

Writing R Extensions

Version 2.3.1 (2006-06-01)

R Development Core Team

Permission is granted to make and distribute verbatim copies of this manual provided the copyright notice and this permission notice are preserved on all copies.

Permission is granted to copy and distribute modified versions of this manual under the conditions for verbatim copying, provided that the entire resulting derived work is distributed under the terms of a permission notice identical to this one.

Permission is granted to copy and distribute translations of this manual into another language, under the above conditions for modified versions, except that this permission notice may be stated in a translation approved by the R Development Core Team.

Copyright © 1999–2005 R Development Core Team

ISBN 3-900051-11-9

Table of Contents

Acknowledgements	1
1 Creating R packages	2
1.1 Package structure	2
1.1.1 The ‘DESCRIPTION’ file	2
1.1.2 The ‘INDEX’ file	5
1.1.3 The ‘install.R’ and ‘R_PROFILE.R’ files	5
1.1.4 Package subdirectories	5
1.1.5 Package bundles	7
1.2 Configure and cleanup	8
1.2.1 Using ‘Makevars’	10
1.2.2 Configure example	11
1.2.3 Using F95 code	13
1.3 Checking and building packages	13
1.3.1 Checking packages	13
1.3.2 Building packages	15
1.3.3 Customizing checking and building	15
1.4 Writing package vignettes	16
1.5 Submitting a package to CRAN	17
1.6 Package name spaces	18
1.6.1 Specifying imports and exports	18
1.6.2 Registering S3 methods	19
1.6.3 Load hooks	19
1.6.4 An example	21
1.6.5 Summary – converting an existing package	22
1.6.6 Name spaces with formal classes and methods	22
1.7 Writing portable packages	23
1.8 Diagnostic messages	23
1.9 Localization	24
1.9.1 C-level messages	25
1.9.2 R messages	25
1.10 Package types	26
1.10.1 Frontend	26
1.10.2 Translation	26
2 Writing R documentation files	27
2.1 Rd format	27
2.1.1 Documenting functions	27
2.1.2 Documenting data sets	31
2.1.3 Documenting S4 classes and methods	31
2.1.4 Documenting packages	32
2.2 Sectioning	32
2.3 Marking text	32
2.4 Lists and tables	34
2.5 Cross-references	34
2.6 Mathematics	35
2.7 Insertions	35

2.8	Indices	36
2.9	Platform-specific documentation	36
2.10	Encoding	36
2.11	Processing Rd format	37
3	Tidying and profiling R code	38
3.1	Tidying R code	38
3.2	Profiling R code	38
3.3	Profiling compiled code	40
3.3.1	Linux	40
3.3.1.1	sprof	40
3.3.1.2	oprofile	40
3.3.2	Solaris	41
3.3.3	MacOS X	41
4	Debugging	42
4.1	Browsing	42
4.2	Debugging R code	43
4.3	Using gctorture and valgrind	47
4.3.1	Using gctorture	47
4.3.2	Using valgrind	47
4.4	Debugging compiled code	48
4.4.1	Finding entry points in dynamically loaded code	49
4.4.2	Inspecting R objects when debugging	50
5	System and foreign language interfaces	52
5.1	Operating system access	52
5.2	Interface functions <code>.C</code> and <code>.Fortran</code>	52
5.3	<code>dyn.load</code> and <code>dyn.unload</code>	54
5.4	Registering native routines	55
5.5	Creating shared objects	57
5.6	Interfacing C++ code	57
5.7	Handling R objects in C	59
5.7.1	Handling the effects of garbage collection	60
5.7.2	Allocating storage	62
5.7.3	Details of R types	62
5.7.4	Attributes	63
5.7.5	Classes	64
5.7.6	Handling lists	65
5.7.7	Finding and setting variables	65
5.7.8	Named objects and copying	66
5.8	Interface functions <code>.Call</code> and <code>.External</code>	67
5.8.1	Calling <code>.Call</code>	67
5.8.2	Calling <code>.External</code>	68
5.8.3	Missing and special values	70
5.9	Evaluating R expressions from C	71
5.9.1	Zero-finding	72
5.9.2	Calculating numerical derivatives	74
5.10	Parsing R code from C	76

6	The R API: entry points for C code	78
6.1	Memory allocation	78
6.1.1	Transient storage allocation	78
6.1.2	User-controlled memory	79
6.2	Error handling	79
6.2.1	Error handling from FORTRAN	79
6.3	Random number generation	80
6.4	Missing and IEEE special values	80
6.5	Printing	81
6.5.1	Printing from FORTRAN	81
6.6	Calling C from FORTRAN and vice versa	81
6.7	Numerical analysis subroutines	82
6.7.1	Distribution functions	82
6.7.2	Mathematical functions	83
6.7.3	Numerical Utilities	84
6.7.4	Mathematical constants	85
6.8	Optimization	86
6.9	Integration	86
6.10	Utility functions	87
6.11	Internationalization	89
6.12	Allowing interrupts	89
6.13	Platform and version information	90
6.14	Inlining C functions	90
6.15	Using these functions in your own C code	90
6.16	Organization of header files	91
7	Generic functions and methods	92
7.1	Adding new generics	93
8	Linking GUIs and other front-ends to R	94
8.1	Embedding R under Unix-alikes	94
8.1.1	Setting R callbacks	95
8.1.2	Meshing event loops	97
8.1.3	Threading issues	98
8.2	Embedding R under Windows	99
8.2.1	Using (D)COM	99
8.2.2	Calling R.dll directly	99
Appendix A	R (internal) programming miscellanea	103
A.1	.Internal and .Primitive	103
A.2	Testing R code	104
Appendix B	R coding standards	105
	Function and variable index	107
	Concept index	110

Acknowledgements

The contributions of Saikat DebRoy (who wrote the first draft of a guide to using `.Call` and `.External`) and of Adrian Trapletti (who provided information on the C++ interface) are gratefully acknowledged.

1 Creating R packages

Packages provide a mechanism for loading optional code and attached documentation as needed. The R distribution provides several packages.

In the following, we assume that you know the `library()` command, including its `lib.loc` argument, and we also assume basic knowledge of the `INSTALL` utility. Otherwise, please look at R's help pages

```
?library
?INSTALL
```

before reading on.

Once a source package is created, it must be installed by the command `R CMD INSTALL`. See [section “Add-on-packages” in *R Installation and Administration*](#), for further details.

Other types of extensions are supported as from R 2.1.0: See [Section 1.10 \[Package types\]](#), page 26.

1.1 Package structure

A package consists of a subdirectory containing a file `DESCRIPTION` and the subdirectories `R`, `data`, `demo`, `exec`, `inst`, `man`, `po`, `src`, and `tests` (some of which can be missing). The package subdirectory may also contain files `INDEX`, `install.R` (deprecated), `R_PROFILE.R` (deprecated), `NAMESPACE`, `configure`, `cleanup`, and `COPYING`. Other files such as `README`, `NEWS` or `ChangeLog` will be ignored by R, but may be useful to end-users.

The `DESCRIPTION`, `INDEX`, `install.R` and `R_PROFILE.R` files are described in the sections below. The `NAMESPACE` file is described in [Section 1.6 \[Package name spaces\]](#), page 18.

The optional files `configure` and `cleanup` are (Bourne shell) script files which are executed before and (provided that option `--clean` was given) after installation on Unix, see [Section 1.2 \[Configure and cleanup\]](#), page 8.

The optional file `COPYING` contains a copy of the license to the package, e.g. a copy of the GNU public license. Whereas you should feel free to include a licence file in your source distribution, please do not arrange to install yet another copy of the GNU `COPYING` or `COPYING.LIB` files but refer to the copies in the R distribution (e.g., in directory `share/licenses` in your own `COPYING` file).

The package subdirectory should be given the same name as the package. Because some file systems (e.g., those on Windows) are not case-sensitive, to maintain portability it is strongly recommended that case distinctions not be used to distinguish different packages. For example, if you have a package named `foo`, do not also create a package named `Foo`.

To ensure that file names are valid across file systems and supported operating system platforms, the ASCII control characters as well as the characters `"`, `*`, `:`, `/`, `<`, `>`, `?`, `\`, and `|` are not allowed in file names. In addition, files with names `con`, `prn`, `aux`, `clock$`, `nul`, `com1` to `com4`, and `lpt1` to `lpt3` after conversion to lower case and stripping possible “extensions”, are disallowed. Also, file names in the same directory must not differ only by case (see the previous paragraph). In addition, the names of `.Rd` files will be used in URLs and so must be ASCII and not contain %.

The R function `package.skeleton` can help to create the structure for a new package: see its help page for details.

1.1.1 The ‘DESCRIPTION’ file

The `DESCRIPTION` file contains basic information about the package in the following format:

```

Package: pkgname
Version: 0.5-1
Date: 2004-01-01
Title: My First Collection of Functions
Author: Joe Developer <Joe.Developer@some.domain.net>, with
       contributions from A. User <A.User@wherever.net>.
Maintainer: Joe Developer <Joe.Developer@some.domain.net>
Depends: R (>= 1.8.0), nlme
Suggests: MASS
Description: A short (one paragraph) description of what
            the package does and why it may be useful.
License: GPL version 2 or newer
URL: http://www.r-project.org, http://www.another.url

```

Continuation lines (for example, for descriptions longer than one line) start with a space or tab. The ‘Package’, ‘Version’, ‘License’, ‘Description’, ‘Title’, ‘Author’, and ‘Maintainer’ fields are mandatory, the remaining fields (‘Date’, ‘Depends’, ‘URL’, ...) are optional.

The ‘DESCRIPTION’ file should be written entirely in ASCII for maximal portability.

The ‘Package’ and ‘Version’ fields give the name and the version of the package, respectively. The name should consist of letters, numbers, and the dot character and start with a letter. The version is a sequence of at least *two* (and usually three) non-negative integers separated by single ‘.’ or ‘-’ characters. The canonical form is as shown in the example, and a version such as ‘0.01’ or ‘0.01.0’ will be handled as if it were ‘0.1-0’. (Translation packages are allowed names of the form Translation-11.)

The ‘License’ field should contain an explicit statement or a well-known abbreviation (such as ‘GPL’, ‘LGPL’, ‘BSD’, or ‘Artistic’), perhaps followed by a reference to the actual license file. It is very important that you include this information! Otherwise, it may not even be legally correct for others to distribute copies of the package.

The ‘Description’ field should give a comprehensive description of what the package does. One can use several (complete) sentences, but only one paragraph.

The ‘Title’ field should give a short description of the package. Some package listings may truncate the title to 65 characters in order to keep the overall size of the listing limited. It should be capitalized, not use any markup, not have any continuation lines, and not end in a period. Older versions of R used a separate file ‘TITLE’ for giving this information; this is now defunct, and the ‘Title’ field in ‘DESCRIPTION’ is required.

The ‘Author’ field describes who wrote the package. It is a plain text field intended for human readers, but not for automatic processing (such as extracting the email addresses of all listed contributors).

The ‘Maintainer’ field should give a *single* name with a *valid* email address in angle brackets (for sending bug reports etc.). It should not end in a period or comma.

The optional ‘Date’ field gives the release date of the current version of the package. It is strongly recommended to use the yyyy-mm-dd format conforming to the ISO standard.

The optional ‘Depends’ field gives a comma-separated list of package names which this package depends on. The package name may be optionally followed by a comparison operator (currently only ‘>=’ and ‘<=’ are supported), whitespace and a valid version number in parentheses. (List package names even if they are part of a bundle.) You can also use the special package name ‘R’ if your package depends on a certain version of R. E.g., if the package works only with R version 1.8.0 or newer, include ‘R (>= 1.8.0)’ in the ‘Depends’ field. Both **library** and the R package checking facilities use this field, hence it is an error to use improper syntax or misuse the ‘Depends’ field for comments on other software that might be needed. Other dependencies (external to the R system) should be listed in the ‘SystemRequirements’ field or a separate ‘README’ file. The R **INSTALL** facilities check if the version of R used is recent enough for the package being installed, and the list of packages which is specified will be attached (after

checking version dependencies) before the current package, both when `library` is called and when saving an image of the package’s code or preparing for lazy-loading.

The optional ‘Imports’ field lists packages whose namespaces are imported from but which do not need to be attached.

The optional ‘Suggests’ field uses the same syntax as ‘Depends’ and lists packages that are not necessarily needed. This includes packages used only in examples or vignettes (see [Section 1.4 \[Writing package vignettes\], page 16](#)), and packages loaded in the body of functions. E.g., suppose an example from package **foo** uses a dataset from package **bar**. Then it is not necessary to have **bar** for routine use of **foo**, unless one wants to execute the examples: it is nice to have **bar**, but not necessary. The general rules are

- Packages whose namespace only is needed to load the package using `library(pkgname)` must be listed in the ‘Imports’ field.
- Packages that need to be attached to successfully load the package using `library(pkgname)` must be listed in the ‘Depends’ field.
- All packages that are needed to successfully run R CMD check on the package must be listed in one of ‘Depends’ or ‘Suggests’ or ‘Imports’.

In particular, large packages providing “only” data for examples or vignettes should be listed in ‘Suggests’ rather than ‘Depends’ in order to make lean installations possible.

The optional ‘URL’ field may give a list of URLs separated by commas or whitespace, for example the homepage of the author or a page where additional material describing the software can be found. These URLs are converted to active hyperlinks on CRAN.

Base and recommended packages (i.e., packages contained in the R source distribution or available from CRAN and recommended to be included in every binary distribution of R) have a ‘Priority’ field with value ‘base’ or ‘recommended’, respectively. These priorities must not be used by “other” packages.

An optional ‘Collate’ field (or OS-specific variants ‘Collate.OStype’, such as e.g. ‘Collate.windows’) can be used for controlling the collation order for the R code files in a package when these are concatenated into a single file upon installation from source. The default is to try collating according to the ‘C’ locale. If present, the collate specification must list *all* R code files in the package (taking possible OS-specific subdirectories into account, see [Section 1.1.4 \[Package subdirectories\], page 5](#)) as a whitespace separated list of file paths relative to the ‘R’ subdirectory. Paths containing white space or quotes need to be quoted. Applicable OS-specific collate specifications take precedence.

The optional ‘LazyLoad’ and ‘LazyData’ fields control whether the R objects and the datasets (respectively) use lazy-loading: set the field’s value to ‘yes’ or ‘true’ for lazy-loading and ‘no’ or ‘false’ for no lazy-loading. (Capitalized values are also accepted.) Note that these values override the command-line options of R CMD INSTALL, ‘--[no-]lazy’ and ‘--[no-]lazy-data’.

The optional ‘SaveImage’ field controls whether the R objects are stored in a saved image see [Section 1.1.3 \[The install.R and R_PROFILE.R files\], page 5](#). (It takes the same values as the field ‘LazyLoad’.)

If the package you are writing uses the **methods** package, specify (preferably) ‘LazyLoad: yes’ or ‘SaveImage: yes’.

The optional ‘ZipData’ field controls whether the automatic Windows build will zip up the data directory or no: set this to ‘no’ if your package will not work with a zipped data directory.

If the ‘DESCRIPTION’ file is not entirely in ASCII it should contain an ‘Encoding’ field specifying an encoding. This is currently used as the encoding of the ‘DESCRIPTION’ file itself, and may in the future be taken as the encoding for other documentation in the package. Only encoding names `latin1`, `latin2` and `UTF-8` are known to be portable.

The optional ‘Type’ field specifies the type of the package: see [Section 1.10 \[Package types\]](#), page 26.

Note: There should be no ‘Built’ or ‘Packaged’ fields, as these are added by the package management tools.

1.1.2 The ‘INDEX’ file

The optional file ‘INDEX’ contains a line for each sufficiently interesting object in the package, giving its name and a description (functions such as print methods not usually called explicitly might not be included). Normally this file is missing, and the corresponding information is automatically generated from the documentation sources (using `Rdindex()` from package **tools**) when installing from source and when using the package builder (see [Section 1.3 \[Checking and building packages\]](#), page 13).

Rather than editing this file, it is preferable to put customized information about the package into an overview man page (see [Section 2.1.4 \[Documenting packages\]](#), page 32) and/or a vignette (see [Section 1.4 \[Writing package vignettes\]](#), page 16).

1.1.3 The ‘install.R’ and ‘R_PROFILE.R’ files

NB: This mechanism is now deprecated and will be removed at R 2.4.0.

The optional file ‘install.R’ serves two purposes. First, its presence tells the **INSTALL** utility to create a binary image of the package workspace. A binary image is created by executing the code in the ‘R’ subdirectory and saving the resulting objects. When the package is attached, the code is not executed again but loaded from the saved image. It is preferable to use the ‘SaveImage’ field in the ‘DESCRIPTION’ file, which takes precedence over the presence of an ‘install.R’ file.

The second purpose for ‘install.R’ is to hold code that needs to be executed each time the package is attached, before the image is loaded. Very few packages have a need for such code.

The optional file ‘R_PROFILE.R’ is executed before the code in the ‘R’ subdirectory when saving an image or preparing for lazy-loading, and should be used to set up an environment needed only to evaluate the code (which is run with the ‘--vanilla’ command-line flag). Packages should not need such code, as the packages in the ‘Depends’ field should suffice.

Both ‘install.R’ and ‘R_PROFILE.R’ should be viewed as experimental; the mechanism to execute code before attaching or installing the package may change in the near future. With the other facilities available as from R 2.0.0 they should be removed if possible.

1.1.4 Package subdirectories

The ‘R’ subdirectory contains R code files, only. The code files to be installed must start with a (lower or upper case) letter or digit¹ and have one of the extensions ‘.R’, ‘.S’, ‘.q’, ‘.r’, or ‘.s’. We recommend using ‘.R’, as this extension seems to be not used by any other software. It should be possible to read in the files using `source()`, so R objects must be created by assignments. Note that there need be no connection between the name of the file and the R objects created by it. The R code files should only create R objects and not call functions with side effects such as `require` and `options`.

Two exceptions are allowed: if the ‘R’ subdirectory contains a file ‘sysdata.rda’ (a saved image of R objects) this will be lazy-loaded into the namespace/package environment – this is intended for system datasets that are not intended to be user-accessible via `data`. Also, files ending in ‘.in’ will be allowed in the ‘R’ directory to allow a ‘configure’ script to generate suitable files,

¹ This is enforced as from R 2.3.0.

Only ASCII characters (and the control characters tab, formfeed, LF and CR) should be used in code files. Other characters are accepted in comments, but then the comments may not be readable in e.g. a UTF-8 locale. Non-ASCII characters in object names will normally² fail when the package is installed. Any byte will be allowed³ in a quoted character string (but `\uxxxx` escapes should not be used), but non-ASCII character strings may not be usable in some locales and may display incorrectly in others.

Various R functions in a package can be used to initialize and clean up. For packages without a name space, these are `.First.lib` and `.Last.lib`. (For packages with a name space, See [Section 1.6.3 \[Load hooks\]](#), page 19.) It is conventional to define these functions in a file called `'zzz.R'`. If `.First.lib` is defined in a package, it is called with arguments `libname` and `pkgname` after the package is loaded and attached. (If a package is installed with version information, the package name includes the version information, e.g. `'ash_1.0.9'`.) A common use is to call `library.dynam()` inside `.First.lib()` to load compiled code: another use is to call those functions with side effects. If `.Last.lib` exists in a package it is called (with argument the full path to the installed package) just before the package is detached. It is uncommon to detach packages and rare to have a `.Last.lib` function: one use is to call `library.dynam.unload` to unload compiled code.

The `'man'` subdirectory should contain (only) documentation files for the objects in the package in *R documentation* (Rd) format. The documentation files to be installed must start with a (lower or upper case ASCII) letter or digit⁴ and have the extension `'Rd'` (the default) or `'rd'`. Further, the names must be valid in `'file://'` URLs, which means⁵ they must be entirely ASCII and not contain `'%'`. See [Chapter 2 \[Writing R documentation files\]](#), page 27, for more information. Note that all user-level objects in a package should be documented; if a package `pkg` contains user-level objects which are for “internal” use only, it should provide a file `'pkg-internal.Rd'` which documents all such objects, and clearly states that these are not meant to be called by the user. See e.g. the sources for package **grid** in the R distribution for an example.

The `'R'` and `'man'` subdirectories may contain OS-specific subdirectories named `'unix'` or `'windows'`.

The C, C++, or FORTRAN⁶ source files for the compiled code are in `'src'`, plus optionally file `'Makevars'` or `'Makefile'`. When a package is installed using `R CMD INSTALL`, Make is used to control compilation and linking into a shared object for loading into R. There are default variables and rules for this (determined when R is configured and recorded in `'R_HOME/etc/Makeconf'`). These rules can be tweaked by setting macros in a file `'src/Makevars'` (see [Section 1.2.1 \[Using Makevars\]](#), page 10). Note that this mechanism should be general enough to eliminate the need for a package-specific `'Makefile'`. If such a file is to be distributed, considerable care is needed to make it general enough to work on all R platforms. In addition, it should have a target `'clean'` which removes all files generated by Make. If necessary, platform-specific files can be used, for example `'Makevars.win'` or `'Makefile.win'` on Windows take precedence over `'Makevars'` or `'Makefile'`.

The `'data'` subdirectory is for additional data files the package makes available for loading using `data()`. Currently, data files can have one of three types as indicated by their extension:

² This is true for OSes which implement the `'C'` locale, unless neither lazy-loading nor saving an image are used, in which case it would fail if loaded in a `'C'` locale. (Windows' idea of the `'C'` locale uses the WinAnsi charset.)

³ It is good practice to encode them as octal or hex escape sequences.

⁴ This is enforced as from R 2.3.0.

⁵ More precisely, they can contain the English alphanumeric characters and the symbols `'$ - _ . + ! ' () , ; = &'`.

⁶ Note that Ratfor is not supported. If you have Ratfor source code, you need to convert it to FORTRAN. Only FORTRAN-77 is supported on all platforms, but some also support FORTRAN-95. If you want to ship Ratfor source files, please do so in a subdirectory of `'src'` and not in the main subdirectory.

plain R code (`.R` or `.r`), tables (`.tab`, `.txt`, or `.csv`), or `save()` images (`.RData` or `.rda`). (All ports of R use the same binary (XDR) format and can read compressed images. Use images saved with `save(, compress = TRUE)`, the default as from R 2.3.0. to save space.) Note that R code should be “self-sufficient” and not make use of extra functionality provided by the package, so that the data file can also be used without having to load the package. It is no longer necessary to provide a `00Index` file in the `data` directory—the corresponding information is generated automatically from the documentation sources when installing from source, or when using the package builder (see [Section 1.3 \[Checking and building packages\]](#), [page 13](#)). If your data files are enormous you can speed up installation by providing a file `datalist` in the `data` subdirectory. This should have one line per topic that `data()` will find, in the format `foo` if `data(foo)` provides `foo`, or `foo: bar bah` if `data(foo)` provides `bar` and `bah`.

The `demo` subdirectory is for R scripts (for running via `demo()`) that demonstrate some of the functionality of the package. Demos may be interactive and are not checked automatically, so if testing is desired use code in the `tests` directory. The script files must start with a (lower or upper case) letter⁷ and have one of the extensions `.R` or `.r`. If present, the `demo` subdirectory should also have a `00Index` file with one line for each demo, giving its name and a description separated by white space. (Note that it is not possible to generate this index file automatically.)

The contents of the `inst` subdirectory will be copied recursively to the installation directory. Subdirectories of `inst` should not interfere with those used by R (currently, `R`, `data`, `demo`, `exec`, `libs`, `man`, `help`, `html`, `latex`, `R-ex`, `chtml`, and `Meta`). The copying of the `inst` happens after `src` is built so its `Makefile` can create files to be installed. Note that with the exceptions of `INDEX` and `COPYING`, information files at the top level of the package will *not* be installed and so not be known to users of Windows and MacOS X compiled packages (and not seen by those who use `R CMD INSTALL` or `install.packages` on the tarball). So any information files you wish an end user to see should be included in `inst`. One thing you might like to add to `inst` is a `CITATION` file for use by the `citation` function.

Subdirectory `tests` is for additional package-specific test code, similar to the specific tests that come with the R distribution. Test code can either be provided directly in a `.R` file, or via a `.Rin` file containing code which in turn creates the corresponding `.R` file (e.g., by collecting all function objects in the package and then calling them with the strangest arguments). The results of running a `.R` file are written to a `.Rout` file. If there is a corresponding `.Rout.save` file, these two are compared, with differences being reported but not causing an error. The whose `tests` is copied to the check area, and the tests are run with the copy as the working directory and with `R_LIBS` set to ensure that the copy of the package installed during testing will be found by `library(pkg_name)`.

Subdirectory `exec` could contain additional executables the package needs, typically scripts for interpreters such as the shell, Perl, or Tcl. This mechanism is currently used only by a very few packages, and still experimental.

Subdirectory `po` is used for files related to *localization*: see [Section 1.9 \[Localization\]](#), [page 24](#).

1.1.5 Package bundles

Sometimes it is convenient to distribute several packages as a *bundle*. (An example is **VR** which contains four packages.) The installation procedures on both Unix and Windows can handle package bundles.

The `DESCRIPTION` file of a bundle has a `Bundle` field and no `Package` field, as in

⁷ This is enforced as from R 2.3.0.

```

Bundle: VR
Priority: recommended
Contains: MASS class nnet spatial
Version: 7.2-12
Date: 2005-01-31
Depends: R (>= 2.0.0), graphics, stats
Suggests: lattice, nlme, survival
Author: S original by Venables & Ripley.
       R port by Brian Ripley <ripley@stats.ox.ac.uk>, following earlier
       work by Kurt Hornik and Albrecht Gebhardt.
Maintainer: Brian Ripley <ripley@stats.ox.ac.uk>
BundleDescription: Functions and datasets to support Venables and
                   Ripley, 'Modern Applied Statistics with S' (4th edition).
License: GPL (version 2 or later) See file LICENCE.
URL: http://www.stats.ox.ac.uk/pub/MASS4/

```

The ‘Contains’ field lists the packages (space separated), which should be contained in separate subdirectories with the names given. During building and installation, packages will be installed in the order specified. Be sure to order this list so that dependencies are met appropriately.

The packages contained in a bundle are standard packages in all respects except that the ‘DESCRIPTION’ file is replaced by a ‘DESCRIPTION.in’ file which just contains fields additional to the ‘DESCRIPTION’ file of the bundle, for example

```

Package: spatial
Description: Functions for kriging and point pattern analysis.
Title: Functions for Kriging and Point Pattern Analysis

```

Any files in the package bundle except the ‘DESCRIPTION’ file and the named packages will be ignored.

The ‘Depends’ field in the bundle’s ‘DESCRIPTION’ file should list the dependencies of all the constituent packages (and similarly for ‘Imports’ and ‘Suggests’), and then ‘DESCRIPTION.in’ files should not contain these fields.

1.2 Configure and cleanup

Note that most of this section is Unix-specific: see the comments later on about the Windows port of R.

If your package needs some system-dependent configuration before installation you can include a (Bourne shell) script ‘configure’ in your package which (if present) is executed by R CMD INSTALL before any other action is performed. This can be a script created by the Autoconf mechanism, but may also be a script written by yourself. Use this to detect if any nonstandard libraries are present such that corresponding code in the package can be disabled at install time rather than giving error messages when the package is compiled or used. To summarize, the full power of Autoconf is available for your extension package (including variable substitution, searching for libraries, etc.).

The (Bourne shell) script ‘cleanup’ is executed as last thing by R CMD INSTALL if present and option ‘--clean’ was given, and by R CMD build when preparing the package for building from its source. It can be used to clean up the package source tree. In particular, it should remove all files created by configure.

As an example consider we want to use functionality provided by a (C or FORTRAN) library `foo`. Using Autoconf, we can create a configure script which checks for the library, sets variable `HAVE_FOO` to `TRUE` if it was found and with `FALSE` otherwise, and then substitutes this value into output files (by replacing instances of ‘@HAVE_FOO@’ in input files with the value of `HAVE_FOO`). For example, if a function named `bar` is to be made available by linking against library `foo` (i.e., using ‘-lfoo’), one could use


```
AC_CHECK_LIB(foo, fun, [HAVE_FOO=TRUE], [HAVE_FOO=FALSE])
AC_SUBST(HAVE_FOO)
.....
AC_CONFIG_FILES([foo.R])
AC_OUTPUT
```

in `configure.ac` (assuming Autoconf 2.50 or better).

The definition of the respective R function in `foo.R.in` could be

```
foo <- function(x) {
  if(!@HAVE_FOO@)
    stop("Sorry, library 'foo' is not available")
  ...
}
```

From this file `configure` creates the actual R source file `foo.R` looking like

```
foo <- function(x) {
  if(!FALSE)
    stop("Sorry, library 'foo' is not available")
  ...
}
```

if library `foo` was not found (with the desired functionality). In this case, the above R code effectively disables the function.

One could also use different file fragments for available and missing functionality, respectively.

You will very likely need to ensure that the same C compiler and compiler flags are used in the `configure` tests as when compiling R or your package. Under Unix, you can achieve this by including the following fragment early in `configure.ac`

```
: ${R_HOME}='R RHOME'
if test -z "${R_HOME}"; then
  echo "could not determine R_HOME"
  exit 1
fi
CC="${R_HOME}/bin/R" CMD config CC
CFLAGS="${R_HOME}/bin/R" CMD config CFLAGS
```

(using `${R_HOME}/bin/R` rather than just `R` is necessary in order to use the ‘right’ version of R when running the script as part of `R CMD INSTALL`.)

Note that earlier versions of this document recommended obtaining the configure information by direct extraction (using `grep` and `sed`) from `R_HOME/etc/Makeconf`, which only works for variables recorded there as literals. R 1.5.0 has added `R CMD config` for getting the value of the basic configuration variables, or the header and library flags necessary for linking against R, see `R CMD config --help` for details.

To check for an external BLAS library using the `ACX_BLAS` macro from the official Autoconf Macro Archive, one can simply do

```
F77="${R_HOME}/bin/R" CMD config F77
AC_PROG_F77
FLIBS="${R_HOME}/bin/R" CMD config FLIBS
ACX_BLAS([], AC_MSG_ERROR([could not find your BLAS library], 1))
```

Note that `FLIBS` as determined by R must be used to ensure that FORTRAN 77 code works on all R platforms. Calls to the Autoconf macro `AC_F77_LIBRARY_LDFLAGS`, which would overwrite `FLIBS`, must not be used (and hence e.g. removed from `ACX_BLAS`). (Recent versions of Autoconf in fact allow an already set `FLIBS` to override the test for the FORTRAN linker flags. Also, recent versions of R can detect external BLAS and LAPACK libraries.)

You should bear in mind that the `configure` script may well not work on Windows systems (this seems normally to be the case for those generated by Autoconf, although simple shell scripts

do work). If your package is to be made publicly available, please give enough information for a user on a non-Unix platform to configure it manually, or provide a ‘`configure.win`’ script to be used on that platform.

In some rare circumstances, the configuration and cleanup scripts need to know the location into which the package is being installed. An example of this is a package that uses C code and creates two shared object/DLLs. Usually, the object that is dynamically loaded by R is linked against the second, dependent, object. On some systems, we can add the location of this dependent object to the object that is dynamically loaded by R. This means that each user does not have to set the value of the `LD_LIBRARY_PATH` (or equivalent) environment variable, but that the secondary object is automatically resolved. Another example is when a package installs support files that are required at run time, and their location is substituted into an R data structure at installation time. (This happens with the Java Archive files in the **SJava** package.) The names of the top-level library directory (i.e., specifiable via the ‘`-l`’ argument) and the directory of the package itself are made available to the installation scripts via the two shell/environment variables `R_LIBRARY_DIR` and `R_PACKAGE_DIR`. Additionally, the name of the package (e.g., ‘`survival`’ or ‘`MASS`’) being installed is available from the shell variable `R_PACKAGE_NAME`.

1.2.1 Using ‘`Makevars`’

Sometimes writing your own ‘`configure`’ script can be avoided by supplying a file ‘`Makevars`’: also one of the commonest uses of a ‘`configure`’ script is to make ‘`Makevars`’ from ‘`Makevars.in`’.

The most common use of a ‘`Makevars`’ file is to set additional preprocessor (for example include paths) flags via `PKG_CPPFLAGS`, and additional compiler flags by setting `PKG_CFLAGS`, `PKG_CXXFLAGS` and `PKG_FFLAGS`, for C, C++, or FORTRAN respectively (see [Section 5.5 \[Creating shared objects\]](#), page 57).

Also, ‘`Makevars`’ can be used to set flags for the linker, for example ‘`-L`’ and ‘`-l`’ options.

When writing a ‘`Makevars`’ file for a package you intend to distribute, take care to ensure that it is not specific to your compiler: flags such as ‘`-O2 -Wall -pedantic`’ are all specific to GCC.

There are some macros which are built whilst configuring the building of R itself, are stored on Unix-alikes in ‘`R_HOME/etc/Makeconf`’ and can be used in ‘`Makevars`’. These include

FLIBS A macro containing the set of libraries need to link FORTRAN code. This may need to be included in `PKG_LIBS`.

BLAS_LIBS A macro containing the BLAS libraries used when building R. This may need to be included in `PKG_LIBS`. Beware that if it is empty then the R executable will contain all the double-precision and double-complex BLAS routines, but no single-precision or complex routines. If `BLAS_LIBS` is included, then `FLIBS` also needs to be, as most BLAS libraries are written in FORTRAN.

LAPACK_LIBS A macro containing the LAPACK libraries (and paths where appropriate) used when building R. This may need to be included in `PKG_LIBS`. This may point to a dynamic library `libRlapack` which contains all the double-precision LAPACK routines as well as those double-complex LAPACK and BLAS routines needed to build R, or it may point to an external LAPACK library, or may be empty if an external BLAS library also contains LAPACK.

[There is no guarantee that the LAPACK library will provide more than all the double-precision and double-complex routines, and some do not provide all the auxiliary routines.]

The macros `BLAS_LIBS` and `FLIBS` should always be included *after* `LAPACK_LIBS`.

`SAFE_FFLAGS`

A macro containing flags which are needed to circumvent over-optimization of Fortran code: it is typically `'-g -O2 -ffloat-store'` on 'ix86' platforms using `g77` or `gfortran`.

Note that `'Makevars'` should not normally contain targets, as it is included before the default `'Makefile'` and `make` is called without an explicit target. To circumvent that, use a suitable phony target before any actual targets: for example **fastICA** has

```
SLAMC_FFLAGS=$(R_XTRA_FFLAGS) $(FPICFLAGS) $(SHLIB_FFLAGS) $(SAFE_FFLAGS)

all: $(SHLIB)

slamc.o: slamc.f
    $(F77) $(SLAMC_FFLAGS) -c -o slamc.o slamc.f
```

to ensure that the LAPACK routines find some constants without infinite looping.

1.2.2 Configure example

It may be helpful to give an extended example of using a `'configure'` script to create a `'src/Makevars'` file: this is based on that in the **RODBC** package.

The `'configure.ac'` file follows: `'configure'` is created from this by running `autoconf` in the top-level package directory (containing `'configure.ac'`).


```

AC_INIT([RODBC], 1.1.5) dnl package name, version

dnl Select an optional include path, from a configure option
dnl or from an environment variable.
AC_ARG_WITH([odbc-include],
             AC_HELP_STRING([--with-odbc-include=INCLUDE_PATH],
                             [the location of ODBC header files]),
             [odbc_include_path=$withval])
if test [ -n "$odbc_include_path" ] ; then
    AC_SUBST([CPPFLAGS],["-I${odbc_include_path} ${CPPFLAGS}"])
else
    if test [ -n "${ODBC_INCLUDE}" ] ; then
        AC_SUBST([CPPFLAGS],["-I${ODBC_INCLUDE} ${CPPFLAGS}"])
    fi
fi

dnl ditto for a library path
AC_ARG_WITH([odbc-lib],
             AC_HELP_STRING([--with-odbc-lib=LIB_PATH],
                             [the location of ODBC libraries]),
             [odbc_lib_path=$withval])
if test [ -n "$odbc_lib_path" ] ; then
    AC_SUBST([LIBS],["-L${odbc_lib_path} ${LIBS}"])
else
    if test [ -n "${ODBC_LIBS}" ] ; then
        AC_SUBST([LIBS],["-L${ODBC_LIBS} ${LIBS}"])
    fi
fi

dnl Another user-specifiable option
AC_ARG_WITH([odbc-manager],
             AC_HELP_STRING([--with-odbc-manager=MGR],
                             [specify the ODBC manager, e.g. odbc or iodbc]),
             [odbc_mgr=$withval])

dnl Now find the compiler and compiler flags to use
: ${R_HOME='R RHOMES'}
if test -z "${R_HOME}"; then
    echo "could not determine R_HOME"
    exit 1
fi
CC="${R_HOME}/bin/R" CMD config CC
CFLAGS="${R_HOME}/bin/R" CMD config CFLAGS

dnl Check the headers can be found
AC_CHECK_HEADERS(sql.h sql.h)
if test "${ac_cv_header_sql_h}" = no ||
    test "${ac_cv_header_sql_h}" = no; then
    AC_MSG_ERROR("ODBC headers sql.h and sql.h not found")
fi

dnl search for a library containing an ODBC function
if test [ -n "${odbc_mgr}" ] ; then
    AC_SEARCH_LIBS(SQLTables, ${odbc_mgr}, ,
                  AC_MSG_ERROR("ODBC driver manager ${odbc_mgr} not found"))
else
    AC_SEARCH_LIBS(SQLTables, odbc iodb, ,
                  AC_MSG_ERROR("no ODBC driver manager found"))
fi

dnl substitute CPPFLAGS and LIBS
AC_SUBST(CPPFLAGS)
AC_SUBST(LIBS)
dnl and do substitution in the src/Makevars.in
AC_OUTPUT(src/Makevars)

```

where ‘src/Makevars.in’ would be simply

```
PKG_CPPFLAGS = @CPPFLAGS@
PKG_LIBS = @LIBS@
```

A user can then be advised to specify the location of the ODBC driver manager files by options like (lines broken for easier reading)

```
R CMD INSTALL
--configure-args='--with-odbc-include=/opt/local/include
--with-odbc-lib=/opt/local/lib --with-odbc-manager=iodbc'
RODBC
```

or by setting the environment variables ODBC_INCLUDE and ODBC_LIBS.

1.2.3 Using F95 code

R currently does not distinguish between FORTRAN 77 and FORTRAN 90/95 code, and assumes all FORTRAN comes in source files with extension `.f`. Commercial Unix systems typically use a F95 compiler, but only since the release of `gcc 4.0.0` in April 2005 have Linux and other non-commercial OSes had much support for F95. The compiler used for R on Windows is a F77 compiler.

This means that portable packages need to be written in correct FORTRAN 77, which will also be valid FORTRAN 95. See <http://developer.r-project.org/Portability.html> for reference resources. In particular, *free source form* F95 code is not portable.

On some systems an alternative F95 compiler is available: from the `gcc` family this might be `gfortran` or `g95`. Configuring R will try to find a compiler which (from its name) appears to be a FORTRAN 90/95 compiler, and set it in macro ‘FC’. Note that it does not check that such a compiler is fully (or even partially) compliant with FORTRAN 90/95. Packages making use of FORTRAN 90/95 features should use file extension ‘`.f90`’ or ‘`.f95`’ for the source files: the variable `PKG_FCFLAGS` specifies any special flags to be used. There is no guarantee that compiled FORTRAN 90/95 code can be mixed with any other type of code, nor that a build of R will have support for such packages.

There is a MinGW build of `gfortran` available from <http://gcc.gnu.org/wiki/GFortranBinaries> and a MinGW build⁸ of `g95` from <http://www.g95.org>. Set F95 in `MkRules` to point to the installed compiler. Then `R CMD SHLIB` and `R CMD INSTALL` work for packages containing FORTRAN 90/95 source code.

1.3 Checking and building packages

Before using these tools, please check that your package can be installed and loaded. `R CMD check` will *inter alia* do this, but you will get more informative error messages doing the checks directly.

1.3.1 Checking packages

Using `R CMD check`, the R package checker, one can test whether *source* R packages work correctly. It can be run on one or more directories, or gzipped package `tar` archives⁹ with extension ‘`.tar.gz`’ or ‘`.tgz`’. This runs a series of checks.

1. The package is installed. This will warn about missing cross-references and duplicate aliases in help files.
2. The file names are checked to be valid across file systems and supported operating system platforms.

⁸ Remember to set `LIBRARY_PATH` to point to your MinGW ‘`lib`’ directory

⁹ This may require GNU `tar`: the command used can be set with environment variable `TAR`.

3. The files and directories are checked for sufficient permissions (Unix only).
4. The ‘DESCRIPTION’ file is checked for completeness, and some of its entries for correctness. Unless installation tests are skipped, checking is aborted if the package dependencies cannot be resolved at run time. One check is that the package name is not that of a standard package, nor of the defunct standard packages (`ctest`, `eda`, `lqs`, `mle`, `modreg`, `mva`, `nls`, `stepfun` and `ts`) which are handled specially by `library`.
5. Available index information (in particular, for demos and vignettes) is checked for completeness.
6. The package subdirectories are checked for suitable file names and for not being empty. The checks on file names are controlled by the option ‘`--check-subdirs=value`’. This defaults to ‘`default`’, which runs the checks only if checking a tarball: the default can be overridden by specifying the value as ‘`yes`’ or ‘`no`’. Further, the check on the ‘`src`’ directory is only run if the package/bundle does not contain a ‘`configure`’ script (which corresponds to the value ‘`yes-maybe`’) and there is no ‘`src/Makefile`’ or ‘`src/Makefile.in`’.
To allow a ‘`configure`’ script to generate suitable files, files ending in ‘`.in`’ will be allowed in the ‘`R`’ directory.
7. The R files are checked for syntax errors.
8. The R files are checked for correct calls to `library.dynam` (with no extension). In addition, it is checked whether methods have all arguments of the corresponding generic, and whether the final argument of replacement functions is called ‘`value`’. All foreign function calls (`.C`, `.Fortran`, `.Call` and `.External` calls) are tested to see if they have a `PACKAGE` argument, and if not, whether the appropriate DLL might be deduced from the namespace of the package. Any other calls are reported. (The check is generous, and users may want to supplement this by examining the output of `tools::checkFF("mypkg", verbose=TRUE)`, especially if the intention were to always use a `PACKAGE` argument)
9. The Rd files are checked for correct syntax and meta data, including the presence of the mandatory (`\name`, `\alias`, `\title`, `\description` and `\keyword`) fields. The Rd name and title are checked for being non-empty, and the keywords found are compared to the standard ones.
10. A check is made for missing documentation entries, such as undocumented user-level objects in the package
11. Documentation for functions, data sets, and S4 classes is checked for consistency with the corresponding code.
12. It is checked whether all function arguments given in `\usage` sections of Rd files are documented in the corresponding `\arguments` section.
13. C, C++ and FORTRAN source and header files are tested for portable (LF-only) line endings.
14. The examples provided by the package’s documentation are run. (see [Chapter 2 \[Writing R documentation files\]](#), [page 27](#), for information on using `\examples` to create executable example code.)
Of course, released packages should be able to run at least their own examples. Each example is run in a ‘clean’ environment (so earlier examples cannot be assumed to have been run), and with the variables `T` and `F` redefined to generate an error unless they are set in the example: See [section “Logical vectors” in *An Introduction to R*](#).
15. If the package sources contain a ‘`tests`’ directory then the tests specified in that directory are run. (Typically they will consist of a set of ‘`.R`’ source files and target output files ‘`.Rout.save`’.)
16. The code in package vignettes (see [Section 1.4 \[Writing package vignettes\]](#), [page 16](#)) is executed.

17. If a working `latex` program is available, the `.dvi` version of the package's manual is created (to check that the `Rd` files can be converted successfully).

Use `R CMD check --help` to obtain more information about the usage of the R package checker. A subset of the checking steps can be selected by adding flags.

1.3.2 Building packages

Using `R CMD build`, the R package builder, one can build R packages from their sources (for example, for subsequent release).

Prior to actually building the package in the common gzipped tar file format, a few diagnostic checks and cleanups are performed. In particular, it is tested whether object indices exist and can be assumed to be up-to-date, and C, C++ and FORTRAN source files are tested and converted to LF line-endings if necessary.

Run-time checks whether the package works correctly should be performed using `R CMD check` prior to invoking the build procedure.

To exclude files from being put into the package, one can specify a list of exclude patterns in file `.Rbuildignore` in the top-level source directory. These patterns should be Perl regexps, one per line, to be matched against the file names relative to the top-level source directory. In addition, directories called `CVS` or `.svn` or `.arch-ids` and files `GNUMakefile` or with base names starting with `.#`, or starting and ending with `#`, or ending in `~`, `.bak` or `.swp`, are excluded by default. In addition, those files in the `R`, `demo` and `man` directories which are flagged by `R CMD check` as having invalid names will be excluded.

Use `R CMD build --help` to obtain more information about the usage of the R package builder.

Unless `R CMD build` is invoked with the `--no-vignettes` option, it will attempt to rebuild the vignettes (see [Section 1.4 \[Writing package vignettes\]](#), page 16) in the package. To do so it installs the current package/bundle into a temporary library tree, but any dependent packages need to be installed in an available library tree (see the Note: below).

One of the checks that `R CMD build` runs is for empty source directories. These are in most cases unintentional, in which case they should be removed and the build re-run.

It can be useful to run `R CMD check --check-subdirs=yes` on the built tarball as a final check on the contents.

`R CMD build` can also build pre-compiled version of packages for binary distributions, but `R CMD INSTALL --build` is preferred (and is considerably more flexible). In particular, Windows users are recommended to use `R CMD INSTALL --build` and install into the main library tree (the default) so that HTML links are resolved.

Note: `R CMD check` and `R CMD build` run R with `--vanilla`, so none of the user's startup files are read. If you need `R_LIBS` set (to find packages in a non-standard library) you will need to set it in the environment.

Note to Windows users: `R CMD check` and `R CMD build` work well under Windows NT4/2000/XP/2003 but may not work correctly on Windows 95/98/ME because of problems with some versions of Perl on those limited OSes. Experiences vary. To use them you will need to have installed the files for building source packages (which is the default).

1.3.3 Customizing checking and building

In addition to the available command line options, `R CMD check` also allows customization by setting (Perl) configuration variables in a configuration file, the location of which can be specified via the `--rcfile` option and defaults to `$HOME/.R/check.conf` provided that the environment variable `HOME` is set.

The following configuration variables are currently available.

`$R_check_use_install_log`

If true, record the output from installing a package as part of its check to a log file ('00install.out' by default), even when running interactively. Default: true.

`$R_check_all_non_ISO_C`

If true, do not ignore compiler (typically GCC) warnings about non ISO C code in *system* headers. Default: false.

`$R_check_weave_vignettes`

If true, weave package vignettes in the process of checking them. Default: true.

`$R_check_subdirs_nocase`

If true, check the case of directories such as 'R' and 'man'. Default: false.

`$R_check_subdirs_strict`

Initial setting for '--check-subdirs'. Default: 'default' (which checks only tarballs, and checks in the 'src' only if there is no 'configure' file).

Values '1' or a string with lower-cased version "yes" or "true" can be used for setting the variables to true; similarly, '0' or strings with lower-cased version "no" or "false" give false.

For example, a configuration file containing

```
$R_check_use_install_log = "TRUE";
$R_check_weave_vignettes = 0;
```

results in using install logs and turning off weaving.

Future versions of R will enhance this customization mechanism, and provide a similar scheme for R CMD build.

There are other internal settings that can be changed via environment variables `_R_CHECK_*_`: see the Perl source code. One that may be interesting is `_R_CHECK_USE_CODETOOLS_` to make use of the **codetools** package available from <http://www.stat.uiowa.edu/~luke/R/codetools>.

1.4 Writing package vignettes

In addition to the help files in Rd format, R packages allow the inclusion of documents in arbitrary other formats. The standard location for these is subdirectory 'inst/doc' of a source package, the contents will be copied to subdirectory 'doc' when the package is installed. Pointers from package help indices to the installed documents are automatically created. Documents in 'inst/doc' can be in arbitrary format, however we strongly recommend to provide them in PDF format, such that users on all platforms can easily read them. To ensure that they can be accessed from a browser, the file names should start with an ASCII letter and be entirely in ASCII letters or digits or minus or underscore

A special case are documents in Sweave format, which we call *package vignettes*. Sweave allows the integration of L^AT_EX documents and R code and is contained in package **utils** which is part of the base R distribution, see the **Sweave** help page for details on the document format. Package vignettes found in directory 'inst/doc' are tested by R CMD check by executing all R code chunks they contain to ensure consistency between code and documentation. Code chunks with option `eval=FALSE` are not tested. The R working directory for all vignette tests in R CMD check is the *installed* version of the 'doc' subdirectory. Make sure all files needed by the vignette (data sets, ...) are accessible by either placing them in the 'inst/doc' hierarchy of the source package, or using calls to `system.file()`.

R CMD build will automatically create PDF versions of the vignettes for distribution with the package sources. By including the PDF version in the package sources it is not necessary that the vignettes can be compiled at install time, i.e., the package author can use private L^AT_EX

extensions which are only available on his machine. Only the R code inside the vignettes is part of the checking procedure, typesetting manuals is not part of the package QC.

By default R CMD build will run Sweave on all files in Sweave format. If no ‘Makefile’ is found in directory ‘inst/doc’, then `texi2dvi --pdf` is run on all vignettes. Whenever a ‘Makefile’ is found, then R CMD build will try to run `make` after the Sweave step, such that PDF manuals can be created from arbitrary source formats (plain L^AT_EX files, ...). The ‘Makefile’ should take care of both creation of PDF files and cleaning up afterwards, i.e., delete all files that shall not appear in the final package archive. Note that the `make` step is executed independently from the presence of any files in Sweave format.

It is no longer necessary to provide a ‘00Index.dcf’ file in the ‘inst/doc’ directory—the corresponding information is generated automatically from the `\VignetteIndexEntry` statements in all Sweave files when installing from source, or when using the package builder (see Section 1.3 [Checking and building packages], page 13). The `\VignetteIndexEntry` statement is best placed in L^AT_EX comment, such that no definition of the command is necessary.

At install time an HTML index for all vignettes is automatically created from the `\VignetteIndexEntry` statements unless a file ‘index.html’ exists in directory ‘inst/doc’. This index is linked into the HTML help system for each package.

1.5 Submitting a package to CRAN

CRAN is a network of WWW sites holding the R distributions and contributed code, especially R packages. Users of R are encouraged to join in the collaborative project and to submit their own packages to CRAN.

Before submitting a package **mypkg**, do run the following steps to test it is complete and will install properly. (Unix procedures only, run from the directory containing ‘mypkg’ as a subdirectory.)

1. Run R CMD check to check that the package will install and will runs its examples, and that the documentation is complete and can be processed. If the package contains code that needs to be compiled, try to enable a reasonable amount of diagnostic messaging (“warnings”) when compiling, such as e.g. ‘`-Wall -pedantic`’ for tools from GCC, the Gnu Compiler Collection. (If R was not configured accordingly, one can achieve this e.g. via `PKG_CFLAGS` and related variables.)
2. Run R CMD build to make the release ‘.tar.gz’ file.

Please ensure that you can run through the complete procedure with only warnings that you understand and have reasons not to eliminate. In principle, packages must pass R CMD check without warnings to be admitted to the main CRAN package area.

When all the testing is done, upload the ‘.tar.gz’ file, using **anonymous** as log-in name and your e-mail address as password, to

`ftp://cran.R-project.org/incoming/`

(note: use **ftp** and not **sftp** to connect to this server) and send a message to **`cran@r-project.org`** about it. The CRAN maintainers will run these tests before putting a submission in the main archive.

Note that the fully qualified name of the ‘.tar.gz’ file must be of the form

`‘package_version[_engine[_type]]’`,

where the ‘[]’ indicates that the enclosed component is optional, *package* and *version* are the corresponding entries in file ‘DESCRIPTION’, *engine* gives the S engine the package is targeted for and defaults to ‘R’, and *type* indicated whether the file contains source or binaries for a certain platform, and defaults to ‘source’. I.e.,


```

OOP_0.1-3.tar.gz
OOP_0.1-3_R.tar.gz
OOP_0.1-3_R_source.tar.gz

```

are all equivalent and indicate an R source package, whereas

```
OOP_0.1-3_Splus6_sparc-sun-solaris.tar.gz
```

is a binary package for installation under Splus6 on the given platform.

This naming scheme has been adopted to ensure usability of code across S engines. R code and utilities operating on package ‘.tar.gz’ files can only be assumed to work provided that this naming scheme is respected. Of course, R CMD **build** automatically creates valid file names.

1.6 Package name spaces

R has a name space management system for packages. This system allows the package writer to specify which variables in the package should be *exported* to make them available to package users, and which variables should be *imported* from other packages.

The current mechanism¹⁰ for specifying a name space for a package is to place a ‘NAMESPACE’ file in the top level package directory. This file contains *name space directives* describing the imports and exports of the name space. Additional directives register any shared objects to be loaded and any S3-style methods that are provided. Note that although the file looks like R code (and often has R-style comments) it is not processed as R code. Only very simple conditional processing of **if** statements is implemented as of R 1.9.0.

Like other packages, packages with name spaces are loaded and attached to the search path by calling **library**. Only the exported variables are placed in the attached frame. Loading a package that imports variables from other packages will cause these other packages to be loaded as well (unless they have already been loaded), but they will *not* be placed on the search path by these implicit loads.

Name spaces are *sealed* once they are loaded. Sealing means that imports and exports cannot be changed and that internal variable bindings cannot be changed. Sealing allows a simpler implementation strategy for the name space mechanism. Sealing also allows code analysis and compilation tools to accurately identify the definition corresponding to a global variable reference in a function body.

1.6.1 Specifying imports and exports

Exports are specified using the **export** directive in the ‘NAMESPACE’ file. A directive of the form

```
export(f, g)
```

specifies that the variables **f** and **g** are to be exported. (Note that variable names may be quoted, and non-standard names such as **[<-.fractions** must be.)

For packages with many variables to export it may be more convenient to specify the names to export with a regular expression using **exportPattern**. The directive

```
exportPattern("[^\\\\.]")
```

exports all variables that do not start with a period.

All packages implicitly import the base name space. Variables from other packages need to be imported explicitly using the directives **import** and **importFrom**. The **import** directive imports all exported variables from the specified package(s). Thus the directives

```
import(foo, bar)
```

specifies that all exported variables in the packages **foo** and **bar** are to be imported. If only some of the variables from a package are needed, then they can be imported using **importFrom**. The directive

¹⁰ Alternative approaches are under consideration and may replace this approach in future R releases.

```
importFrom(foo, f, g)
```

specifies that the exported variables `f` and `g` of the package `foo` are to be imported.

If a package only needs one function from another package it can use a fully qualified variable reference in the code instead of a formal import. A fully qualified reference to the function `f` in package `foo` is of the form `foo::f`. This is less efficient than a formal import and also loses the advantage of recording all dependencies in the ‘NAMESPACE’ file, so this approach is usually not recommended. Evaluating `foo::f` will cause package `foo` to be loaded, but not attached, if it was not loaded already.

1.6.2 Registering S3 methods

The standard method for S3-style `UseMethod` dispatching might fail to locate methods defined in a package that is imported but not attached to the search path. To ensure that these methods are available the packages defining the methods should ensure that the generics are imported and register the methods using `S3method` directives. If a package defines a function `print.foo` intended to be used as a `print` method for class `foo`, then the directive

```
S3method(print, foo)
```

ensures that the method is registered and available for `UseMethod` dispatch. The function `print.foo` does not need to be exported. Since the generic `print` is defined in `base` it does not need to be imported explicitly. This mechanism is intended for use with generics that are defined in a name space. Any methods for a generic defined in a package that does not use a name space should be exported, and the package defining and exporting the methods should be attached to the search path if the methods are to be found.

1.6.3 Load hooks

There are a number of hooks that apply to packages with name spaces. See `help(".onLoad")` for more details.

Packages with name spaces do not use the `.First.lib` function. Since loading and attaching are distinct operations when a name space is used, separate hooks are provided for each. These hook functions are called `.onLoad` and `.onAttach`. They take the same arguments as `.First.lib`; they should be defined in the name space but not exported.

However, packages with name spaces *do* use the `.Last.lib` function. There is also a hook `.onUnload` which is called when the name space is unloaded (via a call to `unloadNamespace`) with argument the full path to the directory in which the package was installed. `.onUnload` should be defined in the name space and not exported, but `.Last.lib` does need to be exported.

Packages are not likely to need `.onAttach` (except perhaps for a start-up banner); code to set options and load shared objects should be placed in a `.onLoad` function, or use made of the `useDynLib` directive described next.

There can be one or more `useDynLib` directives which allow shared objects that need to be loaded to be specified in the ‘NAMESPACE’ file. The directive

```
useDynLib(foo)
```

registers the shared object `foo` for loading with `library.dynam`. Loading of registered object(s) occurs after the package code has been loaded and before running the load hook function. Packages that would only need a load hook function to load a shared object can use the `useDynLib` directive instead.

User-level hooks are also available: see the help on function `setHook`.

The `useDynLib` directive also accepts the names of the native routines that are to be used in R via the `.C`, `.Call`, `.Fortran` and `.External` interface functions. These are given as additional arguments to the directive, for example,


```
useDynLib(foo, myRoutine, myOtherRoutine)
```

By specifying these names in the `useDynLib` directive, the native symbols are resolved when the package is loaded and R variables identifying these symbols are added to the package's name space with these names. These can be used in the `.C`, `.Call`, `.Fortran` and `.External` calls in place of the name of the routine and the `PACKAGE` argument. For instance, we can call the routine `myRoutine` from R with the code

```
.Call(myRoutine, x, y)
```

rather than

```
.Call("myRoutine", x, y, PACKAGE = "foo")
```

There are at least two benefits to this approach. Firstly, the symbol lookup is done just once for each symbol rather than each time it the routine is invoked. Secondly, this removes any ambiguity in resolving symbols that might be present in several compiled libraries. In particular, it allows for correctly resolving routines when different versions of the same package are loaded concurrently in the same R session.

In some circumstances, there will already be an R variable in the package with the same name as a native symbol. For example, we may have an R function in the package named `myRoutine`. In this case, it is necessary to map the native symbol to a different R variable name. This can be done in the `useDynLib` directive by using named arguments. For instance, to map the native symbol name `myRoutine` to the R variable `myRoutine_sym`, we would use

```
useDynLib(foo, myRoutine_sym = myRoutine, myOtherRoutine)
```

We could then call that routine from R using the command

```
.Call(myRoutine_sym, x, y)
```

Symbols without explicit names are assigned to the R variable with that name.

In some cases, it may be preferable not to create R variables in the package's name space that identify the native routines. It may be too costly to compute these for many routines when the package is loaded if many of these routines are not likely to be used. In this case, one can still perform the symbol resolution correctly using the DLL, but do this each time the routine is called. Given a reference to the DLL as an R variable, say `dll`, we can call the routine `myRoutine` using the expression

```
.Call(dll$myRoutine, x, y)
```

The `$` operator resolves the routine with the given name in the DLL using a call to `getNativeSymbol`. This is the same computation as above where we resolve the symbol when the package is loaded. The only difference is that this is done each time in the case of `dll$myRoutine`.

In order to use this dynamic approach (e.g., `dll$myRoutine`), one needs the reference to the DLL as an R variable in the package. The DLL can be assigned to a variable by using the `variable = dllName` format used above for mapping symbols to R variables. For example, if we wanted to assign the DLL reference for the DLL `foo` in the example above to the variable `myDLL`, we would use the following directive in the 'NAMESPACE' file:

```
myDLL = useDynLib(foo, myRoutine_sym = myRoutine, myOtherRoutine)
```

Then, the R variable `myDLL` is in the package's name space and available for calls such as `myDLL$dynRoutine` to access routines that are not explicitly resolved at load time.

If the package has registration information (see [Section 5.4 \[Registering native routines\]](#), [page 55](#)), then we can use that directly rather than specifying the list of symbols again in the `useDynLib` directive in the 'NAMESPACE' file. Each routine in the registration information is specified by giving a name by which the routine is to be specified along with the address of the routine and any information about the number and type of the parameters. Using the `.registration` argument of `useDynLib`, we can instruct the name space mechanism to create R variables for these symbols. For example, suppose we have the following registration information for a DLL named `myDLL`:

```

R_CMethodDef cMethods[] = {
    {"foo", &foo, 4, {REALSXP, INTSXP, STRSXP, LGLSXP}},
    {"bar_sym", &bar, 0},
    {NULL, NULL, 0}
};

R_CallMethodDef callMethods[] = {
    {"R_call_sym", &R_call, 4},
    {"R_version_sym", &R_version, 0},
    {NULL, NULL, 0}
};

```

Then, the directive in the ‘NAMESPACE’ file

```
useDynLib(myDLL, .registration = TRUE)
```

causes the DLL to be loaded and also for the R variables `foo`, `bar_sym`, `R_call_sym` and `R_version_sym` to be defined in the package’s name space.

Note that the names for the R variables are taken from the entry in the registration information and do not need to be the same as the name of the native routine. This allows the creator of the registration information to map the native symbols to non-conflicting variable names in R, e.g. `R_version` to `R_version_sym` for use in an R function such as

```

R_version <- function()
{
    .Call(R_version_sym)
}

```

More information about this symbol lookup, along with some approaches for customizing it, is available from <http://www.omegahat.org/examples/RDotCall>.

1.6.4 An example

As an example consider two packages named **foo** and **bar**. The R code for package **foo** in file ‘foo.R’ is

```

x <- 1
f <- function(y) c(x,y)
foo <- function(x) .Call("foo", x, PACKAGE="foo")
print.foo <- function(x, ...) cat("<a foo>\n")

```

Some C code defines a C function compiled into DLL `foo` (with an appropriate extension). The ‘NAMESPACE’ file for this package is

```

useDynLib(foo)
export(f, foo)
S3method(print, foo)

```

The second package **bar** has code file ‘bar.R’

```

c <- function(...) sum(...)
g <- function(y) f(c(y, 7))
h <- function(y) y+9

```

and ‘NAMESPACE’ file

```
import(foo)
export(g, h)
```

Calling `library(bar)` loads **bar** and attaches its exports to the search path. Package **foo** is also loaded but not attached to the search path. A call to `g` produces

```
> g(6)
[1] 1 13
```

This is consistent with the definitions of `c` in the two settings: in **bar** the function `c` is defined to be equivalent to `sum`, but in **foo** the variable `c` refers to the standard function `c` in **base**.

1.6.5 Summary – converting an existing package

To summarize, converting an existing package to use a name space involves several simple steps:

- Identify the public definitions and place them in `export` directives.
- Identify S3-style method definitions and write corresponding `S3method` declarations.
- Identify dependencies and replace any `require` calls by `import` directives.
- Replace `.First.lib` functions with `.onLoad` functions or `useDynLib` directives.

Some code analysis tools to aid in this process are currently under development.

1.6.6 Name spaces with formal classes and methods

Some additional steps are needed for packages which make use of formal (S4-style) classes and methods (unless these are purely used internally). There needs to be an `.onLoad` action to ensure that the methods package is loaded *and* attached:

```
.onLoad <- function(lib, pkg) require(methods)
```

and any classes and methods which are to be exported need to be declared as such in the ‘NAMESPACE’ file. For example, the now-defunct **mle** package had

```
importFrom(graphics, plot)
importFrom(stats, profile, confint)
exportClasses("mle", "profile.mle", "summary.mle")
exportMethods("confint", "plot", "profile", "summary", "show")
```

All formal classes need to be listed in an `exportClasses` directive. All generics for which formal methods are defined need to be declared in an `exportMethods` directive, and where the generics are formed by taking over existing functions, those functions need to be imported (explicitly unless they are defined in the **base** namespace).

In addition, a package using classes and methods defined in another package needs to import them, with directives

```
importClassesFrom(package, ...)
importMethodsFrom(package, ...)
```

listing the classes and functions with methods respectively. Suppose we had two small packages **A** and **B** with **B** using **A**. Then they could have NAMESPACE files

```
export(f1, ng1)
exportMethods("")
exportClasses(c1)
```

and

```
importFrom(A, ng1)
importClassesFrom(A, c1)
importMethodsFrom(A, f1)
export(f4, f5)
exportMethods(f6, "[")
exportClasses(c1, c2)
```

respectively.

1.7 Writing portable packages

R CMD `check` provides a basic set of checks, but often further problems emerge when people try to install and use packages submitted to CRAN – many of these involve compiled code. Here are some further checks that you can do to make your package more portable.

- If your package has a ‘`configure`’ script, provide a ‘`configure.win`’ script to be used on Windows. The CRAN binary packages for Windows are built automatically, and if your package does not build without intervention it is unlikely to be easily available to a high proportion of R users.
- Make use of the abilities of your compilers to check the standards-conformance of your code. For example, `gcc` can be used with options ‘`-Wall -pedantic`’ to alert you to potential problems. Do not be tempted to assume that these are pure pedantry: for example R is regularly used on platforms where the C compiler does not accept C++ comments.
- Do be very careful with passing arguments between R, C and FORTRAN code. In particular, `long` in C will be 32-bit on most R platforms (including those used by the CRAN maintainers), but 64-bit on many modern Unix and Linux platforms. It is rather unlikely that the use of `long` in C code has been thought through: if you need a longer type than `int` you should use a configure test for a C99 type such as `int_fast64_t` (and failing that, `long long`) and typedef your own type to be `long` or `long long`, or use another suitable type (such as `size_t`). Note that `integer` in FORTRAN corresponds to `int` in C on all R platforms.
- Errors in memory allocation and reading/writing outside arrays are very common causes of crashes (e.g., segfaults) on some machines. See [Section 4.3.2 \[Using valgrind\]](#), page 47 for a tool which can be used to look for this.
- The Mac OS X linker has some restrictions not found on other platforms. Try to ensure that C entry points shared between source files are declared as `extern` in all but one of the files.
- Many platforms will allow unsatisfied entry points in compiled code, but will crash the application (here R) if they are ever used. Some (notably Windows) will not. Looking at the output of

```
nm -pg mypkg.so # or other extension such as '.sl' or '.dylib'
```

and checking if any of the symbols marked `U` is unexpected is a good way to avoid this.

- Conflicts between symbols in DLLs are handled in very platform-specific ways. Good ways to avoid trouble are to make as many symbols as possible static (check with `nm -pg`), and to use unusual names, as well as ensuring you have used the `PACKAGE` argument that R CMD `check` checks for.

1.8 Diagnostic messages

Now that diagnostic messages can be made available for translation, it is important to write them in a consistent style. Using the tools described in the next section to extract all the messages can give a useful overview of your consistency (or lack of it).

Some guidelines follow.

- Messages are sentence fragments, and not viewed in isolation. So it is conventional not to capitalize the first word and not to end with a period (or other punctuation).
- Try not to split up messages into small pieces. In C error messages use a single format string containing all English words in the messages.

In R error messages do not construct a message with `paste` (such messages will not be translated) but via multiple arguments to `stop` or `warning`, or via `gettextf`.

- Do not use colloquialisms such as “can’t” and “don’t”.
- If possible, make quotation marks part of your message as different languages have different conventions. In R messages this means not using `sQuote` or `dQuote` except where the argument is a variable.

Conventionally single quotation marks are used for quotations such as

`'ord'` must be a positive integer, at most the number of knots

and double quotation marks when referring to an R character string such as

`'format'` must be "normal" or "short" - using "normal"

Since ASCII does not contain directional quotation marks, it is best to use ‘’ and let the translator (including automatic translation) use directional quotations where available. The range of quotation styles is immense: unfortunately we cannot reproduce them in a portable `texinfo` document. But as a taster, some languages use ‘up’ and ‘down’ (comma) quotes rather than left or right quotes, and some use guillemets (and some use what Adobe calls ‘guillemotleft’ to start and others use it to end).

- Occasionally messages need to be singular or plural (and in other languages there may be no such concept or several plural forms – Slovenian has four). So avoid constructions such as was once used in `library`

```
if((length(nopkgs) > 0) && !missing(lib.loc)) {
  if(length(nopkgs) > 1)
    warning("libraries ",
            paste(sQuote(nopkgs), collapse = ", "),
            " contain no packages")
  else
    warning("library ", paste(sQuote(nopkgs)),
            " contains no package")
}
```

and was replaced by

```
if((length(nopkgs) > 0) && !missing(lib.loc)) {
  pkglist <- paste(sQuote(nopkgs), collapse = ", ")
  msg <- sprintf(ngettext(length(nopkgs),
                           "library %s contains no packages",
                           "libraries %s contain no packages"),
                 pkglist)
  warning(msg, domain=NA)
}
```

Note that it is much better to have complete clauses as here, since in another language one might need to say ‘There is no package in library %s’ or ‘There are no packages in libraries %s’.

1.9 Localization

R 2.1.0 introduced mechanisms to translate the R- and C-level error and warning messages, and these have been implemented in package **splines** as a demonstration. There are only available if

R is compiled with NLS support (which is requested by `configure` option ‘`--enable-nls`’, the default).

The procedures make use of `msgfmt` and `xgettext` which are part of GNU `gettext` and this will need to be installed: Windows users can find pre-compiled binaries at the GNU archive mirrors and packaged with the `poEdit` package (<http://poedit.sourceforge.net/download.php#win32>).

1.9.1 C-level messages

The process of enabling translations is

- In a header file that will be included in all the C files containing messages that should be translated, declare

```
#include <R.h> /* to include Rconfig.h */

#ifdef ENABLE_NLS
#include <libintl.h>
#define _(String) dgettext ("pkg", String)
/* replace pkg as appropriate */
#else
#define _(String) (String)
#endif
```

- For each message that should be translated, wrap it in `_(...)`, for example

```
error(_("'ord' must be a positive integer"));
```

- In the package’s ‘`src`’ directory run

```
xgettext --keyword=_ -o pkg.pot *.c
```

The file ‘`src/pkg.pot`’ is the template file, and conventionally this is shipped as ‘`po/pkg.pot`’. A translator to another language makes a copy of this file and edits it (see the `gettext` manual) to produce say ‘`ll.po`’, where *ll* is the code for the language in which the translation is to be used. (This file would be shipped in the ‘`po`’ directory.) Next run `msgfmt` on ‘`ll.po`’ to produce ‘`ll.mo`’, and copy that to ‘`inst/po/ll/LC_MESSAGES/pkg.mo`’. Now when the package is loaded after installation it will look for translations of its messages in the ‘`po/lang/LC_MESSAGES/pkg.mo`’ file for any language *lang* that matches the user’s preferences (via the setting of the `LANGUAGE` environment variable or from the locale settings).

1.9.2 R messages

Mechanisms to support the automatic translation of R `stop`, `warning` and `message` messages are in place, provided the package has a namespace. They make use of message catalogs in the same way as C-level messages, but using domain `R-pkg` rather than `pkg`. Translation of character strings inside `stop`, `warning` and `message` calls is automatically enabled, as well as other messages enclosed in calls to `gettext` or `gettextf`. (To suppress this, use argument `domain=NA`.)

Tools to prepare the ‘`R-pkg.pot`’ file are provided in package `tools`: `xgettext2pot` will prepare a file from all strings occurring inside `gettext/gettextf`, `stop`, `warning` and `message` calls. Some of these are likely to be spurious and so the file is likely to need manual editing. `xgettext` extracts the actual calls and so is more useful when tidying up error messages.

Translation of messages which might be singular or plural can be very intricate: languages can have up to four different forms. The R function `ngettext` provides an interface to the C function of the same name, and will choose an appropriate singular or plural form for the selected language depending on the value of its first argument `n`.

Packages without namespaces will need to use `domain="R-pkg"` explicitly in calls to `stop`, `warning`, `message`, `gettext/gettextf` and `ngettext`.

1.10 Package types

The ‘DESCRIPTION’ file has an optional field **Type** which if missing is assumed to be **Package**, the sort of extension discussed so far in this chapter. Currently two other types are recognized, both of which need write permission in the R installation tree.

1.10.1 Frontend

This is a rather general mechanism, designed for adding new front-ends such as the **gnomeGUI** package. If a ‘configure’ file is found in the top-level directory of the package it is executed, and then if a Makefile is found (often generated by ‘configure’), **make** is called. If R CMD INSTALL --clean is used **make clean** is called. No other action is taken.

R CMD build can package up this type of extension, but R CMD check will check the type and skip it.

1.10.2 Translation

Conventionally, a translation package for language *ll* is called **Translation-ll** and has **Type: Translation**. It needs to contain the directories ‘share/locale/ll’ and ‘library/pkgname/po/ll’, or at least those for which translations are available. The files ‘.mo’ are installed in the parallel places in the R installation tree.

For example, a package **Translation-it** might be prepared from an installed (and tested) version of R by

```
mkdir Translation-it
cd Translation-it
(cd $R_HOME; tar cf - share/locale/it library/*/po/it) | tar xf -
# the next step is not needed on Windows
msgfmt -c -o share/locale/it/LC_MESSAGES/RGui.mo $R_SRC_HOME/po/RGui-it.gmo
# create a DESCRIPTION file
cd ..
R CMD build Translation-it
```

It is probably appropriate to give the package a version number based on the version of R which has been translated. So the ‘DESCRIPTION’ file might look like

```
Package: Translation-it
Type: Translation
Version: 2.2.1-1
Title: Italian Translations for R 2.2.1
Description: Italian Translations for R 2.2.1
Author: The translators
Maintainer: Some Body <somebody@some.where.net>
Licence: GPL Version 2 or later.
```


2 Writing R documentation files

2.1 Rd format

R objects are documented in files written in “R documentation” (Rd) format, a simple markup language closely resembling (La)TeX, which can be processed into a variety of formats, including L^AT_EX, HTML and plain text. The translation is carried out by the Perl script `Rdconv` in ‘`R_HOME/bin`’ and by the installation scripts for packages.

The R distribution contains about 1000 such files which can be found in the ‘`src/library/pkg/man`’ directories of the R source tree, where *pkg* stands for the standard packages which are included in the R distribution.

As an example, let us look at the file ‘`src/library/base/man/load.Rd`’ which documents the R function `load`.

```
\name{load}
\alias{load}
\title{Reload Saved Datasets}
\description{
  Reload the datasets written to a file with the function
  \code{save}.
}
\usage{
load(file, envir = parent.frame())
}
\arguments{
  \item{file}{a connection or a character string giving the
    name of the file to load.}
  \item{envir}{the environment where the data should be
    loaded.}
}
\seealso{
  \code{\link{save}}.
}
\examples{
## save all data
save(list = ls(), file= "all.Rdata")

## restore the saved values to the current environment
load("all.Rdata")

## restore the saved values to the workspace
load("all.Rdata", .GlobalEnv)
}
\keyword{file}
```

An Rd file consists of three parts. The header gives basic information about the name of the file, the topics documented, a title, a short textual description and R usage information for the objects documented. The body gives further information (for example, on the function’s arguments and return value, as in the above example). Finally, there is a footer with keyword information. The header and footer are mandatory.

See the “[Guidelines for Rd files](#)” for guidelines for writing documentation in Rd format which should be useful for package writers.

2.1.1 Documenting functions

The basic markup commands used for documenting R objects (in particular, functions) are given in this subsection.

\name{name}

name typically¹ is the basename of the Rd file containing the documentation. It is the “name” of the Rd object represented by the file and has to be unique in a package.

\alias{topic}

The **\alias** entries specify all “topics” the file documents. This information is collected into index data bases for lookup by the on-line (plain text and HTML) help systems.

There may be several **\alias** entries. Quite often it is convenient to document several R objects in one file. For example, file ‘Normal.Rd’ documents the density, distribution function, quantile function and generation of random variates for the normal distribution, and hence starts with

```
\name{Normal}
\alias{Normal}
\alias{dnorm}
\alias{pnorm}
\alias{qnorm}
\alias{rnorm}
```

Also, it is often convenient to have several different ways to refer to an R object, and an **\alias** does not need to be the name of an object.

Note that the **\name** is not necessarily a topic documented, and if so desired it needs to have an explicit **\alias** entry (as in this example).

\title{Title}

Title information for the Rd file. This should be capitalized, not end in a period, and not use any markup (which would cause problems for hypertext search). Use of characters other than English text and punctuation (e.g. <) may limit portability.

\description{...}

A short description of what the function(s) do(es) (one paragraph, a few lines only). (If a description is “too long” and cannot easily be shortened, the file probably tries to document too much at once.)

\usage{fun(arg1, arg2, ...)}

One or more lines showing the synopsis of the function(s) and variables documented in the file. These are set verbatim in typewriter font.

The usage information specified should in general match the function definition *exactly* (such that automatic checking for consistency between code and documentation is possible). Otherwise, include a **\synopsis** section with the actual definition.

For example, **abline** is a function for adding a straight line to a plot which can be used in several different ways, depending on the named arguments specified. Hence, ‘**abline.Rd**’ contains

¹ There can be exceptions: for example Rd files are not allowed to start with a dot, and have to be uniquely named on a case-insensitive file system.

```

\synopsis{
abline(a = NULL, b = NULL, h = NULL, v = NULL, reg = NULL,
       coef = NULL, untf = FALSE, col = par("col"),
       lty = par("lty"), lwd = NULL, ...)
}
\usage{
abline(a, b, untf = FALSE, \dots)
abline(h=, untf = FALSE, \dots)
abline(v=, untf = FALSE, \dots)
...
}

```

Use `\method{generic}{class}` to indicate the name of an S3 method for the generic function *generic* for objects inheriting from class "*class*". In the printed versions, this will come out as *generic* (reflecting the understanding that methods should not be invoked directly but via method dispatch), but `codoc()` and other QC tools always have access to the full name.

For example, 'print.ts.Rd' contains

```

\usage{
\method{print}{ts}(x, calendar, \dots)
}

```

which will print as

Usage:

```

## S3 method for class 'ts':
print(x, calendar, ...)

```

Usage for replacement functions should be given in the style of `dim(x) <- value` rather than explicitly indicating the name of the replacement function ("`dim<=`" in the above). Similarly, one can use `\method{generic}{class}(arglist) <- value` to indicate the usage of an S3 replacement method for the generic replacement function "*generic*<=" for objects inheriting from class "*class*".

`\arguments{...}`

Description of the function's arguments, using an entry of the form

```
\item{arg_i}{Description of arg_i.}
```

for each element of the argument list. There may be optional text before and after these entries.

`\details{...}`

A detailed if possible precise description of the functionality provided, extending the basic information in the `\description` slot.

`\value{...}`

Description of the function's return value.

If a list with multiple values is returned, you can use entries of the form

```
\item{comp_i}{Description of comp_i.}
```

for each component of the list returned. Optional text may precede this list (see the introductory example for `rle`).

`\references{...}`

A section with references to the literature. Use `\url{}` for web pointers.

`\note{...}`

Use this for a special note you want to have pointed out.

For example, ‘pie.Rd’ contains

```
\note{
  Pie charts are a very bad way of displaying information.
  The eye is good at judging linear measures and bad at
  judging relative areas.
  ...
}
```

`\author{...}`

Information about the author(s) of the Rd file. Use `\email{}` without extra delimiters (‘()’ or ‘< >’) to specify email addresses, or `\url{}` for web pointers.

`\seealso{...}`

Pointers to related R objects, using `\code{\link{...}}` to refer to them (`\code` is the correct markup for R object names, and `\link` produces hyperlinks in output formats which support this. See [Section 2.3 \[Marking text\]](#), page 32, and [Section 2.5 \[Cross-references\]](#), page 34).

`\examples{...}`

Examples of how to use the function. These are set as formatted in typewriter font: see [Section 2.7 \[Insertions\]](#), page 35 for when characters need to be escaped. (Markup `\link` and `\var` will be interpreted, but no other.)

Examples are not only useful for documentation purposes, but also provide test code used for diagnostic checking of R. By default, text inside `\examples{}` will be displayed in the output of the help page and run by R CMD `check`. You can use `\dontrun{}` for commands that should only be shown, but not run, and `\dontshow{}` for extra commands for testing that should not be shown to users, but will be run by `example()`. (Previously this was called `\testonly`, and that is still accepted.)

For example,

```
x <- runif(10)      # Shown and run.
\dontrun{plot(x)}   # Only shown.
\dontshow{log(x)}   # Only run.
```

Thus, example code not included in `\dontrun` must be executable! In addition, it should not use any system-specific features or require special facilities (such as Internet access or write permission to specific directories). As from R 1.8.0, code included in `\dontrun` is indicated by comments in the processed help files.

Data needed for making the examples executable can be obtained by random number generation (for example, `x <- rnorm(100)`), or by using standard data sets listed by `data()` (see `?data` for more info).

`\keyword{key}`

Each `\keyword` entry should specify one of the standard keywords (as listed in the file ‘`R_DOC_DIR/KEYWORDS.db`’, default ‘`R_HOME/doc/KEYWORDS.db`’). There must be at least one `\keyword` entry, but can be more than one if the R object being documented falls into more than one category.

The special keyword ‘`internal`’ marks a page of internal objects that are not part of the packages’ API. If the help page for object `foo` has keyword ‘`internal`’, then `help(foo)` gives this help page, but `foo` is excluded from several object indices, like the alphabetical list of objects in the HTML help system.

The R function `prompt` facilitates the construction of files documenting R objects. If `foo` is an R function, then `prompt(foo)` produces file ‘`foo.Rd`’ which already contains the proper function and argument names of `foo`, and a structure which can be filled in with information.

2.1.2 Documenting data sets

The structure of Rd files which document R data sets is slightly different. Whereas sections such as `\arguments` and `\value` are not needed, the format and source of the data should be explained.

As an example, let us look at `'src/library/datasets/man/rivers.Rd'` which documents the standard R data set `rivers`.

```
\name{rivers}
\docType{data}
\alias{rivers}
\title{Lengths of Major North American Rivers}
\description{
  This data set gives the lengths (in miles) of 141 \dQuote{major}
  rivers in North America, as compiled by the US Geological
  Survey.
}
\usage{rivers}
\format{A vector containing 141 observations.}
\source{World Almanac and Book of Facts, 1975, page 406.}
\references{
  McNeil, D. R. (1977) \emph{Interactive Data Analysis}.
  New York: Wiley.
}
\keyword{datasets}
```

This uses the following additional markup commands.

`\docType{...}`

Indicates the “type” of the documentation object. Always `'data'` for data sets.

`\format{...}`

A description of the format of the data set (as a vector, matrix, data frame, time series, ...). For matrices and data frames this should give a description of each column, preferably as a list or table. See [Section 2.4 \[Lists and tables\]](#), page 34, for more information.

`\source{...}`

Details of the original source (a reference or URL). In addition, section `\references` could give secondary sources and usages.

Note also that when documenting data set *bar*,

- The `\usage` entry is always *bar* or (for packages which do not use lazy-loading of data) `data(bar)`. (In particular, only document a *single* data object per Rd file.)
- The `\keyword` entry is always `'datasets'`.

If *bar* is a data frame, documenting it as a data set can be initiated via `prompt(bar)`.

2.1.3 Documenting S4 classes and methods

Structure of and special markup for documenting S4 classes and methods are still under development. There are special ways to use the ‘?’ operator, namely `'class?topic'` and `'methods?topic'`, that depend on conventions for the topic names used in `\alias` entries. The topic names for S4 classes and methods respectively are of the form

```
class-class
generic,signature_list-method
```

where *signature_list* contains the names of the classes in the signature of the method (without quotes) separated by ‘,’ (without whitespace), with ‘ANY’ used for arguments without an explicit

specification. E.g., ‘`genericFunction-class`’ is the topic name for documentation for the S4 class “`genericFunction`”, and ‘`coerce,ANY,NULL-method`’ is the topic name for documentation for the S4 method for `coerce` for signature `c("ANY", "NULL")`.

Skeletons of documentation for S4 classes and methods can be generated by using the functions `promptClass()` and `promptMethods()` from package **methods**. If it is necessary or desired to provide an explicit function declaration (in a `\usage` section) for an S4 method (e.g., if it has “surprising arguments” to be mentioned explicitly), one can use the special markup

```
\S4method{generic}{signature_list}(argument_list)
(e.g., '\S4method{coerce}{ANY,NULL}(from, to)').
```

To allow for making full use of the potential of the on-line documentation system, all user-visible S4 classes and methods in a package should at least have a suitable `\alias` entry in one of the package’s Rd files. If a package has methods for a function defined originally somewhere else, and does not change the underlying default method for the function, the package is responsible for documenting the methods it creates, but not for the function itself or the default method.

See `help("Documentation", package = "methods")` for more information on using and creating on-line documentation for S4 classes and methods.

2.1.4 Documenting packages

Packages may have an overview man page with an `\alias pkgname-package`, e.g. `utils-package` for the `utils` package, when `package?pkgname` will open that help page. If a topic named `pkgname` does not exist in another .Rd file, it is helpful to use this as an additional `\alias`.

Skeletons of documentation for a package can be generated using the function `promptPackage()`. If the `final = TRUE` argument is used, then the Rd file will be generated in final form, containing the information that would be produced by `library(help = pkgname)`. Otherwise (the default) comments will be inserted giving suggestions for content.

The only requirement for this page is that it include a `\docType{package}` statement. All other content is optional. We suggest that it should be a short overview, to give a reader unfamiliar with the package enough information to get started. More extensive documentation is better placed into a package vignette (see [Section 1.4 \[Writing package vignettes\]](#), page 16) and referenced from this page, or into individual man pages for the functions, datasets, or classes.

2.2 Sectioning

To begin a new paragraph or leave a blank line in an example, just insert an empty line (as in (La)TeX). To break a line, use `\cr`.

In addition to the predefined sections (such as `\description{}`, `\value{}`, etc.), you can “define” arbitrary ones by `\section{section_title}{...}`. For example

```
\section{Warning}{You must not call this function unless ...}
```

For consistency with the pre-assigned sections, the section name (the first argument to `\section`) should be capitalized (but not all upper case).

Note that additional named sections are always inserted at a fixed position in the output (before `\note`, `\seealso` and the examples), no matter where they appear in the input (but in the same order as the input).

2.3 Marking text

The following logical markup commands are available for emphasizing or quoting text.

```
\emph{text}
\strong{text}
```

Emphasize *text* using *italic* and **bold** font if possible; `\strong` is stronger.

`\bold{text}`

Set *text* in **bold** font if possible.

`\sQuote{text}`

`\dQuote{text}`

Portably single or double quote *text* (without hard-wiring the quotation marks).

The following logical markup commands are available for indicating specific kinds of text.

`\code{text}`

Indicate text that is a literal example of a piece of a program, e.g., a fragment of R code or the name of an R object, using **typewriter** font if possible. Only percent signs and unpaired braces always need to be escaped (by a backslash) inside `\code`: backslash will need to be escaped if it precedes ‘%’, ‘{’ or ‘}’. The only markup interpreted inside `\code` is `\link` and `\var`.

`\preformatted{text}`

Indicate text that is a literal example of a piece of a program, using **typewriter** font if possible. The same characters need to be escaped as for `\code`. All other formatting, e.g. line breaks, is preserved. The closing brace should be on a line by itself.

`\kbd{keyboard-characters}`

Indicate keyboard input, using *slanted typewriter* font if possible, so users can distinguish the characters they are supposed to type from those the computer outputs.

`\samp{text}`

Indicate text that is a literal example of a sequence of characters.

`\pkg{package-name}`

Indicate the name of an R package.

`\file{file-name}`

Indicate the name of a file. Note that special characters do need to be escaped.

`\email{email-address}`

Indicate an electronic mail address.

`\url{uniform-resource-locator}`

Indicate a uniform resource locator (URL) for the World Wide Web.

`\var{metasyntactic-variable}`

Indicate a metasyntactic variable.

`\env{environment-variable}`

Indicate an environment variable.

`\option{option}`

Indicate a command-line option.

`\command{command-name}`

Indicate the name of a command.

`\dfn{term}`

Indicate the introductory or defining use of a term.

`\cite{reference}`

Indicate a reference without a direct cross-reference via `\link` (see [Section 2.5 \[Cross-references\]](#), page 34), such as the name of a book.

`\acronym{acronym}`

Indicate an acronym (an abbreviation written in all capital letters), such as GNU.

Note that unless explicitly stated otherwise, special characters (see [Section 2.7 \[Insertions\]](#), [page 35](#)) must be escaped inside the above markup commands.

2.4 Lists and tables

The `\itemize` and `\enumerate` commands take a single argument, within which there may be one or more `\item` commands. The text following each `\item` is formatted as one or more paragraphs, suitably indented and with the first paragraph marked with a bullet point (`\itemize`) or a number (`\enumerate`).

`\itemize` and `\enumerate` commands may be nested.

The `\describe` command is similar to `\itemize` but allows initial labels to be specified. The `\items` take two arguments, the label and the body of the item, in exactly the same way as argument and value `\items`. `\describe` commands are mapped to `<DL>` lists in HTML and `\description` lists in L^AT_EX.

The `\tabular` command takes two arguments. The first gives for each of the columns the required alignment ('l' for left-justification, 'r' for right-justification or 'c' for centring.) The second argument consists of an arbitrary number of lines separated by `\cr`, and with fields separated by `\tab`. For example:

```
\tabular{rlll}{
  [,1] \tab Ozone      \tab numeric \tab Ozone (ppb)\cr
  [,2] \tab Solar.R    \tab numeric \tab Solar R (lang)\cr
  [,3] \tab Wind       \tab numeric \tab Wind (mph)\cr
  [,4] \tab Temp       \tab numeric \tab Temperature (degrees F)\cr
  [,5] \tab Month      \tab numeric \tab Month (1--12)\cr
  [,6] \tab Day        \tab numeric \tab Day of month (1--31)
}
```

There must be the same number of fields on each line as there are alignments in the first argument, and they must be non-empty (but can contain only spaces).

2.5 Cross-references

The markup `\link{foo}` (usually in the combination `\code{\link{foo}}`) produces a hyperlink to the help for *foo*. Here *foo* is a *topic*, that is the argument of `\alias` markup in another .Rd file (possibly in another package). Hyperlinks are supported in some of the formats to which .Rd files are converted, for example HTML and PDF, but ignored in others, e.g. the text and S nroff formats.

One main usage of `\link` is in the `\seealso` section of the help page, see [Section 2.1 \[Rd format\]](#), [page 27](#).

You can specify a link to a different topic than its name by `\link[=dest]{name}` which links to topic *dest* with name *name*. This can be used to refer to the documentation for S3/4 classes, for example `\code{"\link[=abc-class]{abc}"}` would be a way to refer to the documentation of an S4 class "abc" defined in your package, and `\code{"\link[=terms.object]{terms}"}` to the S3 "terms" class (in package **stats**). To make these easy to read, `\code{"\linkS4class{abc}"}` expands to the form given above.

There are two other forms of optional argument specified as `\link[pkg]{foo}` and `\link[pkg:bar]{foo}` to link to the package *pkg*, to files '*foo.html*' and '*bar.html*' respectively. These are rarely needed, perhaps to refer to not-yet-installed packages (but there the HTML help system will resolve the link at run time) or in the normally undesirable event that more than one package offers help on a topic² (in which case the present package

² a common example in CRAN packages is `\link[mgcv]{gam}`.

has precedence so this is only needed to refer to other packages). They are only in used in (C)HTML help (and not for hyperlinks in latex nor S sgml conversions of help pages), and link to the file rather than the topic (since there is no way to know which topics are in which files in an uninstalled package).

2.6 Mathematics

Mathematical formulae should be set beautifully for printed documentation yet we still want something useful for text and HTML online help. To this end, the two commands `\eqn{latex}{ascii}` and `\deqn{latex}{ascii}` are used. Where `\eqn` is used for “inline” formulae (corresponding to \TeX ’s $\$...\$$, `\deqn` gives “displayed equations” (as in \LaTeX ’s `displaymath` environment, or \TeX ’s $\$\$...\$\$$).

Both commands can also be used as `\eqn{latexascii}` (only *one* argument) which then is used for both *latex* and *ascii*.

The following example is from ‘Poisson.Rd’:

```
\deqn{p(x) = \frac{\lambda^x e^{-\lambda}}{x!}}{x!}%
      p(x) = lambda^x exp(-lambda)/x!}
for \eqn{x = 0, 1, 2, \ldots}.
```

For the \LaTeX manual, this becomes

$$p(x) = \lambda^x \frac{e^{-\lambda}}{x!}$$

for $x = 0, 1, 2, \dots$

For HTML and text on-line help we get

```
p(x) = lambda^x exp(-lambda)/x!
for x = 0, 1, 2, ....
```

2.7 Insertions

Use `\R` for the R system itself (you don’t need extra ‘`{}`’ or ‘`\`’). Use `\dots` for the dots in function argument lists ‘`...`’, and `\ldots` for ellipsis dots in ordinary text.

After a ‘`%`’, you can put your own comments regarding the help text. The rest of the line will be completely disregarded, normally. Therefore, you can also use it to make part of the “help” invisible.

You can produce a backslash (‘`\`’) by escaping it by another backslash. (Note that `\cr` is used for generating line breaks.)

The “comment” character ‘`%`’ and unpaired braces³ *always* need to be escaped by ‘`\`’, and ‘`\\`’ can be (but rarely needs to be) used for backslash. Inside the verbatim-like commands (`\code`, `\preformatted` and `\examples`), no other characters are special. Note that `\file` is **not** a verbatim-like command.

In “regular” text (no verbatim, no `\eqn`, ...), you currently must escape most \LaTeX special characters, i.e., besides ‘`%`’, ‘`{}`’, and ‘`}`’, the four specials ‘`$`’, ‘`#`’, and ‘`_`’ are produced by preceding each with a ‘`\`’. (‘`&`’ can also be escaped, but need not be.) Further, enter ‘`^`’ as `\eqn{\mbox{\textasciicircum}}{^}`, and ‘`~`’ by `\eqn{\mbox{\textasciitilde}}{~}` or

³ See the examples section in the file ‘Paren.Rd’ for an example.

`\eqn{\sim}{~}` (for a short and long tilde respectively). Also, ‘<’, ‘>’, and ‘|’ must only be used in math mode, i.e., within `\eqn` or `\deqn`.

Text which might need to be represented differently in different encodings should be marked by `\enc`, e.g. `\enc{Jöreskog}{Joreskog}` where the first argument will be used where encodings are allowed and the second should be ASCII (and is used for e.g. the text conversion).

2.8 Indices

The `\alias` command (see [Section 2.1.1 \[Documenting functions\]](#), page 27) is used to specify the “topics” documented, which should include *all* R objects in a package such as functions and variables, data sets, and S4 classes and methods (see [Section 2.1.3 \[Documenting S4 classes and methods\]](#), page 31). The on-line help system searches the index data base consisting of all alias topics.

In addition, it is possible to provide “concept index entries” using `\concept`, which can be used for `help.search()` lookups. E.g., file ‘`cor.test.Rd`’ in the standard package **stats** contains

```
\concept{Kendall correlation coefficient}
\concept{Pearson correlation coefficient}
\concept{Spearman correlation coefficient}
```

so that e.g. `help.search("Spearman")` will succeed in finding the help page for the test for association between paired samples using Spearman’s ρ . (Note that concepts are not currently supported by the HTML search accessed *via* ‘`help.start()`’.)

(Note that `help.search()` only uses “sections” of documentation objects with no additional markup.)

If you want to cross reference such items from other help files via `\link`, you need to use `\alias` and not `\concept`.

2.9 Platform-specific documentation

Sometimes the documentation needs to differ by platform. Currently two OS-specific options are available, ‘`unix`’ and ‘`windows`’, and lines in the help source file can be enclosed in

```
#ifdef OS
...
#endif
```

or

```
#ifndef OS
...
#endif
```

for OS-specific inclusion or exclusion.

If the differences between platforms are extensive or the R objects documented are only relevant to one platform, platform-specific Rd files can be put in a ‘`unix`’ or ‘`windows`’ subdirectory.

2.10 Encoding

Rd files are text files and so it is impossible to deduce the encoding they are written in: ASCII, UTF-8, Latin-1, Latin-9 *etc.* So the `\encoding{}` directive must be used to specify the encoding if it is not ASCII. This is used when creating the header of the HTML conversion (if not present, for back-compatibility the processing to HTML assumes that the file is in Latin-1 (ISO-8859-1)) and to make a comment in the examples file. It is also used to indicate to L^AT_EX how to process the file (see below).

Wherever possible, avoid non-ASCII chars in Rd files, and even symbols such as < > \$ ^ & | ~ * outside verbatim environments (since they may disappear in fonts designed to render text).

For convenience, encoding names `latin1` and `latin2` are always recognized: these and UTF-8 are likely to work fairly widely.

The `\enc` command (see [Section 2.7 \[Insertions\]](#), page 35) can be used to provide transliterations which will be used in conversions that do not support the declared encoding.

The \LaTeX conversion converts an explicit encoding of the file to a

```
\inputencoding{some_encoding}
```

command, and this needs to be matched by a suitable invocation of the `\usepackage{inputenc}` command. The R utility `R CMD Rd2dvi` looks at the converted code and includes the encodings used: it might for example use

```
\usepackage[latin1,latin9,utf8]{inputenc}
```

(Use of `utf8` as an encoding requires \LaTeX dated 2003/12/01 or later.)

Note that this mechanism works best with letters and for example the copyright symbol may be rendered as a subscript and the plus-minus symbol cannot be used in text.

2.11 Processing Rd format

There are several commands to process Rd files from the system command line. All of these need Perl to be installed.

Using `R CMD Rdconv` one can convert R documentation format to other formats, or extract the executable examples for run-time testing. Currently, conversions to plain text, HTML, \LaTeX , and S version 3 or 4 documentation formats are supported.

In addition to this low-level conversion tool, the R distribution provides two user-level programs for processing Rd format. `R CMD Rd2txt` produces “pretty” plain text output from an Rd file, and is particularly useful as a previewer when writing Rd format documentation within Emacs. `R CMD Rd2dvi` generates DVI (or, if option `--pdf` is given, PDF) output from documentation in Rd files, which can be specified either explicitly or by the path to a directory with the sources of a package (or bundle). In the latter case, a reference manual for all documented objects in the package is created, including the information in the ‘DESCRIPTION’ files.

Finally, `R CMD Sd2Rd` converts S version 3 documentation files (which use an extended Nroff format) and S version 4 documentation (which uses SGML markup) to Rd format. This is useful when porting a package originally written for the S system to R. S version 3 files usually have extension `.d`, whereas version 4 ones have extension `.sgml` or `.sgm`.

The exact usage and a detailed list of available options for each of the above commands can be obtained by running `R CMD command --help`, e.g., `R CMD Rdconv --help`. All available commands can be listed using `R --help`.

All of these work under Windows. You will need to have installed the files in the R binary Windows distribution for installing source packages (this is true for a default installation), and for `R CMD Rd2dvi` also the tools to build packages from source as described in the “R Installation and Administration” manual.

3 Tidying and profiling R code

R code which is worth preserving in a package and perhaps making available for others to use is worth documenting, tidying up and perhaps optimizing. The last two of these activities are the subject of this chapter.

3.1 Tidying R code

R treats function code loaded from packages and code entered by users differently. Code entered by users has the source code stored in an attribute, and when the function is listed, the original source is reproduced. Loading code from a package (by default) discards the source code, and the function listing is re-created from the parse tree of the function.

Normally keeping the source code is a good idea, and in particular it avoids comments being moved around in the source. However, we can make use of the ability to re-create a function listing from its parse tree to produce a tidy version of the function, for example with consistent indentation and spaces around operators. This tidied version is much easier to read, not least by other users who are used to the standard format. Although the deparsing cannot do so, we recommend the consistent use of the preferred assignment operator ‘<-’ (rather than ‘=’) for assignment.

We can subvert the keeping of source in two ways.

1. The option `keep.source` can be set to `FALSE` before the code is loaded into R.
2. The stored source code can be removed by removing the `source` attribute, for example by

```
attr(myfun, "source") <- NULL
```

In each case if we then list the function we will get the standard layout.

Suppose we have a file of functions ‘`myfuns.R`’ that we want to tidy up. Create a file ‘`tidy.R`’ containing

```
options(keep.source = FALSE)
source("myfuns.R")
dump(ls(all = TRUE), file = "new.myfuns.R")
```

and run R with this as the source file, for example by `R --vanilla < tidy.R` or by pasting into an R session. Then the file ‘`new.myfuns.R`’ will contain the functions in alphabetical order in the standard layout. You may need to move comments to more appropriate places.

The standard format provides a good starting point for further tidying. Most package authors use a version of Emacs (on Unix or Windows) to edit R code, using the ESS[S] mode of the ESS Emacs package. See [Appendix B \[R coding standards\], page 105](#) for style options within the ESS[S] mode recommended for the source code of R itself.

3.2 Profiling R code

It is possible to profile R code on Windows and most¹ Unix-like versions of R.

The command `Rprof` is used to control profiling, and its help page can be consulted for full details. Profiling works by recording at fixed intervals² (by default every 20 msecs) which R function is being used, and recording the results in a file (default ‘`Rprof.out`’ in the working directory). Then the function `summaryRprof` or the command-line utility `R CMD Rprof Rprof.out` can be used to summarize the activity.

As an example, consider the following code (from Venables & Ripley, 2002).

¹ R has to be built to enable this, but the option ‘`--enable-R-profiling`’ is the default.

² For Unix-alikes these are intervals of CPU time, and for Windows of elapsed time.

```

library(MASS); library(boot)
storm.fm <- nls(Time ~ b*Viscosity/(Wt - c), stormer,
               start = c(b=29.401, c=2.2183))
st <- cbind(stormer, fit=fitted(storm.fm))
storm.bf <- function(rs, i) {
  st$Time <- st$fit + rs[i]
  tmp <- nls(Time ~ (b * Viscosity)/(Wt - c), st,
             start = coef(storm.fm))
  tmp$m$getAllPars()
}
rs <- scale(resid(storm.fm), scale = FALSE) # remove the mean
Rprof("boot.out")
storm.boot <- boot(rs, storm.bf, R = 4999) # pretty slow
Rprof(NULL)

```

Having run this we can summarize the results by

```
R CMD Rprof boot.out
```

```
Each sample represents 0.02 seconds.
Total run time: 80.74 seconds.
```

```
Total seconds: time spent in function and callees.
Self seconds: time spent in function alone.
```

%	total	%	self	
total	seconds	self	seconds	name
100.00	80.74	0.22	0.18	"boot"
99.65	80.46	1.19	0.96	"statistic"
96.33	77.78	2.68	2.16	"nls"
50.21	40.54	1.54	1.24	"<Anonymous>"
47.11	38.04	1.83	1.48	".Call"
23.06	18.62	2.43	1.96	"eval"
19.87	16.04	0.67	0.54	"as.list"
18.97	15.32	0.64	0.52	"switch"
17.88	14.44	0.47	0.38	"model.frame"
17.41	14.06	1.73	1.40	"model.frame.default"
17.41	14.06	2.80	2.26	"nlsModel"
15.43	12.46	1.88	1.52	"qr.qty"
13.40	10.82	3.07	2.48	"assign"
12.73	10.28	2.33	1.88	"storage.mode<-"
12.34	9.96	1.81	1.46	"qr.coef"
10.13	8.18	5.42	4.38	"paste"
...				
%	self	%	total	
self	seconds	total	seconds	name
5.42	4.38	10.13	8.18	"paste"
3.37	2.72	6.71	5.42	"as.integer"
3.29	2.66	5.00	4.04	"as.double"
3.20	2.58	4.29	3.46	"seq.default"
3.07	2.48	13.40	10.82	"assign"
2.92	2.36	5.95	4.80	"names"
2.80	2.26	17.41	14.06	"nlsModel"
2.68	2.16	96.33	77.78	"nls"
2.53	2.04	2.53	2.04	".Fortran"
2.43	1.96	23.06	18.62	"eval"
2.33	1.88	12.73	10.28	"storage.mode<-"
...				

This often produces surprising results and can be used to identify bottlenecks or pieces of R code that could benefit from being replaced by compiled code.

R CMD Rprof uses a Perl script that may be a little faster than `summaryRprof` for large files. On the other hand `summaryRprof` does not require Perl and provides the results as an R object.

Two warnings: profiling does impose a small performance penalty, and the output files can be very large if long runs are profiled.

Profiling short runs can sometimes give misleading results. R from time to time performs *garbage collection* to reclaim unused memory, and this takes an appreciable amount of time which profiling will charge to whichever function happens to provoke it. It may be useful to compare profiling code immediately after a call to `gc()` with a profiling run without a preceding call to `gc`.

3.3 Profiling compiled code

Profiling compiled code is highly system-specific, but this section contains some hints gleaned from various R users. Some methods need to be different for a compiled executable and for dynamic/shared libraries/objects as used by R packages. We know of no good way to profile DLLs on Windows.

3.3.1 Linux

Options include using `sprof` for a shared object, and `oprofile` (see <http://oprofile.sourceforge.net/>) for any executable or shared object.

3.3.1.1 sprof

You can select shared objects to be profiled with `sprof` by setting the environment variable `LD_PROFILE`. For example

```
% setenv LD_PROFILE /path/to/R_HOME/library/stats/libs/stats.so
R
... run the boot example
% sprof /path/to/R_HOME/library/stats/libs/stats.so \
  /var/tmp/path/to/R_HOME/library/stats/libs/stats.so.profile
```

Flat profile:

Each sample counts as 0.01 seconds.

%	cumulative	self		self	total	
time	seconds	seconds	calls	us/call	us/call	name
76.19	0.32	0.32	0	0.00		numeric_deriv
16.67	0.39	0.07	0	0.00		nls_iter
7.14	0.42	0.03	0	0.00		getListElement

```
rm /path/to/R_HOME/library/stats/libs/stats.so.profile
... to clean up ...
```

It is possible that root access is needed to create the directories used for the profile data.

3.3.1.2 oprofile

`oprofile` works by running a daemon which collects information. The daemon must be started as root, e.g.

```
% su
% opcontrol --no-vmlinux
% opcontrol --start
% exit
```

Then as a user

```
% R
```

```

... run the boot example
% opcontrol --dump
% oprofile -l /path/to/R_HOME/library/stats/libs/stats.so
...
samples %      symbol name
1623    75.5939  anonymous symbol from section .plt
349     16.2552  numeric_deriv
113     5.2632   nls_iter
62      2.8878   getListElement
% oprofile -l /path/to/R_HOME/bin/exec/R
...
samples %      symbol name
76052   11.9912  Rf_eval
54670   8.6198   Rf_findVarInFrame3
37814   5.9622   Rf_allocVector
31489   4.9649   Rf_duplicate
28221   4.4496   Rf_protect
26485   4.1759   Rf_cons
23650   3.7289   Rf_matchArgs
21088   3.3250   Rf_findFun
19995   3.1526   findVarLocInFrame
14871   2.3447   Rf_evalList
13794   2.1749   R_Newhashpjw
13522   2.1320   R_gc_internal
...

```

Shutting down the profiler and clearing the records needs to be done as root. You can use `opannotate` to annotate the source code with the times spent in each section, if the appropriate source code was compiled with debugging support.

3.3.2 Solaris

On 64-bit (only) Solaris, the standard profiling tool `gprof` collects information from shared libraries compiled with ‘`-pg`’.

3.3.3 MacOS X

Developers have recommended `sample` (or `Sampler.app`, which is a GUI version) and `Shark` (see <http://developer.apple.com/tools/sharkoptimize.html> and http://developer.apple.com/tools/shark_optimize.html).

4 Debugging

This chapter covers the debugging of R extensions, starting with the ways to get useful error information and moving on to how to deal with errors that crash R. For those who prefer other styles there are contributed packages such as **debug** on CRAN (described in an article in *R-News* 3/3). (There are notes from 2002 provided by Roger Peng at <http://www.biostat.jhsph.edu/~rpeng/docs/R-debug-tools.pdf> which provide complementary examples to those given here.)

4.1 Browsing

Most of the R-level debugging facilities are based around the built-in browser. This can be used directly by inserting a call to `browser()` into the code of a function (for example, using `fix(my_function)`). When code execution reaches that point in the function, control returns to the R console with a special prompt. For example

```
> fix(summary.data.frame) ## insert browser() call after for() loop
> summary(women)
Called from: summary.data.frame(women)
Browse[1]> ls()
[1] "digits" "i"      "lbs"      "lw"      "maxsum" "nm"      "nr"      "nv"
[9] "object" "sms"     "z"
Browse[1]> maxsum
[1] 7
Browse[1]>
      height      weight
Min.   :58.0   Min.   :115.0
1st Qu.:61.5   1st Qu.:124.5
Median :65.0   Median :135.0
Mean   :65.0   Mean   :136.7
3rd Qu.:68.5   3rd Qu.:148.0
Max.   :72.0   Max.   :164.0
> rm(summary.data.frame)
```

At the browser prompt one can enter any R expression, so for example `ls()` lists the objects in the current frame, and entering the name of an object will¹ print it. The following commands are also accepted

- **n**
Enter ‘step-through’ mode. In this mode, hitting return executes the next line of code (more precisely one line and any continuation lines). Typing `c` will continue to the end of the current context, e.g. to the end of the current loop or function.
- **c**
In normal mode, this quits the browser and continues execution, and just return works in the same way. `cont` is a synonym.
- **where**
This prints the call stack. For example


```
> summary(women)
Called from: summary.data.frame(women)
Browse[1]> where
where 1: summary.data.frame(women)
```

¹ With the exceptions of the commands listed below: an object of such a name can be printed *via* an explicit call to `print`.


```
where 2: summary(women)
```

```
Browse[1]>
```

- Q

Quit both the browser and the current expression, and return to the top-level prompt.

Errors in code executed at the browser prompt will normally return control to the browser prompt. Objects can be altered by assignment, and will keep their changed values when the browser is exited. If really necessary, objects can be assigned to the workspace from the browser prompt (by using `<<-` if the name is not already in scope).

4.2 Debugging R code

Suppose your R program gives an error message. The first thing to find out is what R was doing at the time of the error, and the most useful tool is `traceback()`. We suggest that this is run whenever the cause of the error is not immediately obvious. Daily, errors are reported to the R mailing lists as being in some package when `traceback()` would show that the error was being reported by some other package or base R. Here is an example from the regression suite.

```
> success <- c(13,12,11,14,14,11,13,11,12)
> failure <- c(0,0,0,0,0,0,0,2,2)
> resp <- cbind(success, failure)
> predictor <- c(0, 5^(0:7))
> glm(resp ~ 0+predictor, family = binomial(link="log"))
Error: no valid set of coefficients has been found: please supply starting values
> traceback()
3: stop("no valid set of coefficients has been found: please supply
    starting values", call. = FALSE)
2: glm.fit(x = X, y = Y, weights = weights, start = start, etastart = etastart,
    mustart = mustart, offset = offset, family = family, control = control,
    intercept = attr(mt, "intercept") > 0)
1: glm(resp ~ 0 + predictor, family = binomial(link = "log"))
```

The calls to the active frames are given in reverse order (starting with the innermost). So we see the error message comes from an explicit check in `glm.fit`. (`traceback()` shows you all the lines of the function calls, which can be limited by setting option `"deparse.max.lines"`.)

Sometimes the traceback will indicate that the error was detected inside compiled code, for example (from `?nls`)

```
Error in nls(y ~ a + b * x, start = list(a = 0.12345, b = 0.54321), trace = TRUE) :
  step factor 0.000488281 reduced below 'minFactor' of 0.000976563
> traceback()
2: .Call(R_nls_iter, m, ctrl, trace)
1: nls(y ~ a + b * x, start = list(a = 0.12345, b = 0.54321), trace = TRUE)
```

This will be the case if the innermost call is to `.C`, `.Fortran`, `.Call`, `.External` or `.Internal`, but as it is also possible for such code to evaluate R expressions, this need not be the innermost call, as in

```
> traceback()
9: gm(a, b, x)
8: .Call(R_numeric_deriv, expr, theta, rho, dir)
7: numericDeriv(form[[3]], names(ind), env)
6: getRHS()
5: assign("rhs", getRHS(), envir = thisEnv)
4: assign("resid", .swts * (lhs - assign("rhs", getRHS(), envir = thisEnv)),
```

```

        envir = thisEnv)
3: function (newPars)
{
  setPars(newPars)
  assign("resid", .swts * (lhs - assign("rhs", getRHS(), envir = thisEnv)),
        envir = thisEnv)
  assign("dev", sum(resid^2), envir = thisEnv)
  assign("QR", qr(.swts * attr(rhs, "gradient")), envir = thisEnv)
  return(QR$rank < min(dim(QR$qr)))
}(c(-0.00760232418963883, 1.00119632515036))
2: .Call(R_nls_iter, m, ctrl, trace)
1: nls(yeps ~ gm(a, b, x), start = list(a = 0.12345, b = 0.54321))

```

Occasionally `traceback()` does not help, and this can be the case if S4 method dispatch is involved. Consider the following example

```

> xyd <- new("xyloc", x=runif(20), y=runif(20))
Error in as.environment(pkg) : no item called "package:S4nswv"
on the search list
Error in initialize(value, ...) : S language method selection got
an error when called from internal dispatch for function 'initialize'
> traceback()
2: initialize(value, ...)
1: new("xyloc", x = runif(20), y = runif(20))

```

which does not help much, as there is no call to `as.environment` in `initialize` (and the note “called from internal dispatch” tells us so). In this case we searched the R sources for the quoted call, which occurred in only one place, `methods:::asEnvironmentPackage`. So now we knew where the error was occurring. (This was an unusually opaque example.)

The error message

```
evaluation nested too deeply: infinite recursion / options(expressions=)?
```

can be hard to handle with the default value (5000 as from R 2.3.0). Unless you know that there actually is deep recursion going on, it can help to set something like

```
options(expressions=500)
```

and re-run the example showing the error.

Sometimes there is warning that clearly is the precursor to some later error, but it is not obvious where it is coming from. Setting `options(warn = 2)` (which turns warnings into errors) can help here.

Once we have located the error, we have some choices. One way to proceed is to find out more about what was happening at the time of the crash by looking a *post-mortem* dump. To do so, set `options(error=dump.frames)` and run the code again. Then invoke `debugger()` and explore the dump. Continuing our example:

```

> options(error = dump.frames)
> glm(resp ~ 0 + predictor, family = binomial(link = "log"))
Error: no valid set of coefficients has been found: please supply starting values

```

which is the same as before, but an object called `last.dump` has appeared in the workspace. (Such objects can be large, so remove it when it is no longer needed.) We can examine this at a later time by calling the function `debugger`.

```

> debugger()
Message: Error: no valid set of coefficients has been found: please supply starting values
Available environments had calls:
1: glm(resp ~ 0 + predictor, family = binomial(link = "log"))

```

```

2: glm.fit(x = X, y = Y, weights = weights, start = start, etastart = etastart, mus
3: stop("no valid set of coefficients has been found: please supply starting values
Enter an environment number, or 0 to exit Selection:

```

which gives the same sequence of calls as `traceback`, but in outer-first order and with only the first line of the call, truncated to the current width. However, we can now examine in more detail what was happening at the time of the error. Selecting an environment opens the browser in that frame. So we select the function call which spawned the error message, and explore some of the variables (and execute two function calls).

```

Enter an environment number, or 0 to exit Selection: 2
Browsing in the environment with call:
  glm.fit(x = X, y = Y, weights = weights, start = start, etas
Called from: debugger.look(ind)
Browse[1]> ls()
 [1] "aic"          "boundary"    "coefold"     "control"     "conv"
 [6] "dev"          "dev.resids"  "devold"      "EMPTY"       "eta"
[11] "etastart"     "family"     "fit"         "good"        "intercept"
[16] "iter"         "linkinv"     "mu"          "mu.eta"      "mu.eta.val"
[21] "mustart"      "n"           "ngoodobs"    "nobs"        "nvars"
[26] "offset"       "start"       "valideta"    "validmu"     "variance"
[31] "varmu"        "w"           "weights"     "x"           "xnames"
[36] "y"            "ynames"      "z"
Browse[1]> eta
      1          2          3          4          5
0.000000e+00 -2.235357e-06 -1.117679e-05 -5.588393e-05 -2.794197e-04
      6          7          8          9
-1.397098e-03 -6.985492e-03 -3.492746e-02 -1.746373e-01
Browse[1]> valideta(eta)
[1] TRUE
Browse[1]> mu
      1          2          3          4          5          6          7          8
1.0000000 0.9999978 0.9999888 0.9999441 0.9997206 0.9986039 0.9930389 0.9656755
      9
0.8397616
Browse[1]> validmu(mu)
[1] FALSE
Browse[1]> c
Available environments had calls:
1: glm(resp ~ 0 + predictor, family = binomial(link = "log"))
2: glm.fit(x = X, y = Y, weights = weights, start = start, etastart = etastart
3: stop("no valid set of coefficients has been found: please supply starting v

Enter an environment number, or 0 to exit Selection: 0
> rm(last.dump)

```

Because `last.dump` can be looked at later or even in another R session, post-mortem debugging is possible even for batch usage of R. We do need to arrange for the dump to be saved: this can be done either using the command-line flag `--save` to save the workspace at the end of the run, or via a setting such as

```
> options(error = quote({dump.frames(to.file=TRUE); q()}))
```

See the help on `dump.frames` for further options and a worked example.

An alternative error action is to use the function `recover()`:

```

> options(error = recover)
> glm(resp ~ 0 + predictor, family = binomial(link = "log"))
Error: no valid set of coefficients has been found: please supply starting values

Enter a frame number, or 0 to exit

1: glm(resp ~ 0 + predictor, family = binomial(link = "log"))
2: glm.fit(x = X, y = Y, weights = weights, start = start, etastart = etastart)

```

Selection:

which is very similar to `dump.frames`. However, we can examine the state of the program directly, without dumping and re-loading the dump. As its help page says, `recover` can be routinely used as the error action in place of `dump.calls` (used by `traceback`) and `dump.frames`, since it behaves like `dump.frames` in non-interactive use.

Post-mortem debugging is good for finding out exactly what went wrong, but not necessarily why. An alternative approach is to take a closer look at what was happening just before the error, and a good way to do that is to use `debug`. This inserts a call to the browser at the beginning of the function, starting in step-through mode. So in our example we could use

```

> debug(glm.fit)
> glm(resp ~ 0 + predictor, family = binomial(link = "log"))
debugging in: glm.fit(x = X, y = Y, weights = weights, start = start, etastart = etastart,
  mustart = mustart, offset = offset, family = family, control = control,
  intercept = attr(mt, "intercept") > 0)
debug: {
## lists the whole function
Browse[1]>
debug: x <- as.matrix(x)
...
Browse[1]> start
[1] -2.235357e-06
debug: eta <- drop(x %*% start)
Browse[1]> eta
      1      2      3      4      5
0.000000e+00 -2.235357e-06 -1.117679e-05 -5.588393e-05 -2.794197e-04
      6      7      8      9
-1.397098e-03 -6.985492e-03 -3.492746e-02 -1.746373e-01
Browse[1]>
debug: mu <- linkinv(eta <- eta + offset)
Browse[1]> mu
      1      2      3      4      5      6      7      8
1.0000000 0.9999978 0.9999888 0.9999441 0.9997206 0.9986039 0.9930389 0.9656755
      9
0.8397616

```

(The prompt `Browse[1]>` indicates that this is the first level of browsing: it is possible to step into another function that is itself being debugged or contains a call to `browser()`.)

`debug` can be used for hidden functions and S3 methods by e.g. `debug(stats:::predict.Arima)`. (It cannot be used for S4 methods, but an alternative is given on the help page for `debug`.) Sometimes you want to debug a function defined inside another function, e.g. the function `arimafn` defined inside `arima`. To do so, set `debug` on the outer function (here `arima`) and step through it until the inner function has been defined.

Then call `debug` on the inner function (and use `c` to get out of step-through mode in the outer function).

To remove debugging of a function, call `undebug` with the argument previously given to `debug`; debugging otherwise lasts for the rest of the R session (or until the function is edited or otherwise replaced).

`trace` can be used to temporarily insert debugging code into a function, for example to insert a call to `browser()` just before the point of the error. To return to our running example

```
## first get a numbered listing of the expressions of the function
> page(as.list(body(glm.fit)), method="print")
> trace(glm.fit, browser, at=22)
Tracing function "glm.fit" in package "stats"
[1] "glm.fit"
> glm(resp ~ 0 + predictor, family = binomial(link = "log"))
Tracing glm.fit(x = X, y = Y, weights = weights, start = start,
  etastart = etastart, .... step 22
Called from: eval(expr, envir, enclos)
Browse[1]> n
## and single-step from here.
> untrace(glm.fit)
```

For your own functions, it may be as easy to use `fix` to insert temporary code, but `trace` can help with functions in a namespace (as can `fixInNamespace`). Alternatively, use `trace(,edit=TRUE)` to insert code visually.

4.3 Using gctorture and valgrind

Errors in memory allocation and reading/writing outside arrays are very common causes of crashes (e.g., segfaults) on some machines. Often the crash appears long after the invalid memory access: in particular damage to the structures which R itself has allocated may only become apparent at the next garbage collection (or even at later garbage collections after objects have been deleted).

4.3.1 Using gctorture

We can help to detect memory problems earlier by running garbage collection as often as possible. This is achieved by `gctorture(TRUE)`, which as described on its help page

Provokes garbage collection on (nearly) every memory allocation. Intended to ferret out memory protection bugs. Also makes R run *very* slowly, unfortunately.

The reference to ‘memory protection’ is to missing C-level calls to `PROTECT/UNPROTECT` (see [Section 5.7.1 \[Garbage Collection\]](#), page 60) which if missing allow R objects to be garbage-collected when they are still in use. But it can also help with other memory-related errors.

Normally running under `gctorture(TRUE)` will just produce a crash earlier in the R program, hopefully close to the actual cause. See the next section for how to decipher such crashes.

It is possible to run all the examples, tests and vignettes covered by R CMD `check` under `gctorture(TRUE)` by using the option ‘`--use-gct`’.

4.3.2 Using valgrind

If you have access to Linux on an `ix86`, `x86_64` or `ppc32` platform you can use `valgrind` (<http://www.valgrind.org/>, pronounced to rhyme with ‘tinned’) to check for possible problems. To run some examples under `valgrind` use something like

```
R -d valgrind --vanilla < mypkg-Ex.R
R -d "valgrind --tool=memcheck --leak-check=full" --vanilla < mypkg-Ex.R
```

where ‘mypkg-Ex.R’ is a set of examples, e.g. the file created in ‘mypkg.Rcheck’ by R CMD check. Occasionally this reports memory reads of ‘uninitialised values’ that are the result of compiler optimization, so can be worth checking under an unoptimized compile. We know there will be some small memory leaks from `readline` and R itself – these are memory areas that are in use right up to the end of the R session. Expect this to run around 20x slower than without `valgrind`, and in some cases even slower than that. Current versions² of `valgrind` are not happy with many optimized BLASes that use cpu-specific instructions (3D now, SSE, SSE2, SSE3 and similar) so you may need to build a version of R specifically to use with `valgrind`.

On platforms supported by `valgrind` you can build a version of R with extra instrumentation to help `valgrind` detect errors in the use of memory allocated from the R heap. The configure option is ‘--with-valgrind-instrumentation=level’, where `level` is 0, 1, or 2. Level 0 is the default and does not add anything. Level 1 will detect use of uninitialised memory and has little impact on speed. Level 2 will detect many other memory use bugs but makes R much slower when running under `valgrind`. Using this in conjunction with `gctorture` can be even more effective (and even slower).

It is possible to run all the examples, tests and vignettes covered by R CMD check under `valgrind` by using the option ‘--use-valgrind’. If you do this you will need to select the `valgrind` options some other way, for example by having a ‘~/valgrindrc’ file containing

```
--tool=memcheck
--memcheck:leak-check=full
```

or setting the environment variable `VALGRIND_OPTS`.

4.4 Debugging compiled code

Sooner or later programmers will be faced with the need to debug compiled code loaded into R. This section is geared to platforms using `gdb` with code compiled by `gcc`, but similar things are possible with front-ends to `gdb` such as `ddd` and `insight`, and other debuggers such as Sun’s `dbx`.

Consider first ‘crashes’, that is when R terminated unexpectedly with an illegal memory access (a ‘segfault’ or ‘bus error’), illegal instruction or similar. As from R 2.3.0 Unix-alike versions of R use a signal handler which aims to give some basic information. For example

```
*** caught segfault ***
address 0x20000028, cause 'memory not mapped'

Traceback:
 1: .identC(class1[[1]], class2)
 2: possibleExtends(class(sloti), classi, ClassDef2 = getClassDef(classi,
where = where))
 3: validObject(t(cu))
 4: stopifnot(validObject(cu <- as(tu, "dtCMatrix")), validObject(t(cu)),
validObject(t(tu)))

Possible actions:
 1: abort (with core dump)
 2: normal R exit
 3: exit R without saving workspace
 4: exit R saving workspace
Selection: 3
```

² Although this is supposed to have been improved, `valgrind` 3.1.0 still aborts using optimized BLASes on an Opteron.

Since the R process may be damaged, the only really safe option is the first.

Another cause of a ‘crash’ is to overrun the C stack. R tries to track that in its own code, but it may happen in third-party compiled code. For modern POSIX-compliant OSes we can safely catch that and return to the top-level prompt.

```
> .C("aaa")
Error: segfault from C stack overflow
>
```

However, C stack overflows are fatal under Windows and normally defeat attempts at debugging on that platform.

If you have a crash which gives a core dump you can use something like

```
gdb /path/to/R/bin/exec/R core.12345
```

to examine the core dump. If core dumps are disabled or to catch errors that do not generate a dump one can run R directly under a debugger by for example

```
$ R -d gdb --vanilla
...
gdb> run
```

at which point R will run normally, and hopefully the debugger will catch the error and return to its prompt. This can also be used to catch infinite loops or interrupt very long-running code. For a simple example

```
> for(i in 1:1e7) x <- rnorm(100)
[hit Ctrl-C]
Program received signal SIGINT, Interrupt.
0x00397682 in _int_free () from /lib/tls/libc.so.6
(gdb) where
#0 0x00397682 in _int_free () from /lib/tls/libc.so.6
#1 0x00397eba in free () from /lib/tls/libc.so.6
#2 0xb7cf2551 in R_gc_internal (size_needed=313)
    at /users/ripley/R/svn/R-devel/src/main/memory.c:743
#3 0xb7cf3617 in Rf_allocVector (type=13, length=626)
    at /users/ripley/R/svn/R-devel/src/main/memory.c:1906
#4 0xb7c3f6d3 in PutRNGstate ()
    at /users/ripley/R/svn/R-devel/src/main/RNG.c:351
#5 0xb7d6c0a5 in do_random2 (call=0x94bf7d4, op=0x92580e8, args=0x9698f98,
    rho=0x9698f28) at /users/ripley/R/svn/R-devel/src/main/random.c:183
...
```

Some “tricks” are worth knowing.

4.4.1 Finding entry points in dynamically loaded code

Under most compilation environments, compiled code dynamically loaded into R cannot have breakpoints set within it until it is loaded. To use a symbolic debugger on such dynamically loaded code under Unix-alikes use

- Call the debugger on the R executable, for example by `R -d gdb`.
- Start R.
- At the R prompt, use `dyn.load` or `library` to load your shared object.
- Send an interrupt signal. This will put you back to the debugger prompt.
- Set the breakpoints in your code.
- Continue execution of R by typing `signal 0(RET)`.

Under Windows signals may not be able to be used, and if so the procedure is more complicated. See the rw-FAQ and <http://www.stats.uwo.ca/faculty/murdoch/software/debuggingR/gdb.shtml>

4.4.2 Inspecting R objects when debugging

The key to inspecting R objects from compiled code is the function `PrintValue(SEXP s)` which uses the normal R printing mechanisms to print the R object pointed to by `s`, or the safer version `R_PV(SEXP s)` which will only print ‘objects’.

One way to make use of `PrintValue` is to insert suitable calls into the code to be debugged.

Another way is to call `R_PV` from the symbolic debugger. (`PrintValue` is hidden as `Rf_PrintValue`.) For example, from `gdb` we can use

```
(gdb) p R_PV(ab)
```

using the object `ab` from the convolution example, if we have placed a suitable breakpoint in the convolution C code.

To examine an arbitrary R object we need to work a little harder. For example, let

```
R> DF <- data.frame(a = 1:3, b = 4:6)
```

By setting a breakpoint at `do_get` and typing `get("DF")` at the R prompt, one can find out the address in memory of `DF`, for example

```
Value returned is $1 = (SEXP) 0x40583e1c
(gdb) p *$1
$2 = {
  sxpinfo = {type = 19, obj = 1, named = 1, gp = 0,
    mark = 0, debug = 0, trace = 0, = 0},
  attrib = 0x40583e80,
  u = {
    vecsxp = {
      length = 2,
      type = {c = 0x40634700 "O>X0D>X00>X0", i = 0x40634700,
        f = 0x40634700, z = 0x40634700, s = 0x40634700},
      truelength = 1075851272,
    },
    primsxp = {offset = 2},
    symsxp = {pname = 0x2, value = 0x40634700, internal = 0x40203008},
    listsxp = {carval = 0x2, cdrval = 0x40634700, tagval = 0x40203008},
    envsxp = {frame = 0x2, enclos = 0x40634700},
    closxp = {formals = 0x2, body = 0x40634700, env = 0x40203008},
    promsxp = {value = 0x2, expr = 0x40634700, env = 0x40203008}
  }
}
```

(Debugger output reformatted for better legibility).

Using `R_PV()` one can “inspect” the values of the various elements of the SEXP, for example,

```
(gdb) p R_PV($1->attrib)
$names
[1] "a" "b"

$row.names
[1] "1" "2" "3"

$class
[1] "data.frame"

$3 = void
```

To find out where exactly the corresponding information is stored, one needs to go “deeper”:

```
(gdb) set $a = $1->attrib
(gdb) p $a->u.listsxp.tagval->u.symsxp.pname->u.vecsxp.type.c
$4 = 0x405d40e8 "names"
(gdb) p $a->u.listsxp.carval->u.vecsxp.type.s[1]->u.vecsxp.type.c
$5 = 0x40634378 "b"
(gdb) p $1->u.vecsxp.type.s[0]->u.vecsxp.type.i[0]
$6 = 1
(gdb) p $1->u.vecsxp.type.s[1]->u.vecsxp.type.i[1]
$7 = 5
```

5 System and foreign language interfaces

5.1 Operating system access

Access to operating system functions is via the R function `system`. The details will differ by platform (see the on-line help), and about all that can safely be assumed is that the first argument will be a string `command` that will be passed for execution (not necessarily by a shell) and the second argument will be `internal` which if true will collect the output of the command into an R character vector.

The function `system.time` is available for timing (although the information available may be limited on non-Unix-like platforms).

5.2 Interface functions `.C` and `.Fortran`

These two functions provide a standard interface to compiled code that has been linked into R, either at build time or via `dyn.load` (see [Section 5.3 \[dyn.load and dyn.unload\]](#), page 54). They are primarily intended for compiled C and FORTRAN code respectively, but the `.C` function can be used with other languages which can generate C interfaces, for example C++ (see [Section 5.6 \[Interfacing C++ code\]](#), page 57).

The first argument to each function is a character string given the symbol name as known to C or FORTRAN, that is the function or subroutine name. (That the symbol is loaded can be tested by, for example, `is.loaded("cg")`: it is no longer necessary or correct to use `symbol.For.`) (Note that the underscore is not a valid character in a FORTRAN 77 symbol name, and on many platforms `.Fortran` will not correctly translate names containing underscores.)

There can be up to 65 further arguments giving R objects to be passed to compiled code. Normally these are copied before being passed in, and copied again to an R list object when the compiled code returns. If the arguments are given names, these are used as names for the components in the returned list object (but not passed to the compiled code).

The following table gives the mapping between the modes of R vectors and the types of arguments to a C function or FORTRAN subroutine.

R storage mode	C type	FORTRAN type
logical	int *	INTEGER
integer	int *	INTEGER
double	double *	DOUBLE PRECISION
complex	Rcomplex *	DOUBLE COMPLEX
character	char **	CHARACTER*255
raw	char *	none

Do please note the first two. On the 64-bit Unix/Linux platforms, `long` is 64-bit whereas `int` and `INTEGER` are 32-bit. Code ported from S-PLUS (which uses `long *` for `logical` and `integer`) will not work on all 64-bit platforms (although it may appear to work on some). Note too that if your compiled code is a mixture of C functions and FORTRAN subprograms the argument types must match as given in the table above.

C type `Rcomplex` is a structure with `double` members `r` and `i` defined in the header file `'R_ext/Complex.h'` included by `'R.h'`. Only a single character string can be passed to or from FORTRAN, and the success of this is compiler-dependent. Other R objects can be passed to `.C`, but it is better to use one of the other interfaces. An exception is passing an R function for use with `call_R`, when the object can be handled as `void *` en route to `call_R`, but even there `.Call` is to be preferred. Similarly, passing an R list as an argument to a C routine should be done using the `.Call` interface. If one does use the `.C` function to pass a list as an argument,

it is visible to the routine as an array in C of `SEXP` types (i.e., `SEXP *`). The elements of the array correspond directly to the elements of the R list. However, this array must be treated as read-only and one must not assign values to its elements within the C routine. Doing so bypasses R's memory management facilities and will corrupt the object and the R session.

It is possible to pass numeric vectors of storage mode `double` to C as `float *` or FORTRAN as `REAL` by setting the attribute `Csingle`, most conveniently by using the R functions `as.single`, `single` or `storage.mode`. This is intended only to be used to aid interfacing to existing C or FORTRAN code.

Unless formal argument `NAOK` is true, all the other arguments are checked for missing values `NA` and for the IEEE special values `NaN`, `Inf` and `-Inf`, and the presence of any of these generates an error. If it is true, these values are passed unchecked.

Argument `DUP` can be used to suppress copying. It is dangerous: see the on-line help for arguments against its use. It is not possible to pass numeric vectors as `float *` or `REAL` if `DUP=FALSE`.

Argument `PACKAGE` confines the search for the symbol name to a specific shared object (or use `"base"` for code compiled into R). Its use is highly desirable, as there is no way to avoid two package writers using the same symbol name, and such name clashes are normally sufficient to cause R to crash. (If it is not present and the call is from the body of a function defined in a package with a namespace, the shared object loaded by the first (if any) `useDynLib` directive will be used.)

For `.C` only you can specify an `ENCODING` argument: this requests that (unless `DUP = FALSE`) character vectors be re-encoded to the requested encoding before being passed in, and re-encoded from the requested encoding when passed back. Note that encoding names are not standardized, and not all R builds support re-encoding. (The argument is ignored with a warning if re-encoding is not supported at all: R code can test for this *via* `capabilities("iconv")`.) But this can be useful to allow code to work in a UTF-8 locale by specifying `ENCODING = "latin1"`.

Note that the compiled code should not return anything except through its arguments: C functions should be of type `void` and FORTRAN subprograms should be subroutines.

To fix ideas, let us consider a very simple example which convolves two finite sequences. (This is hard to do fast in interpreted R code, but easy in C code.) We could do this using `.C` by

```
void convolve(double *a, int *na, double *b, int *nb, double *ab)
{
    int i, j, nab = *na + *nb - 1;

    for(i = 0; i < nab; i++)
        ab[i] = 0.0;
    for(i = 0; i < *na; i++)
        for(j = 0; j < *nb; j++)
            ab[i + j] += a[i] * b[j];
}
```

called from R by

```
conv <- function(a, b)
.C("convolve",
  as.double(a),
  as.integer(length(a)),
  as.double(b),
  as.integer(length(b)),
  ab = double(length(a) + length(b) - 1))$ab
```

Note that we take care to coerce all the arguments to the correct R storage mode before calling `.C`; mistakes in matching the types can lead to wrong results or hard-to-catch errors.

5.3 `dyn.load` and `dyn.unload`

Compiled code to be used with R is loaded as a shared object (Unix, see [Section 5.5 \[Creating shared objects\]](#), page 57 for more information) or DLL (Windows).

The shared object/DLL is loaded by `dyn.load` and unloaded by `dyn.unload`. Unloading is not normally necessary, but it is needed to allow the DLL to be re-built on some platforms, including Windows.

The first argument to both functions is a character string giving the path to the object. Programmers should not assume a specific file extension for the object/DLL (such as `.so`) but use a construction like

```
file.path(path1, path2, paste("mylib", .Platform$dynlib.ext, sep=""))
```

for platform independence. On Unix systems the path supplied to `dyn.load` can be an absolute path, one relative to the current directory or, if it starts with `~`, relative to the user's home directory.

Loading is most often done via a call to `library.dynam` in the `.First.lib` function of a package. This has the form

```
library.dynam("libname", package, lib.loc)
```

where `libname` is the object/DLL name *with the extension omitted*. Note that the first argument, `chname`, should **not** be `package` since this will not work if the package is installed under another name (as it will be with a versioned install).

Under some Unix systems there is a choice of how the symbols are resolved when the object is loaded, governed by the arguments `local` and `now`. Only use these if really necessary: in particular using `now=FALSE` and then calling an unresolved symbol will terminate R unceremoniously.

R provides a way of executing some code automatically when a object/DLL is either loaded or unloaded. This can be used, for example, to register native routines with R's dynamic symbol mechanism, initialize some data in the native code, or initialize a third party library. On loading a DLL, R will look for a routine within that DLL named `R_init_lib` where `lib` is the name of the DLL file with the extension removed. For example, in the command

```
library.dynam("mylib", package, lib.loc)
```

R looks for the symbol named `R_init_mylib`. Similarly, when unloading the object, R looks for a routine named `R_unload_lib`, e.g., `R_unload_mylib`. In either case, if the routine is present, R will invoke it and pass it a single argument describing the DLL. This is a value of type `DllInfo` which is defined in the `'Rdynload.h'` file in the `'R_ext'` directory.

The following example shows templates for the initialization and unload routines for the `mylib` DLL.

```

#include <R.h>
#include <Rinternals.h>
#include <R_ext/Rdynload.h>

void
R_init_mylib(DllInfo *info)
{
    /* Register routines, allocate resources. */
}

void
R_unload_mylib(DllInfo *info)
{
    /* Release resources. */
}

```

If a shared object/DLL is loaded more than once the most recent version is used. More generally, if the same symbol name appears in several libraries, the most recently loaded occurrence is used. The `PACKAGE` argument provides a good way to avoid any ambiguity in which occurrence is meant.

5.4 Registering native routines

In calls to `.C`, `.Call`, `.Fortran` and `.External`, R must locate the specified native routine by looking in the appropriate shared object/DLL. By default, R uses the operating system-specific dynamic loader to lookup the symbol. Alternatively, the author of the DLL can explicitly register routines with R and use a single, platform-independent mechanism for finding the routines in the DLL. One can use this registration mechanism to provide additional information about a routine, including the number and type of the arguments, and also make it available to S programmers under a different name. In the future, registration may be used to implement a form of “secure” or limited native access.

To register routines with R, one calls the C routine `R_registerRoutines`. This is typically done when the DLL is first loaded within the initialization routine `R_init_dll name` described in [Section 5.3 \[dyn.load and dyn.unload\], page 54](#). `R_registerRoutines` takes 5 arguments. The first is the `DllInfo` object passed by R to the initialization routine. This is where R stores the information about the methods. The remaining 4 arguments are arrays describing the routines for each of the 4 different interfaces: `.C`, `.Call`, `.Fortran` and `.External`. Each argument is a NULL-terminated array of the element types given in the following table:

<code>.C</code>	<code>R_CMethodDef</code>
<code>.Call</code>	<code>R_CallMethodDef</code>
<code>.Fortran</code>	<code>R_FortranMethodDef</code>
<code>.External</code>	<code>R_ExternalMethodDef</code>

Currently, the `R_ExternalMethodDef` is the same as `R_CallMethodDef` type and contains fields for the name of the routine by which it can be accessed in R, a pointer to the actual native symbol (i.e., the routine itself), and the number of arguments the routine expects. For routines with a variable number of arguments invoked via the `.External` interface, one specifies `-1` for the number of arguments which tells R not to check the actual number passed. For example, if we had a routine named `myCall` defined as

```
SEXP myCall(SEXP a, SEXP b, SEXP c);
```

we would describe this as

```
R_CallMethodDef callMethods[] = {
    {"myCall", &myCall, 3},

```

```

    {NULL, NULL, 0}
};

```

along with any other routines for the `.Call` interface.

Routines for use with the `.C` and `.Fortran` interfaces are described with similar data structures, but which have two additional fields for describing the type and “style” of each argument. Each of these can be omitted. However, if specified, each should be an array with the same number of elements as the number of parameters for the routine. The types array should contain the `SEXP` types describing the expected type of the argument. (Technically, the elements of the types array are of type `R_NativePrimitiveArgType` which is just an unsigned integer.) The R types and corresponding type identifiers are provided in the following table:

numeric	REALSXP
integer	INTSXP
logical	LGLSXP
single	SINGLESXP
character	STRSXP
list	VECSXP

Consider a C routine, `myC`, declared as

```
void myC(double *x, int *n, char **names, int *status);
```

We would register it as

```

R_CMethodDef cMethods[] = {
    {"myC", &myC, 4, {REALSXP, INTSXP, STRSXP, LGLSXP}},
    {NULL, NULL, 0}
};

```

One can also specify whether each argument is used simply as input, or as output, or as both input and output. The style field in the description of a method is used for this. The purpose is to allow R to more efficiently transfer values across the R-C/FORTRAN interface by avoiding copying values when it is not necessary. Typically, one omits this information in the registration data.

Having created the arrays describing each routine, the last step is to actually register them with R. We do this by calling `R_registerRoutines`. For example, if we have the descriptions above for the routines accessed by the `.C` and `.Call` we would use the following code:

```

void
R_init_myLib(DllInfo *info)
{
    R_registerRoutines(info, cMethods, callMethods, NULL, NULL);
}

```

This routine will be invoked when R loads the shared object/DLL named `myLib`. The last two arguments in the call to `R_registerRoutines` are for the routines accessed by `.Fortran` and `.External` interfaces. In our example, these are given as `NULL` since we don't have any routines of these types.

When R unloads a shared object/DLL, any registered routines are automatically removed. There is no (direct) facility for unregistering a symbol.

Examples of registering routines can be found in the different packages in the R source tree (e.g., `stats`). Also, there is a brief, high-level introduction in *R News* (volume 1/3, September 2001, pages 20-23).

Additionally, there are (experimental) tools that can be used to automate the generation of the code to register the routines for a collection of C files. See the `GccTranslationUnit` module on the Omegahat Web site at <http://www.omegahat.org/GccTranslationUnit/> for more information.

5.5 Creating shared objects

Under Unix, shared objects for loading into R can be created using `R CMD SHLIB`. This accepts as arguments a list of files which must be object files (with extension `.o`) or C, C++, or FORTRAN sources (with extensions `.c`, `.cc` or `.cpp` or `.C`, and `.f` or `.f90` or `.f95`, respectively). See `R CMD SHLIB --help`, or the on-line help for `SHLIB`, for usage information. If compiling the source files does not work “out of the box”, you can specify additional flags by setting some of the variables `PKG_CPPFLAGS` (for the C preprocessor, typically `-I` flags), `PKG_CFLAGS`, `PKG_CXXFLAGS`, and `PKG_FFLAGS` (for the C, C++, and FORTRAN compilers, respectively) in the file `Makevars` in the compilation directory, or write a `Makefile` in the compilation directory containing the rules required (or, of course, create the object files directly from the command line). Similarly, variable `PKG_LIBS` in `Makevars` can be used for additional `-l` and `-L` flags to be passed to the linker when building the shared object. (On most systems the `PKG_*` flags can be set in the environment instead of in `Makevars`.)

Flags which are set in file `etc/Makeconf` can be overridden by the environment variable `MAKEFLAGS` (at least for systems using GNU `make`), as in (Bourne shell syntax)

```
MAKEFLAGS="CFLAGS=-O3" R CMD SHLIB *.c
```

or by using a `Makefile`.

Note that as `R CMD SHLIB` uses `Make`, it will not remake a shared object just because the flags have changed, and if `test.c` and `test.f` both exist in the current directory

```
R CMD SHLIB test.f
```

will compile `test.c`!

If an add-on package `pkg` contains C, C++, or FORTRAN code in its `src` subdirectory, `R CMD INSTALL` creates a shared object (for loading into R in the `.First.lib` or `.onLoad` function of the package) either automatically using the above `R CMD SHLIB` mechanism, or using `Make` if directory `src` contains a `Makefile`. In both cases, if file `Makevars` exists it is read first when invoking `make`. If a `Makefile` is really needed or provided, if FORTRAN is involved it needs to ensure that the shared object created is linked against all FORTRAN 77 intrinsic and run-time libraries that R was linked against; Make variable `FLIBS` contains this information.

In Windows the same command works, but files `Makevars.win` or `Makefile.win` are used in preference to `Makevars` or `Makefile` if they exist. (This does need the files in the R binary Windows distribution for installing source packages to be installed.) For details of building DLLs with a variety of compilers, see `README.packages` and <http://www.stats.uwo.ca/faculty/murdoch/software/compilingDLLs/>

5.6 Interfacing C++ code

Suppose we have the following hypothetical C++ library, consisting of the two files `X.hh` and `X.cc`, and implementing the two classes `X` and `Y` which we want to use in R.

```
// X.hh

class X {
public: X (); ~X ();
};

class Y {
public: Y (); ~Y ();
};
```

```
// X.cc

#include <iostream>
#include "X.hh"

static Y y;

X::X() { std::cout << "constructor X" << std::endl; }
X::~X() { std::cout << "destructor X" << std::endl; }
Y::Y() { std::cout << "constructor Y" << std::endl; }
Y::~Y() { std::cout << "destructor Y" << std::endl; }
```

To use with R, the only thing we have to do is writing a wrapper function and ensuring that the function is enclosed in

```
extern "C" {

}
```

For example,

```
// X_main.cc:

#include "X.hh"

extern "C" {

void X_main () {
    X x;
}

} // extern "C"
```

Compiling and linking should be done with the C++ compiler-linker (rather than the C compiler-linker or the linker itself); otherwise, the C++ initialization code (and hence the constructor of the static variable Y) are not called. On a properly configured Unix or Windows system, one can simply use

```
R CMD SHLIB X.cc X_main.cc
```

to create the shared object, typically 'X.so' (the file name extension may be different on your platform). Now starting R yields

```
R : Copyright 2000, The R Development Core Team
Version 1.1.0 Under development (unstable) (April 14, 2000)
...
Type      "q()" to quit R.

R> dyn.load(paste("X", .Platform$dynlib.ext, sep = ""))
constructor Y
R> .C("X_main")
constructor X
destructor X
list()
R> q()
Save workspace image? [y/n/c]: y
destructor Y
```

The R for Windows FAQ ([‘rw-FAQ’](#)) contains details of how to compile this example under various Windows compilers.

Using C++ iostreams, as in this example, is best avoided. There is no guarantee that the output will appear in the R console, and indeed it will not on the R for Windows console. Use R code or the C entry points (see [Section 6.5 \[Printing\]](#), [page 81](#)) for all I/O if at all possible.

R header files are C and not C++ header files and should be included within an `extern "C"` block. (This is not always necessary but helps in reading C++ code.)

5.7 Handling R objects in C

Using C code to speed up the execution of an R function is often very fruitful. Traditionally this has been done via the `.C` function in R. One restriction of this interface is that the R objects can not be handled directly in C. This becomes more troublesome when one wishes to call R functions from within the C code. There is a C function provided called `call_R` (also known as `call_S` for compatibility with S) that can do that, but it is cumbersome to use, and the mechanisms documented here are usually simpler to use, as well as more powerful.

If a user really wants to write C code using internal R data structures, then that can be done using the `.Call` and `.External` function. The syntax for the calling function in R in each case is similar to that of `.C`, but the two functions have rather different C interfaces. Generally the `.Call` interface (which is modelled on the interface of the same name in S version 4) is a little simpler to use, but `.External` is a little more general.

A call to `.Call` is very similar to `.C`, for example

```
.Call("convolve2", a, b)
```

The first argument should be a character string giving a C symbol name of code that has already been loaded into R. Up to 65 R objects can be passed as arguments. The C side of the interface is

```
#include <R.h>
#include <Rinternals.h>

SEXP convolve2(SEXP a, SEXP b)
...

```

A call to `.External` is almost identical

```
.External("convolveE", a, b)
```

but the C side of the interface is different, having only one argument

```
#include <R.h>
#include <Rinternals.h>

SEXP convolveE(SEXP args)
...

```

Here `args` is a `LISTSXP`, a Lisp-style list from which the arguments can be extracted.

In each case the R objects are available for manipulation via a set of functions and macros defined in the header file `'Rinternals.h'` or some S4-compatibility macros defined in `'Rdefines.h'`. See [Section 5.8 \[Interface functions `.Call` and `.External`\]](#), [page 67](#) for details on `.Call` and `.External`.

Before you decide to use `.Call` or `.External`, you should look at other alternatives. First, consider working in interpreted R code; if this is fast enough, this is normally the best option. You should also see if using `.C` is enough. If the task to be performed in C is simple enough requiring no call to R, `.C` suffices. The new interfaces are recent additions to S and R, and a great deal of useful code has been written using just `.C` before they were available. The `.Call` and `.External` interfaces allow much more control, but they also impose much greater

responsibilities so need to be used with care. Neither `.Call` nor `.External` copy their arguments. You should treat arguments you receive through these interfaces as read-only.

There are two approaches that can be taken to handling R objects from within C code. The first (historically) is to use the macros and functions that have been used to implement the core parts of R through `.Internal` calls. A public subset of these is defined in the header file `'Rinternals.h'` in the directory `'R_INCLUDE_DIR'` (default `'R_HOME/include'`) that should be available on any R installation.

Another approach is to use R versions of the macros and functions defined for the S version 4 interface `.Call`, which are defined in the header file `'Rdefines.h'`. This is a somewhat simpler approach, and is to be preferred if the code is intended to be shared with S. However, it is less well documented and even less tested. Note too that some idionatic S4 constructions with these macros (such as assigning elements of character vectors or lists) are invalid in R since version 1.2.0.

A substantial amount of R is implemented using the functions and macros described here, so the R source code provides a rich source of examples and “how to do it”: indeed many of the examples here were developed by examining closely R system functions for similar tasks. Do make use of the source code for inspirational examples.

It is necessary to know something about how R objects are handled in C code. All the R objects you will deal with will be handled with the type `SEXP`¹, which is a pointer to a structure with typedef `SEXP`. Think of this structure as a *variant type* that can handle all the usual types of R objects, that is vectors of various modes, functions, environments, language objects and so on. The details are given later in this section, but for most purposes the programmer does not need to know them. Think rather of a model such as that used by Visual Basic, in which R objects are handed around in C code (as they are in interpreted R code) as the variant type, and the appropriate part is extracted for, for example, numerical calculations, only when it is needed. As in interpreted R code, much use is made of coercion to force the variant object to the right type.

5.7.1 Handling the effects of garbage collection

We need to know a little about the way R handles memory allocation. The memory allocated for R objects is not freed by the user; instead, the memory is from time to time *garbage collected*. That is, some or all of the allocated memory not being used is freed. (Prior to R 1.2, objects could be moved, too.)

The R object types are represented by a C structure defined by a typedef `SEXP` in `'Rinternals.h'`. It contains several things among which are pointers to data blocks and to other `SEXP`s. A `SEXP` is simply a pointer to a `SEXP`.

If you create an R object in your C code, you must tell R that you are using the object by using the `PROTECT` macro on a pointer to the object. This tells R that the object is in use so it is not destroyed. Notice that it is the object which is protected, not the pointer variable. It is a common mistake to believe that if you invoked `PROTECT(p)` at some point then `p` is protected from then on, but that is not true once a new object is assigned to `p`.

Protecting an R object automatically protects all the R objects pointed to in the corresponding `SEXP`.

The programmer is solely responsible for housekeeping the calls to `PROTECT`. There is a corresponding macro `UNPROTECT` that takes as argument an `int` giving the number of objects to unprotect when they are no longer needed. The protection mechanism is stack-based, so `UNPROTECT(n)` unprotects the last `n` objects which were protected. The calls to `PROTECT` and

¹ `SEXP` is an acronym for *S*imple *EXP*ression, common in LISP-like language syntaxes.

UNPROTECT must balance when the user's code returns. R will warn about "stack imbalance in .Call" (or .External) if the housekeeping is wrong.

Here is a small example of creating an R numeric vector in C code. First we use the macros in 'Rinternals.h':

```
#include <R.h>
#include <Rinternals.h>

SEXP ab;

....
PROTECT(ab = allocVector(REALSXP, 2));
REAL(ab)[0] = 123.45;
REAL(ab)[1] = 67.89;
UNPROTECT(1);
```

and then those in 'Rdefines.h':

```
#include <R.h>
#include <Rdefines.h>

SEXP ab;

....
PROTECT(ab = NEW_NUMERIC(2));
NUMERIC_POINTER(ab)[0] = 123.45;
NUMERIC_POINTER(ab)[1] = 67.89;
UNPROTECT(1);
```

Now, the reader may ask how the R object could possibly get removed during those manipulations, as it is just our C code that is running. As it happens, we can do without the protection in this example, but in general we do not know (nor want to know) what is hiding behind the R macros and functions we use, and any of them might cause memory to be allocated, hence garbage collection and hence our object `ab` to be removed. It is usually wise to err on the side of caution and assume that any of the R macros and functions might remove the object.

In some cases it is necessary to keep better track of whether protection is really needed. Be particularly aware of situations where a large number of objects are generated. The pointer protection stack has a fixed size (default 10,000) and can become full. It is not a good idea then to just PROTECT everything in sight and UNPROTECT several thousand objects at the end. It will almost invariably be possible to either assign the objects as part of another object (which automatically protects them) or unprotect them immediately after use.

Protection is not needed for objects which R already knows are in use. In particular, this applies to function arguments.

There is a less-used macro UNPROTECT_PTR(*s*) that unprotects the object pointed to by the SEXP *s*, even if it is not the top item on the pointer protection stack. This is rarely needed outside the parser (the R sources have one example, in 'src/main/plot3d.c').

Sometimes an object is changed (for example duplicated, coerced or grown) yet the current value needs to be protected. For these cases PROTECT_WITH_INDEX saves an index of the protection location that can be used to replace the protected value using REPROTECT. For example (from the internal code for `optim`)

```
PROTECT_INDEX ipx;

....
PROTECT_WITH_INDEX(s = eval(OS->R_fcall, OS->R_env), &ipx);
REPROTECT(s = coerceVector(s, REALSXP), ipx);
```

5.7.2 Allocating storage

For many purposes it is sufficient to allocate R objects and manipulate those. There are quite a few `allocXxx` functions defined in ‘`Rinternals.h`’—you may want to explore them. These allocate R objects of various types, and for the standard vector types there are equivalent `NEW_XXX` macros defined in ‘`Rdefines.h`’.

If storage is required for C objects during the calculations this is best allocating by calling `R_alloc`; see [Section 6.1 \[Memory allocation\]](#), page 78. All of these memory allocation routines do their own error-checking, so the programmer may assume that they will raise an error and not return if the memory cannot be allocated.

5.7.3 Details of R types

Users of the ‘`Rinternals.h`’ macros will need to know how the R types are known internally: if the ‘`Rdefines.h`’ macros are used then S4-compatible names are used.

The different R data types are represented in C by `SEXPTYPE`. Some of these are familiar from R and some are internal data types. The usual R object modes are given in the table.

SEXPTYPE	R equivalent
<code>REALSXP</code>	numeric with storage mode <code>double</code>
<code>INTSXP</code>	integer
<code>CPLXSXP</code>	complex
<code>LGLSXP</code>	logical
<code>STRSXP</code>	character
<code>VECSXP</code>	list (generic vector)
<code>LISTSXP</code>	“dotted-pair” list
<code>DOTSXP</code>	a ‘ <code>...</code> ’ object
<code>NILSXP</code>	<code>NULL</code>
<code>SYMSXP</code>	name/symbol
<code>CLOSXP</code>	function or function closure
<code>ENVSXP</code>	environment

Among the important internal `SEXPTYPE`s are `LANGSXP`, `CHARSXP`, `PROMSXP`, etc. (**Note:** although it is possible to return objects of internal types, it is unsafe to do so as assumptions are made about how they are handled which may be violated at user-level evaluation.)

Unless you are very sure about the type of the arguments, the code should check the data types. Sometimes it may also be necessary to check data types of objects created by evaluating an R expression in the C code. You can use functions like `isReal`, `isInteger` and `isString` to do type checking. See the header file ‘`Rinternals.h`’ for definitions of other such functions. All of these take a `SEXP` as argument and return 1 or 0 to indicate `TRUE` or `FALSE`. Once again there are two ways to do this, and ‘`Rdefines.h`’ has macros such as `IS_NUMERIC`.

What happens if the `SEXP` is not of the correct type? Sometimes you have no other option except to generate an error. You can use the function `error` for this. It is usually better to coerce the object to the correct type. For example, if you find that an `SEXP` is of the type `INTEGER`, but you need a `REAL` object, you can change the type by using, equivalently,

```
PROTECT(newSexp = coerceVector(oldSexp, REALSXP));
```

or

```
PROTECT(newSexp = AS_NUMERIC(oldSexp));
```

Protection is needed as a new *object* is created; the object formerly pointed to by the `SEXP` is still protected but now unused.

All the coercion functions do their own error-checking, and generate NAs with a warning or stop with an error as appropriate.

So far we have only seen how to create and coerce R objects from C code, and how to extract the numeric data from numeric R vectors. These can suffice to take us a long way in interfacing R objects to numerical algorithms, but we may need to know a little more to create useful return objects.

5.7.4 Attributes

Many R objects have attributes: some of the most useful are `classes` and the `dim` and `dimnames` that mark objects as matrices or arrays. It can also be helpful to work with the `names` attribute of vectors.

To illustrate this, let us write code to take the outer product of two vectors (which `outer` and `%o%` already do). As usual the R code is simple

```
out <- function(x, y)
{
  storage.mode(x) <- storage.mode(y) <- "double"
  .Call("out", x, y)
}
```

where we expect `x` and `y` to be numeric vectors (possibly integer), possibly with names. This time we do the coercion in the calling R code.

C code to do the computations is

```
#include <R.h>
#include <Rinternals.h>

SEXP out(SEXP x, SEXP y)
{
  int i, j, nx, ny;
  double tmp;
  SEXP ans;

  nx = length(x); ny = length(y);
  PROTECT(ans = allocMatrix(REALSXP, nx, ny));
  for(i = 0; i < nx; i++) {
    tmp = REAL(x)[i];
    for(j = 0; j < ny; j++)
      REAL(ans)[i + nx*j] = tmp * REAL(y)[j];
  }
  UNPROTECT(1);
  return(ans);
}
```

but we would like to set the `dimnames` of the result. Although `allocMatrix` provides a short cut, we will show how to set the `dim` attribute directly.

```
#include <R.h>
#include <Rinternals.h>

SEXP out(SEXP x, SEXP y)
{
  R_len_t i, j, nx, ny;
  double tmp;
  SEXP ans, dim, dimnames;
```



```

    nx = length(x); ny = length(y);
    PROTECT(ans = allocVector(REALSXP, nx*ny));
    for(i = 0; i < nx; i++) {
        tmp = REAL(x)[i];
        for(j = 0; j < ny; j++)
            REAL(ans)[i + nx*j] = tmp * REAL(y)[j];
    }

    PROTECT(dim = allocVector(INTSXP, 2));
    INTEGER(dim)[0] = nx; INTEGER(dim)[1] = ny;
    setAttrib(ans, R_DimSymbol, dim);

    PROTECT(dimnames = allocVector(VECSXP, 2));
    SET_VECTOR_ELT(dimnames, 0, getAttrib(x, R_NamesSymbol));
    SET_VECTOR_ELT(dimnames, 1, getAttrib(y, R_NamesSymbol));
    setAttrib(ans, R_DimNamesSymbol, dimnames);

    UNPROTECT(3);
    return(ans);
}

```

This example introduces several new features. The `getAttrib` and `setAttrib` functions get and set individual attributes. Their second argument is a `SEXP` defining the name in the symbol table of the attribute we want; these and many such symbols are defined in the header file `'Rinternals.h'`.

There are shortcuts here too: the functions `namesgets`, `dimgets` and `dimnamesgets` are the internal versions of `names<-`, `dim<-` and `dimnames<-`, and there are functions such as `GetMatrixDimnames` and `GetArrayDimnames`.

What happens if we want to add an attribute that is not pre-defined? We need to add a symbol for it *via* a call to `install`. Suppose for illustration we wanted to add an attribute `"version"` with value 3.0. We could use

```

SEXP version;
PROTECT(version = allocVector(REALSXP, 1));
REAL(version)[0] = 3.0;
setAttrib(ans, install("version"), version);
UNPROTECT(1);

```

Using `install` when it is not needed is harmless and provides a simple way to retrieve the symbol from the symbol table if it is already installed.

5.7.5 Classes

In R the class is just the attribute named `"class"` so it can be handled as such, but there is a shortcut `classgets`. Suppose we want to give the return value in our example the class `"mat"`. We can use

```

#include <R.h>
#include <Rdefines.h>

....
SEXP ans, dim, dimnames, class;

....
PROTECT(class = allocVector(STRSXP, 1));
SET_STRING_ELT(class, 0, mkChar("mat"));
classgets(ans, class);
UNPROTECT(4);
return(ans);
}

```

As the value is a character vector, we have to know how to create that from a C character array, which we do using the function `mkChar`.

5.7.6 Handling lists

Some care is needed with lists, as R has moved from using LISP-like lists (now called “pairlists”) to S-like generic vectors. As a result, the appropriate test for an object of mode `list` is `isNewList`, and we need `allocVector(VECSXP, n)` and *not* `allocList(n)`.

List elements can be retrieved or set by direct access to the elements of the generic vector. Suppose we have a list object

```
a <- list(f=1, g=2, h=3)
```

Then we can access `a$g` as `a[[2]]` by

```

double g;
....
g = REAL(VECTOR_ELT(a, 1))[0];

```

This can rapidly become tedious, and the following function (based on one in package `stats`) is very useful:

```

/* get the list element named str, or return NULL */

SEXP getListElement(SEXP list, char *str)
{
    SEXP elmt = R_NilValue, names = getAttrib(list, R_NamesSymbol);
    int i;

    for (i = 0; i < length(list); i++)
        if(strcmp(CHAR(STRING_ELT(names, i)), str) == 0) {
            elmt = VECTOR_ELT(list, i);
            break;
        }
    return elmt;
}

```

and enables us to say

```

double g;
g = REAL(getListElement(a, "g"))[0];

```

5.7.7 Finding and setting variables

It will be usual that all the R objects needed in our C computations are passed as arguments to `.Call` or `.External`, but it is possible to find the values of R objects from within the C given their names. The following code is the equivalent of `get(name, envir = rho)`.

```

SEXP getvar(SEXP name, SEXP rho)
{
    SEXP ans;

    if(!isString(name) || length(name) != 1)
        error("name is not a single string");
    if(!isEnvironment(rho))
        error("rho should be an environment");
    ans = findVar(install(CHAR(STRING_ELT(name, 0))), rho);
    printf("first value is %f\n", REAL(ans)[0]);
    return(R_NilValue);
}

```

The main work is done by `findVar`, but to use it we need to install `name` as a name in the symbol table. As we wanted the value for internal use, we return `NULL`.

Similar functions with syntax

```

void defineVar(SEXP symbol, SEXP value, SEXP rho)
void setVar(SEXP symbol, SEXP value, SEXP rho)

```

can be used to assign values to R variables. `defineVar` creates a new binding or changes the value of an existing binding in the specified environment frame; it is the analogue of `assign(symbol, value, envir = rho, inherits = FALSE)`, but unlike `assign`, `defineVar` does not make a copy of the object `value`.² `setVar` searches for an existing binding for `symbol` in `rho` or its enclosing environments. If a binding is found, its value is changed to `value`. Otherwise, a new binding with the specified value is created in the global environment. This corresponds to `assign(symbol, value, envir = rho, inherits = TRUE)`.

5.7.8 Named objects and copying

When assignments are done in R such as

```

x <- 1:10
y <- x

```

the named object is not necessarily copied, so after those two assignments `y` and `x` are bound to the same `SEXP` (the structure a `SEXP` points to). This means that any code which alters one of them has to make a copy before modifying the copy if the usual R semantics are to apply. Note that whereas `.C` and `.Fortran` do copy their arguments (unless the dangerous `dup = FALSE` is used), `.Call` and `.External` do not. So `duplicate` is commonly called on arguments to `.Call` before modifying them.

However, at least some of this copying is unneeded. In the first assignment shown, `x <- 1:10`, R first creates an object with value `1:10` and then assigns it to `x` but if `x` is modified no copy is necessary as the temporary object with value `1:10` cannot be referred to again. R distinguishes between named and unnamed objects *via* a field in a `SEXP` that can be accessed via the macros `NAMED` and `SET_NAMED`. This can take values

- 0 The object is not bound to any symbol
- 1 The object has been bound to exactly one symbol
- 2 The object has potentially been bound to two or more symbols, and one should act as if another variable is currently bound to this value.

Note the past tenses: R does not do full reference counting and there may currently be fewer bindings.

² You can assign a *copy* of the object in the environment frame `rho` using `defineVar(symbol, duplicate(value), rho)`.

It is safe to modify the value of any `SEXP` for which `NAMED(foo)` is zero, and if `NAMED(foo)` is two, the value should be duplicated (via a call to `duplicate`) before any modification. Note that it is the responsibility of the author of the code making the modification to do the duplication, even if it is `x` whose value is being modified after `y <- x`.

The case `NAMED(foo) == 1` allows some optimization, but it can be ignored (and duplication done whenever `NAMED(foo) > 0`). This optimization is unlikely to be worthwhile in user code. It is intended for use within assignment functions. Suppose we used

```
x <- 1:10
foo(x) <- 3
```

which is computed as

```
x <- 1:10
x <- "foo<-"(x)
```

Then inside `"foo<-"` the object pointing to the current value of `x` will have `NAMED(foo)` as one, and it would be safe to modify it as the only symbol bound to it is `x` and that will be rebound immediately. (Provided the remaining code in `"foo<-"` make no reference to `x`, and no one is going to attempt a direct call such as `y <- "foo<-"(x)`.)

Another possible example of optimizing `NAMED(foo) == 1` is inside `.Call` code which alters an argument and the calling R code is written to be aware that the argument will be altered.

5.8 Interface functions `.Call` and `.External`

In this section we consider the details of the R/C interfaces.

These two interfaces have almost the same functionality. `.Call` is based on the interface of the same name in S version 4, and `.External` is based on `.Internal`. `.External` is more complex but allows a variable number of arguments.

5.8.1 Calling `.Call`

Let us convert our finite convolution example to use `.Call`, first using the `'Rdefines.h'` macros. The calling function in R is

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

which could hardly be simpler, but as we shall see all the type checking must be transferred to the C code, which is

```

#include <R.h>
#include <Rdefines.h>

SEXP convolve2(SEXP a, SEXP b)
{
    int i, j, na, nb, nab;
    double *xa, *xb, *xab;
    SEXP ab;

    PROTECT(a = AS_NUMERIC(a));
    PROTECT(b = AS_NUMERIC(b));
    na = LENGTH(a); nb = LENGTH(b); nab = na + nb - 1;
    PROTECT(ab = NEW_NUMERIC(nab));
    xa = NUMERIC_POINTER(a); xb = NUMERIC_POINTER(b);
    xab = NUMERIC_POINTER(ab);
    for(i = 0; i < nab; i++) xab[i] = 0.0;
    for(i = 0; i < na; i++)
        for(j = 0; j < nb; j++) xab[i + j] += xa[i] * xb[j];
    UNPROTECT(3);
    return(ab);
}

```

Note that unlike the macros in S version 4, the R versions of these macros do check that coercion can be done and raise an error if it fails. They will raise warnings if missing values are introduced by coercion. Although we illustrate doing the coercion in the C code here, it often is simpler to do the necessary coercions in the R code.

Now for the version in R-internal style. Only the C code changes.

```

#include <R.h>
#include <Rinternals.h>

SEXP convolve2(SEXP a, SEXP b)
{
    R_len_t i, j, na, nb, nab;
    double *xa, *xb, *xab;
    SEXP ab;

    PROTECT(a = coerceVector(a, REALSXP));
    PROTECT(b = coerceVector(b, REALSXP));
    na = length(a); nb = length(b); nab = na + nb - 1;
    PROTECT(ab = allocVector(REALSXP, nab));
    xa = REAL(a); xb = REAL(b);
    xab = REAL(ab);
    for(i = 0; i < nab; i++) xab[i] = 0.0;
    for(i = 0; i < na; i++)
        for(j = 0; j < nb; j++) xab[i + j] += xa[i] * xb[j];
    UNPROTECT(3);
    return(ab);
}

```

This is called in exactly the same way.

5.8.2 Calling `.External`

We can use the same example to illustrate `.External`. The R code changes only by replacing `.Call` by `.External`

```
conv <- function(a, b) .External("convolveE", a, b)
```

but the main change is how the arguments are passed to the C code, this time as a single SEXP. The only change to the C code is how we handle the arguments.

```
#include <R.h>
#include <Rinternals.h>

SEXP convolveE(SEXP args)
{
    int i, j, na, nb, nab;
    double *xa, *xb, *xab;
    SEXP a, b, ab;

    PROTECT(a = coerceVector(CADR(args), REALSXP));
    PROTECT(b = coerceVector(CADDR(args), REALSXP));
    ...
}
```

Once again we do not need to protect the arguments, as in the R side of the interface they are objects that are already in use. The macros

```
first = CADR(args);
second = CADDR(args);
third = CADDDR(args);
fourth = CAD4R(args);
```

provide convenient ways to access the first four arguments. More generally we can use the `CDR` and `CAR` macros as in

```
args = CDR(args); a = CAR(args);
args = CDR(args); b = CAR(args);
```

which clearly allows us to extract an unlimited number of arguments (whereas `.Call` has a limit, albeit at 65 not a small one).

More usefully, the `.External` interface provides an easy way to handle calls with a variable number of arguments, as `length(args)` will give the number of arguments supplied (of which the first is ignored). We may need to know the names ('tags') given to the actual arguments, which we can by using the `TAG` macro and using something like the following example, that prints the names and the first value of its arguments if they are vector types.

```
#include <R_ext/PrtUtil.h>

SEXP showArgs(SEXP args)
{
    int i, nargs;
    Rcomplex cpl;
    char *name;
```

```

if((nargs = length(args) - 1) > 0) {
  for(i = 0; i < nargs; i++) {
    args = CDR(args);
    name = CHAR(PRINTNAME(TAG(args)));
    switch(TYPEOF(CAR(args))) {
      case REALSXP:
        Rprintf("[%d] '%s' %f\n", i+1, name, REAL(CAR(args))[0]);
        break;
      case LGLSXP:
      case INTSXP:
        Rprintf("[%d] '%s' %d\n", i+1, name, INTEGER(CAR(args))[0]);
        break;
      case CPLXSXP:
        cpl = COMPLEX(CAR(args))[0];
        Rprintf("[%d] '%s' %f + %fi\n", i+1, name, cpl.r, cpl.i);
        break;
      case STRSXP:
        Rprintf("[%d] '%s' %s\n", i+1, name,
                  CHAR(String_ELT(CAR(args), 0)));
        break;
      default:
        Rprintf("[%d] '%s' R type\n", i+1, name);
    }
  }
}
return(R_NilValue);
}

```

This can be called by the wrapper function

```
showArgs <- function(...) .External("showArgs", ...)
```

Note that this style of programming is convenient but not necessary, as an alternative style is

```
showArgs <- function(...) .Call("showArgs1", list(...))
```

5.8.3 Missing and special values

One piece of error-checking the `.C` call does (unless `NAOK` is true) is to check for missing (`NA`) and IEEE special values (`Inf`, `-Inf` and `NaN`) and give an error if any are found. With the `.Call` interface these will be passed to our code. In this example the special values are no problem, as IEEE arithmetic will handle them correctly. In the current implementation this is also true of `NA` as it is a type of `NaN`, but it is unwise to rely on such details. Thus we will re-write the code to handle `NA`s using macros defined in `'R_exts/Arith.h'` included by `'R.h'`.

The code changes are the same in any of the versions of `convolve2` or `convolveE`:

```

...
for(i = 0; i < na; i++)
  for(j = 0; j < nb; j++)
    if(ISNA(xa[i]) || ISNA(xb[j]) || ISNA(xab[i + j]))
      xab[i + j] = NA_REAL;
    else
      xab[i + j] += xa[i] * xb[j];
...

```

Note that the `ISNA` macro, and the similar macros `ISNAN` (which checks for `NaN` or `NA`) and `R_FINITE` (which is false for `NA` and all the special values), only apply to numeric values of type

double. Missingness of integers, logicals and character strings can be tested by equality to the constants `NA_INTEGER`, `NA_LOGICAL` and `NA_STRING`. These and `NA_REAL` can be used to set elements of R vectors to NA.

The constants `R_NaN`, `R_PosInf`, `R_NegInf` and `R_NaReal` can be used to set doubles to the special values.

5.9 Evaluating R expressions from C

We noted that the `call_R` interface could be used to evaluate R expressions from C code, but the current interfaces are much more convenient to use. The main function we will use is

```
SEXP eval(SEXP expr, SEXP rho);
```

the equivalent of the interpreted R code `eval(expr, envir = rho)`, although we can also make use of `findVar`, `defineVar` and `findFun` (which restricts the search to functions).

To see how this might be applied, here is a simplified internal version of `lapply` for expressions, used as

```
a <- list(a = 1:5, b = rnorm(10), test = runif(100))
.Call("lapply", a, quote(sum(x)), new.env())
```

with C code

```
SEXP lapply(SEXP list, SEXP expr, SEXP rho)
{
    R_len_t i, n = length(list);
    SEXP ans;

    if(!isNewList(list)) error("'list' must be a list");
    if(!isEnvironment(rho)) error("'rho' should be an environment");
    PROTECT(ans = allocVector(VECSXP, n));
    for(i = 0; i < n; i++) {
        defineVar(install("x"), VECTOR_ELT(list, i), rho);
        SET_VECTOR_ELT(ans, i, eval(expr, rho));
    }
    setAttrib(ans, R_NamesSymbol, getAttrib(list, R_NamesSymbol));
    UNPROTECT(1);
    return(ans);
}
```

It would be closer to `lapply` if we could pass in a function rather than an expression. One way to do this is *via* interpreted R code as in the next example, but it is possible (if somewhat obscure) to do this in C code. The following is based on the code in `'src/main/optimize.c'`.

```

SEXP lapply2(SEXP list, SEXP fn, SEXP rho)
{
  R_len_t i, n = length(list);
  SEXP R_fcall, ans;

  if(!isNewList(list)) error("'list' must be a list");
  if(!isFunction(fn)) error("'fn' must be a function");
  if(!isEnvironment(rho)) error("'rho' should be an environment");
  PROTECT(R_fcall = lang2(fn, R_NilValue));
  PROTECT(ans = allocVector(VECSXP, n));
  for(i = 0; i < n; i++) {
    SETCADR(R_fcall, VECTOR_ELT(list, i));
    SET_VECTOR_ELT(ans, i, eval(R_fcall, rho));
  }
  setAttrib(ans, R_NamesSymbol, getAttrib(list, R_NamesSymbol));
  UNPROTECT(2);
  return(ans);
}

```

used by

```
.Call("lapply2", a, sum, new.env())
```

Function `lang2` creates an executable ‘list’ of two elements, but this will only be clear to those with a knowledge of a LISP-like language.

As a more comprehensive example of constructing an R call in C code and evaluating, consider the following fragment of `printAttributes` in ‘`src/main/print.c`’.

```

/* Need to construct a call to
   print(CAR(a), digits=digits)
   based on the R_print structure, then eval(call, env).
   See do_docall for the template for this sort of thing.
*/
SEXP s, t;
PROTECT(t = s = allocList(3));
SET_TYPEOF(s, LANGSXP);
CAR(t) = install("print"); t = CDR(t);
CAR(t) = CAR(a); t = CDR(t);
CAR(t) = allocVector(INTSXP, 1);
INTEGER(CAR(t))[0] = digits;
SET_TAG(t, install("digits"));
eval(s, env);
UNPROTECT(1);

```

At this point `CAR(a)` is the R object to be printed, the current attribute. There are three steps: the call is constructed as a pairlist of length 3, the list is filled in, and the expression represented by the pairlist is evaluated.

A pairlist is quite distinct from a generic vector list, the only user-visible form of list in R. A pairlist is a linked list (with `CDR(t)` computing the next entry), with items (accessed by `CAR(t)`) and names or tags (set by `SET_TAG`). In this call there are to be three items, a symbol (pointing to the function to be called) and two argument values, the first unnamed and the second named. Setting the type makes this a call which can be evaluated.

5.9.1 Zero-finding

In this section we re-work the example of `call_S` in Becker, Chambers & Wilks (1988) on finding a zero of a univariate function, which used to be used as an example for `call_R` in the now defunct `demo(dynload)`. The R code and an example are

```
zero <- function(f, guesses, tol = 1e-7) {
  f.check <- function(x) {
    x <- f(x)
    if(!is.numeric(x)) stop("Need a numeric result")
    as.double(x)
  }
  .Call("zero", body(f.check), as.double(guesses), as.double(tol),
        new.env())
}

cube1 <- function(x) (x^2 + 1) * (x - 1.5)
zero(cube1, c(0, 5))
```

where this time we do the coercion and error-checking in the R code. The C code is

```
SEXP mkans(double x)
{
  SEXP ans;
  PROTECT(ans = allocVector(REALSXP, 1));
  REAL(ans)[0] = x;
  UNPROTECT(1);
  return ans;
}

double feval(double x, SEXP f, SEXP rho)
{
  defineVar(install("x"), mkans(x), rho);
  return(REAL(eval(f, rho))[0]);
}

SEXP zero(SEXP f, SEXP guesses, SEXP stol, SEXP rho)
{
  double x0 = REAL(guesses)[0], x1 = REAL(guesses)[1],
        tol = REAL(stol)[0];
  double f0, f1, fc, xc;

  if(tol <= 0.0) error("non-positive tol value");
  f0 = feval(x0, f, rho); f1 = feval(x1, f, rho);
  if(f0 == 0.0) return mkans(x0);
  if(f1 == 0.0) return mkans(x1);
  if(f0*f1 > 0.0) error("x[0] and x[1] have the same sign");
```

```

    for(;;) {
        xc = 0.5*(x0+x1);
        if(fabs(x0-x1) < tol) return mkans(xc);
        fc = feval(xc, f, rho);
        if(fc == 0) return mkans(xc);
        if(f0*fc > 0.0) {
            x0 = xc; f0 = fc;
        } else {
            x1 = xc; f1 = fc;
        }
    }
}

```

The C code is essentially unchanged from the `call_R` version, just using a couple of functions to convert from `double` to `SEXP` and to evaluate `f.check`.

5.9.2 Calculating numerical derivatives

We will use a longer example (by Saikat DebRoy) to illustrate the use of evaluation and `.External`. This calculates numerical derivatives, something that could be done as effectively in interpreted R code but may be needed as part of a larger C calculation.

An interpreted R version and an example are

```

numeric.deriv <- function(expr, theta, rho=sys.frame(sys.parent()))
{
    eps <- sqrt(.Machine$double.eps)
    ans <- eval(substitute(expr), rho)
    grad <- matrix(,length(ans), length(theta),
                  dimnames=list(NULL, theta))
    for (i in seq(along=theta)) {
        old <- get(theta[i], envir=rho)
        delta <- eps * min(1, abs(old))
        assign(theta[i], old+delta, envir=rho)
        ans1 <- eval(substitute(expr), rho)
        assign(theta[i], old, envir=rho)
        grad[, i] <- (ans1 - ans)/delta
    }
    attr(ans, "gradient") <- grad
    ans
}
omega <- 1:5; x <- 1; y <- 2
numeric.deriv(sin(omega*x*y), c("x", "y"))

```

where `expr` is an expression, `theta` a character vector of variable names and `rho` the environment to be used.

For the compiled version the call from R will be

```
.External("numeric_deriv", expr, theta, rho)
```

with example usage

```
.External("numeric_deriv", quote(sin(omega*x*y)),
        c("x", "y"), .GlobalEnv)
```

Note the need to quote the expression to stop it being evaluated.

Here is the complete C code which we will explain section by section.

```

#include <R.h> /* for DOUBLE_EPS */
#include <Rinternals.h>

SEXP numeric_deriv(SEXP args)
{
  SEXP theta, expr, rho, ans, ans1, gradient, par, dimnames;
  double tt, xx, delta, eps = sqrt(DOUBLE_EPS);
  int start, i, j;

  expr = CADR(args);
  if(!isString(theta = CADDR(args)))
    error("theta should be of type character");
  if(!isEnvironment(rho = CADDDR(args)))
    error("rho should be an environment");

  PROTECT(ans = coerceVector(eval(expr, rho), REALSXP));
  PROTECT(gradient = allocMatrix(REALSXP, LENGTH(ans), LENGTH(theta)));

  for(i = 0, start = 0; i < LENGTH(theta); i++, start += LENGTH(ans)) {
    PROTECT(par = findVar(install(CHAR(STRING_ELT(theta, i))), rho));
    tt = REAL(par)[0];
    xx = fabs(tt);
    delta = (xx < 1) ? eps : xx*eps;
    REAL(par)[0] += delta;
    PROTECT(ans1 = coerceVector(eval(expr, rho), REALSXP));
    for(j = 0; j < LENGTH(ans); j++)
      REAL(gradient)[j + start] =
        (REAL(ans1)[j] - REAL(ans)[j])/delta;
    REAL(par)[0] = tt;
    UNPROTECT(2); /* par, ans1 */
  }

  PROTECT(dimnames = allocVector(VECSXP, 2));
  SET_VECTOR_ELT(dimnames, 1, theta);
  dimnamesgets(gradient, dimnames);
  setAttrib(ans, install("gradient"), gradient);
  UNPROTECT(3); /* ans gradient dimnames */
  return ans;
}

```

The code to handle the arguments is

```

expr = CADR(args);
if(!isString(theta = CADDR(args)))
  error("theta should be of type character");
if(!isEnvironment(rho = CADDDR(args)))
  error("rho should be an environment");

```

Note that we check for correct types of `theta` and `rho` but do not check the type of `expr`. That is because `eval` can handle many types of R objects other than `EXPRSXP`. There is no useful coercion we can do, so we stop with an error message if the arguments are not of the correct mode.

The first step in the code is to evaluate the expression in the environment `rho`, by

```

PROTECT(ans = coerceVector(eval(expr, rho), REALSXP));

```

We then allocate space for the calculated derivative by

```
PROTECT(gradient = allocMatrix(REALSXP, LENGTH(ans), LENGTH(theta)));
```

The first argument to `allocMatrix` gives the `SEXPTYPE` of the matrix: here we want it to be `REALSXP`. The other two arguments are the numbers of rows and columns.

```
for(i = 0, start = 0; i < LENGTH(theta); i++, start += LENGTH(ans)) {
  PROTECT(par = findVar(install(CHAR(String_ELT(theta, i))), rho));
```

Here, we are entering a for loop. We loop through each of the variables. In the for loop, we first create a symbol corresponding to the *i*'th element of the `STRSXP` `theta`. Here, `String_ELT(theta, i)` accesses the *i*'th element of the `STRSXP` `theta`. Macro `CHAR()` extracts the actual character representation of it: it returns a pointer. We then install the name and use `findVar` to find its value.

```
  tt = REAL(par)[0];
  xx = fabs(tt);
  delta = (xx < 1) ? eps : xx*eps;
  REAL(par)[0] += delta;
  PROTECT(ans1 = coerceVector(eval(expr, rho), REALSXP));
```

We first extract the real value of the parameter, then calculate `delta`, the increment to be used for approximating the numerical derivative. Then we change the value stored in `par` (in environment `rho`) by `delta` and evaluate `expr` in environment `rho` again. Because we are directly dealing with original R memory locations here, R does the evaluation for the changed parameter value.

```
  for(j = 0; j < LENGTH(ans); j++)
    REAL(gradient)[j + start] =
      (REAL(ans1)[j] - REAL(ans)[j])/delta;
  REAL(par)[0] = tt;
  UNPROTECT(2);
}
```

Now, we compute the *i*'th column of the gradient matrix. Note how it is accessed: R stores matrices by column (like FORTRAN).

```
PROTECT(dimnames = allocVector(VECSXP, 2));
SET_VECTOR_ELT(dimnames, 1, theta);
dimnamesgets(gradient, dimnames);
setAttrib(ans, install("gradient"), gradient);
UNPROTECT(3);
return ans;
}
```

First we add column names to the gradient matrix. This is done by allocating a list (a `VECSXP`) whose first element, the row names, is `NULL` (the default) and the second element, the column names, is set as `theta`. This list is then assigned as the attribute having the symbol `R_DimNamesSymbol`. Finally we set the gradient matrix as the gradient attribute of `ans`, unprotect the remaining protected locations and return the answer `ans`.

5.10 Parsing R code from C

Suppose an R extension want to accept an R expression from the user and evaluate it. The previous section covered evaluation, but the expression will be entered as text and needs to be parsed first. A small part of R's parse interface is declared in header file `'R_ext/Parse.h'`³.

An example of the usage can be found in the (example) Windows package `windlgs` included in the R source tree. The essential part is

³ This is only guaranteed to show the current interface: it is liable to change.

```

#include <R.h>
#include <Rinternals.h>
#include <R_ext/Parse.h>

SEXP menu_ttest3()
{
    char cmd[256];
    SEXP cmdSexp, cmdexpr, ans = R_NilValue;
    int i;
    ParseStatus status;
    ...
    if(done == 1) {
PROTECT(cmdSexp = allocVector(STRSXP, 1));
SET_STRING_ELT(cmdSexp, 0, mkChar(cmd));
cmdexpr = PROTECT(R_ParseVector(cmdSexp, -1, &status));
if (status != PARSE_OK) {
    UNPROTECT(2);
    error("invalid call %s", cmd);
}
/* Loop is needed here as EXPSEXP will be of length > 1 */
for(i = 0; i < length(cmdexpr); i++)
    ans = eval(VECTOR_ELT(cmdexpr, i), R_GlobalEnv);
UNPROTECT(2);
    }
    return ans;
}

```

Note that a single line of text may give rise to more than one R expression.

6 The R API: entry points for C code

There are a large number of entry points in the R executable/DLL that can be called from C code (and some that can be called from FORTRAN code). Only those documented here are stable enough that they will only be changed with considerable notice.

The recommended procedure to use these is to include the header file ‘R.h’ in your C code by

```
#include <R.h>
```

This will include several other header files from the directory ‘R_INCLUDE_DIR/R_ext’, and there are other header files there that can be included too, but many of the features they contain should be regarded as undocumented and unstable.

An alternative is to include the header file ‘S.h’, which may be useful when porting code from S. This includes rather less than ‘R.h’, and has extra some compatibility definitions (for example the `S_complex` type from S).

The defines used for compatibility with S sometimes causes conflicts (notably with Windows headers), and the known problematic defines can be removed by defining `STRICT_R_HEADERS`.

Most of these header files, including all those included by ‘R.h’, can be used from C++ code. Some others need to be included within an `extern "C"` declaration, and for clarity this is advisable for all R header files.

Note: Because R re-maps many of its external names to avoid clashes with user code, it is *essential* to include the appropriate header files when using these entry points.

This remapping can cause problems¹, and can be eliminated by defining `R_NO_REMAP` and prepending `Rf_` to *all* the function names used from ‘Rinternals.h’ and ‘R_ext/Error.h’.

6.1 Memory allocation

There are two types of memory allocation available to the C programmer, one in which R manages the clean-up and the other in which user has full control (and responsibility).

6.1.1 Transient storage allocation

Here R will reclaim the memory at the end of the call to `.C`. Use

```
char* R_alloc(long n, int size)
```

which allocates *n* units of *size* bytes each. A typical usage (from package `stats`) is

```
x = (int *) R_alloc(nrows(merge)+2, sizeof(int));
```

There is a similar calls, `S_alloc` (for compatibility with older versions of S) which zeroes the memory allocated,

```
char* S_alloc(long n, int size)
```

and

```
char* S_realloc(char *p, long new, long old, int size)
```

which changes the allocation size from *old* to *new* units, and zeroes the additional units.

For compatibility with current versions of S, header ‘S.h’ (only) defines wrapper macros equivalent to

```
type* Salloc(long n, int type)
type* Srealloc(char *p, long new, long old, int type)
```

¹ Known problems are redefining `error`, `length`, `vector` and `warning`

This memory is taken from the heap, and released at the end of the `.C`, `.Call` or `.External` call. Users can also manage it, by noting the current position with a call to `vmaxget` and clearing memory allocated subsequently by a call to `vmaxset`. This is only recommended for experts.

Note that this memory will be freed on error or user interrupt (if allowed: see [Section 6.12 \[Allowing interrupts\]](#), page 89).

Note that although `n` is `long`, there are limits imposed by R's internal allocation mechanism. These will only come into play on 64-bit systems, where the current limit for `n` is just under 16Gb.

6.1.2 User-controlled memory

The other form of memory allocation is an interface to `malloc`, the interface providing R error handling. This memory lasts until freed by the user and is additional to the memory allocated for the R workspace.

The interface functions are

```
type* Calloc(size_t n, type)
type* Realloc(any *p, size_t n, type)
void Free(any *p)
```

providing analogues of `calloc`, `realloc` and `free`. If there is an error during allocation it is handled by R, so if these routines return the memory has been successfully allocated or freed. `Free` will set the pointer `p` to `NULL`. (Some but not all versions of S do so.)

Users should arrange to `Free` this memory when no longer needed, including on error or user interrupt. This can often be done most conveniently from an `on.exit` action in the calling R function – see `pwilcox` for an example.

Do not assume that memory allocated by `Calloc/Realloc` comes from the same pool as used by `malloc`: in particular do not use `free` or `strdup` with it.

6.2 Error handling

The basic error handling routines are the equivalents of `stop` and `warning` in R code, and use the same interface.

```
void error(const char * format, ...);
void warning(const char * format, ...);
```

These have the same call sequences as calls to `printf`, but in the simplest case can be called with a single character string argument giving the error message. (Don't do this if the string contains `'%'` or might otherwise be interpreted as a format.)

If `STRICT_R_HEADERS` is not defined there is also an S-compatibility interface which uses calls of the form

```
PROBLEM ..... ERROR
MESSAGE ..... WARN
PROBLEM ..... RECOVER(NULL_ENTRY)
MESSAGE ..... WARNING(NULL_ENTRY)
```

the last two being the forms available in all S versions. Here `'.....'` is a set of arguments to `printf`, so can be a string or a format string followed by arguments separated by commas.

6.2.1 Error handling from FORTRAN

There are two interface function provided to call `error` and `warning` from FORTRAN code, in each case with a simple character string argument. They are defined as

```
subroutine rexit(message)
subroutine rwarn(message)
```

Messages of more than 255 characters are truncated, with a warning.

6.3 Random number generation

The interface to R's internal random number generation routines is

```
double unif_rand();
double norm_rand();
double exp_rand();
```

giving one uniform, normal or exponential pseudo-random variate. However, before these are used, the user must call

```
GetRNGstate();
```

and after all the required variates have been generated, call

```
PutRNGstate();
```

These essentially read in (or create) `.Random.seed` and write it out after use.

File 'S.h' defines `seed_in` and `seed_out` for S-compatibility rather than `GetRNGstate` and `PutRNGstate`. These take a `long *` argument which is ignored.

The random number generator is private to R; there is no way to select the kind of RNG or set the seed except by evaluating calls to the R functions.

The C code behind R's `rxxx` functions can be accessed by including the header file 'Rmath.h'; See [Section 6.7.1 \[Distribution functions\]](#), page 82. Those calls generate a single variate and should also be enclosed in calls to `GetRNGstate` and `PutRNGstate`.

In addition, there is an interface (defined in header 'R_ext/Applic.h') to the generation of random 2-dimensional tables with given row and column totals using Patefield's algorithm.

```
void rcont2 (int* nrow, int* ncol, int* nrowt, int* ncolt, int* ntotal,      [Function]
            double* fact, int* jwork, int* matrix)
```

Here, `nrow` and `ncol` give the numbers `nr` and `nc` of rows and columns and `nrowt` and `ncolt` the corresponding row and column totals, respectively, `ntotal` gives the sum of the row (or columns) totals, `jwork` is a workspace of length `nc`, and on output `matrix` contains the `nr * nc` generated random counts in the usual column-major order.

6.4 Missing and IEEE special values

A set of functions is provided to test for NA, Inf, -Inf and NaN. These functions are accessed via macros:

<code>ISNA(x)</code>	True for R's NA only
<code>ISNAN(x)</code>	True for R's NA and IEEE NaN
<code>R_FINITE(x)</code>	False for Inf, -Inf, NA, NaN

and via function `R_IsNaN` which is true for NaN but not NA.

Do use `R_FINITE` rather than `isfinite` or `finite`; the latter is often mendacious and `isfinite` is only available on a few platforms, on which `R_FINITE` is a macro expanding to `isfinite`.

Currently `ISNAN` is a macro calling `isnan`, although this was not necessarily true before R 2.0.0.

You can check for Inf or -Inf by testing equality to `R_PosInf` or `R_NegInf`, and set (but not test) an NA as `NA_REAL`.

All of the above apply to *double* variables only. For integer variables there is a variable accessed by the macro `NA_INTEGER` which can be used to set or test for missingness.

6.5 Printing

The most useful function for printing from a C routine compiled into R is `Rprintf`. This is used in exactly the same way as `printf`, but is guaranteed to write to R's output (which might be a GUI console rather than a file). It is wise to write complete lines (including the `"\n"`) before returning to R.

The function `REprintf` is similar but writes on the error stream (`stderr`) which may or may not be different from the standard output stream. Functions `Rvprintf` and `REvprintf` are analogues using the `vprintf` interface.

6.5.1 Printing from FORTRAN

On many systems FORTRAN `write` and `print` statements can be used, but the output may not interleave well with that of C, and will be invisible on GUI interfaces. They are not portable and best avoided.

Three subroutines are provided to ease the output of information from FORTRAN code.

```
subroutine dblepr(label, nchar, data, ndata)
subroutine realpr(label, nchar, data, ndata)
subroutine intpr (label, nchar, data, ndata)
```

Here *label* is a character label of up to 255 characters, *nchar* is its length (which can be -1 if the whole label is to be used), and *data* is an array of length at least *ndata* of the appropriate type (`double precision`, `real` and `integer` respectively). These routines print the label on one line and then print *data* as if it were an R vector on subsequent line(s). They work with zero *ndata*, and so can be used to print a label alone.

6.6 Calling C from FORTRAN and vice versa

Naming conventions for symbols generated by FORTRAN differ by platform: it is not safe to assume that FORTRAN names appear to C with a trailing underscore. To help cover up the platform-specific differences there is a set of macros that should be used.

`F77_SUB(name)`

to define a function in C to be called from FORTRAN

`F77_NAME(name)`

to declare a FORTRAN routine in C before use

`F77_CALL(name)`

to call a FORTRAN routine from C

`F77_COMDECL(name)`

to declare a FORTRAN common block in C

`F77_COM(name)`

to access a FORTRAN common block from C

On most current platforms these are all the same, but it is unwise to rely on this. Note that names with underscores are not legal in FORTRAN 77, and are not portably handled by the above macros. (Also, all FORTRAN names for use by R are lower case, but this is not enforced by the macros.)

For example, suppose we want to call R's normal random numbers from FORTRAN. We need a C wrapper along the lines of

```
#include <R.h>

void F77_SUB(rndstart)(void) { GetRNGstate(); }
void F77_SUB(rndend)(void) { PutRNGstate(); }
double F77_SUB(normrnd)(void) { return norm_rand(); }
```

to be called from FORTRAN as in

```
subroutine testit()
double precision normrnd, x
call rndstart()
x = normrnd()
call dblepr("X was", 5, x, 1)
call rndend()
end
```

Note that this is not guaranteed to be portable, for the return conventions might not be compatible between the C and FORTRAN compilers used. (Passing values via arguments is safer.)

The standard packages, for example **stats**, are a rich source of further examples.

6.7 Numerical analysis subroutines

R contains a large number of mathematical functions for its own use, for example numerical linear algebra computations and special functions.

The header files ‘**R_ext/BLAS.h**’, ‘**R_ext/Lapack.h**’ and ‘**R_ext/Linpack.h**’ contains declarations of the BLAS, LAPACK and LINPACK/EISPACK linear algebra functions included in R. These are expressed as calls to FORTRAN subroutines, and they will also be usable from users’ FORTRAN code. Although not part of the official API, this set of subroutines is unlikely to change (but might be supplemented).

The header file ‘**Rmath.h**’ lists many other functions that are available and documented in the following subsections. Many of these are C interfaces to the code behind R functions, so the R function documentation may give further details.

6.7.1 Distribution functions

The routines used to calculate densities, cumulative distribution functions and quantile functions for the standard statistical distributions are available as entry points.

The arguments for the entry points follow the pattern of those for the normal distribution:

```
double dnorm(double x, double mu, double sigma, int give_log);
double pnorm(double x, double mu, double sigma, int lower_tail,
             int give_log);
double qnorm(double p, double mu, double sigma, int lower_tail,
             int log_p);
double rnorm(double mu, double sigma);
```

That is, the first argument gives the position for the density and CDF and probability for the quantile function, followed by the distribution’s parameters. Argument *lower_tail* should be **TRUE** (or 1) for normal use, but can be **FALSE** (or 0) if the probability of the upper tail is desired or specified.

Finally, *give_log* should be non-zero if the result is required on log scale, and *log-p* should be non-zero if *p* has been specified on log scale.

Note that you directly get the cumulative (or “integrated”) *hazard* function, $H(t) = -\log(1 - F(t))$, by using

```
- pdist(t, ..., /*lower_tail = */ FALSE, /* give_log = */ TRUE)
```

or shorter (and more cryptic) `- pdist(t, ..., 0, 1)`.

The random-variate generation routine **rnorm** returns one normal variate. See [Section 6.3 \[Random numbers\]](#), page 80, for the protocol in using the random-variate routines.

Note that these argument sequences are (apart from the names and that **rnorm** has no *n*) exactly the same as the corresponding R functions of the same name, so the documentation of the R functions can be used.

For reference, the following table gives the basic name (to be prefixed by ‘d’, ‘p’, ‘q’ or ‘r’ apart from the exceptions noted) and distribution-specific arguments for the complete set of distributions.

beta	beta	a, b
non-central beta	nbeta	a, b, lambda
binomial	binom	n, p
Cauchy	cauchy	location, scale
chi-squared	chisq	df
non-central chi-squared	nchisq	df, lambda
exponential	exp	scale
F	f	n1, n2
non-central F	nf (*)	n1, n2, ncp
gamma	gamma	shape, scale
geometric	geom	p
hypergeometric	hyper	NR, NB, n
logistic	logis	location, scale
lognormal	lnorm	logmean, logsd
negative binomial	nbinom	n, p
normal	norm	mu, sigma
Poisson	pois	lambda
Student’s t	t	n
non-central t	nt	df, delta
Studentized range	tukey (*)	rr, cc, df
uniform	unif	a, b
Weibull	weibull	shape, scale
Wilcoxon rank sum	wilcox	m, n
Wilcoxon signed rank	signrank	n

Entries marked with an asterisk only have ‘p’ and ‘q’ functions available, and none of the non-central distributions have ‘r’ functions. After a call to `dwilcox`, `pwilcox` or `qwilcox` the function `wilcox_free()` should be called, and similarly for the signed rank functions.

The argument names are not all quite the same as the R ones.

6.7.2 Mathematical functions

<code>double gammafn (double x)</code>	[Function]
<code>double lgammafn (double x)</code>	[Function]
<code>double digamma (double x)</code>	[Function]
<code>double trigamma (double x)</code>	[Function]
<code>double tetragamma (double x)</code>	[Function]
<code>double pentagamma (double x)</code>	[Function]
<code>double psigamma (double x, double deriv)</code>	[Function]

The Gamma function, its natural logarithm and first four derivatives and the n-th derivative of Psi, the digamma function.

<code>double beta (double a, double b)</code>	[Function]
<code>double lbeta (double a, double b)</code>	[Function]

The (complete) Beta function and its natural logarithm.

<code>double choose (double n, double k)</code>	[Function]
<code>double lchoose (double n, double k)</code>	[Function]

The number of combinations of k items chosen from n and its natural logarithm. n and k are rounded to the nearest integer.


```
double bessel_i (double x, double nu, double expo) [Function]
double bessel_j (double x, double nu) [Function]
double bessel_k (double x, double nu, double expo) [Function]
double bessel_y (double x, double nu) [Function]
```

Bessel functions of types I, J, K and Y with index *nu*. For **bessel_i** and **bessel_k** there is the option to return $\exp(-x)$ $I(x; nu)$ or $\exp(x)$ $K(x; nu)$ if *expo* is 2. (Use *expo* == 1 for unscaled values.)

6.7.3 Numerical Utilities

There are a few other numerical utility functions available as entry points.

```
double R_pow (double x, double y) [Function]
double R_pow_di (double x, int i) [Function]
```

R_pow(*x*, *y*) and **R_pow_di**(*x*, *i*) compute x^y and x^i , respectively using **R_FINITE** checks and returning the proper result (the same as R) for the cases where *x*, *y* or *i* are 0 or missing or infinite or NaN.

```
double pythag (double a, double b) [Function]
pythag(a, b) computes  $\sqrt{a^2 + b^2}$  without overflow or destructive underflow: for example it still works when both a and b are between 1e200 and 1e300 (in IEEE double precision).
```

```
double log1p (x) [Function]
log1p(x) Computes  $\log(1 + x)$  (log 1 plus x), accurately even for small x, i.e.,  $|x| \ll 1$ .
This may be provided by your platform, in which case it is not included in ‘Rmath.h’, but is (probably) in ‘math.h’ which ‘Rmath.h’ includes. For backwards compatibility with R versions prior to 1.5.0, the entry point Rf_log1p is still provided.
```

```
double log1pmx (x) [Function]
log1pmx(x) Computes  $\log(1 + x) - x$  (log 1 plus x minus x), accurately even for small x, i.e.,  $|x| \ll 1$ .
```

```
double expm1 (x) [Function]
expm1(x) Computes  $\exp(x) - 1$  (exp x minus 1), accurately even for small x, i.e.,  $|x| \ll 1$ .
This may be provided by your platform, in which case it is not included in ‘Rmath.h’, but is (probably) in ‘math.h’ which ‘Rmath.h’ includes.
```

```
double lgamma1p (x) [Function]
lgamma1p(x) Computes  $\log(\Gamma(x + 1))$  (log(gamma(1 plus x))), accurately even for small x, i.e.,  $0 < x < 0.5$ .
```

```
double logspace_add (double logx, double logy) [Function]
double logspace_sub (double logx, double logy) [Function]
```

Compute the log of a sum or difference from logs of terms, i.e., “*x + y*” as $\log(\exp(\log x) + \exp(\log y))$ and “*x - y*” as $\log(\exp(\log x) - \exp(\log y))$, without causing overflows or throwing away too much accuracy.

```
int imax2 (int x, int y) [Function]
int imin2 (int x, int y) [Function]
double fmax2 (double x, double y) [Function]
double fmin2 (double x, double y) [Function]
```

Return the larger (**max**) or smaller (**min**) of two integer or double numbers, respectively.

```
double sign (double x) [Function]
sign(x) Compute the signum function, where sign(x) is 1, 0, or -1, when x is positive, 0, or negative, respectively.
```


double fsign (*double x, double y*) [Function]
 Performs “transfer of sign” and is defined as $|x| * \text{sign}(y)$.

double fprec (*double x, double digits*) [Function]
 Returns the value of x rounded to *digits* decimal digits (after the decimal point).
 This is the function used by R’s `round()`.

double fround (*double x, double digits*) [Function]
 Returns the value of x rounded to *digits significant* decimal digits.
 This is the function used by R’s `signif()`.

double ftrunc (*double x*) [Function]
 Returns the value of x truncated (to an integer value) towards zero.

6.7.4 Mathematical constants

R has a set of commonly used mathematical constants encompassing constants usually found ‘`math.h`’ and contains further ones that are used in statistical computations. All these are defined to (at least) 30 digits accuracy in ‘`Rmath.h`’. The following definitions use `ln(x)` for the natural logarithm (`log(x)` in R).

Name	Definition ($\ln = \log$)	<code>round(value, 7)</code>
<code>M_E</code>	e	2.7182818
<code>M_LOG2E</code>	$\log_2(e)$	1.4426950
<code>M_LOG10E</code>	$\log_{10}(e)$	0.4342945
<code>M_LN2</code>	$\ln(2)$	0.6931472
<code>M_LN10</code>	$\ln(10)$	2.3025851
<code>M_PI</code>	π	3.1415927
<code>M_PI_2</code>	$\pi/2$	1.5707963
<code>M_PI_4</code>	$\pi/4$	0.7853982
<code>M_1_PI</code>	$1/\pi$	0.3183099
<code>M_2_PI</code>	$2/\pi$	0.6366198
<code>M_2_SQRTPI</code>	$2/\sqrt{\pi}$	1.1283792
<code>M_SQRT2</code>	$\sqrt{2}$	1.4142136
<code>M_SQRT1_2</code>	$1/\sqrt{2}$	0.7071068
<code>M_SQRT_3</code>	$\sqrt{3}$	1.7320508
<code>M_SQRT_32</code>	$\sqrt{32}$	5.6568542
<code>M_LOG10_2</code>	$\log_{10}(2)$	0.3010300
<code>M_2PI</code>	2π	6.2831853
<code>M_SQRT_PI</code>	$\sqrt{\pi}$	1.7724539
<code>M_1_SQRT_2PI</code>	$1/\sqrt{2\pi}$	0.3989423
<code>M_SQRT_2dPI</code>	$\sqrt{2/\pi}$	0.7978846
<code>M_LN_SQRT_PI</code>	$\ln(\sqrt{\pi})$	0.5723649
<code>M_LN_SQRT_2PI</code>	$\ln(\sqrt{2\pi})$	0.9189385
<code>M_LN_SQRT_PId2</code>	$\ln(\sqrt{\pi/2})$	0.2257914

There are a set of constants (`PI`, `DOUBLE_EPS`) (and so on) defined (unless `STRICT_R_HEADERS` is defined) in the included header ‘`R_ext/Constants.h`’, mainly for compatibility with S.

Further, the included header ‘`R_ext/Boolean.h`’ has constants `TRUE` and `FALSE = 0` of type `Rboolean` in order to provide a way of using “logical” variables in C consistently.

6.8 Optimization

The C code underlying `optim` can be accessed directly. The user needs to supply a function to compute the function to be minimized, of the type

```
typedef double optimfn(int n, double *par, void *ex);
```

where the first argument is the number of parameters in the second argument. The third argument is a pointer passed down from the calling routine, normally used to carry auxiliary information.

Some of the methods also require a gradient function

```
typedef void optimgr(int n, double *par, double *gr, void *ex);
```

which passes back the gradient in the `gr` argument. No function is provided for finite-differencing, nor for approximating the Hessian at the result.

The interfaces (defined in header ‘R_ext/Applic.h’) are

- Nelder Mead:

```
void nmmin(int n, double *xin, double *x, double *Fmin, optimfn fn,
           int *fail, double abstol, double intol, void *ex,
           double alpha, double beta, double gamma, int trace,
           int *fnccount, int maxit);
```

- BFGS:

```
void vmmin(int n, double *x, double *Fmin,
           optimfn fn, optimgr gr, int maxit, int trace,
           int *mask, double abstol, double reltol, int nREPORT,
           void *ex, int *fnccount, int *grcount, int *fail);
```

- Conjugate gradients:

```
void cgmin(int n, double *xin, double *x, double *Fmin,
           optimfn fn, optimgr gr, int *fail, double abstol,
           double intol, void *ex, int type, int trace,
           int *fnccount, int *grcount, int maxit);
```

- Limited-memory BFGS with bounds:

```
void lbfgsb(int n, int lmm, double *x, double *lower,
            double *upper, int *nbd, double *Fmin, optimfn fn,
            optimgr gr, int *fail, void *ex, double factr,
            double pgtol, int *fnccount, int *grcount,
            int maxit, char *msg, int trace, int nREPORT);
```

- Simulated annealing:

```
void sammin(int n, double *x, double *Fmin, optimfn fn, int maxit,
            int tmax, double temp, int trace, void *ex);
```

Many of the arguments are common to the various methods. `n` is the number of parameters, `x` or `xin` is the starting parameters on entry and `x` the final parameters on exit, with final value returned in `Fmin`. Most of the other parameters can be found from the help page for `optim`: see the source code ‘src/appl/lbfgsb.c’ for the values of `nbd`, which specifies which bounds are to be used.

6.9 Integration

The C code underlying `integrate` can be accessed directly. The user needs to supply a *vectorizing* C function to compute the function to be integrated, of the type

```
typedef void integr_fn(double *x, int n, void *ex);
```

where `x[]` is both input and output and has length `n`, i.e., a C function, say `fn`, of type `integr_fn` must basically do `for(i in 1:n) x[i] := f(x[i], ex)`. The vectorization requirement can be used to speed up the integrand instead of calling it `n` times. Note that in the current implementation built on QUADPACK, `n` will be either 15 or 21. The `ex` argument is a pointer passed down from the calling routine, normally used to carry auxiliary information.

There are interfaces (defined in header ‘`R_ext/Applic.h`’) for definite and for indefinite integrals. ‘Indefinite’ means that at least one of the integration boundaries is not finite.

- Finite:

```
void Rdqags(integr_fn f, void *ex, double *a, double *b,
            double *epsabs, double *epsrel,
            double *result, double *abserr, int *neval, int *ier,
            int *limit, int *lenw, int *last,
            int *iwork, double *work);
```

- Indefinite:

```
void Rdqagi(integr_fn f, void *ex, double *bound, int *inf,
            double *epsabs, double *epsrel,
            double *result, double *abserr, int *neval, int *ier,
            int *limit, int *lenw, int *last,
            int *iwork, double *work);
```

Only the 3rd and 4th argument differ for the two integrators; for the definite integral, using `Rdqags`, `a` and `b` are the integration interval bounds, whereas for an indefinite integral, using `Rdqagi`, `bound` is the finite bound of the integration (if the integral is not doubly-infinite) and `inf` is a code indicating the kind of integration range,

`inf = 1` corresponds to (bound, +Inf),

`inf = -1` corresponds to (-Inf, bound),

`inf = 2` corresponds to (-Inf, +Inf),

`f` and `ex` define the integrand function, see above; `epsabs` and `epsrel` specify the absolute and relative accuracy requested, `result`, `abserr` and `last` are the output components `value`, `abs.err` and `subdivisions` of the R function `integrate`, where `neval` gives the number of integrand function evaluations, and the error code `ier` is translated to R’s `integrate()` \$ `message`, look at that function definition. `limit` corresponds to `integrate(..., subdivisions = *)`. It seems you should always define the two work arrays and the length of the second one as

```
lenw = 4 * limit;
iwork = (int *) R_alloc(limit, sizeof(int));
work = (double *) R_alloc(lenw, sizeof(double));
```

The comments in the source code in ‘`src/appl/integrate.c`’ give more details, particularly about reasons for failure (`ier >= 1`).

6.10 Utility functions

R has a fairly comprehensive set of sort routines which are made available to users’ C code. These are declared in header file ‘`R_ext/Uutils.h`’ (included by ‘`R.h`’) and include the following.

```
void R_isort (int* x, int n) [Function]
void R_rsort (double* x, int n) [Function]
void R_csort (Rcomplex* x, int n) [Function]
void rsort_with_index (double* x, int* index, int n) [Function]
```

The first three sort integer, real (double) and complex data respectively. (Complex numbers are sorted by the real part first then the imaginary part.) NAs are sorted last.

`rsort_with_index` sorts on `x`, and applies the same permutation to `index`. NAs are sorted last.

`void revsort (double* x, int* index, int n)` [Function]

Is similar to `rsort_with_index` but sorts into decreasing order, and NAs are not handled.

`void iPsort (int* x, int n, int k)` [Function]

`void rPsort (double* x, int n, int k)` [Function]

`void cPsort (Rcomplex* x, int n, int k)` [Function]

These all provide (very) partial sorting: they permute `x` so that `x[k]` is in the correct place with smaller values to the left, larger ones to the right.

`void R_qsort (double *v, int i, int j)` [Function]

`void R_qsort_I (double *v, int *I, int i, int j)` [Function]

`void R_qsort_int (int *iv, int i, int j)` [Function]

`void R_qsort_int_I (int *iv, int *I, int i, int j)` [Function]

These routines sort `v[i:j]` or `iv[i:j]` (using 1-indexing, i.e., `v[1]` is the first element) calling the quicksort algorithm as used by R's `sort(v, method = "quick")` and documented on the help page for the R function `sort`. The `..._I()` versions also return the `sort.index()` vector in `I`. Note that the ordering is *not* stable, so tied values may be permuted.

Note that NAs are not handled (explicitly) and you should use different sorting functions if NAs can be present.

`subroutine qsort4 (double precision v, integer indx, integer ii, integer jj)` [Function]

`subroutine qsort3 (double precision v, integer ii, integer jj)` [Function]

The FORTRAN interface routines for sorting double precision vectors are `qsort3` and `qsort4`, equivalent to `R_qsort` and `R_qsort_I`, respectively.

`void R_max_col (double* matrix, int* nr, int* nc, int* maxes, int* ties_meth)` [Function]

Given the `nr` by `nc` matrix `matrix` in column-major ("FORTRAN") order, `R_max_col()` returns in `maxes[i-1]` the column number of the maximal element in the *i*-th row (the same as R's `max.col()` function). In the case of ties (multiple maxima), `*ties_meth` is an integer code in 1:3 determining the method: 1 = "random", 2 = "first" and 3 = "last". See R's help page `?max.col`.

`int findInterval (double* xt, int n, double x, Rboolean rightmost_closed, Rboolean all_inside, int ilo, int* mflag)` [Function]

Given the ordered vector `xt` of length `n`, return the interval or index of `x` in `xt[]`, typically `max(i; 1 ≤ i ≤ n & xt[i] ≤ x)` where we use 1-indexing as in R and FORTRAN (but not C). If `rightmost_closed` is true, also returns `n - 1` if `x` equals `xt[n]`. If `all_inside` is not 0, the result is coerced to lie in 1:(`n-1`) even when `x` is outside the `xt[]` range. On return, `*mflag` equals `-1` if `x < xt[1]`, `+1` if `x ≥ xt[n]`, and 0 otherwise.

The algorithm is particularly fast when `ilo` is set to the last result of `findInterval()` and `x` is a value of a sequence which is increasing or decreasing for subsequent calls.

There is also an `F77_CALL(interv)()` version of `findInterval()` with the same arguments, but all pointers.

The following two functions do *numerical* colorspace conversion from HSV to RGB and back. Note that all colours must be in `[0,1]`.

`void hsv2rgb (double h, double s, double v, double *r, double *g, double *b)` [Function]

`void rgb2hsv (double r, double g, double b, double *h, double *s, double *v)` [Function]

A system-independent interface to produce the name of a temporary file is provided as

char * R_tmpnam (const char* *prefix*) [Function]
 Return a pathname for a temporary file with name beginning with *prefix*. A NULL prefix is replaced by "".

There is also the internal function used to expand file names in several R functions, and called directly by `path.expand`.

char * R_ExpandFileName (char* *fn*) [Function]
 Expand a path name *fn* by replacing a leading tilde by the user's home directory (if defined). The precise meaning is platform-specific; it will usually be taken from the environment variable HOME if this is defined.

6.11 Internationalization

R has its own C-level interface to the the encoding conversion capabilities provided by `iconv`, for the following reasons

- These wrapper routines do error-handling when no usable implementation of `iconv` was available at configure time.
- Under Windows they arrange to load the '`iconv.dll`' at first use.
- There are incompatibilities between the declarations in different implementations of `iconv`.

These are declared in header file '`R_ext/Riconv.h`'.

void *Riconv_open (char* *to*, char* *from*) [Function]
 Set up a pointer to an encoding object to be used to convert between two encodings: "" indicates the current locale.

size_t Riconv (void **cd*, char *inbuf*, size_t **inbytesleft*, char ***outbuf*, size_t **outbytesleft*)** [Function]

Convert as much as possible of *inbuf* to *outbuf*. Initially the `int` variables indicate the number of bytes available in the buffers, and they are updated (and the `char` pointers are updated to point to the next free byte in the buffer). The return value is the number of characters converted, or `(size_t)-1` (beware: `size_t` is usually an unsigned type). It should be safe to assume that an error condition sets `errno` to one of `E2BIG` (the output buffer is full), `EILSEQ` (the input cannot be converted, and might be invalid in the encoding specified) or `EINVAL` (the input does not end with a complete multi-byte character).

int Riconv_close (void * *cd*) [Function]
 Free the resources of an encoding object.

6.12 Allowing interrupts

As from R 1.8.0 no part of R can be interrupted whilst running long computations in compiled code, so programmers should make provision for the code to be interrupted at suitable points by calling from C

```
#include <R_ext/Utils.h>
```

```
void R_CheckUserInterrupt(void);
```

and from FORTRAN

```
subroutine rchkusr()
```

These check if the user has requested an interrupt, and if so branch to R's error handling functions.

Note that it is possible that the code behind one of the entry points defined here if called from your C or FORTRAN code could be interruptible or generate an error and so not return to your code.

6.13 Platform and version information

The header files define `USING_R`, which should be used to test if the code is indeed being used with R.

Header file `'Rconfig.h'` (included by `'R.h'`) is used to define platform-specific macros that are mainly for use in other header files. The macro `WORDS_BIGENDIAN` is defined on big-endian systems (e.g. `sparc-sun-solaris2.6`) and not on little-endian systems (such as `i686` under Linux or Windows). It can be useful when manipulating binary files.

Header file `'Rversion.h'` (**not** included by `'R.h'`) defines a macro `R_VERSION` giving the version number encoded as an integer, plus a macro `R_Version` to do the encoding. This can be used to test if the version of R is late enough, or to include back-compatibility features. For protection against earlier versions of R which did not have this macro, use a construction such as

```
#if defined(R_VERSION) && R_VERSION >= R_Version(1, 9, 0)
...
#endif
```

More detailed information is available in the macros `R_MAJOR`, `R_MINOR`, `R_YEAR`, `R_MONTH` and `R_DAY`: see the header file `'Rversion.h'` for their format. Note that the minor version includes the patchlevel (as in `'9.0'`).

6.14 Inlining C functions

The C99 keyword `inline` is recognized by some compilers used to build R whereas others need `__inline__` or do not support inlining. Portable code can be written using the macro `R_INLINE` (defined in file `'Rconfig.h'` included by `'R.h'`), as for example from package **cluster**

```
#include <R.h>

static R_INLINE int ind_2(int l, int j)
{
...
}
```

Be aware that using inlining with functions in more than one compilation unit is almost impossible to do portably: see <http://www.greenend.org.uk/rjk/2003/03/inline.html>. All the R configure code has checked is that `R_INLINE` can be used in a single C file with the compiler used to build R. We recommend that packages making extensive use of inlining include their own configure code.

6.15 Using these functions in your own C code

It is possible to build `Mathlib`, the R set of mathematical functions documented in `'Rmath.h'`, as a standalone library `'libRmath'` under Unix and Windows. (This includes the functions documented in [Section 6.7 \[Numerical analysis subroutines\]](#), page 82 as from that header file.)

The library is not built automatically when R is installed, but can be built in the directory `'src/nmath/standalone'` in the R sources: see the file `'README'` there. To use the code in your own C program include

```
#define MATHLIB_STANDALONE
#include <Rmath.h>
```

and link against `'-lRmath'`. There is an example file `'test.c'`.

A little care is needed to use the random-number routines. You will need to supply the uniform random number generator

```
double unif_rand(void)
```


or use the one supplied (and with a dynamic library or DLL you will have to use the one supplied, which is the Marsaglia-multicarry with an entry points

```
set_seed(unsigned int, unsigned int)
```

to set its seeds and

```
get_seed(unsigned int *, unsigned int *)
```

to read the seeds).

6.16 Organization of header files

The header files which R installs are in directory ‘R_INCLUDE_DIR’ (default ‘R_HOME/include’). This currently contains

‘R.h’	includes many other files
‘S.h’	different version for code ported from S
‘Rinternals.h’	definitions for using R’s internal structures
‘Rdefines.h’	macros for an S-like interface to the above
‘Rmath.h’	standalone math library
‘Rversion.h’	R version information
‘Rinterface.h’	for add-on front-ends (Unix-alikes only)
‘R_ext/RStartup.h’	ditto
‘R_ext/eventloop.h’	for add-on front-ends and for packages that need to share in the R event loops (on all platforms)
‘R_ext/Applic.h’	optimization and integration
‘R_ext/BLAS.h’	C definitions for BLAS routines
‘R_ext/Lapack.h’	C definitions for some LAPACK routines
‘R_ext/Linpack.h’	C definitions for some LINPACK routines, not all of which are included in R
‘R_ext/GetX11Image.h’	X11Image interface used by package trkplot
‘R_ext/Parse.h’	a small part of R’s parse interface
‘R_ext/Callbacks.h’	C (and R function) top-level task handlers
‘R_ext/RConvertors.h’	
‘R_ext/Rdynload.h’	needed to register compiled code in packages
‘R_ext/R-ftp-http.h’	interface to internal method of <code>download.file</code>
‘R_ext/Riconv.h’	interface to <code>iconv</code>

The following headers are included by ‘R.h’:

‘Rconfig.h’	configuration info that is made available
‘R_ext/Arith.h’	handling for NAs, NaNs, Inf/−Inf
‘R_ext/Boolean.h’	TRUE/FALSE type
‘R_ext/Complex.h’	C typedefs for R’s complex
‘R_ext/Constants.h’	constants
‘R_ext/Error.h’	error handling
‘R_ext/Memory.h’	memory allocation
‘R_ext/Print.h’	Rprintf and variations.
‘R_ext/Random.h’	random number generation
‘R_ext/RS.h’	definitions common to ‘R.h’ and ‘S.h’, including F77_CALL etc.
‘R_ext/Utils.h’	sorting and other utilities
‘R_ext/libextern.h’	definitions for exports from ‘R.dll’ on Windows.

The graphics systems are exposed in headers ‘Rdevices.h’ (for writing graphics devices), ‘Rgraphics.h’ and ‘R_ext/Graphics{Base,Device,Engine}.h’.

7 Generic functions and methods

R programmers will often want to add methods for existing generic functions, and may want to add new generic functions or make existing functions generic. In this chapter we give guidelines for doing so, with examples of the problems caused by not adhering to them.

This chapter only covers the ‘informal’ class system copied from S3, and not with the formal methods of package **methods** of R 1.4.0 and later.

The key function for methods is `NextMethod`, which dispatches the next method. It is quite typical for a method function to make a few changes to its arguments, dispatch to the next method, receive the results and modify them a little. An example is

```
t.data.frame <- function(x)
{
  x <- as.matrix(x)
  NextMethod("t")
}
```

Also consider `predict.glm`: it happens that in R for historical reasons it calls `predict.lm` directly, but in principle (and in S originally and currently) it could use `NextMethod`. (`NextMethod` seems under-used in the R sources. Do be aware that there are S/R differences in this area, and the example above works because there is a *next* method, the default method, not that a new method is selected when the class is changed.)

Any method a programmer writes may be invoked from another method by `NextMethod`, with the arguments appropriate to the previous method. Further, the programmer cannot predict which method `NextMethod` will pick (it might be one not yet dreamt of), and the end user calling the generic needs to be able to pass arguments to the next method. For this to work

A method must have all the arguments of the generic, including ... if the generic does.

It is a grave misunderstanding to think that a method needs only to accept the arguments it needs. The original S version of `predict.lm` did not have a `...` argument, although `predict` did. It soon became clear that `predict.glm` needed an argument `dispersion` to handle over-dispersion. As `predict.lm` had neither a `dispersion` nor a `...` argument, `NextMethod` could no longer be used. (The legacy, two direct calls to `predict.lm`, lives on in `predict.glm` in R, which is based on the workaround for S3 written by Venables & Ripley.)

Further, the user is entitled to use positional matching when calling the generic, and the arguments to a method called by `UseMethod` are those of the call to the generic. Thus

A method must have arguments in exactly the same order as the generic.

To see the scale of this problem, consider the generic function `scale`, defined (in R 1.4.0) as

```
scale <- function (x, center = TRUE, scale = TRUE)
  UseMethod("scale")
```

Suppose an unthinking package writer created methods such as

```
scale.foo <- function(x, scale = FALSE, ...) { }
```

Then for `x` of class "foo" the calls

```
scale(x, , TRUE)
scale(x, scale = TRUE)
```

would most likely do different things, to the justifiable consternation of the end user.

To add a further twist, which default is used when a user calls `scale(x)` in our example? What if

```
scale.bar <- function(x, center, scale = TRUE) NextMethod("scale")
```

and `x` has class `c("bar", "foo")`? We are not going to give you the answers because it is unreasonable that a user should be expected to anticipate such behaviour. This leads to the recommendation:

A method should use the same defaults as the generic.

Here there might be justifiable exceptions, which will need careful documentation.

7.1 Adding new generics

When creating a new generic function, bear in mind that its argument list will be the maximal set of arguments for methods, including those written elsewhere years later. So choosing a good set of arguments may well be an important design issue, and there need to be good arguments *not* to include a `...` argument.

If a `...` argument is supplied, some thought should be given to its position in the argument sequence. Arguments which follow `...` must be named in calls to the function, and they must be named in full (partial matching is suppressed after `...`). Formal arguments before `...` can be partially matched, and so may ‘swallow’ actual arguments intended for `...`. Although it is commonplace to make the `...` argument the last one, that is not always the right choice.

Sometimes package writers want to make generic a function in the base package, and request a change in R. This may be justifiable, but making a function generic with the old definition as the default method does have a small performance cost. It is never necessary, as a package can take over a function in the base package and make it generic by

```
foo <- function(object, ...) UseMethod("foo")
foo.default <- base::foo
```

(If the thus defined default method needs a ‘`...`’ added to its argument list, one can e.g. use `formals(foo.default) <- c(formals(foo.default), alist(... =))`.)

The same idea can be applied for functions in other packages with namespaces.

8 Linking GUIs and other front-ends to R

8.1 Embedding R under Unix-alikes

R can be built as a shared library¹ if configured with ‘`--enable-R-shlib`’. This shared library can be used to run R from alternative front-end programs. We will assume this has been done for the rest of this section.

The command-line R front-end, ‘`R_HOME/bin/exec/R`’ is one such example, and the unbundled GNOME and MacOS X consoles are others. The source for ‘`R_HOME/bin/exec/R`’ is in file ‘`src/main/Rmain.c`’ and is very simple

```
int Rf_initialize_R(int ac, char **av); /* in ../unix/system.c */
void Rf_mainloop();                  /* in main.c */

int main(int ac, char **av)
{
    Rf_initialize_R(ac, av);
    Rf_mainloop(); /* does not return */
    return 0;
}
```

indeed, misleadingly simple. Remember that ‘`R_HOME/bin/exec/R`’ is run from a shell script ‘`R_HOME/bin/R`’ which sets up the environment for the executable, and this is used for

- Setting `R_HOME` and checking it is valid, as well as the path `R_SHARE_DIR` and `R_DOC_DIR` to the installed ‘`share`’ and ‘`doc`’ directory trees.
- Setting `LD_LIBRARY_PATH` to include the directories used in linking R. This is recorded as the default setting of `R_LD_LIBRARY_PATH` in the shell script.
- Creating the per-session temporary directory used by `tempdir()`, exported as the environment variable `R_SESSION_TMPDIR`.
- Processing some of the arguments, for example to run R under a debugger and to launch alternative front-ends to provide GUIs.

The first two of these can be achieved for your front-end by running it *via* `R CMD`, and the third is optional, for if `R_SESSION_TMPDIR` is not set, the R startup code will create a suitable per-session directory. So, for example

```
R CMD /usr/local/lib/R/bin/exec/R
R CMD exec/R
```

will both work in a standard R installation. (`R CMD` looks first for executables in ‘`R_HOME/bin`’.) If you do not want to run your front-end in this way, you need to ensure that `R_HOME` is set and `LD_LIBRARY_PATH` is suitable. (The latter might well be, but modern Unix/Linux systems do not normally include ‘`/usr/local/lib`’, and R does look there for system components.)

The other senses in which this example is too simple are that all the internal defaults are used and that control is handed over to the R main loop. There are a number of small examples² in the ‘`tests/Embedding`’ directory. These make use of `Rf_initEmbeddedR` in ‘`src/main/Rembedded.c`’, and essentially use

```
int Rf_initEmbeddedR(int argc, char **argv)
{ /* This is already compiled into R */
    Rf_initialize_R(argc, argv);
    setup_Rmainloop();
}
```

¹ In the parlance of MacOS X this is a *dynamic* library, and is the normal way to build R on that platform.

² but these are not part of the automated test procedures and so little tested.

```

    return(1);
}

int main(int ac, char **av)
{
    /* do some setup */
    Rf_initEmbeddedR(argc, argv);
    /* do some more setup */

    /* submit some code to R, which is done interactively via
       run_Rmainloop();
    */
    end_Rmainloop();
    return 0;
}

```

If you don't want to pass R arguments, you can fake an `argv` array, for example by

```

char *argv[] = {"REmbeddedPostgres", "--gui=none", "--silent"};
Rf_initEmbeddedR(sizeof(argv)/sizeof(argv[0]), argv);

```

However, to make a GUI we do want to run `run_Rmainloop` after setting up various parts of R to talk to our GUI, and arranging for our GUI callbacks to be called during the R mainloop.

8.1.1 Setting R callbacks

As from R version 2.1.0 there is a public header file ‘`Rinterface.h`’ that makes it possible to change the standard callbacks used by R in a documented way. This defines pointers (if `R_INTERFACE_PTRS` is defined)

```

extern void (*ptr_R_Suicide)(char *);
extern void (*ptr_R_ShowMessage)(char *);
extern int  (*ptr_R_ReadConsole)(char *, unsigned char *, int, int);
extern void (*ptr_R_WriteConsole)(char *, int);
extern void (*ptr_R_ResetConsole)();
extern void (*ptr_R_FlushConsole)();
extern void (*ptr_R_ClearerrConsole)();
extern void (*ptr_R_Busy)(int);
extern void (*ptr_R_CleanUp)(SA_TYPE, int, int);
extern int  (*ptr_R_ShowFiles)(int, char **, char **, char *,
                               Rboolean, char *);
extern int  (*ptr_R_ChooseFile)(int, char *, int);
extern int  (*ptr_R_EditFile)(char *);
extern void (*ptr_R_loadhistory)(SEXP, SEXP, SEXP, SEXP);
extern void (*ptr_R_savehistory)(SEXP, SEXP, SEXP, SEXP);
extern void (*ptr_R_addhistory)(SEXP, SEXP, SEXP, SEXP);

```

which allow standard R callbacks to be redirected to your GUI. What these do is generally documented in the file ‘`src/unix/system.txt`’.

`void R_ShowMessage (char *message)` [Function]

This should display the message, which may have multiple lines: it should be brought to the user's attention immediately.

`void R_Busy (int which)` [Function]

This function invokes actions (such as change of cursor) when R embarks on an extended computation (*which*=1) and when such a state terminates (*which*=0).

```

int R_ReadConsole (char *prompt, char *buf, int buflen, int hist)      [Function]
void R_WriteConsole (char *buf, int buflen)                            [Function]
void R_ResetConsole ()                                                [Function]
void R_FlushConsole ()                                                [Function]
void R_ClearErrConsole ()                                             [Function]

```

These functions interact with a console.

`R_ReadConsole` prints the given prompt at the console and then does a `gets(3)`-like operation, transferring up to `buflen` characters into the buffer `buf`. The last two characters should be set to `"\n\0"` to preserve sanity. If `hist` is non-zero, then the line should be added to any command history which is being maintained. The return value is 0 on success and >0 on failure.

`R_WriteConsole` writes the given buffer to the console.

`R_ResetConsole` is called when the system is reset after an error. `R_FlushConsole` is called to flush any pending output to the system console. `R_ClearerrConsole` clears any errors associated with reading from the console.

```

void R_ShowFiles (int nfile, char **file, char **headers, char *wtitle,  [Function]
                  Rboolean del, char *pager)

```

This function is used to display the contents of files.

```

int R_ChooseFile (int new, char *buf, int len)                        [Function]

```

Choose a file and return its name in `buf` of length `len`. Return value is 0 for success, > 0 otherwise.

```

int R_EditFile (char *buf)                                           [Function]

```

Send a file to an editor window.

```

SEXP R_loadhistory (SEXP, SEXP, SEXP, SEXP);                        [Function]
SEXP R_savehistory (SEXP, SEXP, SEXP, SEXP);                        [Function]
SEXP R_addhistory (SEXP, SEXP, SEXP, SEXP);                         [Function]

```

.Internal functions for `loadhistory`, `savehistory` and `timestamp`: these are called after checking the number of arguments.

If the console has no history mechanism these can be as simple as

```

SEXP R_loadhistory (SEXP call, SEXP op, SEXP args, SEXP env)
{
    errorcall(call, "loadhistory is not implemented");
    return R_NilValue;
}
SEXP R_savehistory (SEXP call, SEXP op, SEXP args, SEXP env)
{
    errorcall(call, "savehistory is not implemented");
    return R_NilValue;
}
SEXP R_addhistory (SEXP call, SEXP op, SEXP args, SEXP env)
{
    return R_NilValue;
}

```

The `R_addhistory` function should return silently if no history mechanism is present, as a user may be calling `timestamp` purely to write the time stamp to the console.

```

void R_Suicide (char *message)                                       [Function]

```

This should abort R as rapidly as possible, displaying the message. A possible implementation is

```

void R_Suicide (char *message)
{
    char pp[1024];
    snprintf(pp, 1024, "Fatal error: %s\n", s);
    R_ShowMessage(pp);
    R_CleanUp(SA_SUICIDE, 2, 0);
}

```

`void R_CleanUp (SA_TYPE saveact, int status, int RunLast)` [Function]

This function invokes any actions which occur at system termination. It needs to be quite complex:

```

#include <Rinterface.h>
#include <Rdevices.h>    /* for KillAllDevices */

void R_CleanUp (SA_TYPE saveact, int status, int RunLast)
{
    char *tmpdir;

    if(saveact == SA_DEFAULT) saveact = SaveAction;
    if(saveact == SA_SAVEASK) {
        /* ask what to do and set saveact */
    }
    switch (saveact) {
    case SA_SAVE:
        if(runLast) R_dot_Last();
        if(R_DirtyImage) R_SaveGlobalEnv();
        /* save the console history in R_HistoryFile */
        break;
    case SA_NOSAVE:
        if(runLast) R_dot_Last();
        break;
    case SA_SUICIDE:
    default:
        break;
    }

    R_RunExitFinalizers();
    /* clean up after the editor e.g. CleanEd() */

    if((tmpdir = getenv("R_SESSION_TMPDIR"))) {
        char buf[1024]
        snprintf((char *)buf, 1024, "rm -rf %s", tmpdir);
        R_system((char *)buf);
    }

    /* close all the graphics devices */
    if(saveact != SA_SUICIDE) KillAllDevices();
    fpu_setup(FALSE);

    exit(status);
}

```

8.1.2 Meshing event loops

One of the most difficult issues in interfacing R to a front-end is the handling of event loops, at least if a single thread is used. R uses events and timers for

- Running X11 windows such as the graphics device and data editor, and interacting with them (e.g., using `locator()`).
- Supporting Tcl/Tk events for the **tcltk** package.
- Preparing input.
- Timing operations, for example for profiling R code and `Sys.sleep()`.
- Interrupts, where permitted.

Specifically, the Unix command-line version of R runs separate event loops for

- Preparing input at the console command-line, in file `'src/unix/sys-unix.c'`.
- Waiting for a response from a socket in the internal functions underlying FTP and HTTP transfers in `download.file()` and for direct socket access, in files `'src/modules/internet/nanoftp.c'`, `'src/modules/internet/nanohttp.c'` and `'src/modules/internet/Rsock.c'`
- Mouse and window events when displaying the X11-based dataentry window, in file `'src/modules/X11/dataentry.c'`. This is regarded as *modal*, and no other events are serviced whilst it is active.

There is a protocol for adding event handlers to the first two types of event loops, using types and functions declared in the header `'R_ext/eventloop.h'` and described in comments in file `'src/unix/sys-std.c'`. It is possible to add (or remove) an input handler for events on a particular file descriptor, or to set a polling interval (*via* `R_wait_usec`) and a function to be called periodically via `R_PolledEvents`: the polling mechanism is used by the **tcltk** package.

An alternative front-end needs both to make provision for other R events whilst waiting for input, and to ensure that it is not frozen out during events of the second type. This is not handled very well in the existing examples. The GNOME front-end can run a own handler for polled events by setting

```
extern int (*R_timeout_handler)();
extern long R_timeout_val;

if (R_timeout_handler && R_timeout_val)
    gtk_timeout_add(R_timeout_val, R_timeout_handler, NULL);
gtk_main ();
```

whilst it is waiting for console input. This obviously handles events for Gtk windows (such as the graphics device in the **gtkDevice** package), but not X11 events (such as the `X11()` device) or for other event handlers that might have been registered with R. It does not attempt to keep itself alive whilst R is waiting on sockets. The ability to add a polled handler as `R_timeout_handler` is used by the **tcltk** package.

8.1.3 Threading issues

Embedded R is designed to be run in the main thread, and all the testing is done in that context. There is a potential issue with the stack-checking mechanism where threads are involved. This uses two variables declared in `'Rinterface.h'` (if `CSTACK_DEFN`s is defined) as

```
extern uintptr_t R_CStackLimit; /* C stack limit */
extern uintptr_t R_CStackStart; /* Initial stack address */
```

Note that `uintptr_t` is a C99 type for which a substitute is defined in R, so your code needs to define `HAVE_UINTPTR_T` appropriately.

These will be set³ when `R_initialize_R` is called, to values appropriate to the main thread. Stack-checking can be disabled by setting `R_CStackLimit = (uintptr_t)-1`, but it is better to if possible set appropriate values. (What these are and how to determine them are OS-specific, and the stack size limit may differ for secondary threads. If you have a choice of stack size, at least 8Mb is recommended.)

You may also want to consider how signals are handled: R sets signal handlers for several signals, including `SIGINT`, `SIGSEGV`, `SIGPIPE`, `SIGUSR1` and `SIGUSR2`, but these can all be suppressed by setting the variable `R_SignalHandlers` (declared in `Rinterface.h`) to 0.

8.2 Embedding R under Windows

All Windows interfaces to R call entry points in the DLL `'R.dll'`, directly or indirectly. Simpler applications may find it easier to use the indirect route via (D)COM.

8.2.1 Using (D)COM

(D)COM is a standard Windows mechanism used for communication between Windows applications. One application (here R) is run as COM server which offers services to clients, here the front-end calling application. The services are described in a 'Type Library' and are (more or less) language-independent, so the calling application can be written in C or C++ or Visual Basic or Perl or Python and so on. The 'D' in (D)COM refers to 'distributed', as the client and server can be running on different machines.

The basic R distribution is not a (D)COM server, but two addons are currently available that interface directly with R and provide a (D)COM server:

- There is a (D)COM server called `StatConnector` written by Thomas Baier available on CRAN (<http://cran.r-project.org/other-software.html>) which works with `'Rproxy.dll'` (in the R distribution) and `'R.dll'` to support transfer of data to and from R and remote execution of R commands, as well as embedding of an R graphics window. The `rcom` package on CRAN provides a (D)COM server in a running R session.
- Another (D)COM server, `RDCOMServer`, is available from <http://www.omegahat.org/>. Its philosophy is discussed in <http://www.omegahat.org/RDCOMServer/Docs/Paradigm.html> and is very different from the purpose of this section.

8.2.2 Calling R.dll directly

The R DLL is mainly written in C and has `_cdecl` entry points. Calling it directly will be tricky except from C code (or C++ with a little care).

Examples of calling `'R.dll'` directly are provided in the directory `'src/gnuwin32/front-ends'`, including `'Rproxy.dll'` used by `StatConnector` and a simple command-line front end `'rtest.c'` whose code is

```
#define Win32
#include <windows.h>
#include <stdio.h>
#include <Rversion.h>
#include <R_ext/RStartup.h>
/* for askok and askyesnocancel */
#include "graphapp/graphapp.h"

/* for signal-handling code */
#include <psignal.h>
```

³ at least on platforms where the values are available, that is having `getrlimit` and on Linux or having `sysctl` supporting `KERN_USRSTACK`, including FreeBSD and MacOS X.

```

void R_Suicide(char*); /* In Rinterface.h */
void R_CleanUp(SA_TYPE, int, int); /* from Startup.h */

/* one way to allow user interrupts: called in ProcessEvents */
__declspec(dllimport) int UserBreak;

/* calls into the R DLL */
extern char *getDLLVersion(), *getRUser(), *get_R_HOME();
extern void R_DefParams(Rstart), R_SetParams(Rstart), R_setStartTime();
extern void setup_term_ui(void), ProcessEvents(void);
extern void run_Rmainloop(void), end_Rmainloop(void), R_ReplDLLinit(void);
extern int R_ReplDLLdo1();

/* simple input, simple output */

/* This version blocks all events: a real one needs to call ProcessEvents
   frequently. See rterm.c and ../system.c for one approach using
   a separate thread for input.
*/
int myReadConsole(char *prompt, char *buf, int len, int addtohistory)
{
    fputs(prompt, stdout);
    fflush(stdout);
    if(fgets(buf, len, stdin)) return 1; else return 0;
}

void myWriteConsole(char *buf, int len)
{
    printf("%s", buf);
}

void myCallBack()
{
    /* called during i/o, eval, graphics in ProcessEvents */
}

void myBusy(int which)
{
    /* set a busy cursor ... if which = 1, unset if which = 0 */
}

static void my_onintr(int sig) { UserBreak = 1; }

int main (int argc, char **argv)
{
    structRstart rp;
    Rstart Rp = &rp;
    char Rversion[25], *RHome;

    sprintf(Rversion, "%s.%s", R_MAJOR, R_MINOR);
    if(strcmp(getDLLVersion(), Rversion) != 0) {
        fprintf(stderr, "Error: R.DLL version does not match\n");
        exit(1);
    }

    R_setStartTime();
    R_DefParams(Rp);
    if((RHome = get_R_HOME()) == NULL) {
        fprintf(stderr, "R_HOME must be set in the environment or Registry\n");
        exit(1);
    }
    Rp->rhome = RHome;
    Rp->home = getRUser();
    Rp->CharacterMode = LinkDLL;
    Rp->ReadConsole = myReadConsole;

```

```

Rp->WriteConsole = myWriteConsole;
Rp->CallBack = myCallBack;
Rp->ShowMessage = askok;
Rp->YesNoCancel = askyesnocancel;
Rp->Busy = myBusy;

Rp->R_Quiet = TRUE;
Rp->R_Interactive = FALSE;
Rp->RestoreAction = SA_RESTORE;
Rp->SaveAction = SA_NOSAVE;
R_SetParams(Rp); /* so R_ShowMessage is set */
R_SizeFromEnv(Rp);
R_SetParams(Rp);
R_set_command_line_arguments(argc, argv);

FlushConsoleInputBuffer(GetStdHandle(STD_INPUT_HANDLE));

signal(SIGBREAK, my_onintr);
setup_term_ui(); /* initialize graphapp, eventloop, read Rconsole */
setup_Rmainloop();
#ifdef SIMPLE_CASE
    run_Rmainloop();
    end_Rmainloop();
#else
    R_ReplDLLinit();
    while(R_ReplDLLdo1() > 0) {
/* add user actions here if desired */
    }
/* only get here on EOF (not q()) */
    R_CleanUp(SA_DEFAULT, 0, 1);
#endif
    end_Rmainloop();
    return 0;
}

```

The ideas are

- Check that the front-end and the linked ‘R.dll’ match – other front-ends may allow a looser match.
- Find and set the R home directory and the user’s home directory. The former may be available from the Windows Registry: it will normally be in HKEY_LOCAL_MACHINE\Software\R-core\R\InstallPath and can be set there by running the program ‘R_HOME\bin\RSetReg.exe’.
- Define startup conditions and callbacks via the `Rstart` structure. `R_DefParams` sets the defaults, and `R_SetParams` sets updated values.
- Record the command-line arguments used by `R_set_command_line_arguments` for use by the R function `commandArgs()`.
- Set up the signal handler and the basic user interface.
- Run the main R loop, possibly with our actions intermeshed.
- Arrange to clean up.

An underlying theme is the need to keep the GUI ‘alive’, and this has not been done in this example. The R callback `R_ProcessEvents` needs to be called frequently to ensure that Windows events in R windows are handled expeditiously. Conversely, R needs to allow the GUI code (which is running in the same process) to update itself as needed – two ways are provided to allow this:

- `R_ProcessEvents` calls the callback registered by `Rp->callback`. A version of this is used to run package Tcl/Tk for `tcltk` under Windows, for the code is

```
void R_ProcessEvents(void)
```

```

{
    while (peekevent()) doevent(); /* Windows events for GraphApp */
    if (UserBreak) { UserBreak = FALSE; onintr(); }
    R_CallBackHook();
    if(R_tcldo) R_tcldo();
}

```

- The mainloop can be split up to allow the calling application to take some action after each line of input has been dealt with: see the alternative code below `#ifdef SIMPLE_CASE`.

It may be that no R GraphApp windows need to be considered, although these include pagers, the `windows()` graphics device, the R data and script editors and various popups such as `choose.file()` and `select.list()`. It would be possible to replace all of these, but it seems easier to allow GraphApp to handle most of them.

It is possible to run R in a GUI in a single thread (as `RGui.exe` shows) but it will normally be easier⁴ to use multiple threads.

Note that R's own front ends use a stack size of 10Mb, whereas MinGW executables default to 2Mb, and Visual C++ ones to 1Mb. The latter stack sizes are too small for a number of R applications, so general-purpose front-ends should use a larger stack size.

⁴ An attempt to use only threads in the late 1990s failed to work correctly under Windows 95, the predominant version of Windows at that time.

Appendix A R (internal) programming miscellanea

A.1 .Internal and .Primitive

C code compiled into R at build time can be called “directly” or via the `.Internal` interface, which is very similar to the `.External` interface except in syntax. More precisely, R maintains a table of R function names and corresponding C functions to call, which by convention all start with ‘do_’ and return a SEXP. Via this table (`R_FunTab` in file ‘`src/main/names.c`’) one can also specify how many arguments to a function are required or allowed, whether the arguments are to be evaluated before calling or not, and whether the function is “internal” in the sense that it must be accessed via the `.Internal` interface, or directly accessible in which case it is printed in R as `.Primitive`.

R’s functionality can also be extended by providing corresponding C code and adding to this function table.

In general, all such functions use `.Internal()` as this is safer and in particular allows for transparent handling of named and default arguments. For example, `axis` is defined as

```
axis <- function(side, at = NULL, labels = NULL, ...)
  .Internal(axis(side, at, labels, ...))
```

However, for reasons of convenience and also efficiency (as there is some overhead in using the `.Internal` interface), there are exceptions which can be accessed directly. Note that these functions make no use of R code, and hence are very different from the usual interpreted functions. In particular, `args` and `body` return `NULL` for such objects.

The list of these “primitive” functions is subject to change: currently, it includes the following.

1. “Special functions” which really are *language* elements, however exist as “primitive” functions in R:

```
{      (      if      for      while repeat break next
return function quote on.exit
```

2. Basic *operators* (i.e., functions usually *not* called as `foo(a, b, ...)`) for subsetting, assignment, arithmetic and logic. These are the following 1-, 2-, and N -argument functions:

```
[      [[      $      <-      <<-      =      [<-      [[<-      $<-

+      -      *      /      ^      %%      %*%      %/%
<      <=      ==      !=      >=      >
|      ||      &      &&      !
```

3. “Low level” 0- and 1-argument functions which belong to one of the following groups of functions:

- a. Basic mathematical functions with a single argument, i.e.,

```
sign      abs
floor     ceiling trunc

sqrt      exp
cos       sin      tan
acos      asin     atan
cosh      sinh     tanh
acosh     asinh    atanh

cumsum    cumprod
cummax    cummin
```

```

      Im      Re
      Arg     Conj      Mod

```

Note however that the R function `log` has an optional named argument `base`, and therefore is defined as

```

log <- function(x, base = exp(1)) {
  if(missing(base))
    .Internal(log(x))
  else
    .Internal(log(x, base))
}

```

in order to ensure that `log(x = pi, base = 2)` is identical to `log(base = 2, x = pi)`.

- b. Functions rarely used outside of “programming” (i.e., mostly used inside other functions), such as

```

nargs      missing
interactive is.xxx
.Primitive .Internal
symbol.C    symbol.For
globalenv   baseenv      emptyenv      pos.to.env
unclass

```

(where *xxx* stands for almost 30 different notions, such as `function`, `vector`, `numeric`, and so forth, but not `is.loaded`).

- c. The programming and session management utilities

```

debug      undebg      trace  untrace
browser    proc.time   gc.time

```

4. The following basic assignment and extractor functions

```

length      length<-
class       class<-
oldClass    oldClass<-
attr        attr<-
attributes  attributes<-
dim         dim<-
dimnames    dimnames<-
environment<-

```

5. The following few *N*-argument functions are “primitive” for efficiency reasons:

```

:          ~          c          list
call       as.call    as.character
expression substitute as.environment
UseMethod  invisible  standardGeneric
.C         .Fortran   .Call      .External
.Call.graphics .External.graphics
.subset    .subset2   .primTrace .primUntrace
reg.finalizer

```

A.2 Testing R code

When you (as R developer) add new functions to the R base (all the packages distributed with R), be careful to check if *make test-Specific* or particularly, *cd tests; make no-segfault.Rout* still works (without interactive user intervention, and on a standalone computer). If the new function, for example, accesses the Internet, or requires GUI interaction, please add its name to the “stop list” in ‘`tests/no-segfault.Rin`’.

Appendix B R coding standards

R is meant to run on a wide variety of platforms, including Linux and most variants of Unix as well as 32-bit Windows versions and on MacOS X. Therefore, when extending R by either adding to the R base distribution or by providing an add-on package, one should not rely on features specific to only a few supported platforms, if this can be avoided. In particular, although most R developers use GNU tools, they should not employ the GNU extensions to standard tools. Whereas some other software packages explicitly rely on e.g. GNU make or the GNU C++ compiler, R does not. Nevertheless, R is a GNU project, and the spirit of the *GNU Coding Standards* should be followed if possible.

The following tools can “safely be assumed” for R extensions.

- As from R 2.3.0 an ISO C99 C compiler. Note that extensions such as POSIX 1003.1 must be tested for, typically using Autoconf (see [Section 1.2 \[Configure and cleanup\]](#), page 8), unless you are sure they are supported on all mainstream R platforms (including Windows and MacOS X). Packages will be more portable if written assuming only C89, but this should not be done where using C99 features will make for cleaner or more robust code.
- A FORTRAN 77 compiler (but not FORTRAN 9x).
- A simple `make`, considering the features of `make` in 4.2 BSD systems as a baseline.

GNU or other extensions, including pattern rules using ‘%’, the automatic variable ‘\$^’, the ‘+=’ syntax to append to the value of a variable, the (“safe”) inclusion of makefiles with no error, conditional execution, and many more, must not be used (see Chapter “Features” in the *GNU Make Manual* for more information). On the other hand, building R in a separate directory (not containing the sources) should work provided that `make` supports the `VPATH` mechanism.

Windows-specific makefiles can assume GNU `make` 3.75 or later, as no other `make` is viable on that platform.

- A Bourne shell and the “traditional” Unix programming tools, including `grep`, `sed`, and `awk`.

There are POSIX standards for these tools, but these may not fully be supported. Baseline features could be determined from a book such as *The UNIX Programming Environment* by Brian W. Kernighan & Rob Pike. Note in particular that ‘|’ in a regexp is an extended regexp, and is not supported by all versions of `grep` or `sed`. The Open Group Base Specifications, Issue 6, which is technically identical to ISO/IEC 9945 and IEEE Std 1003.1 (POSIX), 2004, are available at <http://www.opengroup.org/onlinepubs/009695399/mindex.html>.

Under Windows, most users will not have these tools installed, and you should not require their presence for the operation of your package. However, users who install your package from source will have them, as they can be assumed to have followed the instructions in “the Windows toolset” appendix of the “R Installation and Administration” manual to obtain them. Redirection cannot be assumed to be available via `system` as this does not use a standard shell (let alone a Bourne shell).

In addition, the following tools are needed for certain tasks.

- Perl version 5 is needed for converting documentation written in Rd format to plain text, HTML, \LaTeX , and to extract the examples. In addition, several other tools, in particular `check` and `build` (see [Section 1.3 \[Checking and building packages\]](#), page 13), require Perl. The R Core Team has decided that Perl (version 5) can safely be assumed for building R from source, building and checking add-on packages, and for installing add-on packages from source. On the other hand, Perl cannot be assumed at all for installing *binary* (pre-built) versions of add-on packages, or at run time.

- Makeinfo version 4.7 is needed to build the Info files for the R manuals written in the GNU Texinfo system. (Future distributions of R may contain the Info files.)

It is also important that code is written in a way that allows others to understand it. This is particularly helpful for fixing problems, and includes using self-descriptive variable names, commenting the code, and also formatting it properly. The R Core Team recommends to use a basic indentation of 4 for R and C (and most likely also Perl) code, and 2 for documentation in Rd format. Emacs users can implement this indentation style by putting the following in one of their startup files. (For GNU Emacs 20: for GNU Emacs 21 use customization to set the `c-default-style` to "bsd" and `c-basic-offset` to 4.)

```
;;; C
(add-hook 'c-mode-hook
  (lambda () (c-set-style "bsd")))

;;; ESS
(add-hook 'ess-mode-hook
  (lambda ()
    (ess-set-style 'C++)
    ;; Because
    ;;
    

|                                   | DEF | GNU | BSD | K&R | C++ |
|-----------------------------------|-----|-----|-----|-----|-----|
| ;; ess-indent-level               | 2   | 2   | 8   | 5   | 4   |
| ;; ess-continued-statement-offset | 2   | 2   | 8   | 5   | 4   |
| ;; ess-brace-offset               | 0   | 0   | -8  | -5  | -4  |
| ;; ess-arg-function-offset        | 2   | 4   | 0   | 0   | 0   |
| ;; ess-expression-offset          | 4   | 2   | 8   | 5   | 4   |
| ;; ess-else-offset                | 0   | 0   | 0   | 0   | 0   |
| ;; ess-close-brace-offset         | 0   | 0   | 0   | 0   | 0   |


    (add-hook 'local-write-file-hooks
      (lambda ()
        (ess-nyuke-trailing-whitespace))))
  (setq ess-nyuke-trailing-whitespace-p 'ask)
  ;; or even
  ;; (setq ess-nyuke-trailing-whitespace-p t)
  ;;; Perl
  (add-hook 'perl-mode-hook
    (lambda () (setq perl-indent-level 4))))
```

(The 'GNU' styles for Emacs' C and R modes use a basic indentation of 2, which has been determined not to display the structure clearly enough when using narrow fonts.)

Function and variable index

*

*Riconv_open..... 89

.

.C..... 52
 .Call..... 59, 67
 .External..... 59, 69
 .First.lib..... 6
 .Fortran..... 52
 .Internal..... 103
 .Last.lib..... 6, 19
 .onAttach..... 19
 .onLoad..... 19
 .onUnload..... 19
 .Primitive..... 103
 .Random.seed..... 80

\

\acronym..... 34
 \alias..... 28
 \arguments..... 29
 \author..... 30
 \bold..... 33
 \cite..... 33
 \code..... 33
 \command..... 33
 \concept..... 36
 \cr..... 32
 \deqn..... 35
 \describe..... 34
 \description..... 28
 \details..... 29
 \dfn..... 33
 \dontrun..... 30
 \dontshow..... 30
 \dots..... 35
 \dQuote..... 33
 \email..... 33
 \emph..... 32
 \enc..... 36
 \enumerate..... 34
 \env..... 33
 \eqn..... 35
 \examples..... 30
 \file..... 33
 \format..... 31
 \itemize..... 34
 \kbd..... 33
 \keyword..... 30
 \ldots..... 35
 \link..... 34
 \method..... 29
 \name..... 28
 \note..... 29
 \option..... 33
 \pkg..... 33
 \preformatted..... 33
 \R..... 35
 \references..... 29

\samp..... 33
 \section..... 32
 \seealso..... 30
 \source..... 31
 \sQuote..... 33
 \strong..... 32
 \synopsis..... 28
 \tabular..... 34
 \title..... 28
 \url..... 33
 \usage..... 28
 \value..... 29
 \var..... 33

B

bessel_i..... 84
 bessel_j..... 84
 bessel_k..... 84
 bessel_y..... 84
 beta..... 83
 BLAS_LIBS..... 10
 browser..... 42

C

Calloc..... 79
 CAR..... 69
 CDR..... 69
 cgmin..... 86
 choose..... 83
 cPsort..... 88

D

debug..... 46
 debugger..... 44
 defineVar..... 66
 digamma..... 83
 dump.frames..... 44
 duplicate..... 66
 dyn.load..... 54
 dyn.unload..... 54

E

exp_rand..... 80
 expm1..... 84
 export..... 18
 exportClasses..... 22
 exportMethods..... 22
 exportPattern..... 18

F

FALSE..... 85
 findInterval..... 88
 findVar..... 66
 FLIBS..... 10
 fmax2..... 84
 fmin2..... 84

fprec	85
Free	79
fround	85
fsign	85
ftrunc	85

G

gammafn	83
gctorture	47
getAttrib	64
GetRNGstate	80

H

hsv2rgb	88
---------------	----

I

imax2	84
imin2	84
import	18
importClassesFrom	22
importFrom	18
importMethodsFrom	22
install	64
iPsort	88
ISNA	70, 80
ISNAN	70, 80

L

LAPACK_LIBS	10
lbeta	83
lbfgsb	86
lchoose	83
lgamma1p	84
lgammafn	83
library.dynam	6, 54
log1p	84
log1pmx	84
logspace_add	84
logspace_sub	84

M

M_E	85
M_PI	85

N

NA_REAL	80
nmmin	86
norm_rand	80

P

pentagamma	83
PKG_CFLAGS	57
PKG_CPPFLAGS	57
PKG_CXXFLAGS	57
PKG_FFLAGS	57
PKG_LIBS	57
prompt	30

PROTECT	60
PROTECT_WITH_INDEX	61
psigamma	83
PutRNGstate	80
pythag	84

Q

qsort3	88
qsort4	88

R

R CMD build	15
R CMD check	13
R CMD config	9
R CMD Rd2dvi	37
R CMD Rd2txt	37
R CMD Rdconv	37
R CMD Sd2Rd	37
R CMD SHLIB	57
R_addhistory	96
R_alloc	78
R_Busy	95
R_ChooseFile	96
R_CleanUp	97
R_ClearErrConsole	96
R_csort	87
R_EditFile	96
R_ExpandFileName	89
R_FINITE	80
R_FlushConsole	96
R_INLINE	90
R_IsNaN	80
R_isort	87
R_LIBRARY_DIR	10
R_loadhistory	96
R_max_col	88
R_NegInf	80
R_PACKAGE_DIR	10
R_PosInf	80
R_pow	84
R_pow_di	84
R_qsort	88
R_qsort_I	88
R_qsort_int	88
R_qsort_int_I	88
R_ReadConsole	96
R_ResetConsole	96
R_rsort	87
R_savehistory	96
R_ShowFiles	96
R_ShowMessage	95
R_Suicide	96
R_tmpnam	89
R_Version	90
R_WriteConsole	96
rcont2	80
Rdqagi	87
Rdqags	87
Realloc	79
recover	45
REprintf	81
REPROTECT	61

REvprintf.....	81
revsort.....	88
rgb2hsv.....	88
Riconv.....	89
Riconv_close.....	89
Rprintf.....	81
Rprof.....	38
rPsort.....	88
rsort_with_index.....	87
Rvprintf.....	81

S

S_alloc.....	78
S_realloc.....	78
S3method.....	19
SAFE_FFLAGS.....	11
samin.....	86
seed_in.....	80
seed_out.....	80
setAttrib.....	64
setVar.....	66
sign.....	84
system.....	52

system.time.....	52
------------------	----

T

tetragamma.....	83
trace.....	47
traceback.....	43
trigamma.....	83
TRUE.....	85

U

undebug.....	47
unif_rand.....	80
UNPROTECT.....	60
UNPROTECT_PTR.....	61
useDynLib.....	19

V

vmxget.....	78
vmxset.....	78
vmmin.....	86

Concept index

\
 \linkS4class 34

A

Allocating storage 62
 Attributes 63

B

Bessel functions 84
 Beta function 83
 Building packages 15

C

C++ code, interfacing 57
 Calling C from FORTRAN and vice versa 81
 Checking packages 13
 Classes 64
 cleanup file 2
 configure file 2
 COPYING file 2
 Copying objects 66
 CRAN 17
 CRAN submission 17
 Creating packages 2
 Creating shared objects 57
 Cross-references in documentation 34
 cumulative hazard 82

D

Debugging 48
 DESCRIPTION file 2
 Details of R types 62
 Distribution functions from C 82
 Documentation, writing 27
 Dynamic loading 54

E

encoding 36
 Error handling from C 79
 Error handling from FORTRAN 79
 Evaluating R expressions from C 71

F

Finding variables 65

G

Gamma function 83
 Garbage collection 60
 Generic functions 92

H

Handling lists 65

Handling R objects in C 59

I

IEEE special values 70, 80
 INDEX file 5
 Indices 36
 Inspecting R objects when debugging 50
 install.R file 5
 integration 86
 Interfaces to compiled code 52, 67
 Interfacing C++ code 57
 Interrupts 89

L

Lists and tables in documentation 34

M

Marking text in documentation 32
 Mathematics in documentation 35
 Memory allocation from C 78
 Method functions 92
 Missing values 70, 80

N

name spaces 18
 Numerical analysis subroutines from C 82
 Numerical derivatives 74

O

Operating system access 52
 optimization 86

P

Package builder 15
 Package bundles 7
 Package structure 2
 Package subdirectories 5
 Packages 2
 Parsing R code from C 76
 Platform-specific documentation 36
 Printing from C 81
 Printing from FORTRAN 81
 Processing Rd format 37
 Profiling 38, 40

R

R_PROFILE.R file 5
 Random numbers in C 80, 82
 Random numbers in FORTRAN 81
 Registering native routines 55

S

Setting variables	65
Sort functions from C	87
Submitting to CRAN	17
Sweave	16

T

Tidying R code	38
----------------------	----

V

Version information from C	90
vignettes	16

Z

Zero-finding	73
--------------------	----