



Interfacing R and C

(IN WINDOWS)

A BASIC INTRODUCTION
FROM A BASIC USER

Brock Stewart, PhD
jnn6@cdc.gov

CDC R User Group
August 27, 2015

Outline

Introduction

C: A minimal intro just big enough for use with R

C in R: Basic layout using `.C()`

C in R: Advanced interfaces (`.Call`, `.External`, `Rcpp`, `inline`)

Postscript

Introduction

Why C in R? Interpreted languages (such as R) are slower than compiled code. Almost always unnecessary to create user-defined source code: many R functions already use compiled code; extra time involved is only justified in extreme situations. However, when we feel it is worth the effort, we can load compiled programs into R. Furthermore, learning C gives you a greater understanding of R, since much of R is written in C.

Why pointers? Memory management is fundamental to C programming; is a powerful feature of C. Even if using packages, wrappers, or other simplifying R tools, is useful to understand this aspect of C. R interfaces with C programs via pointers. Also, you'll likely need dynamic memory allocation, for which pointers are necessary.

C: A minimal intro just big enough for use with
R

About C

- ▶ Dennis Ritchie, AT&T Bell Labs, released 1972
- ▶ *The C Programming Language*: Kernighan & Ritchie, 1978
- ▶ Still one of the most – if not the most – widely used programming language (esp. including its variants)
- ▶ Was used to write OS's UNIX, Windows, Linux
- ▶ Was used to write R, MySQL, ...
- ▶ Can be executed by nearly any microprocessor you can find; portable to (almost?) every digital computing architecture that has ever existed: computers, phones, calculators, appliances, robots, machinery ... jet airplanes
- ▶ Allows programs to directly access and manipulate memory

C Basics

C

```
1  /* my C program */
2  #include <stdio.h>
3
4  int main()
5  {
6      printf("Hi RUG \n");
7      return 0;
8  }
```

Preprocessor directives

- ▶ header files
- ▶ constants
- ▶ macros

C program \approx collection of functions

- ▶ Preprocessor directives
- ▶ Functions
- ▶ Variables

Statements end in semicolons

C is case-sensitive

Base C is considered “small” – header files are similar to R libraries

Most header files you need will be included with your compiler

Can even include R header files to use R functions, e.g. `rnorm()`

C – Preprocessor Directives

C

```
1 #include <stdio.h>
2 #include myHeader.h
3 #define PI 3.141592654f
```

The statement `#include <stdio.h>` merges the contents of “stdio.h” into the program before compilation.

You can include your own header files by omitting the angle brackets.

Other things can be done with preprocessor directives; don't worry too much about them when first learning about C.

C – Variables

- ▶ Variable = a name given to an accessible storage area
- ▶ Define: `type variable_list;`
- ▶ Data types: `char`, `int`, `float`, `double`, `void`
- ▶ `variable_list` is one or more names separated by commas
- ▶ Initialize (Assign an initial value):

C

```
1 int ack, foo=4, bar;  
2 ack=3;
```

- ▶ Uninitialized variables have values either NULL or undefined, depending on type.

C – Variables & Memory

- ▶ `int a;`

Assigns the name “a” to a block of memory with exactly the size to store an int

- ▶ `a = 3;`

Assigns the value “3” to the storage block named “a”

- ▶ `int a=3;` does both

- ▶ Each memory block has:

- ▶ a size, which depends on the data type and system
- ▶ an address, which is assigned by the system

C – Functions

C

```
1 type name( param list )
2 {
3     body of the function
4 }
```

- ▶ type If the function returns a value, then the data type of the value. If the function does not return a value, then void
- ▶ name The function's name
- ▶ param list optional; comma-separated

Self-contained C programs must contain a function named “main”

C programs written for R should not contain a “main” function

C – Pointers

- ▶ Pointer = a variable whose value is the address of another variable
- ▶ Define: `type *name;` Same as other variables, with “*”
- ▶ The address of any variable can be obtained using “&”
- ▶ We also use “*” to access the value of the variable that a pointer is pointing to

C

```
1 int ack = 20;
2 int *ptr;
3 ptr = &ack;
4 printf("%x",&ack); /* 12fee2 */
5 printf("%x", ptr); /* 12fee2 */
6 printf("%d",*ptr); /* 20 */
```

C – Arrays

- ▶ A data structure storing a fixed-size collection of elements (variables) of the same type
- ▶ Consists of contiguous memory locations
- ▶ Declare: `type name[size];` where size is an `int > 0`
- ▶ Initialize:

```
C
1  int arr[6] = { 2, 20, 40, 12, 16, 17 };
2  /* or */
3  arr[0] = 2;
4  arr[1] = 20; /* etc. */
5  printf("%i",arr[3]); /* 12 */
```

- ▶ Indexing begins at 0
- ▶ Access elements with square brackets

C – Arrays & Pointers

- ▶ The name of an array is a pointer to the first element in the array.
- ▶ Pointers can be indexed like arrays.

C – Dynamic Memory Allocation

Memory Management Functions (from cplusplus.com)

- ▶ These functions are found in the header file `<stdlib.h>`.

`malloc(size)` Allocates a block of *size* bytes of memory, returning a pointer to the beginning of the block. The content is not initialized; indeterminate values.

`calloc(num, size)` Allocates a block of memory for an array of *num* elements, each of *size* bytes long, initialized to zero. I.e. allocates a zero-initialized memory block of ($num * size$) bytes.

`realloc(ptr, size)` Changes the size of the memory block pointed to by *ptr*. Existing content is preserved. Value of newly allocated portion is indeterminate.

`free(ptr)` Deallocates a block of memory previously allocated by a call to one of the above functions – the one pointed to by *ptr* – making it available again for further allocations.

C – Dynamic Memory Allocation

C

```
1  int conv=10, sum=3, i=2, j;  
2  int *arr = calloc(i,sizeof(int));  
3  arr[0] = 1;  
4  arr[1] = 2;  
5  while( sum < conv )  
6  {  
7      i++;  
8      arr = realloc(arr,i*sizeof(int));  
9      arr[i-1] = i;  
10     sum += arr[i-1];  
11 }
```

Can't do `int arr[i];` since we want to use `realloc`

Can replace `calloc` line with `int *arr = malloc(i*sizeof(int));`

C in R: Basic layout using .C()

Step 0: Install Rtools

- ▶ Why? In Linux, all the tools needed to build packages – including compiling source code – are included in the R installation. In Windows, these must be added separately.
- ▶ Get Rtools here:
<https://cran.r-project.org/bin/windows/Rtools/>
- ▶ Will install to C:\Rtools by default.
- ▶ Add Rtools to the Windows system PATH variable
- ▶ Authoritative info:
<https://cran.r-project.org/doc/manuals/R-admin.html#The-Windows-toolset>

Step 1: Write C code

- ▶ Should be composed of functions
- ▶ Do not write a “main” function
- ▶ All function arguments should be pointers
- ▶ Functions should not return any value

Step 2: Compile C code

Command Prompt

```
> cd filepath\to\your\code  
> R CMD SHLIB My_C_code.c
```

Can use `pushd` instead of `cd` for network drives

Step 3: Use Compiled code in R

R

```
1 dyn.load()  
2 .C("Fun",...)
```

list of parameters should be “typed”

Step 3: Use Compiled code in R

R

```
1 dyn.load()  
2 .C("Fun",...)
```

list of parameters should be “typed”

What happens if you don't type the parameters?

Nobody knows for sure, but don't provoke the gremlins



Tiny Example

C: foo01.c

```
1 void foo(int *n, double *x)
2 {
3     int i;
4     for (i=0; i<*n; i++)
5         x[i] = x[i]*x[i];
6 }
```

Command Prompt

```
> pushd \\cdc.gov\private\L132\JNN6\fooProj
> R CMD SHLIB foo01.c
```

R

```
1 dyn.load("foo01.dll")
2 .C("foo", n=as.integer(5), x=as.double(rnorm(5)))
```

Tiny Example

R

```
1 .C("foo", n=as.integer(3), x=as.double(rnorm(3)))
2 $n
3 [1] 3
4
5 $x
6 [1] 1.1859348 0.0226387 1.6485473
```

Output is a list, a named list when inputs are given names

Tiny Example

R

```
1 foo = function(y){
2     .C("foo",
3         n=as.integer(length(y)),
4         x=as.double(y)
5     )
6 }
7 foo(rnorm(3))
8 $n
9 [1] 3
10
11 $x
12 [1] 1.692649 2.530549 5.003614
```

Can wrap in an R function ...

Tiny Example

R

```
1 foo = function(y){
2     .C("foo",
3         n=as.integer(length(y)),
4         x=as.double(y)
5     )$x
6 }
7 foo(rnorm(3))
8 [1] 0.1241827 0.0109811 0.3376740
```

... and return only part of the outputted list ...

Tiny Example

R

```
1  foo = function(y){
2      if(!is.numeric(y))
3          stop("Hey! y must be numeric!")
4      val = .C("foo",
5              n=as.integer(length(y)),
6              x=as.double(y)
7          )
8      return(val$x)
9  }
10 foo(rnorm(3))
11 [1] 0.01478424 0.17330274 0.03974468
12 foo("bar")
13 Error in foo("bar") : Hey! y must be numeric!
```

... and include error-checking

Running-time Example

C

```
1 void convolve(double *a, int *na, double *b, int
   *nb, double *ab)
2 {
3     int nab = *na + *nb - 1;
4
5     for(int i = 0; i < nab; i++)
6         ab[i] = 0.0;
7
8     for(int i = 0; i < *na; i++)
9         for(int j = 0; j < *nb; j++)
10            ab[i + j] += a[i] * b[j];
11 }
```

From the Writing R Extensions manual on CRAN

Running-time Example

R

```
1 dyn.load("convolve.dll")
2 conv = function(a,b)
3   .C( "convolve",
4     as.double(a),
5     as.integer(length(a)),
6     as.double(b),
7     as.integer(length(b)),
8     ab = double(length(a) + length(b) - 1)
9   )$ab
```

Running-time Example

R

```
1 convR = function(a,b){
2   an=length(a)
3   bn=length(b)
4   nab=an+bn
5   ab=rep(0,nab)
6   for(i in 1:an)
7     for(j in 1:bn)
8       ab[i + j] = ab[i+j] + a[i] * b[j]
9   ab[-1]
10 }
```

Running-time Example

R

```
1 library(rbenchmark)
2 A=rnorm(1000)
3 B=rnorm(5000)
4 benchmark(
5     conv(A,B),
6     convR(A,B),
7     replications=10,order="relative"
8 )
9     test replications elapsed relative
10 conv(A, B)           10    0.05      1.0
11 convR(A, B)          10  112.09  2241.8
```

Further info on R CMD SHLIB

R CMD SHLIB --help

- ▶ Build a shared object for dynamic loading from the specified source or object files (which are automatically made from their sources) or linker options. If not given via '-output', the name for the shared object is determined from the first source or object file.
- ▶ -h: print short help message and exit
- ▶ -o=LIB: use LIB as (full) name for the built library
- ▶ -c: remove files created during compilation

Also see:

- ▶ [?SHLIB](#)
- ▶ [Writing R Extensions Section 5.5](#)

Further info on .C()

Writing R Extensions Section 5.2

- ▶ Provides interface to compiled code that's been linked into R either at build time or via `dyn.load()`

Further info on .C()

Writing R Extensions Section 5.2

- ▶ Provides interface to compiled code that's been linked into R either at build time or via `dyn.load()`
- ▶ Primarily intended for compiled C, but can be used with other languages which can generate C interfaces, e.g. C++

Further info on .C()

Writing R Extensions Section 5.2

- ▶ Provides interface to compiled code that's been linked into R either at build time or via `dyn.load()`
- ▶ Primarily intended for compiled C, but can be used with other languages which can generate C interfaces, e.g. C++
- ▶ 1st arg is char string specifying the function in a loaded DLL

Further info on .C()

Writing R Extensions Section 5.2

- ▶ Provides interface to compiled code that's been linked into R either at build time or via `dyn.load()`
- ▶ Primarily intended for compiled C, but can be used with other languages which can generate C interfaces, e.g. C++
- ▶ 1st arg is char string specifying the function in a loaded DLL
- ▶ Can test if a function is loaded using `is.loaded()`

Further info on .C()

Writing R Extensions Section 5.2

- ▶ Provides interface to compiled code that's been linked into R either at build time or via `dyn.load()`
- ▶ Primarily intended for compiled C, but can be used with other languages which can generate C interfaces, e.g. C++
- ▶ 1st arg is char string specifying the function in a loaded DLL
- ▶ Can test if a function is loaded using `is.loaded()`
- ▶ Up to 65 further args giving R objects to be passed to compiled code

Further info on .C()

Writing R Extensions Section 5.2

- ▶ Provides interface to compiled code that's been linked into R either at build time or via `dyn.load()`
- ▶ Primarily intended for compiled C, but can be used with other languages which can generate C interfaces, e.g. C++
- ▶ 1st arg is char string specifying the function in a loaded DLL
- ▶ Can test if a function is loaded using `is.loaded()`
- ▶ Up to 65 further args giving R objects to be passed to compiled code
- ▶ These are **copied** before being passed in, and again to an R list when the compiled code returns – that is *copied twice*

Further info on .C()

Writing R Extensions Section 5.2

- ▶ Provides interface to compiled code that's been linked into R either at build time or via `dyn.load()`
- ▶ Primarily intended for compiled C, but can be used with other languages which can generate C interfaces, e.g. C++
- ▶ 1st arg is char string specifying the function in a loaded DLL
- ▶ Can test if a function is loaded using `is.loaded()`
- ▶ Up to 65 further args giving R objects to be passed to compiled code
- ▶ These are **copied** before being passed in, and again to an R list when the compiled code returns – that is *copied twice*
- ▶ Look-up speed consideration: see PACKAGE argument

Further info on .C()

Writing R Extensions Section 5.2

- ▶ Provides interface to compiled code that's been linked into R either at build time or via `dyn.load()`
- ▶ Primarily intended for compiled C, but can be used with other languages which can generate C interfaces, e.g. C++
- ▶ 1st arg is char string specifying the function in a loaded DLL
- ▶ Can test if a function is loaded using `is.loaded()`
- ▶ Up to 65 further args giving R objects to be passed to compiled code
- ▶ These are **copied** before being passed in, and again to an R list when the compiled code returns – that is *copied twice*
- ▶ Look-up speed consideration: see PACKAGE argument
- ▶ The compiled code should not return anything except through its args; C funcs should be of type `void`

Further info on .C()

Writing R Extensions Section 5.2

- ▶ Provides interface to compiled code that's been linked into R either at build time or via `dyn.load()`
- ▶ Primarily intended for compiled C, but can be used with other languages which can generate C interfaces, e.g. C++
- ▶ 1st arg is char string specifying the function in a loaded DLL
- ▶ Can test if a function is loaded using `is.loaded()`
- ▶ Up to 65 further args giving R objects to be passed to compiled code
- ▶ These are **copied** before being passed in, and again to an R list when the compiled code returns – that is *copied twice*
- ▶ Look-up speed consideration: see PACKAGE argument
- ▶ The compiled code should not return anything except through its args; C funcs should be of type `void`
- ▶ Coerce all args to an R storage mode before calling .C

Further info on dyn.load

Writing R Extensions Section 5.3

- ▶ Compiled code to be used with R is loaded as a shared object, a DLL in Windows
- ▶ `dyn.unload()` to unload; not normally necessary, but is in Windows for, e.g., developing the C code
- ▶ See further info in this section regarding package building

C in R: Advanced interfaces (.Call, .External,
Rcpp, inline)

.Call vs .C

From <http://mazamascience.com/WorkingWithData/?p=1099>

`.C()`

- ▶ allows for simple C code that doesn't know R
- ▶ only simple data types can be passed
- ▶ all argument type conversion and checking must be done in R
- ▶ all arguments are copied locally before being passed to the C function (memory bloat)

`.Call()`

- ▶ allows for simple R code
- ▶ allows for complex data types
- ▶ allows for a C function return value
- ▶ does not require wasteful argument copying
- ▶ requires **much** more knowledge of R internals
- ▶ is the recommended, modern approach for serious C-in-R programmers

.Call vs .External and Rcpp

- ▶ Primary difference between .Call and .External: args are specified in .Call as in .C, i.e. as comma-separated list. .External has only 2 args: function name, and one argument holding all args sent to the compiled code.
- ▶ .External allows a variable number of arguments to be supplied to its args argument
- ▶ Rcpp: *Very* extensive development in recent years. Uses the .Call interface, so requires more knowledge of R internals than the simple use of .C
- ▶ My favorite Rcpp starting point:
<http://mazamascience.com/WorkingWithData/?p=1125>
- ▶ Rcpp, especially some of its extensions (sugar), requires less code than .Call

Some CRAN manual info on `.Call` (and `.External`)

Writing R Extensions Section 5.9

- ▶ Traditionally, speed up using C code done with `.C`. Use `.Call` or `.External` to utilize internal R structures. The R-side is similar between `.C` and `.Call`; the C-side changes substantially.
- ▶ `.Call` – 1st arg is also name of loaded function
- ▶ `.Call` – also takes up to 65 args to be passed to the compiled code
- ▶ Ask yourself: Do I really need it? Consider standard R first. If only atomic vectors, then consider `.C`.
- ▶ `.Call` and `.External` allow much more control, but impose greater responsibility and knowledge
- ▶ **Important:** With `.Call` (and `.External`) we use C structures, data types, functions, and macros available in header files, esp. [`Rinternals.h`](#) or [`Rdefines.h`](#)

Example C code for .Call – convolve2.c

C: convolve2.c

```
1  #include <R.h>
2  #include <Rinternals.h>
3  SEXP convolve2(SEXP a, SEXP b){
4  int na, nb, nab;
5  double *xa, *xb, *xab;
6  SEXP ab;
7  a = PROTECT(coerceVector(a, REALSXP));
8  b = PROTECT(coerceVector(b, REALSXP));
9  na = length(a); nb = length(b); nab = na + nb - 1;
10 ab = PROTECT(allocVector(REALSXP, nab));
11 xa = REAL(a); xb = REAL(b); xab = REAL(ab);
12 for(int i = 0; i < nab; i++) xab[i] = 0.0;
13 for(int i = 0; i < na; i++)
14 for(int j=0;j<nb;j++) xab[i+j] += xa[i]*xb[j];
15 UNPROTECT(3);
16 return ab;
17 }
```

Example C code for .Call – convolve2.c

Command Prompt

```
> pushd \\cdc.gov\private\L132\JNN6\convProj  
> R CMD SHLIB convolve2.c
```

R

```
1 dyn.load("convolve2.dll")  
2 conv2 = function(a, b) .Call("convolve2", a, b)
```

New C-side code

- ▶ Notice the use of R.h and Rinternals.h
- ▶ SEXP: a pointer to a structure with typedef SEXPREC
- ▶ SEXP = Simple *EX*pression (from LISP language)
- ▶ SEXPREC = Simple *EX*pression *REC*ord
- ▶ C func and its args are always type SEXP
- ▶ SEXP is kind of a “catch all” data type used to handle R objects; converted to int, double, etc. in C code
- ▶ Requires handling garbage collection: PROTECT, UNPROTECT (must balance)
- ▶ Details on SEXP and all of its types (e.g. REALSXP) found in [R Internals Section 1](#)
- ▶ Other functions from Rinternals.h, e.g.: coerceVector, allocVector

Postscript

Summary

.C is easy. Normal C code. Used with atomic R objects. Can use R-analogs Calloc and Realloc for dynamic memory allocation ([Writing R Extensions 6.1.2](#)). Quite a bit of R built with .C, you can do quite a bit with it too, although .Call considered more appropriate for “modern” usage by enthusiasts.

.Call has a *much* higher learning curve, due to use of utilities in Rinternals.h and other header files. Unfortunately, these aren't necessarily well-documented – might have to RTFS. However, allows much more utility, e.g. easily handles all types of R objects, including multi-dimensional, functions, expressions, and object attributes (names, class, etc.).

Rcpp *et al.* (which I didn't really cover) improve on .Call, allowing you the same (or more) power with simpler code than .Call. I strongly suggest learning and going through examples in the order presented here: C, .C, .Call, then Rcpp.

Additional References & Resources

C

- ▶ <http://www.tutorialspoint.com/cprogramming/index.htm>
- ▶ https://en.wikibooks.org/wiki/A_Little_C_Primer
- ▶ https://en.wikibooks.org/wiki/C_Programming
- ▶ <http://www.cprogramming.com/>
- ▶ <http://en.cppreference.com/w/>
- ▶ <http://www.cplusplus.com/>
- ▶ <http://forum.codecall.net/forum/161-cc-tutorials/>

R

- ▶ <http://adv-r.had.co.nz/C-interface.html>
- ▶ <http://users.stat.umn.edu/~geyer/rc/>
- ▶ <http://mcglinn.web.unc.edu/blog/linking-c-with-r-in-windows/>
- ▶ http://wiki.math.yorku.ca/index.php/R:_C_code:_including_it_in_R
- ▶ <https://nugaemeae.wordpress.com/2012/08/30/the-why-and-the-how-of-installing-rtools/>

References

- ▶ *Writing R Extensions*, Chapters 5 & 6
<https://cran.r-project.org/doc/manuals/r-release/R-exts.html>
- ▶ *R Internals*, Chapter 1
<https://cran.r-project.org/doc/manuals/R-ints.html>
- ▶ ?SHLIB
- ▶ ?C
- ▶ ?Call
- ▶ ?External