,r23):
  tst:=unapply(%,xi): # Owen's function for the test data
> remDigits:=Digits: Digits:=14;
plot(tst, 0..0.6); #plot(tst, 0.353685..0.353686);
mid:=0.3536856;
tst(mid): evalf(%);
Digits:=remDigits:
```

Digits := 14



```
r23 := 0.999506560365732
```

```
> # new values  
x1 := -0.400000000;  
x2 := 1.000000000;  
x3 := 0.330000000;
```

```
x1 := -0.400000000
```

```
x2 := 1.000000000
```

```
x3 := 0.330000000
```

For the DLL that change does not matter:

```
> st:=time():  
if 0 <= check_coefficients(r12,r13,r23) then  
  print(`cdfN3 =`);  
  cdfN3_lcc(x1,x2,x3,r12,r13,r23);  
  #evalf(%,16);  
end if;  
new_result_from_DLL:=%:  
`seconds`=time()-st;
```

```
valid correlation matrix, determinant = 0.973766 10-12
```

```
cdfN3 =
```


```
0.004670924202417529607877661369138761508077131238829783499498108536734169803673820850\  
69678355927915376347992
```

```
seconds = 4.786
```

but Owen's function now looks better to handle:

```
> theF(xi,x2,x3, r12,r13,r23):  
  tst:=unapply(%,xi):  
> remDigits:=Digits: Digits:=14;  
  plot(tst, -2.. x1);  
  Digits:=remDigits:
```

```
Digits := 14
```



```
>  
> TheIntegrand:='eval(theIntegrand, L)';  
TheIntegral:= 'Int(TheIntegrand,xi = -23.0 .. x1, method = _Gquad)';
```

```

TheIntegrand := eval(theIntegrand, L)
TheIntegral := Int(TheIntegrand, ξ = -23.0 .. x1, method = _Gquad)
> 'TheIntegral'=evalf[105](TheIntegral):
`eval(cdfN2_lcc(a1+b1*x1, a2+b2*x1,r)*cdfN(x1), L) - TheIntegral =`;
evalf(eval(cdfN2_lcc(a1+b1*x1, a2+b2*x1,r)*cdfN(x1), L) -TheIntegral,105);
eval(cdfN2_lcc(a1+b1*x1, a2+b2*x1,r)*cdfN(x1), L) - TheIntegral =
0.004670924202417529607877661369138761508077131238829783499498108536734169803673820850\
6967835592791537634770

```

and up to the last places that approach gives the same

#### Test 4

```

> theta1 := 1e-3;
theta2 := 1 - 1e-6; #1 - 1e-6;
theta3 := -1e-3;

R:='R': A:='A': `R`= A*A^t;
`A`=Matrix('[[1,0,0],
[cos(theta1*Pi),sin(theta1*Pi),0],
[cos(theta2*Pi)*cos(theta3*Pi),cos(theta2*Pi)*sin(theta3*Pi),sin(theta2*Pi)
]]'):
rhs(%):
A:=evalf(%): #evalf(evalhf(%),15);
LinearAlgebra:-Multiply(A, LinearAlgebra:-Transpose(A)):
R:=evalf(evalhf(%),15); A:='A':
``;
r12:=R[1,2];
r13:=R[1,3];
r23:=R[2,3];

check_coefficients(r12,r13,r23):
theta1 := 0.001
theta2 := 0.999999
theta3 := -0.001
R = A^t
R := [
1. 0.999995065201858 -0.999995065196923
0.999995065201858 1. -0.999980260851202
-0.999995065196923 -0.999980260851202 1.
]
r12 := 0.999995065201858
r13 := -0.999995065196923
r23 := -0.999980260851202
valid correlation matrix, determinant = 0.974223 10^-16

> x1 := 1.000000000;
x2 := 0.330000000;
x3 := -0.400000000;
x1 := 1.000000000
x2 := 0.330000000
x3 := -0.400000000

```

For that case Owen's function can not be used, it would evaluate to zero:

```
> theF(xi,x2,x3, r12,r13,r23):
  tst:=unapply(%,xi):
  tst(xi): op(%) : evalf(%,8) : cdfN2_lcc(%);
          cdfN2_lcc(105.04244 - 318.30884  $\xi$ , -127.32410 + 318.30868  $\xi$ , 0.99999950)
```

Generally the DLL uses sorting on the correlation coefficients and takes care that the integral is provided with a 'positive integrand' as follows:

```
> arrangeXR:=proc(x,r)
  local R,X,swapR,swapX;

  X := [x[1],x[2],x[3]];
  R := [r[1],r[2],r[3]];

  if R[2] < R[1] then # swapR them
    swapR:= R[2];
    R[2]:=R[1];
    R[1]:=swapR;

    swapX:= X[2];
    X[2]:=X[3];
    X[3]:=swapX;
  end if;
  #print(R);

  if R[3] < R[2] then # swapR them
    swapR:= R[3];
    R[3]:=R[2];
    R[2]:=swapR;

    swapX:=X[1];
    X[1]:=X[2];
    X[2]:=swapX;
  end if;
  #print(R);

  if R[2] < R[1] then # swapR them
    swapR:=R[2];
    R[2]:=R[1];
    R[1]:=swapR;

    swapX:= X[2];
    X[2]:=X[3];
    X[3]:=swapX;
  end if;
  #print([X,R]);

  if ( R[3] < 0 or 0 < R[1] or R[1] = 0 or R[2] = 0) then
  else
    if 0 < R[2] then
  print("we do it ", [X,R]);
    swapR:=R[1];
    R[1]:=R[2];
    R[2]:=R[3];
    R[3]:=swapR;

    swapX:=X[1];
    X[1]:=X[3];
    X[3]:=X[2];
    X[2]:=swapX;
    end if;
  end if;
```

```

end if;

return [X,R];
end proc;
arrangeXR := proc(x, r)
local R, X, swapR, swapX;
  X := [x[1], x[2], x[3]];
  R := [r[1], r[2], r[3]];
  if R[2] < R[1] then
    swapR := R[2]; R[2] := R[1]; R[1] := swapR; swapX := X[2]; X[2] := X[3]; X[3] := swapX
  end if;
  if R[3] < R[2] then
    swapR := R[3]; R[3] := R[2]; R[2] := swapR; swapX := X[1]; X[1] := X[2]; X[2] := swapX
  end if;
  if R[2] < R[1] then
    swapR := R[2]; R[2] := R[1]; R[1] := swapR; swapX := X[2]; X[2] := X[3]; X[3] := swapX
  end if;
  if R[3] < 0 or 0 < R[1] or R[1] = 0 or R[2] = 0 then
  else
    if 0 < R[2] then
      print("we do it ", [X, R]);
      swapR := R[1];
      R[1] := R[2];
      R[2] := R[3];
      R[3] := swapR;
      swapX := X[1];
      X[1] := X[3];
      X[3] := X[2];
      X[2] := swapX
    end if
  end if;
  return [X, R]
end proc

```

I simply use that for the test

```

> X:= [x1,x2,x3]: R:= [r12,r13,r23]:
LL:=arrangeXR(X,R):

# new values
r12 := LL[2][1];
r13 := LL[2][2];
r23 := LL[2][3];

x1 := LL[1][1];
x2 := LL[1][2];
x3 := LL[1][3];

r12 := -0.999995065196923
r13 := -0.999980260851202

```

```

r23 := 0.999995065201858
x1 := -0.400000000
x2 := 1.000000000
x3 := 0.330000000
> st:=time():
if 0 <= check_coefficients(r12,r13,r23) then
  print(`cdfN3 =`);
  cdfN3_lcc(x1,x2,x3,r12,r13,r23);
  #evalf(%,16);
end if;
new_result_from_DLL:=%;
`seconds`=time()-st;
      valid correlation matrix, determinant = 0.974223 10-16
      cdfN3 =
0.822162192527598196788847962130947249442660639960626051658136981277493942265067765874\
044067167190879386083 10-32
      seconds = 4.376

```


where Owen's function becomes somewhat nicer at least:

```

> theF(xi,x2,x3, r12,r13,r23):
tst:=unapply(%,xi):
tst(xi): op(%) : evalf(%,8) : cdfN2_lcc(%);

remDigits:=Digits: Digits:=14:
plot(tst, -0.41.. x1);
Digits:=remDigits:
      cdfN2_lcc(318.31025 + 318.30868 ξ, 52.521470 + 159.15283 ξ, 0.99999987)

```



Again compare both results:

```

> TheIntegrand:='eval(theIntegrand, L)';
TheIntegral:= 'Int(TheIntegrand,xi = -23.0 .. x1, method = _Gquad)';
``;
'TheIntegral'=evalf[105](TheIntegral):
`eval(cdfN2_lcc(a1+b1*x1, a2+b2*x1,r)*cdfN(x1), L) - TheIntegral =`;
evalf(eval(cdfN2_lcc(a1+b1*x1, a2+b2*x1,r)*cdfN(x1), L)

```

```

-TheIntegral,105);
`computational error` = % -new_result_from_DLL;
                    TheIntegrand := eval(theIntegrand, L)
                    TheIntegral := Int(TheIntegrand, ξ = -23.0 .. x1, method = _Gquad)

                    eval(cdfN2_lcc(a1+b1*x1, a2+b2*x1,r)*cdfN(x1), L) - TheIntegral =
0.822162192527598196788847962130947249442660639960626051658136981277493941993302336397\
84224873905886620 10-32
                    computational error = -0.271765429476201818428132013186083 10-104

```

The computational error is less than 1E-104 which is the exactness of qfloat (it is not in digits, but decimal places).

## Test 5

The correlation coefficients need not be close to 1 to give an ugly matrix neither the determinant need to be extremely small to cause difficulties:

```

> r12:=1/2;
r13:=1/2;
r23:=-1/2 + 10(-3); r23 := evalf(r23):

check_coefficients(r12,r13,r23):

                    r12 :=  $\frac{1}{2}$ 
                    r13 :=  $\frac{1}{2}$ 
                    r23 :=  $\frac{-499}{1000}$ 

                    valid correlation matrix, determinant = 0.00149900

> x1 := 1.000000000;
x2 := 0.330000000;
x3 := -0.400000000;

                    x1 := 1.000000000
                    x2 := 0.330000000
                    x3 := -0.400000000

```

The DLL asserts an result in moderate time:

```

> st:=time():
if 0 <= check_coefficients(r12,r13,r23) then
  print(`cdfN3 =`);
  cdfN3_lcc(x1,x2,x3,r12,r13,r23);
  #evalf(%,16);
end if;
new_result_from_DLL:=%;
`seconds`=time()-st;

                    valid correlation matrix, determinant = 0.00149900
                    cdfN3 =
0.142296394003137257703313606337041820472803102143364302364135641820820384071866324589\
730288702970003275120

                    seconds = 8.416

```




```
> TheIntegrand:='eval(theIntegrand, L)';
TheIntegral:= 'Int(evalf(TheIntegrand),xi = -23.0 .. x1)';
TheIntegrand := eval(theIntegrand, L)
```

$$\text{TheIntegral} := \int_{-23.0}^{x1} \text{evalf}(\text{TheIntegrand}) d\xi$$

And looking at the integrand Maple should be able to compute that (otherwise split it into pieces, for r23 towards -1/2 it becomes quite ugly ... I lost patients with Maple to integrate then and did not work out an example for that):

```
> remDigits:=Digits: Digits:=14:
plot(TheIntegrand,xi = -3.0 .. x1);
DI:=evalf(diff(TheIntegrand,xi)): expand(%): Tryhard(%): combine(% ,exp):
evalf(% ,14):
fsolve(%=0, xi=-1..0);
Digits:=remDigits:
```



Again have patients to let Maple compute it up to 105 Digits and then compare

```
> evalf(TheIntegrand): normal(%): evalf(%):
TheIntegral:= 'Int(% ,xi = -23.0 .. x1)':
TheI_105:=evalf(% ,105);
```

```
TheI_105 := -0.1422963940031372577033136063370418204728031021433643023641356418208203840\
71866324589730288702970003275121
```

```
> `eval(cdfN2_lcc(a1+b1*x1, a2+b2*x1,r)*cdfN(x1), L) - TheIntegral =`;
evalf(eval(cdfN2_lcc(a1+b1*x1, a2+b2*x1,r)*cdfN(x1), L) -TheI_105,105);
`computational error` = % -new_result_from_DLL;
```

```
eval(cdfN2_lcc(a1+b1*x1, a2+b2*x1,r)*cdfN(x1), L) - TheIntegral =
0.142296394003137257703313606337041820472803102143364302364135641820820384071866324589\
730288702970003275121
```

computational error = 0.1 10<sup>-104</sup>

which shows that within the exactness of qfloat the result is correct.

[ >