

# Seamless R and C++ Integration: Rcpp and RInside

Dirk Eddebuettel  
Debian Project

*Joint work with Romain François*

Invited seminar presentation  
Institute for Statistics and Mathematics  
Wirtschaftsuniversität Wien  
20 May 2010



# Preliminaries

- We assume a recent version of R so that `install.packages(c("Rcpp", "RInside", "inline"))` gets us current versions of the packages
- All examples shown should work 'as is' on Linux, OS X and Windows *provided a complete R development environment*
- The *R Installation and Administration* manual is an excellent start if you need to address the preceding point
- In particular, one must use the same compilers used to build R in order to extend or embed the R engine
- However, there is a known issue with the current RInside / Rcpp on Windows; but releases 0.2.1 and 0.7.1 *do* work




# Outline

- 1 Extending R
  - Why ?
  - The standard API
  - Inline
- 2 Rcpp
  - Overview
  - New API
  - Examples
- 3 Rcpp Usage Examples
  - RInside
  - RProtoBuf
  - Others
- 4 Summary
  - Key points
  - Resources



# Motivation



Chambers. *Software for Data Analysis: Programming with R*. Springer, 2008

Chambers (2008) opens chapter 11 (*Interfaces I: Using C and Fortran*) with these words:

*Since the core of R is in fact a program written in the C language, it's not surprising that the most direct interface to non-R software is for code written in C, or directly callable from C. All the same, including additional C code is a serious step, with some added dangers and often a substantial amount of programming and debugging required. You should have a good reason.*



# Motivation

Chambers (2008) then proceeds with this rough map of the road ahead:

## Against:

- It's more work
- Bugs will bite
- Potential platform dependency
- Less readable software

## In Favor:

- New and trusted computations
- Speed
- Object references

So is the deck stacked against us?



# Outline

- 1 Extending R
  - Why ?
  - The standard API
  - Inline
- 2 Rcpp
  - Overview
  - New API
  - Examples
- 3 Rcpp Usage Examples
  - RInside
  - RProtoBuf
  - Others
- 4 Summary
  - Key points
  - Resources



# Compiled Code: The Basics

R offers several functions to access compiled code: we focus on `.C` and `.Call` here. (*R Extensions*, sections 5.2 and 5.9; *Software for Data Analysis*).

The canonical example is the convolution function:

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



# Compiled Code: The Basics cont.

The convolution function is called from R by

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

As stated in the manual, one must take care to coerce all the arguments to the correct R storage mode before calling `.C` as mistakes in matching the types can lead to wrong results or hard-to-catch errors.





# Example: Running the convolution code via .C

All these files are at <http://dirk.eddelbuettel.com/code/rcppTut>

- Turn the C source file into a dynamic library using  
`R CMD SHLIB convolve.C.c`
- Load it inside an R script or session using  
`dyn.load("convolve.C.so")`
- Use it via the `.C()` interface as shown on previous slide
- All together in a helper file `convolve.C.sh`

```
#!/bin/sh
```

```
R CMD SHLIB convolve.C.c
```

```
cat convolve.C.call.R | R --no-save
```



# Compiled Code: The Basics cont.

Using `.Call`, the example becomes

```
1 #include <R.h>
2 #include <Rdefines.h>
3
4 extern "C" SEXP convolve2(SEXP a, SEXP b)
5 {
6     int i, j, na, nb, nab;
7     double *xa, *xb, *xab;
8     SEXP ab;
9
10    PROTECT(a = AS_NUMERIC(a));
11    PROTECT(b = AS_NUMERIC(b));
12    na = LENGTH(a); nb = LENGTH(b); nab = na + nb - 1;
13    PROTECT(ab = NEW_NUMERIC(nab));
14    xa = NUMERIC_POINTER(a); xb = NUMERIC_POINTER(b);
15    xab = NUMERIC_POINTER(ab);
16    for(i = 0; i < nab; i++) xab[i] = 0.0;
17    for(i = 0; i < na; i++)
18        for(j = 0; j < nb; j++) xab[i + j] += xa[i] * xb[j];
19    UNPROTECT(3);
20    return(ab);
21 }
```



# Compiled Code: The Basics cont.

Now the call simplifies to just the function name and the vector arguments—all other handling is done at the C/C++ level:

```
1 conv <- function(a, b) .Call("convolve2", a, b)
```

In summary, we see that

- there are different entry points
- using different calling conventions
- leading to code that may need to do more work at the lower level.



# Example: Running the convolution code via .Call

- Turn the C source file into a dynamic library using
- Load it inside an R script or session using
- Use it via the `.Call()` interface as shown previously
- All together in a helper file `convolve.Call.sh`

```
R CMD SHLIB convolve.Call.c
```

```
dyn.load("convolve.Call.so")
```

```
#!/bin/sh
```

```
R CMD SHLIB convolve.Call.c
```

```
cat convolve.Call.call.R | R --no-save
```



# Outline

- 1 Extending R
  - Why ?
  - The standard API
  - **Inline**
- 2 Rcpp
  - Overview
  - New API
  - Examples
- 3 Rcpp Usage Examples
  - RInside
  - RProtoBuf
  - Others
- 4 Summary
  - Key points
  - Resources



# Compiled Code: inline

`inline` is a package by Oleg Sklyar et al that provides the function `cfunction` which can wrap Fortran, C or C++ code.

```
1 ## A simple Fortran example
2 code <- "
3     integer i
4     do 1 i=1, n(1)
5     1 x(i) = x(i)**3
6 "
7 cubefn <- cfunction(signature(n="integer", x="numeric"),
8                     code, convention=".Fortran")
9 x <- as.numeric(1:10)
10 n <- as.integer(10)
11 cubefn(n, x)$x
```

`cfunction` takes care of compiling, linking, loading, ... by placing the resulting dynamically-loadable object code in the per-session temporary directory used by R.



# Example: Convolution via .C with inline

Using the file `convolve.C.inline.R`

```
1 require(inline)
2
3 code <- "int i, j, nab = *na + *nb - 1;
4
5     for(i = 0; i < nab; i++)
6         ab[i] = 0.0;
7
8     for(i = 0; i < *na; i++) {
9         for(j = 0; j < *nb; j++)
10            ab[i + j] += a[i] * b[j];
11     }
12 "
13
14 fun <- cfunction(signature(a="numeric", na="numeric",
15                          b="numeric", nb="numeric",
16                          ab="numeric"),
17                 code, language="C", convention=".C")
18
19 fun(1:10, 10, 10:1, 10, numeric(19))$ab
```



# Example: Convolution via .Call with inline

Using the file `convolve.Call.inline.R`

```
1 require(inline)
2 code <- "int i, j, na, nb, nab;
3     double *xa, *xb, *xab;
4     SEXP ab;
5
6     PROTECT(a = AS_NUMERIC(a)); PROTECT(b = AS_NUMERIC(b));
7     na = LENGTH(a); nb = LENGTH(b); nab = na + nb - 1;
8     PROTECT(ab = NEW_NUMERIC(nab));
9
10    xa = NUMERIC_POINTER(a); xb = NUMERIC_POINTER(b);
11    xab = NUMERIC_POINTER(ab);
12    for(i = 0; i < nab; i++) xab[i] = 0.0;
13
14    for(i = 0; i < na; i++)
15        for(j = 0; j < nb; j++)
16            xab[i + j] += xa[i] * xb[j];
17
18    UNPROTECT(3);
19    return(ab); "
20
21 fun <- cfunction(signature(a="numeric", b="numeric"),
22                 code, language="C")
23
24 fun(1:10, 10:1)
```





# Outline

- 1 Extending R
  - Why ?
  - The standard API
  - Inline
- 2 Rcpp
  - **Overview**
  - New API
  - Examples
- 3 Rcpp Usage Examples
  - RInside
  - RProtoBuf
  - Others
- 4 Summary
  - Key points
  - Resources



# Compiled Code: Rcpp

In a nutshell:

- `Rcpp` makes it easier to interface C++ and R code.
- Using the `.Call` interface, we can use features of the C++ language to automate the tedious bits of the macro-based C-level interface to R.
- One major advantage of using `.Call` is that richer R objects (vectors, matrices, lists, ... in fact most SEXP types incl functions, environments etc) can be passed directly between R and C++ without the need for explicit passing of dimension arguments.
- By using the C++ class layers, we do not need to manipulate the SEXP objects using any of the old-school C macros.
- `inline` eases usage, development and testing.



# Example: Convolution using classic Rcpp

Using the file `convolve.Call.Rcpp.classic.R`

```
1 require(inline)
2 code <- '
3   RcppVector<double> xa(a);
4   RcppVector<double> xb(b);
5
6   int nab = xa.size() + xb.size() - 1;
7   RcppVector<double> xab(nab);
8   for (int i = 0; i < nab; i++) xab(i) = 0.0;
9
10  for (int i = 0; i < xa.size(); i++)
11    for (int j = 0; j < xb.size(); j++)
12      xab(i + j) += xa(i) * xb(j);
13
14  RcppResultSet rs;
15  rs.add("ab", xab);
16  return rs.getReturnList();
17 '
18
19 fun <- cppfunction(signature(a="numeric", b="numeric"), code)
20 fun(1:10, 10:1)
```



# Outline

- 1 Extending R
  - Why ?
  - The standard API
  - Inline
- 2 Rcpp
  - Overview
  - **New API**
  - Examples
- 3 Rcpp Usage Examples
  - RInside
  - RProtoBuf
  - Others
- 4 Summary
  - Key points
  - Resources



# Rcpp: The 'New API'

Rcpp was significantly extended over the last few months to permit more natural expressions. Consider this comparison between the R API and the new Rcpp API:

```

1 SEXP ab;
2 PROTECT(ab = allocVector(STRSXP, 2));
3 SET_STRING_ELT( ab, 0, mkChar("foo") );
4 SET_STRING_ELT( ab, 1, mkChar("bar") );
5 UNPROTECT(1);

```

```

1 CharacterVector ab(2) ;
2 ab[0] = "foo" ;
3 ab[1] = "bar" ;

```

Data types, including STL containers and iterators, can be nested and other niceties. Implicit converters allow us to combine types:

```

1 std::vector<double> vec;
2 [...]
3 List x(3);
4 x[0] = vec;
5 x[1] = "some text";
6 x[2] = 42;

```

```

1 // With Rcpp 0.7.11 or later we can do:
2 std::vector<double> vec;
3 [...]
4 List x = List::create(vec,
5                       "some text",
6                       42);

```

# Functional programming in both languages

In R, functional programming is easy:

```

1 R> data(faithful); lapply(faithful, summary)
2 $eruptions
3   Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
4   1.60   2.16   4.00   3.49   4.45   5.10
5
6 $waiting
7   Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
8   43.0   58.0   76.0   70.9   82.0   96.0

```

We can do that in C++ as well and pass the R function down to the data that we let the STL `transform` function iterate over:

```

1 src <- 'Rcpp::List input(data);
2       Rcpp::Function f(fun) ;
3       Rcpp::List output(input.size());
4       std::transform(input.begin(), input.end(), output.begin(), f);
5       output.names() = input.names();
6       return output; '
7 cpp_lapply <- cppfunction(signature(data="list", fun = "function"), src)

```



# Exception handling

## Automatic catching and conversion of C++ exceptions:

```
R> library(Rcpp); library(inline)
R> cpp <- '
+     Rcpp::NumericVector x(xs); // automatic conversion from SEXP
+     for (int i=0; i<x.size(); i++) {
+         if (x[i] < 0)
+             throw std::range_error("Non-negative values required");
+         x[i] = log(x[i]);
+     }
+     return x; // automatic conversion to SEXP
+ '
R> fun <- cppfunction(signature(xs="numeric"), cpp)
R> fun( seq(2, 5) )

[1] 0.6931 1.0986 1.3863 1.6094

R> fun( seq(5, -2) )

Error in fun(seq(5, -2)) : Non-negative values required

R> fun( LETTERS[1:5] )

Error in fun(LETTERS[1:5]) : not compatible with INTSEXP
R>
```



# Exception handling: Usage

- We attempted to automate forwarding of exceptions from the C++ layer to the R layer.
- This works (thanks to some `gcc` magic) on operating system with an X in their name, but not on Windows.
- We therefore once again recommend to wrap code with

```
try {
```

and

```
    } catch( std::exception &ex) {  
        forward_exception_to_r(ex);  
    } catch(...) {  
        ::Rf_error("c++ exception (unknown reason)");  
    }  
}
```

- Because this is invariant, we provide macros `BEGIN_RCPP` and `END_RCPP`.
- We provide a variant `cppfunction` of `inline::cfunction` which automatically inserts these at the beginning and end of the code snippets.





# Outline

- 1 Extending R
  - Why ?
  - The standard API
  - Inline
- 2 Rcpp
  - Overview
  - New API
  - **Examples**
- 3 Rcpp Usage Examples
  - RInside
  - RProtoBuf
  - Others
- 4 Summary
  - Key points
  - Resources



# Example: Convolution using new Rcpp

Using the file `convolve.Call.Rcpp.new.R`

```
1 require(inline)
2
3 code <- '
4   Rcpp::NumericVector xa(a); // automatic conversion from SEXP
5   Rcpp::NumericVector xb(b);
6
7   int n_xa = xa.size();
8   int n_xb = xb.size();
9   int nab = n_xa + n_xb - 1;
10
11  Rcpp::NumericVector xab(nab);
12
13  for (int i = 0; i < n_xa; i++)
14    for (int j = 0; j < n_xb; j++)
15      xab[i + j] += xa[i] * xb[j];
16
17  return xab; // automatic conversion to SEXP
18 '
19
20 fun <- cppfunction(signature(a="numeric", b="numeric"), code)
21
22 fun(1:10, 10:1)
```



# Speed comparison

See the directory `Rcpp/examples/ConvolveBenchmarks`

In a recently-submitted paper, the following table summarises the performance of convolution examples:

Implementation	Time in millisec	Relative to R API
R API (as benchmark)	32	
<code>RcppVector&lt;double&gt;</code>	354	11.1
<code>NumericVector::operator[]</code>	52	1.6
<code>NumericVector::begin</code>	33	1.0

Table 1: Performance for convolution example

We averaged 1000 replications with two 100-element vectors – see `examples/ConvolveBenchmarks/` in `Rcpp` for details.



# Another Speed Comparison Example

- Regression is a key component of many studies. In simulations, we often want to run a very large number of regressions.
- R has `lm()` as the general purposes function. It is very powerful and returns a rich object—but it is not *lightweight*.
- For this purpose, R has `lm.fit()`. But, this does not provide all relevant auxiliary data as *e.g.* the standard error of the estimate.
- For the most recent *Introduction to High-Performance Computing with R* tutorial, I had written a hybrid R/C/C++ solution using the GNU GSL.
- We complement this with a new C++ implementation around the Armadillo linear algebra classes.



# Linear regression via GSL: lmGSL()

See the directory `Rcpp/examples/FastLM`

```

1  lmGSL <- function() {
2    src <- '
3
4    RcppVectorView<double> Yr(Ysexp);
5    RcppMatrixView<double> Xr(Xsexp);
6
7    int i, j, n = Xr.dim1(), k = Xr.dim2();
8    double chi2;
9
10   gsl_matrix *X = gsl_matrix_alloc(n, k);
11   gsl_vector *y = gsl_vector_alloc(n);
12   gsl_vector *c = gsl_vector_alloc(k);
13   gsl_matrix *cov = gsl_matrix_alloc(k, k);
14
15   for (i = 0; i < n; i++) {
16     for (j = 0; j < k; j++) {
17       gsl_matrix_set(X, i, j, Xr(i, j));
18     }
19     gsl_vector_set(y, i, Yr(i));
20   }
21
22   gsl_multifit_linear_workspace *wk =
23     gsl_multifit_linear_alloc(n, k);
24   gsl_multifit_linear(X, y, c, cov, &chi2, wk);
25   gsl_multifit_linear_free(wk);
26   RcppVector<double> StdErr(k);
27   RcppVector<double> Coef(k);

```

```

28   for (i = 0; i < k; i++) {
29     Coef(i) = gsl_vector_get(c, i);
30     StdErr(i) =
31       sqrt(gsl_matrix_get(cov, i, i));
32   }
33
34   gsl_matrix_free(X);
35   gsl_vector_free(y);
36   gsl_vector_free(c);
37   gsl_matrix_free(cov);
38
39   RcppResultSet rs;
40   rs.add("coef", Coef);
41   rs.add("stderr", StdErr);
42
43   return = rs.getReturnList();
44   '
45   ## turn into a function that R can call
46   ## args redundant on Debian/Ubuntu
47   fun <-
48     cppfunction(signature(Ysexp="numeric",
49       Xsexp="numeric"), src,
50     includes=
51       "#include <gsl/gsl_multifit.h>",
52     cppargs="-I/usr/include",
53     libargs="-lgsl -lgslcblas")
54 }

```

# Linear regression via Armadillo: ImArmadillo example

Also see the directory `Rcpp/examples/FastLM`

```

1  ImArmadillo <- function() {
2      src <- '
3      Rcpp::NumericVector yr(Ysexp);
4      Rcpp::NumericVector Xr(Xsexp);          // actually an n x k matrix
5      std::vector<int> dims = Xr.attr("dim");
6      int n = dims[0], k = dims[1];
7      arma::mat X(Xr.begin(), n, k, false);    // use advanced armadillo constructors
8      arma::colvec y(yr.begin(), yr.size());
9      arma::colvec coef = solve(X, y);        // model fit
10     arma::colvec resid = y - X*coef;         // to comp. std.errr of the coefficients
11     arma::mat covmat = trans(resid)*resid/(n-k) * arma::inv(arma::trans(X)*X);
12
13     Rcpp::NumericVector coefr(k), stderrestr(k);
14     for (int i=0; i<k; i++) {                // with RcppArmadillo template converters
15         coefr[i] = coef[i];                 // this would not be needed but we only
16         stderrestr[i] = sqrt(covmat(i,i));  // have Rcpp.h here
17     }
18
19     return Rcpp::List::create( Rcpp::Named( "coefficients", coefr), // Rcpp 0.7.11
20                               Rcpp::Named( "stderr", stderrestr));
21
22
23     ## turn into a function that R can call
24     fun <- cppfunction(signature(Ysexp="numeric", Xsexp="numeric"),
25                         src, includes="#include <armadillo>",
26                         cppargs="-I/usr/include", libargs="-larmadillo")
27 }

```

# Linear regression via Armadillo: RcppArmadillo

See `fastLm` in the RcppArmadillo package

`fastLm` in the new RcppArmadillo release does even better:

```

1 #include <RcppArmadillo.h>
2 extern "C" SEXP fastLm(SEXP ys, SEXP Xs) {
3   try {
4     Rcpp::NumericVector yr(ys);           // creates Rcpp vector from SEXP
5     Rcpp::NumericMatrix Xr(Xs);          // creates Rcpp matrix from SEXP
6     int n = Xr.nrow(), k = Xr.ncol();
7
8     arma::mat X(Xr.begin(), n, k, false); // reuses memory and avoids extra copy
9     arma::colvec y(yr.begin(), yr.size(), false);
10
11     arma::colvec coef = arma::solve(X, y); // fit model y ~ X
12     arma::colvec res = y - X*coef;        // residuals
13
14     double s2 = std::inner_product(res.begin(), res.end(), res.begin(), double()) / (n-k);
15                                     // std.errors of coefficients
16     arma::colvec stderr = arma::sqrt(s2*arma::diagvec(arma::inv(arma::trans(X)*X)));
17
18     return Rcpp::List::create(Rcpp::Named("coefficients") = coef,
19                               Rcpp::Named("stderr") = stderr,
20                               Rcpp::Named("df") = n - k);
21   } catch( std::exception &ex ) {
22     forward_exception_to_r( ex );
23   } catch (...) {
24     ::Rf_error( "c++ exception (unknown reason)" );
25   }
26   return R_NilValue; // -Wall
27 }

```



# Linear regression via GNU GSL: RcppGSL

See `fastLm` in the RcppGSL package (on R-Forge)

We also wrote `fastLm` in a new package RcppGSL:

```


1 extern "C" SEXP fastLm(SEXP ys, SEXP Xs) {
2   BEGIN_RCPP
3   RcppGSL::vector<double> y = ys;      // create gsl data structures from SEXP
4   RcppGSL::matrix<double> X = Xs;
5   int n = X.nrow(), k = X.ncol();
6   double chisq;
7   RcppGSL::vector<double> coef(k);     // to hold the coefficient vector
8   RcppGSL::matrix<double> cov(k,k);   // and the covariance matrix
9   // the actual fit requires working memory we allocate and free
10  gsl_multifit_linear_workspace *work = gsl_multifit_linear_alloc (n, k);
11  gsl_multifit_linear (X, y, coef, cov, &chisq, work);
12  gsl_multifit_linear_free (work);
13  // extract the diagonal as a vector view
14  gsl_vector_view diag = gsl_matrix_diagonal(cov) ;
15  // currently there is not a more direct interface in Rcpp::NumericVector
16  // that takes advantage of wrap, so we have to do it in two steps
17  Rcpp::NumericVector stderr ; stderr = diag;
18  std::transform( stderr.begin(), stderr.end(), stderr.begin(), sqrt );
19  Rcpp::List res = Rcpp::List::create(Rcpp::Named("coefficients") = coef,
20                                     Rcpp::Named("stderr") = stderr,
21                                     Rcpp::Named("df") = n - k);
22  // free all the GSL vectors and matrices — as these are really C data structures
23  // we cannot take advantage of automatic memory management
24  coef.free(); cov.free(); y.free(); X.free();
25  return res;    // return the result list to R
26  END_RCPP
27 }

```





# Rcpp Example: Regression timings




The small `longley` example exhibits less variability between methods, but the larger data set shows the gains more clearly.

For the small data set, all three appear to improve similarly on `lm`.

Source: Our calculations, see `examples/FastLM/` in `Rcpp`.

# Another Rcpp example (cont.)

Comparison of R and linear model fit routines



By dividing the `lm` time by the respective times, we obtain the 'possible gains' from switching.

One caveat, measurements depends critically on the size of the data as well as the cpu and libraries that are used.

Source: Our calculations, see `examples/FastLM/` in Rcpp.

# Possible gains from template meta-programming

Armadillo uses delayed evaluation (via recursive template and template meta-programming) to combine several operations into one expression reducing / eliminating temporary objects.

Operation	Relative performance improvement for			
	small matrices		medium to large	
	IT++	Newmat	IT++	Newmat
<code>A + B</code>	15.0	10.0	3.5	1.0
<code>A + B + C + D</code>	15.0	10.0	6.0	1.5
<code>A * B * C * D</code>	2.5	10.0	2.5	20.0
<code>B.row(size-1) = A.row(0)</code>	16.0	44.0	2.0	4.5
<code>trans(p) * inv(diagmat(q)) * r</code>	77.0	23.0	1086.0	5.0

Table 2: Gains from C++ template programming

See <http://arma.sourceforge.net/speed.html> for details.



# Outline

- 1 Extending R
  - Why ?
  - The standard API
  - Inline
- 2 Rcpp
  - Overview
  - New API
  - Examples
- 3 Rcpp Usage Examples
  - RInside
  - RProtoBuf
  - Others
- 4 Summary
  - Key points
  - Resources



# From RApache to littler to RInside

See the file `RInside/standard/rinside_sample0.cpp`

Jeff Horner's work on `RApache` lead to joint work in `littler`, a scripting / cmdline front-end. As it embeds `R` and simply 'feeds' the REPL loop, the next step was to embed `R` in proper `C++` classes: `RInside`.

```
1 #include <RInside.h> // for the embedded R via RInside
2
3 int main(int argc, char *argv[]) {
4     RInside R(argc, argv); // create an embedded R instance
5
6     R["txt"] = "Hello, world!\n"; // assign a char* (string) to 'txt'
7
8     R.parseEvalQ("cat(txt)"); // eval the init string, ignoring any returns
9
10    exit(0);
11
12 }
```



# Another simple example

See `RInside/standard/rinside_sample8.cpp` (in SVN, older version in pkg)

This example shows some of the new assignment and converter code:

```
1 #include <RInside.h> // for the embedded R via RInside
2
3
4 int main(int argc, char *argv[]) {
5
6     RInside R(argc, argv); // create an embedded R instance
7
8     R["x"] = 10 ;
9     R["y"] = 20 ;
10
11     R.parseEvalQ("z <- x + y") ;
12
13     int sum = R["z"];
14
15     std::cout << "10 + 20 = " << sum << std::endl ;
16     exit(0);
17 }
```



# A finance example

See the file `RInside/standard/rinside_sample4.cpp` (edited)

```

1  #include <RInside.h>                // for the embedded R via RInside
2  #include <iomanip>
3  int main(int argc, char *argv[]) {
4      RInside R(argc, argv);          // create an embedded R instance
5      SEXP ans;
6      R.parseEvalQ("suppressMessages(library(fPortfolio))");
7      txt = "lppData <- 100 * LPP2005.RET[, 1:6]; "
8           "ewSpec <- portfolioSpec(); nAssets <- ncol(lppData); ";
9      R.parseEval(txt, ans);          // prepare problem
10     const double dvec[6] = { 0.1, 0.1, 0.1, 0.1, 0.3, 0.3 }; // weights
11     const std::vector<double> w(dvec, &dvec[6]);
12     R.assign( w, "weightsvec");      // assign STL vec to R's 'weightsvec'
13
14     R.parseEvalQ("setWeights(ewSpec) <- weightsvec");
15     txt = "ewPortfolio <- feasiblePortfolio(data = lppData, spec = ewSpec, "
16         "constraints = \"LongOnly\"); "
17         "print(ewPortfolio); "
18         "vec <- getCovRiskBudgets(ewPortfolio@portfolio)";
19     ans = R.parseEval(txt);          // assign covRiskBudget weights to ans
20     Rcpp::NumericVector V(ans);      // convert SEXP variable to an RcppVector
21
22     ans = R.parseEval("names(vec)"); // assign columns names to ans
23     Rcpp::CharacterVector n(ans);
24
25     for (int i=0; i<names.size(); i++) {
26         std::cout << std::setw(16) << n[i] << "\t" << std::setw(11) << V[i] << "\n";
27     }
28     exit(0);
29 }

```



# And another *parallel* example

See the file `RInside/mpi/rinside_mpi_sample2.cpp`

```
1 // MPI C++ API version of file contributed by Jianping Hua
2
3 #include <mpi.h> // mpi header
4 #include <RInside.h> // for the embedded R via RInside
5
6 int main(int argc, char *argv[]) {
7
8     MPI::Init(argc, argv); // mpi initialization
9     int myrank = MPI::COMM_WORLD.Get_rank(); // obtain current node rank
10    int nodesize = MPI::COMM_WORLD.Get_size(); // obtain total nodes running.
11
12    RInside R(argc, argv); // create an embedded R instance
13
14    std::stringstream txt;
15    txt << "Hello from node " << myrank // node information
16        << " of " << nodesize << " nodes!" << std::endl;
17    R.assign( txt.str(), "txt"); // assign string to R variable 'txt'
18
19    std::string evalstr = "cat(txt)"; // show node information
20    R.parseEvalQ(evalstr); // eval the string, ign. any returns
21
22    MPI::Finalize(); // mpi finalization
23
24    exit(0);
25 }
```



# RInside workflow

- C++ programs compute, gather or aggregate raw data.
- Data is saved and analysed before a new 'run' is launched.
- With `RInside` we now skip a step:
  - collect data in a vector or matrix
  - pass data to `R` — easy thanks to `Rcpp` wrappers
  - pass one or more short 'scripts' as strings to `R` to evaluate
  - pass data back to C++ programm — easy thanks to `Rcpp` converters
  - resume main execution based on new results
- A number of simple examples ship with `RInside`
  - *nine* different examples in `examples/standard`
  - *four* different examples in `examples/mpi`



# Outline

- 1 Extending R
  - Why ?
  - The standard API
  - Inline
- 2 Rcpp
  - Overview
  - New API
  - Examples
- 3 Rcpp Usage Examples
  - RInside
  - **RProtoBuf**
  - Others
- 4 Summary
  - Key points
  - Resources



# About Google ProtoBuf

Quoting from the page at Google Code:

*Protocol buffers are a flexible, efficient, automated mechanism for serializing structured data—think XML, but smaller, faster, and simpler.*

*You define how you want your data to be structured once, then you can use special generated source code to easily write and read your structured data to and from a variety of data streams and using a variety of languages.*

*You can even update your data structure without breaking deployed programs that are compiled against the "old" format.*

Google provides native bindings for C++, Java and Python.



# Google ProtoBuf

```
1 R> library( RProtoBuf )           ## load the package
2 R> readProtoFiles( "addressbook.proto" ) ## acquire protobuf information
3 R> bob <- new( tutorial.Person,    ## create new object
4 +   email = "bob@example.com",
5 +   name = "Bob",
6 +   id = 123 )
7 R> writeLines( bob$toString() )    ## serialize to stdout
8 name: "Bob"
9 id: 123
10 email: "bob@example.com"
11
12 R> bob$email                       ## access and/or override
13 [1] "bob@example.com"
14 R> bob$id <- 5
15 R> bob$id
16 [1] 5
17
18 R> serialize( bob, "person.pb" )   ## serialize to compact binary format
```

Under the hood, `Rcpp` is used and works very well in conjunction with the rich C++ API provided by Google.



# Outline

- 1 Extending R
  - Why ?
  - The standard API
  - Inline
- 2 Rcpp
  - Overview
  - New API
  - Examples
- 3 Rcpp Usage Examples
  - RInside
  - RProtoBuf
  - **Others**
- 4 Summary
  - Key points
  - Resources



# Users of Rcpp

- RInside uses Rcpp for object transfer and more
- RcppArmadillo and RcppGSL (which contain fastLm())
- RcppExamples is a 'this is how you can do it' stanza
- RProtoBuf is what got Romain and me here, it may get rewritten to take more advantage of Rcpp
- RQuantLib is where Rcpp originally started
- highlight is Romain's first re-use of Rcpp
- mvabund, sdcTable, bifactorial, minqa, pcaMethods (BioC), phylobase are truly external users which are all on CRAN
- upcoming: possibly lme4a
- *Your package here next?*



# Outline

- 1 Extending R
  - Why ?
  - The standard API
  - Inline
- 2 Rcpp
  - Overview
  - New API
  - Examples
- 3 Rcpp Usage Examples
  - RInside
  - RProtoBuf
  - Others
- 4 Summary
  - **Key points**
  - Resources



# Wrapping up

This tutorial has tried to show you that

- While the deck way be stacked against you (when adding C/C++ to R), you can still pick where to play
- R can be extended in many ways; we focus on something that allows us write extensions
  - that are efficient: we want speed and features
  - that correspond to the R object model
  - that also allow us to embed R inside C++
- And all this while retaining 'high-level' STL-alike semantics, templates and other goodies in C++
- Using C++ abstractions wisely can keep the code both clean and readable – yet very efficient





# Outline

- 1 Extending R
  - Why ?
  - The standard API
  - Inline
- 2 Rcpp
  - Overview
  - New API
  - Examples
- 3 Rcpp Usage Examples
  - RInside
  - RProtoBuf
  - Others
- 4 **Summary**
  - Key points
  - **Resources**



# Some pointers

- <http://dirk.eddelbuettel.com/code/rcpp.html>
- <http://dirk.eddelbuettel.com/code/rcppTut/>
- <http://romainfrancois.blog.free.fr/index.php?category/R-package/Rcpp>
- <http://cran.r-project.org/package=Rcpp>
- <http://r-forge.r-project.org/projects/rcpp/>
- and likewise for RInside, RProtoBuf and more.



# The end

*Thank you!*

