

Interfacing C code from R

Søren Højsgaard
Aalborg University, Denmark

November 9, 2012

Contents

1	Introduction	1
2	Preparation	1
3	Programming in C	2
3.1	Example: The exponential function	2
3.2	Example: Roots of 2nd degree polynomial	4
4	Further reading	5
5	Interfacing C-code with .C() and .Call()	5
5.1	Using .C()	5
5.2	Using .Call()	6
5.3	Using linear algebra libraries	7
5.4	Benchmarking - I	9
6	Calling R from C	10
7	Creating C code on the fly using inline	11
7.1	Example: The exponential function – using .C()	11
7.2	Example: Square the numbers $1, \dots, n$ – using .C()	13
7.3	Example: Multiplying matrices – using .Call()	13
7.4	Benchmarking - II	14
8	EXERCISES: C programming	14

1 Introduction

In the following we describe how to call C or C++ functions from R.

- Calling `C` (or `C++`) functions from `R` is usually done to gain speed.
- Knowledge of `C` / `C++` is not a prerequisite for the course. We will show some examples for inspiration and refer you to studying the languages on your own.
- A typical workflow is to first make a pure `R` implementation of your method; then identify possible bottlenecks wrt. computing time. These would then be a candidate for a `C` implementation.
- However, some of the `C` and `Fortran` code underlying `R` is very efficient so there is not necessarily much (if anything) to be gained from a `C` implementation
- The exposition is based on working on a Windows platform.

2 Preparation

For Linux users:

- Things seem to be working out of the box (on the most recent Ubuntu)

For Windows users:

- Download the `Rtools` from a CRAN mirror near you.
- Install `Rtools` and allow for the suggested change of the computers `PATH` variable.
- It is a good idea to read the file `Rtools.txt` in the `Rtools` installation directory; especially if you have `cygwin` installed on your computer.

For Mac users:

- Unknown – sorry

3 Programming in C

3.1 Example: The exponential function

Recall our implementation of the exponential function in `R`:

```

R> expf <- function(x, eps=.Machine$double.eps){
+   if (x<0) {
+     1/expf(-x)
+   } else {
+     n <- 0; term <- 1; ans <- 1
+     while(abs(term)> eps) {
+       n = n + 1
+       term = term * x / n
+       ans <- ans + term
+     }
+     ans
+   }
+ }

```

Various C-implementations are shown below.

```

1  /* Compile with gcc -o expfun1 expfun1.c */
2  #include <stdio.h>
3  #include <math.h>
4
5  int main (void){
6    char str[8];
7    double x, ans=1.0;
8    printf("---_expfun1_---\n");
9    printf("enter_x:_");
10   fgets(str, sizeof(str), stdin);
11   sscanf(str, "%lf", &x);
12
13   double term=1.0, eps=1e-16;
14   int n=0;
15   while (fabs(term)>eps){
16     n++;
17     term = term * x / n;
18     ans = ans + term;
19   }
20
21   double expx = exp(x);
22   double rerr = (ans-expx)/expx;
23   printf("input=%f_ ans=%f_ expx=%f_ rerr=%f\n", x, ans, expx, rerr);
24 }

```

```

1  /* Compile with gcc -o expfun2 expfun2.c */
2  #include <stdio.h>
3  #include <math.h>
4
5  double C_expfun2 (double x){
6    double ans=1.0, term=1.0, eps=1e-16;
7    int n=0;
8    while (fabs(term)>eps){
9      n++;
10     term = term * x / n;
11     ans = ans + term;
12   }
13   return(ans);
14 }
15
16 int main (void){
17   char str[8];
18   double x, ans;
19   printf("---_expfun2_---\n");
20   printf("enter_x:_");
21   fgets(str, sizeof(str), stdin);
22   sscanf(str, "%lf", &x);
23 }

```

```

24  if (x<0)
25      ans = 1/C_expfun2(-x);
26  else
27      ans = C_expfun2(x);
28
29  double expx = exp(x);
30  double rerr = (ans-expx)/expx;
31  printf("input=%f ans=%f expx=%f rerr=%f\n", x, ans, expx, rerr);
32  }

```

```

1  /* Compile with gcc -o expfun3 expfun3.c */
2  #include <stdio.h>
3  #include <math.h>
4
5  void C_expfun3 (double x, double *ans){
6      double term=1.0, eps=1e-16;
7      int n=0;
8      while (fabs(term)>eps){
9          n++;
10         term = term * x / n;
11         *ans = *ans + term;
12     }
13 }
14
15 int main (void){
16     char str[8];
17     double x, ans=1;
18     printf("---expfun2---\n");
19     printf("enter x:");
20     fgets(str, sizeof(str), stdin);
21     sscanf(str, "%lf", &x);
22
23     if (x<0){
24         C_expfun3(-x, &ans);
25         ans = 1/ans;
26     } else {
27         C_expfun3(x, &ans);
28     }
29
30     double expx = exp(x);
31     double rerr = (ans-expx)/expx;
32     printf("input=%f ans=%f expx=%f rerr=%f\n", x, ans, expx, rerr);
33 }

```

The workhorse for calculating the exponential is simple; everything else in the programs is about getting data in and out.

```

1
2 void C_expfun3 (double x, double *ans){
3     double term=1.0, eps=1e-16;
4     int n=0;
5     while (fabs(term)>eps){
6         n++;
7         term = term * x / n;
8         *ans = *ans + term;
9     }
10 }

```

This simple workhorse can be called directly from R. More about this later.

3.2 Example: Roots of 2nd degree polynomial

```
1 /* Compile into findroots2 with gcc -o findroots2 findroots2.c */
2 #include <stdio.h>
3 #include <math.h>
4 # define nroot 2
5 FILE *output;
6
7 void getroots (double parm[3], double roots[nroot]) {
8     roots[0] = (-parm[1] + sqrt(pow(parm[1],2) - 4*parm[0]*parm[2]))/(2*parm[0]);
9     roots[1] = (-parm[1] - sqrt(pow(parm[1],2) - 4*parm[0]*parm[2]))/(2*parm[0]);
10    return;
11 }
12
13 int main(void){
14     int i;
15     double coef[3], roots[nroot];
16     char line[256];
17
18     printf("Solving a x^2 + b x + c = 0\n");
19     printf("set a:"); fgets(line, sizeof(line), stdin);
20     sscanf(line, "%lf", &coef[0]);
21     printf("set b:"); fgets(line, sizeof(line), stdin);
22     sscanf(line, "%lf", &coef[1]);
23     printf("set c:"); fgets(line, sizeof(line), stdin);
24     sscanf(line, "%lf", &coef[2]);
25
26     sprintf(line, "a=%f b=%f c=%f\n", coef[0], coef[1], coef[2]);
27     printf(line);
28
29     getroots(coef, roots);
30     printf("Solution: %f %f", roots[0], roots[1]);
31
32     output = fopen("getroots.output.csv", "w");
33     for(i=0; i<nroot; i++) {
34         fprintf(output, "%f, ", roots[i]);
35     }
36     fprintf(output, "\n");
37     fclose(output);
38     return;
39 }
```

4 Further reading

- In general: The manual “Writing R Extensions”
- C programming: *The C Programming Language* by Kernighan and Ritchie. A classic.
- C programming: <http://www.cprogramming.com/>

5 Interfacing C-code with .C() and .Call()

C functions can be called from R using .C() and .Call().

- The .C() function is the basic way of calling C functions from R.

- With `.C()` all arguments must be “primitive” (essentially vectors of doubles or integers).
- `.C()` “returns” only primitive values.
- The `.Call()` interface is more difficult but offers several advantages
- `.Call()` accepts R data structures such as matrices and lists.
- `.Call()` returns R data structures.
- (Fortran code can be called in a way similar to `.C()` using `.Fortran()`. We shall not discuss calling Fortran in detail.)

We illustrate the mechanisms by considering matrix multiplication (which is something you do not want to write your own C functions for; those in R are very good).

```
R> A <- matrix(1:12, nc=6)
R> B <- matrix(1:18, nr=6)
R> A %*% B

      [,1] [,2] [,3]
[1,] 161  377  593
[2,] 182  434  686
```

5.1 Using `.C()`

A C function for matrix multiplication to be called with `.C()` is:

```
1  /* File: matprod1.c: Calculate the product of matrices X and Y */
2  void matprod1(double *X, int *dimX, double *Y, int *dimY, double *ans){
3      double sum;
4      int ii, jj, kk;
5      int nrX=dimX[0], ncX=dimX[1], nrY=dimY[0], ncY=dimY[1];
6
7      for (ii=0; ii<nrX; ii++){
8          for (jj=0; jj<ncY; jj++){
9              sum = 0;
10             for (kk=0; kk<ncX; kk++){
11                 sum = sum + X[ii+nrX*kk]*Y[kk+nrY*jj];
12             }
13             ans[ii+nrX*jj] = sum;
14         }
15     }
16 }
```

To compile the C function, issue the following command at the command prompt/in a shell:

```
R --vanilla CMD SHLIB src/matprod1.c
```

On Windows this creates `matprod1.dll`; on Linux this creates `matprod1.so`.

The function is called with:

```

R> mprod1_C_C <- function(A,B){ # Using .C and own C code
+   ans1 <- .C("matprod1", as.numeric(A), dim(A), as.numeric(B), dim(B),
+             result=numeric(nrow(A)*ncol(B)))$result
+   dim(ans1) <- c(nrow(A), ncol(B))
+   ans1
+ }
R> dyn.load("src/matprod1.dll")
R> mprod1_C_C(A, B)

      [,1] [,2] [,3]
[1,] 161  377  593
[2,] 182  434  686

R> dyn.unload("src/matprod1.dll")

```

Notice the following:

- Matrices are converted to vectors (of the right type) and dimensions of the matrices are supplied explicitly.
- A vector designated to hold the answer is created and supplied as well.

5.2 Using `.Call()`

It would be much more convenient to be able to simply pass the matrices A and B to the matrix multiplication function without having to pass the dimensions explicitly etc.

This can be done using SEXPs (S expressions).

- R objects are SEXPs.
- The SEXPs are passed to the C code where they are “digested”.
- Advantage: calling the C-function from R is simpler and faster code is often obtained.
- Disadvantage: The C code gets more involved.

A C function for matrix multiplication to be called with `.Call()` is:

```

1  /* File: matprod2.c: Calculates the product of matrices X and Y */
2  #include <Rdefines.h>
3  #include "matprod1.h"
4
5  SEXP matprod2(SEXP X, SEXP Y) {
6     int nprot=0;
7     PROTECT(X = AS_NUMERIC(X)); nprot++;      /* Digest SEXPs from R */
8     PROTECT(Y = AS_NUMERIC(Y)); nprot++;
9     double *xptr; xptr = REAL(X);
10    double *yptr; yptr = REAL(Y);
11    int *dimX; dimX = INTEGER(GET_DIM(X));
12    int *dimY; dimY = INTEGER(GET_DIM(Y));
13    SEXP ans;                                  /* Create SEXP to hold result */
14    PROTECT(ans = allocMatrix(REALSXP, dimX[0], dimY[1])); nprot++;
15    double *ansptr; ansptr = REAL(ans);
16    matprod1(xptr, dimX, yptr, dimY, ansptr); /* Calculate product */

```

```

17 UNPROTECT(nprot);          /* Wrap up; */
18 return(ans);              /* Return the result to R */
19 }

```

Above we used the header file `matprod1.h`:

```

1 void matprod1(double *X, int *dimX, double *Y, int *dimY, double *ans);

```

To compile the file do:

```

R --vanilla CMD SHLIB src/matprod2.c src/matprod1.c

```

Load the file with:

```

R> mprod2_Call_C <- function(A,B){ # Using .Call and own C code
+   .Call("matprod2", A, B)
+ }
R> dyn.load("src/matprod2.dll")
R> mprod2_Call_C(A, B)

      [,1] [,2] [,3]
[1,]  161  377  593
[2,]  182  434  686

R> dyn.unload("src/matprod2.dll")

```

5.3 Using linear algebra libraries

There are various libraries with linear algebra routines available online, e.g.

- blas (<http://www.netlib.org/blas/>) or
- lapack (<http://www.netlib.org/lapack/>) or
- linpack (<http://www.netlib.org/linpack/>) or

These routines are implemented in Fortran and the libraries are available in R.

Matrix multiplication can be made using the routine `dgemm` in blas as:

```

1  /* File: matprod3.c */
2  #include <R.h>
3  #include <R_ext/BLAS.h>
4
5  void matprod3(double *X, int *dimX, double *Y, int *dimY, double *ans){
6
7     int nrX=dimX[0], ncX=dimX[1], nrY=dimY[0], ncY=dimY[1];
8
9     char *transa = "N", *transb = "N";
10    double one = 1.0, zero = 0.0;
11    F77_CALL(dgemm)(transa, transb, &nrX, &ncY, &ncX, &one,
12                   X, &nrX, Y, &nrY, &zero, ans, &nrX);
13 }

```


We may alternatively use the `.Call()` interface:

```
1  /* File: matprod4.c */
2  #include <R.h>
3  #include <Rdefines.h>
4  #include <R_ext/BLAS.h>
5
6  SEXP matprod4(SEXP X, SEXP Y) {
7      int nprot=0;
8      PROTECT(X = AS_NUMERIC(X)); nprot++;
9      PROTECT(Y = AS_NUMERIC(Y)); nprot++;
10     double *xptra; xptra = REAL(X);
11     double *yptra; yptra = REAL(Y);
12     int *dimX; dimX = INTEGER(GET_DIM(X));
13     int *dimY; dimY = INTEGER(GET_DIM(Y));
14
15     SEXP ans; PROTECT(ans = allocMatrix(REALSXP, dimX[0], dimY[1])); nprot++;
16     double *ansptra; ansptra = REAL(ans);
17
18     char *transa = "N", *transb = "N";
19     double one = 1.0, zero = 0.0;
20     F77_CALL(dgemm)(transa, transb, &dimX[0], &dimY[1], &dimX[1], &one,
21                    xptra, &dimX[0], yptra, &dimY[0], &zero, ansptra, &dimX[0]);
22     UNPROTECT(nprot);
23     return(ans);
24 }
```

To compile these files, first create a file named `Makevars` with the contents

```
PKG_LIBS=$(BLAS_LIBS)
```

and then compile as described above – or compile both using

```
R --vanilla CMD SHLIB src/matprod3.c src/matprod4.c
```

which creates `matprod3.dll`.

The calls are as before

```

R> mprod3_C_For <- function(A,B){ # Using .C and a Fortran routine
+   ans1 <- .C("matprod3", as.numeric(A), dim(A), as.numeric(B), dim(B),
+             result=numeric(nrow(A)*ncol(B)))$result
+   dim(ans1) <- c(nrow(A), ncol(B))
+   ans1
+ }
R> mprod4_Call_For <- function(A,B){ # Using .Call and a Fortran routine
+   .Call("matprod4", A, B)
+ }
R> dyn.load("src/matprod3.dll")
R> mprod3_C_For(A, B)

      [,1] [,2] [,3]
[1,] 161  377  593
[2,] 182  434  686

R> mprod4_Call_For(A, B)

      [,1] [,2] [,3]
[1,] 161  377  593
[2,] 182  434  686

R> dyn.unload("src/matprod3.dll")

```

5.4 Benchmarking - I

Compile everything:

```
R CMD SHLIB -o src/matprod.dll src/matprod4.c src/matprod3.c src/matprod2.c src/matprod1.c
```

Compare performances:

```

R> library(rbenchmark)
R> cols <- c("test", "replications", "elapsed", "relative")

```

```

R> A <- matrix(rnorm(100), nr=10); B<-A[,1:5]
R> dyn.load("src/matprod.dll")
R> N <- 50000
R> benchmark(A %*% B, mprod1_C_C(A,B), mprod2_Call_C(A,B), mprod3_C_For(A,B),
+           mprod4_Call_For(A,B), columns=cols, order="relative", replications=N)

      test replications elapsed relative
3  mprod2_Call_C(A, B)      50000  0.56  1.000
1      A %*% B              50000  0.70  1.250
5  mprod4_Call_For(A, B)   50000  0.86  1.536
4  mprod3_C_For(A, B)     50000  1.61  2.875
2  mprod1_C_C(A, B)       50000  2.02  3.607

R> dyn.unload("src/matprod.dll")

```

Notice:

- Timewise, `.C()` performs poor relative to `.Call()`. There are two (at least) reasons:

- On the R side, as `.numeric()`, `numeric()` etc. takes time.
- The arguments are copied when calling the C code, and that takes time.

6 Calling R from C

It is possible to call R functions from within C. Consider the following (not very elegant) way of generating observations from a χ_f^2 -distribution.

```

1 #include <R.h>
2 #include <Rmath.h>
3 #include <math.h>
4
5 void myrchisq1(int *n, int *df, double *ans){
6     double chisq, tmp;
7     GetRNGstate();          /* set random number seed */
8     for (int jj=0; jj<*n; jj++){ //Rprintf("jj=%i\n", jj);
9         chisq=0;
10        for (int ii=0; ii<*df; ii++){
11            tmp = rnorm(0.0,1.0); /* call R function */
12            //Rprintf("%2ith draw= %8.4f \n", (ii+1), tmp);
13            chisq=chisq + tmp*tmp;
14        }
15        ans[jj] = chisq;
16    }
17    PutRNGstate();          /* write back ran number seed */
18 }

```

An alternative interface to the function is provided as follows:

```

1 #include <R.h>
2 #include <Rmath.h>
3 #include <math.h>
4 #include <Rdefines.h>
5
6 SEXP myrchisq2(SEXP n, SEXP df){
7     int nprot=0;
8     PROTECT(n); nprot++;
9     PROTECT(df); nprot++;
10    int _df = INTEGER_VALUE(df);
11    int _n  = INTEGER_VALUE(n);
12
13    SEXP ans;
14    PROTECT(ans=NEW_NUMERIC(_n)); nprot++;
15    double *ansptr; ansptr = REAL(ans);
16    myrchisq1(&_n, &_df, ansptr);
17    UNPROTECT(nprot);
18    return ans;
19 }

```

```
R --vanilla CMD SHLIB -o src/callRfromC.dll src/callRfromC2.c src/callRfromC1.c
```

```

R> nn <- 10
R> nu <- 7
R> dyn.load("src/callRfromC.dll")
R> .C("myrchisq1", as.integer(nn), as.integer(nu), ans=numeric(nn))$ans

[1] 14.438440  3.110139  4.501082  3.949214  5.574821 12.023367  7.682628
[8] 11.469489  6.076906  9.359699

R> .Call("myrchisq2", nn, nu)

[1] 1.922361  1.126529  6.489263  5.867347  5.965098  9.346552  4.489872
[8] 3.928444 11.865221 14.109356

R> dyn.unload("src/callRfromC.dll")

```

7 Creating C code on the fly using **inline**

- The **inline** package allows to define R functions with in-lined C (or C++) code.
- The **inline** package is very good for prototyping: getting code up and running quickly.
- Both `.C()` and `.Call()` may be used as calling mechanisms.
- The C (or C++) code is given as a text string.

7.1 Example: The exponential function – using `.C()`

```

R> expfun.src <- '
+ double term=1.0, eps=1e-16, ans=1.0;
+ int n=0;
+ while (fabs(term)>eps){
+   n++;
+   term = term * *x / n;
+   ans = ans + term;
+ }
+ *res = ans;
+ '
R> library(inline)
R> expfun <- cfunction(signature(x="numeric", res="numeric"), expfun.src, convention=".C")
R> expfun

An object of class "CFunc"
function (x, res)
  .Primitive(".C")(<pointer: 0x0000000067e814b0>, x = as.double(x),
    res = as.double(res))
<environment: 0x00000000aa85fe0>
Slot "code":
[1] "#include <R.h>\n\nnextern \"C\" {\n void filec34479062e4 ( double * x, double * res );\n}\n\nvoid filec34479062e4 ( dou

```

The function:

```

R> cat(expfun@code)

#include <R.h>

extern "C" {
  void filec34479062e4 ( double * x, double * res );
}

void filec34479062e4 ( double * x, double * res ) {

  double term=1.0, eps=1e-16, ans=1.0;
  int n=0;
  while (fabs(term)>eps){
    n++;
    term = term * *x / n;
    ans = ans + term;
  }
  *res = ans;
}

```

```

R> x <- 1
R> expfun(x, numeric(1))$res

[1] 2.718282

R> (expfun(x, numeric(1))$res - exp(x)) / exp(x)

[1] 1.633713e-16

R> x <- 20
R> expfun(x, numeric(1))$res

[1] 485165195

R> (expfun(x, numeric(1))$res - exp(x)) / exp(x)

[1] -1.228543e-16

R> x<- -20
R> expfun(x, numeric(1))$res

[1] 5.621884e-09

R> (expfun(x, numeric(1))$res - exp(x)) / exp(x)

[1] 1.727543

```

7.2 Example: Square the numbers $1, \dots, n$ – using `.C()`

```
R> codeSq <- "  
+   for (int i=0; i < *n; i++) {  
+     x[i] = x[i]*x[i];  
+   }"  
R> library("inline")  
R> sqfn <- cfunction(signature(n="integer", x="numeric"), codeSq, convention=".C")  
R> x <- as.numeric(1:10)  
R> n <- as.integer(10)  
R> sqfn(n, x)$x  
  
[1] 1 4 9 16 25 36 49 64 81 100
```

7.3 Example: Multiplying matrices – using `.Call()`

```
R> codeMatProd <- '  
+ /* Calculates the product of matrices X and Y */  
+   int nprot=0;  
+   PROTECT(X = AS_NUMERIC(X)); nprot++;  
+   PROTECT(Y = AS_NUMERIC(Y)); nprot++;  
+   double *xptr;   xptr = REAL(X);  
+   double *yptr;   yptr = REAL(Y);  
+   int nrX = INTEGER(GET_DIM(X))[0], ncX = INTEGER(GET_DIM(X))[1];  
+   int nrY = INTEGER(GET_DIM(Y))[0], ncY = INTEGER(GET_DIM(Y))[1];  
+   SEXP ans; PROTECT(ans = allocMatrix(REALSXP, nrX, ncY)); nprot++;  
+   double *ansptr; ansptr = REAL(ans);  
  
+   int ii, jj, kk;  
+   for (ii=0; ii<nrX; ii++){  
+     for (jj=0; jj<ncY; jj++){  
+       double sum = 0;  
+       for (kk=0; kk<ncX; kk++){  
+         sum += xptr[ii+nrX*kk]*yptr[kk+nrY*jj];  
+       }  
+       ansptr[ii+nrX*jj] = sum;  
+     }  
+   }  
+   UNPROTECT(nprot);  
+   return ans; '
```

```
R> library(inline)  
R> mprod5_inline_Call <- cfunction(signature(X = "matrix", Y="matrix"),  
+                                 body=codeMatProd,  
+                                 convention=".Call")
```

7.4 Benchmarking - II

```
R> A <- matrix(rnorm(100), nr=10); B<-A[,1:5]
R> dyn.load("src/matprod.dll")
R> N <- 2000
R> benchmark(A %*%B, mprod1_C_C(A,B),mprod2_Call_C(A,B),mprod3_C_For(A,B),
+             mprod4_Call_For(A,B), mprod5_inline_Call(A,B),
+             columns=cols, order="relative", replications=N)

      test replications elapsed relative
1      A %*% B           2000    0.03    1.000
3  mprod2_Call_C(A, B)     2000    0.03    1.000
5  mprod4_Call_For(A, B)   2000    0.03    1.000
6  mprod5_inline_Call(A, B) 2000    0.03    1.000
4  mprod3_C_For(A, B)     2000    0.11    3.667
2  mprod1_C_C(A, B)       2000    0.13    4.333
```

8 EXERCISES: C programming

1. Write a C-program that will find the roots of a 2nd degree polynomial
2. Interface this C-program from R.
3. Can you beat R's built-in function for matrix multiplication `%*%` in terms of computing time? In this quest you are allowed to use A^T as input when computing AB . You are also allowed to use parallel computing.