

R Installation and Administration

Version 4.0.3 (2020-10-10)

R Core Team

This manual is for R, version 4.0.3 (2020-10-10).

Copyright © 2001–2020 R 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 Core Team.

Table of Contents

1	Obtaining R	1
1.1	Getting and unpacking the sources	1
1.2	Getting patched and development versions	1
1.2.1	Using Subversion and rsync	1
2	Installing R under Unix-alikes	3
2.1	Simple compilation	3
2.2	Help options	4
2.3	Making the manuals	5
2.4	Installation	7
2.5	Uninstallation	9
2.6	Sub-architectures	9
2.6.1	Multilib	10
2.7	Other Options	11
2.7.1	Debugging Symbols	11
2.7.2	OpenMP Support	11
2.7.3	C++ Support	12
2.7.4	Link-Time Optimization	13
2.7.4.1	LTO with GCC	14
2.7.4.2	LTO with LLVM	14
2.7.4.3	LTO for package checking	15
2.8	Testing an Installation	15
3	Installing R under Windows	17
3.1	Building from source	17
3.1.1	The Windows toolset	17
3.1.2	L ^A T _E X	17
3.2	Checking the build	18
3.3	Testing an Installation	18
4	Installing R under macOS	19
4.1	Running R under macOS	19
4.2	Uninstalling under macOS	20
4.3	Multiple versions	21
5	Running R	22

6	Add-on packages	23
6.1	Default packages	23
6.2	Managing libraries	23
6.3	Installing packages	24
6.3.1	Windows	25
6.3.2	macOS	26
6.3.3	Customizing package compilation	27
6.3.4	Multiple sub-architectures	28
6.3.5	Byte-compilation	29
6.3.6	External software	29
6.4	Updating packages	30
6.5	Removing packages	30
6.6	Setting up a package repository	30
6.7	Checking installed source packages	31
7	Internationalization and Localization	33
7.1	Locales	33
7.1.1	Locales under Unix-alikes	33
7.1.2	Locales under Windows	34
7.1.3	Locales under macOS	34
7.2	Localization of messages	34
8	Choosing between 32- and 64-bit builds	36
9	The standalone Rmath library	38
9.1	Unix-alikes	38
9.2	Windows	39
Appendix A Essential and useful other		
	programs under a Unix-alike	41
A.1	Essential programs and libraries	41
A.2	Useful libraries and programs	43
A.2.1	Tcl/Tk	45
A.2.2	Java support	46
A.2.3	Other compiled languages	47
A.3	Linear algebra	47
A.3.1	BLAS	47
A.3.1.1	ATLAS	48
A.3.1.2	OpenBLAS	49
A.3.1.3	Intel MKL	50
A.3.1.4	Shared BLAS	51
A.3.2	LAPACK	52
A.3.3	Caveats	53

Appendix B Configuration on a Unix-alike 54

B.1	Configuration options	54
B.2	Internationalization support	55
B.3	Configuration variables	55
B.3.1	Setting paper size	56
B.3.2	Setting the browsers	56
B.3.3	Compilation flags	56
B.3.4	Making manuals	56
B.4	Setting the shell	56
B.5	Using make	57
B.6	Using Fortran	57
B.7	Compile and load flags	57
B.8	Maintainer mode	59

Appendix C Platform notes 60

C.1	X11 issues	60
C.2	Linux	61
C.2.1	Clang	63
C.2.2	Intel compilers	64
C.3	macOS	65
C.3.1	Prerequisites	65
C.3.1.1	Note for Catalina users	67
C.3.2	Other C/C++ compilers	67
C.3.3	Other libraries	68
C.3.4	Tcl/Tk headers and libraries	69
C.3.5	Java	70
C.3.6	Frameworks	71
C.3.7	Building R.app	71
C.3.8	Building binary packages	71
C.4	Solaris	72
C.4.1	64-bit builds	74
C.4.2	Using gcc	75
C.5	FreeBSD	75
C.6	OpenBSD	75
C.7	Cygwin	76
C.8	New platforms	76

Function and variable index 77

Concept index 78

Environment variable index 79

1 Obtaining R

Sources, binaries and documentation for R can be obtained via CRAN, the “Comprehensive R Archive Network” whose current members are listed at <https://CRAN.R-project.org/mirrors.html>.

1.1 Getting and unpacking the sources

The simplest way is to download the most recent `R-x.y.z.tar.gz` file, and unpack it with

```
tar -xf R-x.y.z.tar.gz
```

on systems that have a suitable¹ `tar` installed. On other systems you need to have the `gzip` program installed, when you can use

```
gzip -dc R-x.y.z.tar.gz | tar -xf -
```

The pathname of the directory into which the sources are unpacked should not contain spaces, as most `make` programs (and specifically GNU `make`) do not expect spaces.

If you want the build to be usable by a group of users, set `umask` before unpacking so that the files will be readable by the target group (e.g., `umask 022` to be usable by all users). Keep this setting of `umask` whilst building and installing.

If you use a fairly recent GNU version of `tar` and do this as a root account (which on Windows includes accounts with administrator privileges) you may see many warnings about changing ownership. In which case you can use

```
tar --no-same-owner -xf R-x.y.z.tar.gz
```

and perhaps also include the option `--no-same-permissions`. (These options can also be set in the `TAR_OPTIONS` environment variable: if more than one option is included they should be separated by spaces.)

1.2 Getting patched and development versions

A patched version of the current release, ‘`r-patched`’, and the current development version, ‘`r-devel`’, are available as daily tarballs and via access to the R Subversion repository. (For the two weeks prior to the release of a minor (4.x.0) version, ‘`r-patched`’ tarballs may refer to beta/release candidates of the upcoming release, the patched version of the current release being available via Subversion.)

The tarballs are available from <https://stat.ethz.ch/R/daily/>. Download `R-patched.tar.gz` or `R-devel.tar.gz` (or the `.tar.bz2` versions) and unpack as described in the previous section. They are built in exactly the same way as distributions of R releases.

1.2.1 Using Subversion and rsync

Sources are also available via <https://svn.R-project.org/R/>, the R Subversion repository. If you have a Subversion client (see <https://subversion.apache.org/>), you can check out and update the current ‘`r-devel`’ from <https://svn.r-project.org/R/trunk/> and the current ‘`r-patched`’ from ‘<https://svn.r-project.org/R/branches/>’

¹ e.g. GNU `tar` version 1.15 or later, or that from the ‘`libarchive`’ (as used on macOS) or ‘`Heirloom Toolchest`’ distributions.

R-x-y-branch/' (where x and y are the major and minor number of the current released version of R). E.g., use

```
svn checkout https://svn.r-project.org/R/trunk/ path
```

to check out 'r-devel' into directory *path* (which will be created if necessary). The alpha, beta and RC versions of an upcoming x.y.0 release are available from 'https://svn.r-project.org/R/branches/R-x-y-branch/' in the four-week period prior to the release.

Note that 'https:' is required², and that the SSL certificate for the Subversion server of the R project should be recognized as from a trusted source.

Note that retrieving the sources by e.g. `wget -r` or `svn export` from that URL will not work (and will give a error early in the `make` process): the Subversion information is needed to build R.

The Subversion repository does not contain the current sources for the recommended packages, which can be obtained by `rsync` or downloaded from CRAN. To use `rsync` to install the appropriate sources for the recommended packages, run `./tools/rsync-recommended` from the top-level directory of the R sources.

If downloading manually from CRAN, do ensure that you have the correct versions of the recommended packages: if the number in the file `VERSION` is 'x.y.z' you need to download the contents of 'https://CRAN.R-project.org/src/contrib/dir', where *dir* is 'x.y.z/Recommended' for r-devel or x.y-patched/Recommended for r-patched, respectively, to directory `src/library/Recommended` in the sources you have unpacked. After downloading manually you need to execute `tools/link-recommended` from the top level of the sources to make the requisite links in `src/library/Recommended`. A suitable incantation from the top level of the R sources using `wget` might be (for the correct value of *dir*)

```
wget -r -ll --no-parent -A*.gz -nd -P src/library/Recommended \
  https://CRAN.R-project.org/src/contrib/dir
./tools/link-recommended
```

² for some Subversion clients 'http:' may appear to work, but requires continual redirection.

2 Installing R under Unix-alikes

R will configure and build under most common Unix and Unix-alike platforms including ‘`cpu-*-linux-gnu`’ for the ‘`alpha`’, ‘`arm64`’, ‘`hppa`’, ‘`ix86`’, ‘`m68k`’, ‘`mips`’, ‘`mipsel`’, ‘`ppc64`’, ‘`s390x`’, ‘`sparc64`’, and ‘`x86_64`’ CPUs, ‘`x86_64-apple-darwin`’ and ‘`i386-sun-solaris`’ as well as perhaps (it is tested less frequently on these platforms) ‘`i386-*-freebsd`’, ‘`x86_64-*-freebsd`’, ‘`i386-*-netbsd`’, ‘`x86_64-*-openbsd`’ and ‘`powerpc-ibm-aix6*`’

In addition, binary distributions are available for some common Linux distributions and for macOS (formerly OS X and Mac OS X). See the FAQ for current details. These are installed in platform-specific ways, so for the rest of this chapter we consider only building from the sources.

Cross-building is not possible: installing R builds a minimal version of R and then runs many R scripts to complete the build.

2.1 Simple compilation

First review the essential and useful tools and libraries in Appendix A [Essential and useful other programs under a Unix-alike], page 41, and install those you want or need. Ensure that either the environment variable `TMPDIR` is either unset (and `/tmp` exists and can be written in and scripts can be executed from) or points to the absolute path to a valid temporary directory (one from which execution of scripts is allowed) which does not contain spaces.¹

Choose a directory to install the R tree (R is not just a binary, but has additional data sets, help files, font metrics etc). Let us call this place *R_HOME*. Untar the source code. This should create directories `src`, `doc`, and several more under a top-level directory: change to that top-level directory (At this point North American readers should consult Section B.3.1 [Setting paper size], page 56.) Issue the following commands:

```
./configure
make
```

(See Section B.5 [Using make], page 57, if your make is not called ‘`make`’.) Users of Debian-based 64-bit systems² may need

```
./configure LIBnn=lib
make
```

Then check the built system works correctly by

```
make check
```

Failures are not necessarily problems as they might be caused by missing functionality, but you should look carefully at any reported discrepancies. (Some non-fatal errors are expected in locales that do not support Latin-1, in particular in true C locales and non-UTF-8 non-Western-European locales.) A failure in `tests/ok-errors.R` may indicate inadequate resource limits (see Chapter 5 [Running R], page 22).

More comprehensive testing can be done by

```
make check-devel
```

¹ Most aspects will work with paths containing spaces, but external software used by R may not.

² which use `lib` rather than `lib64` for their primary 64-bit library directories: attempts are made to detect such systems.

or

```
make check-all
```

see file `tests/README` and Section 2.8 [Testing a Unix-alike Installation], page 15, for the possibilities of doing this in parallel. Note that these checks are only run completely if the recommended packages are installed.

If the `configure` and `make` commands execute successfully, a shell-script front-end called `R` will be created and copied to `R_HOME/bin`. You can link or copy this script to a place where users can invoke it, for example to `/usr/local/bin/R`. You could also copy the man page `R.1` to a place where your `man` reader finds it, such as `/usr/local/man/man1`. If you want to install the complete R tree to, e.g., `/usr/local/lib/R`, see Section 2.4 [Installation], page 7. Note: you do not *need* to install R: you can run it from where it was built.

You do not necessarily have to build R in the top-level source directory (say, `TOP_SRCDIR`). To build in `BUILDDIR`, run

```
cd BUILDDIR
TOP_SRCDIR/configure
make
```

and so on, as described further below. This has the advantage of always keeping your source tree clean and is particularly recommended when you work with a version of R from Subversion. (You may need GNU `make` to allow this, and you will need no spaces in the path to the build directory. It is unlikely to work if the source directory has previously been used for a build.)

There are many settings which can be customized when building R and most are described in the file `config.site` in the top-level source directory. This can be edited, but for an installation using `BUILDDIR` it is better to put the changed settings in a newly-created file `config.site` in the build directory.

Now `rehash` if necessary, type `R`, and read the R manuals and the R FAQ (files `FAQ` or `doc/manual/R-FAQ.html`, or <https://CRAN.R-project.org/doc/FAQ/R-FAQ.html> which always has the version for the latest release of R).

Note: if you already have R installed, check that where you installed R replaces or comes earlier in your path than the previous installation. Some systems are set up to have `/usr/bin` (the standard place for a system installation) ahead of `/usr/local/bin` (the default place for installation of R) in their default path, and some do not have `/usr/local/bin` on the default path.

2.2 Help options

R by default provides help pages as plain text displayed in a pager, with the options (see the help for `help`) of displaying help as HTML or PDF.

By default HTML help pages are created when needed rather than being built at install time.

If you need to disable the server and want HTML help, there is the option to build HTML pages when packages are installed (including those installed with R). This is enabled by the `configure` option `--enable-prebuilt-html`. Whether `R CMD INSTALL` (and hence `install.packages`) pre-builds HTML pages is determined by looking at the R installation

and is reported by `R CMD INSTALL --help`: it can be overridden by specifying one of the `INSTALL` options `--html` or `--no-html`.

The server is disabled by setting the environment variable `R_DISABLE_HTTPD` to a non-empty value, either before R is started or within the R session before HTML help (including `help.start`) is used. It is also possible that system security measures will prevent the server from being started, for example if the loopback interface has been disabled. See `?tools::startDynamicHelp` for more details.

2.3 Making the manuals

There is a set of manuals that can be built from the sources,

<code>'fullrefman'</code>	Printed versions of all the help pages for base and recommended packages (around 3750 pages).
<code>'refman'</code>	Printed versions of the help pages for selected base packages (around 2200 pages)
<code>'R-FAQ'</code>	R FAQ
<code>'R-intro'</code>	"An Introduction to R".
<code>'R-data'</code>	"R Data Import/Export".
<code>'R-admin'</code>	"R Installation and Administration", this manual.
<code>'R-exts'</code>	"Writing R Extensions".
<code>'R-lang'</code>	"The R Language Definition".

To make these (with `'fullrefman'` rather than `'refman'`), use

```
make pdf      to create PDF versions
make info     to create info files (not 'refman' nor 'fullrefman').
```

You will not be able to build any of these unless you have `texi2any` version 5.1 or later installed, and for PDF you must have `texi2dvi` and `texinfo.tex` installed (which are part of the GNU `texinfo` distribution but are, especially `texinfo.tex`, often made part of the `TeX` package in re-distributions). The path to `texi2any` can be set by macro `'TEXI2ANY'` in `config.site`. NB: `texi2any` require `perl`.

The PDF versions can be viewed using any recent PDF viewer: they have hyperlinks that can be followed. The info files are suitable for reading online with Emacs or the standalone GNU `info` program. The PDF versions will be created using the paper size selected at configuration (default ISO a4): this can be overridden by setting `R_PAPERSIZE` on the `make` command line, or setting `R_PAPERSIZE` in the environment and using `make -e`. (If re-making the manuals for a different paper size, you should first delete the file `doc/manual/version.texi`. The usual value for North America would be `'letter'`.)

There are some issues with making the PDF reference manual, `fullrefman.pdf` or `refman.pdf`. The help files contain both ISO Latin1 characters (e.g. in `text.Rd`) and upright quotes, neither of which are contained in the standard `LATEX` Computer Modern fonts. We have provided four alternatives:

times (The default.) Using standard PostScript fonts, Times Roman, Helvetica and Courier. This works well both for on-screen viewing and for printing. One disadvantage is that the Usage and Examples sections may come out rather wide: this can be overcome by using *in addition* either of the options **inconsolata** (on a Unix-alike only if found by **configure**) or **beramono**, which replace the Courier monospaced font by Inconsolata or Bera Sans mono respectively. (You will need a recent version of the appropriate L^AT_EX package **inconsolata**³ or **bera** installed.)

Note that in most L^AT_EX installations this will not actually use the standard fonts for PDF, but rather embed the URW clones NimbusRom, NimbusSans and (for Courier, if used) NimbusMon.

This needs L^AT_EX packages **times**, **helvetic** and (if used) **courier** installed.

lm Using the *Latin Modern* fonts. These are not often installed as part of a T_EX distribution, but can be obtained from <https://www.ctan.org/tex-archive/fonts/ps-type1/lm/> and mirrors. This uses fonts rather similar to Computer Modern, but is not so good on-screen as **times**.

cm-super Using type-1 versions of the Computer Modern fonts by Vladimir Volovich. This is a large installation, obtainable from <https://www.ctan.org/tex-archive/fonts/ps-type1/cm-super/> and its mirrors. These type-1 fonts have poor hinting and so are nowhere near as readable on-screen as the other three options.

ae A package to use composites of Computer Modern fonts. This works well most of the time, and its PDF is more readable on-screen than the previous two options. There are three fonts for which it will need to use bitmapped fonts, **tctt0900.600pk**, **tctt1000.600pk** and **tcrm1000.600pk**. Unfortunately, if those files are not available, Acrobat Reader will substitute completely incorrect glyphs so you need to examine the logs carefully.

The default can be overridden by setting the environment variable **R_RD4PDF**. (On Unix-alikes, this will be picked up at install time and stored in **etc/Renviron**, but can still be overridden when the manuals are built, using **make -e**.) The usual⁴ default value for **R_RD4PDF** is **'times,inconsolata,hyper'**: omit **'hyper'** if you do not want hyperlinks (e.g. for printing the manual) or do not have L^AT_EX package **hyperref**, and omit **'inconsolata'** if you do not have L^AT_EX package **inconsolata** installed.

Further options, e.g. for **hyperref**, can be included in a file **Rd.cfg** somewhere on your L^AT_EX search path. For example, if you prefer to hyperlink the text and not the page number in the table of contents use

```
\ifthenelse{\boolean{Rd@use@hyper}}{\hypersetup{linktoc=section}}{}
```

or

```
\ifthenelse{\boolean{Rd@use@hyper}}{\hypersetup{linktoc=all}}{}
```

to hyperlink both text and page number.

³ Instructions on how to install the latest version are at <https://www.ctan.org/tex-archive/fonts/inconsolata/>.

⁴ on a Unix-alike, **'inconsolata'** is omitted if not found by **configure**.

Ebook versions of most of the manuals in one or both of `.epub` and `.mobi` formats can be made by running in `doc/manual` one of

```
make ebooks
make epub
make mobi
```

This requires `ebook-convert` from Calibre (<https://calibre-ebook.com/download>), or from most Linux distributions. If necessary the path to `ebook-convert` can be set as make macro `EBOOK` to by editing `doc/manual/Makefile` (which contains a commented value suitable for macOS) or using `make -e`.

2.4 Installation

To ensure that the installed tree is usable by the right group of users, set `umask` appropriately (perhaps to `'022'`) before unpacking the sources and throughout the build process.

After

```
./configure
make
make check
```

(or, when building outside the source, `TOP_SRCDIR/configure`, etc) have been completed successfully, you can install the complete R tree to your system by typing

```
make install
```

A parallel make can be used (but run `make` before `make install`). Those using GNU `make` 4.0 or later may want to use `make -j n -O` to avoid interleaving of output.

This will install to the following directories:

`prefix/bin` or `bindir`

the front-end shell script and other scripts and executables

`prefix/man/man1` or `mandir/man1`

the man page

`prefix/LIBnn/R` or `libdir/R`

all the rest (libraries, on-line help system, ...). Here `LIBnn` is usually `'lib'`, but may be `'lib64'` on some 64-bit Linux systems. This is known as the R home directory.

where `prefix` is determined during configuration (typically `/usr/local`) and can be set by running `configure` with the option `--prefix`, as in

```
./configure --prefix=/where/you/want/R/to/go
```

where the value should be an absolute path. This causes `make install` to install the R script to `/where/you/want/R/to/go/bin`, and so on. The prefix of the installation directories can be seen in the status message that is displayed at the end of `configure`. The installation may need to be done by the owner of `prefix`, often a root account.

There is the option of using `make install-strip` (see Section 2.7.1 [Debugging Symbols], page 11).

You can install into another directory tree by using

```
make prefix=/path/to/here install
```

at least with GNU **make** (but not some other Unix makes).

More precise control is available at configure time via options: see **configure --help** for details. (However, most of the ‘Fine tuning of the installation directories’ options are not used by R.)

Configure options **--bindir** and **--mandir** are supported and govern where a copy of the R script and the **man** page are installed.

The configure option **--libdir** controls where the main R files are installed: the default is **‘eprefix/LIBnn’**, where *eprefix* is the prefix used for installing architecture-dependent files, defaults to *prefix*, and can be set via the configure option **--exec-prefix**.

Each of **bindir**, **mandir** and **libdir** can also be specified on the **make install** command line (at least for GNU **make**).

The **configure** or **make** variables **rdocdir** and **rsharedir** can be used to install the system-independent **doc** and **share** directories to somewhere other than **libdir**. The C header files can be installed to the value of **rincludedir**: note that as the headers are not installed into a subdirectory you probably want something like **rincludedir=/usr/local/include/R-4.0.3**.

If you want the R home to be something other than **libdir/R**, use **rhome**: for example

```
make install rhome=/usr/local/lib64/R-4.0.3
```

will use a version-specific R home on a non-Debian Linux 64-bit system.

If you have made R as a shared/static library you can install it in your system’s library directory by

```
make prefix=/path/to/here install-libR
```

where **prefix** is optional, and **libdir** will give more precise control.⁵ However, you should not install to a directory mentioned in **LDPATHS** (e.g. **/usr/local/lib64**) if you intend to work with multiple versions of R, since that directory may be given precedence over the **lib** directory of other R installations.

```
make install-strip
```

will install stripped executables, and on platforms where this is supported, stripped libraries in directories **lib** and **modules** and in the standard packages.

Note that installing R into a directory whose path contains spaces is not supported, and some aspects (such as installing source packages) will not work.

To install info and PDF versions of the manuals, use one or both of

```
make install-info
make install-pdf
```

Once again, it is optional to specify **prefix**, **libdir** or **rhome** (the PDF manuals are installed under the R home directory).

More precise control is possible. For info, the setting used is that of **infodir** (default **prefix/info**, set by configure option **--infodir**). The PDF files are installed into the R **doc** tree, set by the **make** variable **rdocdir**.

A staged installation is possible, that it is installing R into a temporary directory in order to move the installed tree to its final destination. In this case **prefix** (and so on)

⁵ This will be needed if more than one sub-architecture is to be installed.

should reflect the final destination, and `DESTDIR` should be used: see https://www.gnu.org/prep/standards/html_node/DESTDIR.html.

You can optionally install the run-time tests that are part of `make check-all` by

```
make install-tests
```

which populates a `tests` directory in the installation.

2.5 Uninstallation

You can uninstall R by

```
make uninstall
```

optionally specifying `prefix` etc in the same way as specified for installation.

This will also uninstall any installed manuals. There are specific targets to uninstall info and PDF manuals in file `doc/manual/Makefile`.

Target `uninstall-tests` will uninstall any installed tests, as well as removing the directory `tests` containing the test results.

An installed shared/static `libR` can be uninstalled by

```
make prefix=/path/to/here uninstall-libR
```

2.6 Sub-architectures

Some platforms can support closely related builds of R which can share all but the executables and dynamic objects. Examples include builds under Linux and Solaris for different CPUs or 32- and 64-bit builds.

R supports the idea of architecture-specific builds, specified by adding `'r_arch=name'` to the `configure` line. Here *name* can be anything non-empty, and is used to name subdirectories of `lib`, `etc`, `include` and the package `libs` subdirectories. Example names from other software are the use of `sparcv9` on Sparc Solaris and `32` by `gcc` on `'x86_64'` Linux.

If you have two or more such builds you can install them over each other (and for 32/64-bit builds on one architecture, one build can be done without `'r_arch'`). The space savings can be considerable: on `'x86_64'` Linux a basic install (without debugging symbols) took 74Mb, and adding a 32-bit build added 6Mb. If you have installed multiple builds you can select which build to run by

```
R --arch=name
```

and just running `'R'` will run the last build that was installed.

`R CMD INSTALL` will detect if more than one build is installed and try to install packages with the appropriate library objects for each. This will not be done if the package has an executable `configure` script or a `src/Makefile` file. In such cases you can install for extra builds by

```
R --arch=name CMD INSTALL --libs-only pkg1 pkg2 ...
```

If you want to mix sub-architectures compiled on different platforms (for example `'x86_64'` Linux and `'i686'` Linux), it is wise to use explicit names for each, and you may also need to set `libdir` to ensure that they install into the same place.

When sub-architectures are used the version of `Rscript` in e.g. `/usr/bin` will be the last installed, but architecture-specific versions will be available in e.g.

`/usr/lib64/R/bin/exec${R_ARCH}`. Normally all installed architectures will run on the platform so the architecture of `Rscript` itself does not matter. The executable `Rscript` will run the R script, and at that time the setting of the `R_ARCH` environment variable determines the architecture which is run.

When running post-install tests with sub-architectures, use

```
R --arch=name CMD make check[-devel|all]
```

to select a sub-architecture to check.

Sub-architectures are also used on Windows, but by selecting executables within the appropriate `bin` directory, `R_HOME/bin/i386` or `R_HOME/bin/x64`. For backwards compatibility there are executables `R_HOME/bin/R.exe` and `R_HOME/bin/Rscript.exe`: these will run an executable from one of the subdirectories, which one being taken first from the `R_ARCH` environment variable, then from the `--arch` command-line option⁶ and finally from the installation default (which is 32-bit for a combined 32/64 bit R installation).

2.6.1 Multilib

For some Linux distributions⁷, there is an alternative mechanism for mixing 32-bit and 64-bit libraries known as *multilib*. If the Linux distribution supports multilib, then parallel builds of R may be installed in the sub-directories `lib` (32-bit) and `lib64` (64-bit). The build to be run may then be selected using the `setarch` command. For example, a 32-bit build may be run by

```
setarch i686 R
```

The `setarch` command is only operational if both 32-bit and 64-bit builds are installed. If there is only one installation of R, then this will always be run regardless of the architecture specified by the `setarch` command.

There can be problems with installing packages on the non-native architecture. It is a good idea to run e.g. `setarch i686 R` for sessions in which packages are to be installed, even if that is the only version of R installed (since this tells the package installation code the architecture needed).

There is a potential problem with packages using Java, as the post-install for a ‘i686’ RPM on ‘x86_64’ Linux reconfigures Java and will find the ‘x86_64’ Java. If you know where a 32-bit Java is installed you may be able to run (as root)

```
export JAVA_HOME=<path to jre directory of 32-bit Java>
setarch i686 R CMD javareconf
```

to get a suitable setting.

When this mechanism is used, the version of `Rscript` in e.g. `/usr/bin` will be the last installed, but an architecture-specific version will be available in e.g. `/usr/lib64/R/bin`. Normally all installed architectures will run on the platform so the architecture of `Rscript` does not matter.

⁶ with possible values ‘i386’, ‘x64’, ‘32’ and ‘64’.

⁷ mainly on RedHat and Fedora, whose layout is described here.

2.7 Other Options

There are many other installation options, most of which are listed by `configure --help`. Almost all of those not listed elsewhere in this manual are either standard `autoconf` options not relevant to R or intended for specialist uses by the R developers.

One that may be useful when working on R itself is the option `--disable-byte-compiled-packages`, which ensures that the base and recommended packages are not byte-compiled. (Alternatively the (make or environment) variable `R_NO_BASE_COMPILE` can be set to a non-empty value for the duration of the build.)

Option `--with-internal-tzcode` makes use of R's own code and copy of the IANA database for managing timezones. This will be preferred where there are issues with the system implementation, usually involving times after 2037 or before 1916. An alternative time-zone directory⁸ can be used, pointed to by environment variable `TZDIR`: this should contain files such as `Europe/London`. On all tested OSes the system timezone was deduced correctly, but if necessary it can be set as the value of environment variable `TZ`.

2.7.1 Debugging Symbols

By default, `configure` adds a flag (usually `-g`) to the compilation flags for C, Fortran and CXX sources. This will slow down compilation and increase object sizes of both R and packages, so it may be a good idea to change those flags (set `'CFLAGS'` etc in `config.site` before configuring, or edit files `Makeconf` and `etc/Makeconf` between running `configure` and `make`).

Having debugging symbols available is useful both when running R under a debugger (e.g., `R -d gdb`) and when using sanitizers and `valgrind`, all things intended for experts.

Debugging symbols (and some others) can be 'stripped' on installation by using

```
make install-strip
```

How well this is supported depends on the platform: it works best on those using GNU `binutils`. On 'x86_64' Linux a typical reduction in overall size was from 92MB to 66MB. On macOS debugging symbols are not by default included in `.dylib` and `.so` files, so there is negligible difference.

2.7.2 OpenMP Support

By default `configure` searches for suitable flags⁹ for OpenMP support for the C, C++ (default standard) and Fortran compilers.

Only the C result is currently used for R itself, and only if `MAIN_LD/DYLIB_LD` were not specified. This can be overridden by specifying

```
R_OPENMP_CFLAGS
```

Use for packages has similar restrictions (involving `SHLIB_LD` and similar: note that as Fortran code is by default linked by the C (or C++) compiler, both need to support OpenMP) and can be overridden by specifying some of

```
SHLIB_OPENMP_CFLAGS
```

⁸ How to prepare such a directory is described in file `src/extra/tzone/Notes` in the R sources.

⁹ for example, `-fopenmp`, `-xopenmp` or `-qopenmp`. This includes for `clang` and the Intel and Oracle compilers.


```
SHLIB_OPENMP_CXXFLAGS
SHLIB_OPENMP_FFLAGS
```

Setting these to an empty value will disable OpenMP for that compiler (and configuring with `--disable-openmp` will disable all detection¹⁰ of OpenMP). The `configure` detection test is to compile and link a standalone OpenMP program, which is not the same as compiling a shared object and loading it into the C program of R's executable. Note that overridden values are not tested.

2.7.3 C++ Support

C++ is not used by R itself, but support is provided for installing packages with C++ code via `make` macros defined in file `etc/Makeconf` (and with explanations in file `config.site`):

```
CXX
CXXFLAGS
CXXPICFLAGS
CXXSTD

CXX11
CXX11STD
CXX11FLAGS
CXX11PICFLAGS

CXX14
CXX14STD
CXX14FLAGS
CXX14PICFLAGS

CXX17
CXX17STD
CXX17FLAGS
CXX17PICFLAGS

CXX20
CXX20STD
CXX20FLAGS
CXX20PICFLAGS
```

The macros `CXX` etc are those used by default for C++ code. `configure` will attempt to set the rest suitably, choosing for `CXXSTD` and `CXX11STD` a suitable flag such as `-std=c++11` for C++11 support (which is required if C++ is to be supported at all). Similarly, `configure` will if possible choose for `CXX14STD` a flag¹¹ such as `-std=c++14` for C++14 support, `-std=c++17` or `-std=c++1z` for support for C++17 and `-std=c++20` or `-std=c++2a` for support for

¹⁰ This does not necessarily disable *use* of OpenMP – the `configure` code allows for platforms where OpenMP is used without a flag. For the `flang` compiler in late 2017, the Fortran runtime always used OpenMP.

¹¹ This was a valid option for `g++` 5 and later and 2016 versions of the Intel and Solaris compilers. For earlier versions of `g++` one could try `-std=c++1y`.

C++20.. The inferred values can be overridden in file `config.site` or on the `configure` command line: user-supplied values will be tested by compiling some C++11/14/17/20 code.

It may be that there is no suitable flag for C++14/17/20 support with the default compiler, in which case a different compiler could be selected for CXX14 or CXX17 or CXX20 with its corresponding flags.

The `-std` flag is supported by the GCC, `clang++`, Intel and Solaris compilers (the latter from version 12.4). Currently accepted values are (plus some synonyms)

```
g++:      c++11 gnu+11 c++14 gnu++14 c++17 gnu++17 c++2a gnu++2a (from 8)
          c++20 gnu++20 (from 10)
Intel:    c++11 c++14 (from 16.0) c++17 (from 17.0)
Solaris:  c++11 c++14 (from 12.5)
```

(Those for `clang++` are documented at https://clang.llvm.org/cxx_status.html, and follow `g++`: `-std=c++20` is supported from Clang 10.)

‘Standards’ for `g++` starting with ‘gnu’ enable ‘GNU extensions’: what those are is hard to track down.

For the use of C++11 and later in R packages see the ‘Writing R Extensions’ manual. Prior to R 3.6.0 the default C++ standard was that of the compiler used: currently it is C++11: this can be overridden by setting ‘CXXSTD’ when R is configured, for example to ‘`-std=gnu++14`’.

https://en.cppreference.com/w/cpp/compiler_support indicates which versions of common compilers support (parts of) which C++ standards.

2.7.4 Link-Time Optimization

There is support for using link-time optimization (LTO) if the toolchain supports it: configure with flag `--enable-lto`. When LTO is enabled it is used for compiled code in add-on packages. To use LTO for R itself but not add-on packages (including recommended packages) use the flag `--enable-lto=R`.

The main benefit seen to date from LTO has been detecting long-standing bugs in the ways packages pass arguments to compiled code and between compilation units. Benchmarking in 2020 with `gcc/gfortran` 10 showed gains of a few percent in increased performance and reduction in installed size for builds without debug symbols, but large size reductions for some packages¹² with debug symbols. (Performance and size gains are said to be most often seen in complex C++ builds.)

Whether toolchains support LTO is often unclear: all of the C compiler, the Fortran compiler¹³ and linker have to support it, and support it by the same mechanism (so mixing compiler families may not work and a non-default linker may be needed). It has been supported by the GCC and LLVM projects for some years with diverging implementations.

LTO support was added in 2011 for GCC 4.5 on Linux but was little used before 2019: compiler support has steadily improved over those years and `--enable-lto=R` is nowadays used for some routine CRAN checking.

Unfortunately `--enable-lto` may be accepted but silently do nothing useful if some of the toolchain does not support LTO: this is less common than it once was.

¹² A complete CRAN installation reduced from 50 to 35GB.

¹³ although there is the possibility to exclude Fortran but that misses some of the benefits.

Various macros can be set in file `config.site` to customize how LTO is used. If the Fortran compiler is not of the same family as the C/C++ compilers, set macro ‘`LTO_FC`’ (probably to empty). Macro ‘`LTO_LD`’ can be used to select an alternative linker should that be needed.

2.7.4.1 LTO with GCC

This has been tested on Linux with `gcc/gfortran` 8 and later: that needed setting (e.g. in `config.site`)

```
AR=gcc-ar
RANLIB=gcc-ranlib
```

For non-system compilers or if those wrappers have not been installed one may need something like

```
AR="ar --plugin=/path/to/liblto_plugin.so"
RANLIB="ranlib --plugin=/path/to/liblto_plugin.so"
```

and `NM` may be needed to be set analogously. (If using an LTO-enabled build to check packages, set environment variable `UserNM`¹⁴ to ‘`gcc-nm`’.)

With GCC 5 and later it is possible to parallelize parts of the LTO linking process: set the make macro ‘`LTO`’ to something like ‘`LTO=-flto=8`’ (to use 8 threads), for example in file `config.site`.

Under some circumstances and for a few packages, the PIC flags have needed overriding on Linux with GCC 9 and later: e.g use in `config.site`:

```
CPICFLAGS=-fPIC
CXXPICFLAGS=-fPIC
CXX11PICFLAGS=-fPIC
CXX14PICFLAGS=-fPIC
CXX17PICFLAGS=-fPIC
CXX20PICFLAGS=-fPIC
FPICFLAGS=-fPIC
```

We suggest only using these if the problem is encountered (it was not seen on CRAN with GCC 10 at the time of writing).

Note that R may need to be re-compiled after even a minor update to the compiler (e.g. from 10.1 to 10.2) but this may not be clear from confused compiler messages.

2.7.4.2 LTO with LLVM

LLVM supports another type of LTO called ‘Thin LTO’ as well as a similar implementation to GCC, sometimes called ‘Full LTO’. (See <https://clang.llvm.org/docs/ThinLTO.html>.) Currently the only LLVM compiler relevant to R is `clang` for which this can be selected by setting macro ‘`LTO=-flto=thin`’. LLVM has

```
AR=llvm-ar
RANLIB=llvm-ranlib
```

(but macOS does not, and these are not needed there). Where the linker supports a parallel backend for Thin LTO this can be specified *via* the macro ‘`LTO_LD`’: see the URL above for per-linker settings and further linking optimizations.)

¹⁴ not `NM` as we found `make` overriding that.

For example, on macOS one might use

```
LTO=-flto=thin
LTO_FC=
LTO_LD=-Wl,-mllvm,-threads=4
```

to use Thin LTO with 4 threads for C/C++ code, but skip LTO for Fortran code compiled with `gfortran`.

It is said to be particularly beneficial to use `-O3` for `clang` in conjunction with LTO.

2.7.4.3 LTO for package checking

LTO effectively compiles all the source code in a package as a single compilation unit and so allows the compiler (with sufficient diagnostic flags such as `-Wall`) to check consistency between what are normally separate compilation units.

With `gcc/gfortran` 9.x and later¹⁵ LTO will flag inconsistencies in calls to Fortran subroutines/functions, both between Fortran source files and between Fortran and C/C++. `gfortran` 8.4, 9.2 and later can help understanding these by extracting C prototypes from Fortran source files with its `-fc-prototypes-external` option, e.g. that (at the time of writing) Fortran `LOGICAL` corresponds to `int_least32_t *` in C.

On some systems it is possible to build the BLAS, LINPACK and LAPACK support as static libraries containing intermediate-level code that LTO will compile for all objects linked against these libraries, by configuring R with `--enable-lto=check`. This checks the consistency of calls to BLAS/LINPACK/LAPACK in any packages installed using the build. **NB:** as its name suggests, this option is intended only for checking installation of R and packages: it includes these library routines (those called directly and all they depend on) in each package. This is unlikely to work in conjunction with non-default options for BLAS and LAPACK, and ‘linking’ against these libraries may be very slow.

2.8 Testing an Installation

Full post-installation testing is possible only if the test files have been installed with

```
make install-tests
```

which populates a `tests` directory in the installation.

If this has been done, two testing routes are available. The first is to move to the home directory of the R installation (as given by `R.RHOME` or from R as `R.home()`) and run

```
cd tests
## followed by one of
../bin/R CMD make check
../bin/R CMD make check-devel
../bin/R CMD make check-all
```

and other useful targets are `test-BasePackages` and `test-Recommended` to run tests of the standard and recommended packages (if installed) respectively.

This re-runs all the tests relevant to the installed R (including for example the code in the package vignettes), but not for example the ones checking the example code in the manuals nor making the standalone Rmath library. This can occasionally be useful when

¹⁵ probably also 8.4 and later.

the operating environment has been changed, for example by OS updates or by substituting the BLAS (see Section A.3.1.4 [Shared BLAS], page 51).

Parallel checking of packages may be possible: set the environment variable `TEST_MC_CORES` to the maximum number of processes to be run in parallel. This affects both checking the package examples (part of `make check`) and package sources (part of `make check-devel` and `make check-recommended`). It does require a `make` command which supports the `make -j n` option: most do but on Solaris you need to select GNU `make` or `dmake`.

Alternatively, the installed R can be run, preferably with `--vanilla`. Then

```
pdf("tests.pdf") ## optional, but prevents flashing graphics windows
Sys.setenv(LC_COLLATE = "C", LC_TIME = "C", LANGUAGE = "en")
tools::testInstalledBasic("both")
tools::testInstalledPackages(scope = "base")
tools::testInstalledPackages(scope = "recommended")
```

runs the basic tests and then all the tests on the standard and recommended packages. These tests can be run from anywhere: the basic tests write their results in the `tests` folder of the R home directory and run fewer tests than the first approach: in particular they do not test things which need Internet access—that can be tested by

```
tools::testInstalledBasic("internet")
```

These tests work best if `diff` (in `Rtools*.exe` for Windows users) is in the path.

It is possible to test the installed packages (but not their package-specific tests) by `testInstalledPackages` even if `make install-tests` was not run.

Note that the results may depend on the language set for times and messages: for maximal similarity to reference results you may want to try setting (before starting the R session)

```
LANGUAGE=en
```

and use a UTF-8 or Latin-1 locale.

3 Installing R under Windows

The `bin/windows` directory of a CRAN site contains binaries for a base distribution and a large number of add-on packages from CRAN to run on 32- or 64-bit Windows (Windows 7 and later are tested; XP is known to fail some tests) on ‘`ix86`’ and ‘`x86_64`’ CPUs.

Your file system must allow long file names (as is likely except perhaps for some network-mounted systems). If it does not also support conversion to short name equivalents (a.k.a. DOS 8.3 names), then R *must* be installed in a path that does not contain spaces.

Installation is *via* the installer `R-4.0.3-win.exe`. Just double-click on the icon and follow the instructions. When installing on a 64-bit version of Windows the options will include 32- or 64-bit versions of R (and the default is to install both). You can uninstall R from the Control Panel.

You will be asked to choose a language for installation: that choice applies to both installation and un-installation but not to running R itself.

See the R Windows FAQ (<https://CRAN.R-project.org/bin/windows/base/rw-FAQ.html>) for more details on the binary installer.

3.1 Building from source

It is possible to use other toolchains (including that used for R 3.6.3, ‘`MSYS2`’ and ‘`MXE`’ (M cross environment)) to build R, but this manual only documents that used for binary distributions of R 4.0.x.

3.1.1 The Windows toolset

The binary distribution of R is currently built with tools distributed in `rtools40-x86_64.exe`. The webpage <https://cran.r-project.org/bin/windows/Rtools/> describes how to install the toolset and has links to example scripts to build and package R.

This toolset was modified (by Jeroen Ooms) from the ‘`MSYS2`’ project (<https://www.msys2.org/>), which uses compilers (GCC version 8.3.0) and runtime libraries from the ‘`Mingw-w64`’ project (<http://mingw-w64.org/>)¹ and provides a package manager (‘`pacman`’) for components and additional libraries.

The toolsets used for 64-bit Windows from 2008–20 were based on Mingw-w64. The assistance of Yu Gong at a crucial step in porting R to Mingw-w64 is gratefully acknowledged, as well as help from Kai Tietz, the lead developer of the Mingw-w64 project.

3.1.2 L^AT_EX

Both building R and checking packages need a distribution of L^AT_EX installed, with the directory containing `pdflatex` on the path.

The ‘`MiKTeX`’ (<https://miktex.org/>) distribution of L^AT_EX is that used on CRAN. This can be set up to install extra packages ‘on the fly’ (without asking), which is the simplest way to use it. The ‘`basic`’ version of ‘`MiKTeX`’ will need to add some packages.² In any case

¹ Capitalization of this project’s name is inconsistent even on its own webpages.

² There are reports of segfaults when ‘`MiKTeX`’ installs additional packages when making `NEWS.pdf`: re-running `make` seems to solve this.

ensure that the **inconsolata** package is installed—you can check with the ‘MiKTeX’ Package Manager.

It is also possible to use the TeX Live distribution from <https://www.tug.org/texlive/>. (The CRAN package **tinytex** (<https://CRAN.R-project.org/package=tinytex>) can install and manage a subset of TeX Live.)

3.2 Checking the build

You can test a build by running

```
make check
```

The recommended packages can be checked by

```
make check-recommended
```

Other levels of checking are

```
make check-devel
```

for a more thorough check of the R functionality, and

```
make check-all
```

for both **check-devel** and **check-recommended**.

If a test fails, there will almost always be a **.Rout.fail** file in the directory being checked (often **tests/Examples** or **tests**): examine the file to help pinpoint the problem.

Parallel checking of package sources (part of **make check-devel** and **make check-recommended**) is possible: see the environment variable **TEST_MC_CORES** to the maximum number of processes to be run in parallel.

3.3 Testing an Installation

The Windows installer contains a set of test files used when building R.

The toolset is not needed to run these tests, but more comprehensive analysis of errors will be given if **diff** is in the path (and **errorsAreFatal = FALSE** is then not needed below).

Launch either **Rgui** or **Rterm**, preferably with **--vanilla**. Then run

```
Sys.setenv(LC_COLLATE = "C", LANGUAGE = "en")
library("tools")
testInstalledBasic("both")
testInstalledPackages(scope = "base", errorsAreFatal = FALSE)
testInstalledPackages(scope = "recommended", errorsAreFatal = FALSE)
```

runs the basic tests and then all the tests on the standard and recommended packages. These tests can be run from anywhere: they write some of their results in the **tests** folder of the R home directory (as given by **R.home()**), and hence may need to be run under the account used to install R.

The results of **example(md5sums)** when testing **tools** will differ from the reference output as some files are installed with Windows’ CRLF line endings.

4 Installing R under macOS

(‘macOS’ was known as ‘OS X’ from 2012–2016 and as ‘Mac OS X’ before that.)

The front page of a CRAN site has a link ‘Download R for (Mac) OS X’. Click on that, then download the file `R-4.0.3.pkg` and install it. This runs on macOS 10.13 and later on Intel CPU¹ (High Sierra, Mojave, Catalina, Big Sur, ...).

Installers for R-patched and R-devel are usually available from <https://mac.R-project.org>. (Some of these packages are unsigned/not notarized: to install those Control/right/two-finger click, select ‘Open With’ and ‘Installer’.)

For some older versions of the OS you can in principle (it is little tested) install R from the sources (see Section C.3 [macOS], page 65).

It is important that if you use a binary installer package that your OS is fully updated: look at ‘Updates’ from the ‘App Store’ (‘Software Update’ in ‘System Preferences’ on Catalina) to be sure. (If using XQuartz, check that is current.)

To install, just double-click on the icon of the file you downloaded. At the ‘Installation Type’ stage, note the option to ‘Customize’. This currently shows four components: everyone will need the ‘R Framework’ component: the remaining components are optional. (The ‘Tcl/Tk’ component is needed to use package `tcltk`. The ‘Texinfo’ component is only needed by those installing source packages or R from its sources.)

This is an Apple Installer package. If you encounter any problem during the installation, please check the Installer log by clicking on the “Window” menu and item “Installer Log”. The full output (select “Show All Log”) is useful for tracking down problems. Note the the installer is clever enough to try to upgrade the last-installed version of the application where you installed it (which may not be where you want this time ...).

Various parts of the build require XQuartz to be installed: see <https://xquartz.macosforge.org/>. These include the `tcltk` package and the X11 device: attempting to use these without XQuartz will remind you. Also for the cairographics-based devices (which are not often used on macOS) such as `png(type = "cairo")`.

If you update your macOS version, you should re-install R (and perhaps XQuartz): the installer may tailor the installation to the current version of the OS.

For building R from source, see Section C.3 [macOS], page 65.

4.1 Running R under macOS

There are two ways to run R on macOS from a CRAN binary distribution.

There is a GUI console normally installed with the R icon in `/Applications` which you can run by double-clicking (e.g. from Launchpad or Finder). (If you cannot find it there it was possibly installed elsewhere so try searching for it in Spotlight.) This is usually referred to as R.APP to distinguish it from command-line R: its user manual is currently part of the macOS FAQ at <https://cran.r-project.org/bin/macosx/RMacOSX-FAQ.html> and can be viewed from R.APP’s ‘Help’ menu.

¹ macOS 11 has been a little tested on ‘Apple Silicon’ using ‘Rosetta’ emulation.

You can run command-line R and `Rscript` from a Terminal² so these can be typed as commands like any other Unix-alike: see the next chapter of this manual. There are some small differences which may surprise users of R on other platforms, notably the default location of the personal library directory (under `~/Library/R`, e.g. `~/Library/R/4.0/library`), and that warnings, messages and other output to `stderr` are highlighted in bold.

Those using the `zsh` shell (the default for new user accounts as from Catalina) might find the command `R` being masked by the `zsh` builtin `r` (which recalls commands). One can use a full path to R in an alias, or add `disable r` to `~/.zshrc`.

It has been reported that running `R.APP` may fail if no preferences are stored, so if it fails when launched for the very first time, try it again (the first attempt will store some preferences).

Users of `R.APP` need to be aware of the ‘App Nap’ feature (https://developer.apple.com/library/archive/releasenotes/MacOSX/WhatsNewInOSX/Articles/MacOSX10_9.html) which can cause R tasks to appear to run very slowly when not producing output in the console. Here are ways to avoid it:

- Ensure that the console is completely visible (or at least the activity indicator at the top right corner is visible).
- In a Terminal, run

```
defaults write org.R-project.R NSAppSleepDisabled -bool YES
```

(see https://developer.apple.com/library/archive/releasenotes/MacOSX/WhatsNewInOSX/Articles/MacOSX10_9.html).

Using the X11 device or the X11-based versions of `View()` and `edit()` for data frames and matrices (the latter are the default for command-line R but not `R.APP`) requires an X sub-system to be installed: see Section C.3 [macOS], page 65. So do the `tcltk` package and some third-party packages.

4.2 Uninstalling under macOS

R for macOS consists of two parts: the GUI (`R.APP`) and the R framework. Un-installation is as simple as removing those folders (e.g. by dragging them onto the Trash aka Bin). The typical installation will install the GUI into the `/Applications/R.app` folder and the R framework into the `/Library/Frameworks/R.framework` folder. The links to R and `Rscript` in `/usr/local/bin` should also be removed.

If you want to get rid of R more completely using a Terminal, simply run:

```
sudo rm -Rf /Library/Frameworks/R.framework /Applications/R.app \
/usr/local/bin/R /usr/local/bin/Rscript
```

The installation consists of up to four Apple packages:³ `org.r-project.R.fw.pkg`, `org.r-project.R.GUI.pkg`, `org.r-project.x86_64.tcltk.x11` and `org.r-project.x86_64.texinfo`. You can use `pkgutil --forget` if you want the Apple Installer to forget

² The installer puts links to R and `Rscript` in `/usr/local/bin`. If these are missing, you can run directly the copies in `/Library/Frameworks/R.framework/Resources/`.

³ At the time of writing: use `pkgutil --pkgs | grep org.r-project` to check..

about the package without deleting its files (useful for the R framework when installing multiple R versions in parallel), or after you have deleted the files.

Uninstalling the Tcl/Tk or Texinfo components (which are installed under `/usr/local`) is not as simple. You can list the files they installed in a Terminal by

```
pkgutil --files org.r-project.x86_64.tcltk.x11
pkgutil --files org.r-project.x86_64.texinfo
```

These are paths relative to `/`, the root of the file system.

4.3 Multiple versions

The installer will remove any previous version⁴ of the R framework which it finds installed. This can be avoided by using `pkgutil --forget` (see the previous section). However, note that different versions are installed under `/Library/Frameworks/R.framework/Versions` as 4.0, 3.6 and so on, so it is not possible to have different ‘3.x.y’ or ‘4.x.y’ versions installed for the same ‘x’.

A version of R can be run directly from the command-line as e.g.

```
/Library/Frameworks/R.framework/Versions/4.0/Resources/bin/R
```

However, R.APP will always run the ‘current’ version, that is the last installed version. A small utility, `Rswitch.app` (available at <https://mac.R-project.org/#other>: it is 32-bit so not usable on Catalina), can be used to change the ‘current’ version. This is of limited use as R.APP is compiled against a particular version of R and will likely crash if switched to an earlier version. This may allow you to install a development version of R (de-selecting R.APP) and then switch back to the release version.

⁴ More precisely, of the Apple package of the same name: this means that installing a package for 3.6.x does not remove an installation for 4.0 or later.

5 Running R

How to start R and what command-line options are available is discussed in Section “Invoking R” in *An Introduction to R*.

You should ensure that the shell has set adequate resource limits: R expects a stack size of at least 8MB and to be able to open at least 256 file descriptors. (Any modern OS should have default limits at least as large as these, but apparently NetBSD may not. Use the shell command `ulimit (sh/bash)` or `limit (csh/tcsh)` to check.) For some compilers¹ and packages a larger stack size has been needed: 20-25MB has sufficed to date.

R makes use of a number of environment variables, the default values of many of which are set in file `R_HOME/etc/Renviron` (there are none set by default on Windows and hence no such file). These are set at `configure` time, and you would not normally want to change them – a possible exception is `R_PAPERSIZE` (see Section B.3.1 [Setting paper size], page 56). The paper size will be deduced from the ‘`LC_PAPER`’ locale category if it exists and `R_PAPERSIZE` is unset, and this will normally produce the right choice from ‘`a4`’ and ‘`letter`’ on modern Unix-alikes (but can always be overridden by setting `R_PAPERSIZE`).

Various environment variables can be set to determine where R creates its per-session temporary directory. The environment variables `TMPDIR`, `TMP` and `TEMP` are searched in turn and the first one which is set and points to a writable area is used. If none do, the final default is `/tmp` on Unix-alikes and the value of `R_USER` on Windows. The path should be an absolute path not containing spaces (and it is best to avoid non-alphanumeric characters such as `+`).

Some Unix-alike systems are set up to remove files and directories periodically from `/tmp`, for example by a `cron` job running `tmpwatch`. Set `TMPDIR` to another directory before starting long-running jobs on such a system.

Note that `TMPDIR` will be used to execute `configure` scripts when installing packages, so if `/tmp` has been mounted as ‘`noexec`’, `TMPDIR` needs to be set to a directory from which execution is allowed.

¹ The Oracle compilers on Solaris (where the issue is parsing very complex R expressions) and GCC 9 on Linux.

6 Add-on packages

It is helpful to use the correct terminology. A *package* is loaded from a *library* by the function `library()`. Thus a library is a directory containing installed packages; the main library is `R_HOME/library`, but others can be used, for example by setting the environment variable `R_LIBS` or using the R function `.libPaths()`. To avoid any confusion you will often see a library directory referred to as a ‘library tree’.

6.1 Default packages

The set of packages loaded on startup is by default

```
> getOption("defaultPackages")
[1] "datasets" "utils"      "grDevices" "graphics"  "stats"      "methods"
```

(plus, of course, **base**) and this can be changed by setting the option in startup code (e.g. in `~/.Rprofile`). It is initially set to the value of the environment variable `R_DEFAULT_PACKAGES` if set (as a comma-separated list). Setting `R_DEFAULT_PACKAGES=NULL` ensures that only package **base** is loaded.

Changing the set of default packages is normally used to reduce the set for speed when scripting: in particular not using **methods** will reduce the start-up time by a factor of up to two. But it can also be used to customize R, e.g. for class use. **Rscript** also checks the environment variable `R_SCRIPT_DEFAULT_PACKAGES`; if set, this takes precedence over `R_DEFAULT_PACKAGES`.

6.2 Managing libraries

R packages are installed into *libraries*, which are directories in the file system containing a subdirectory for each package installed there.

R comes with a single library, `R_HOME/library` which is the value of the R object ‘`.Library`’ containing the standard and recommended¹ packages. Both sites and users can create others and make use of them (or not) in an R session. At the lowest level ‘`.libPaths()`’ can be used to add paths to the collection of libraries or to report the current collection.

R will automatically make use of a site-specific library `R_HOME/site-library` if this exists (it does not in a vanilla R installation). This location can be overridden by setting² ‘`.Library.site`’ in `R_HOME/etc/Rprofile.site`, or (not recommended) by setting the environment variable `R_LIBS_SITE`. Like ‘`.Library`’, the site libraries are always included by ‘`.libPaths()`’.

Users can have one or more libraries, normally specified by the environment variable `R_LIBS_USER`. This has a default value (to see it, use ‘`Sys.getenv("R_LIBS_USER")`’ within an R session), but that is only used if the corresponding directory actually exists (which by default it will not).

Both `R_LIBS_USER` and `R_LIBS_SITE` can specify multiple library paths, separated by colons (semicolons on Windows).

¹ unless they were excluded in the build.

² its binding is locked once the startup files have been read, so users cannot easily change it.

6.3 Installing packages

Packages may be distributed in source form or compiled binary form. Installing source packages which contain C/C++/Fortran code requires that compilers and related tools be installed. Binary packages are platform-specific and generally need no special tools to install, but see the documentation for your platform for details.

Note that you may need to specify implicitly or explicitly the library to which the package is to be installed. This is only an issue if you have more than one library, of course.

Ensure that the environment variable `TMPDIR` is either unset (and `/tmp` exists and can be written in and executed from) or is the absolute path to a valid temporary directory, not containing spaces.

For most users it suffices to call `'install.packages(pkgname)'` or its GUI equivalent if the intention is to install a CRAN package and internet access is available.³ On most systems `'install.packages()'` will allow packages to be selected from a list box (typically with thousands of items).

To install packages from source on a Unix-alike use in a terminal

```
R CMD INSTALL -l /path/to/library pkg1 pkg2 ...
```

The part `'-l /path/to/library'` can be omitted, in which case the first library of a normal R session is used (that shown by `.libPaths()[1]`).

There are a number of options available: use `R CMD INSTALL --help` to see the current list.

Alternatively, packages can be downloaded and installed from within R. First choose your nearest CRAN mirror using `chooseCRANmirror()`. Then download and install packages **pkg1** and **pkg2** by

```
> install.packages(c("pkg1", "pkg2"))
```

The essential dependencies of the specified packages will also be fetched. Unless the library is specified (argument `lib`) the first library in the library search path is used: if this is not writable, R will ask the user (in an interactive session) if the default personal library should be created, and if allowed to will install the packages there.

If you want to fetch a package and all those it depends on (in any way) that are not already installed, use e.g.

```
> install.packages("Rcmdr", dependencies = TRUE)
```

`install.packages` can install a source package from a local `.tar.gz` file (or a URL to such a file) by setting argument `repos` to `NULL`: this will be selected automatically if the name given is a single `.tar.gz` file.

`install.packages` can look in several repositories, specified as a character vector by the argument `repos`: these can include a CRAN mirror, Bioconductor, R-forge, rforge.net, local archives, local files, ...). Function `setRepositories()` can select amongst those repositories that the R installation is aware of.

Something which sometimes puzzles users is that `install.packages()` may report that a package which they believe should be available is not found. Some possible reasons:

- The package, such as **grid** or **tcltk**, is part of R itself and not otherwise available.

³ If a proxy needs to be set, see `?download.file`.

- The package is not in the available repositories, so check which have been selected by `getOption("repos")`
- The package is available, but not for the current version of R or for the type of OS (Unix/Windows). To retrieve the information on available versions of package `pkg`, use


```
av <- available.packages(filters=list())
av[av[, "Package"] == pkg, ]
```

in your R session, and look at the ‘Depends’ and ‘OS_type’ fields (there may be more than one matching entry). If the package depends on a version of R later than the one in use, it is possible that an earlier version is available which will work with your version of R: for CRAN look for ‘Old sources’ on the package’s CRAN landing page and manually retrieve an appropriate version (of comparable age to your version of R).

Naive users sometimes forget that as well as installing a package, they have to use `library` to make its functionality available.

6.3.1 Windows

What `install.packages` does by default is different on Unix-alikes (except macOS) and Windows. On Unix-alikes it consults the list of available *source* packages on CRAN (or other repository/ies), downloads the latest version of the package sources, and installs them (via R CMD INSTALL). On Windows it looks (by default) first at the list of *binary* versions of packages available for your version of R and downloads the latest versions (if any). If no binary version is available or the source version is newer, it will install the source versions of packages without compiled C/C++/Fortran code, and offer to do so for those with, if `make` is available (and this can be tuned by option `"install.packages.compile.from.source"`).

On Windows `install.packages` can also install a binary package from a local zip file (or the URL of such a file) by setting argument `repos` to `NULL`. `Rgui.exe` has a menu `Packages` with a GUI interface to `install.packages`, `update.packages` and `library`.

Windows binary packages for R are distributed as a single binary containing either or both architectures (32- and 64-bit).

R CMD INSTALL works in Windows to install source packages. No additional tools are needed if the package does not contain compiled code, and `install.packages(type="source")` will work for such packages. Those with compiled code need the tools (see Section 3.1.1 [The Windows toolset], page 17) to be on the path: this can be set from within an R session by something like

```
Sys.setenv(PATH=paste("C:\\rtools40\\usr\\bin", Sys.getenv("PATH"), sep=";"))
```

Occasional permission problems after unpacking source packages have been seen on some systems: these have been circumvented by setting the environment variable `R_INSTALL_TAR` to ‘`tar.exe`’.

If you have only a source package that is known to work with current R and just want a binary Windows build of it, you could make use of the building service offered at <https://win-builder.r-project.org/>.

For almost all packages R CMD INSTALL will attempt to install both 32- and 64-bit builds of a package if run from a 32/64-bit install of R. It will report success if the installation of the architecture of the running R succeeded, whether or not the other architecture was successfully installed. The exceptions are packages with a non-empty `configure.win` script

or which make use of `src/Makefile.win`. If `configure.win` does something appropriate to both architectures use⁴ option `--force-biarch`: otherwise `R CMD INSTALL --merge-multiarch` can be applied to a source tarball to merge separate 32- and 64-bit installs. (This can only be applied to a tarball, and will only succeed if both installs succeed.)

If you have a package without compiled code and no Windows-specific help, you can zip up an installation on another OS and install from that zip file on Windows. However, such a package can be installed from the sources on Windows without any additional tools.

6.3.2 macOS

On macOS `install.packages` works as it does on other Unix-alike systems, but there is an additional type `mac.binary` (available for the CRAN distribution but not when compiling R from source) which can be passed to `install.packages` in order to download and install binary packages from a suitable repository. These binary package files for macOS have the extension `'.tgz'`. The R.APP GUI provides menus for installation of either binary or source packages, from CRAN or local files.

On R builds using binary packages, the default is type `both`: this looks first at the list of binary packages available for your version of R and installs the latest versions (if any). If no binary version is available or the source version is newer, it will install the source versions of packages without compiled C/C++/Fortran code and offer to do so for those with, if `make` is available.

Note that most binary packages which include compiled code are tied to a particular series (e.g. R 4.0.x or 3.6.x) of R.

Installing source packages which do not contain compiled code should work with no additional tools. For others you will need the ‘Command Line Tools’ for `Xcode` and compilers which match those used to build R: see Section C.3 [macOS], page 65.

Package **rJava** (<https://CRAN.R-project.org/package=rJava>) and those which depend on it need a Java runtime installed and several packages need X11 installed, including those using Tk. See Section C.3 [macOS], page 65, and Section C.3.5 [Java (macOS)], page 70.

Tcl/Tk extension **BWidget** used to be distributed with R but no longer is; **Tktable** has been distributed with most versions of R (but not 4.0.0).

A few of the binary packages need other software to be installed on your system. In particular packages using Gtk+ (**RGtk2** (<https://CRAN.R-project.org/package=RGtk2>), **cairoDevice** (<https://CRAN.R-project.org/package=cairoDevice>) and those that depend on them) need the GTK framework installed from <https://mac.R-project.org/libs/>: the appropriate version at the time of writing was https://mac.R-project.org/libs/GTK_2.24.17-X11.pkg

The default compilers specified are shown in file `/Library/Frameworks/R.framework/Resources/etc/Makeconf`. At the time of writing those settings assumed that the C, Fortran and C++ compilers were on the path, using `gfortran` 8.2.0 (see Section C.3 [macOS], page 65). The settings can be changed, either by editing that file or in a file such as

⁴ for a small number of CRAN packages where this is known to be safe and is needed by the autobuilder this is the default. Look at the source of `tools:::install_packages` for the list. It can also be specified in the package’s `DESCRIPTION` file.

`~/.R/Makevars` (see the next section). Entries which may need to be changed include ‘CC’, ‘CXX’, ‘FC’, ‘FLIBS’ and the corresponding flags, and perhaps ‘CXXCPP’, ‘DYLIB_LD’, ‘MAIN_LD’, ‘SHLIB_CXXLD’ and ‘SHLIB_LD’, as well as the ‘CXX11’, ‘CXX14’, ‘CXX17’ and ‘CXX20’ variants

So for example on Mojave or Catalina you could select a specific non-Apple build of `clang` for both C and C++ with extensive checking by having in `~/.R/Makevars`

```
CC=/usr/local/clang/bin/clang -isysroot
/Library/Developer/CommandLineTools/SDKs/MacOSX.sdk
CXX=/usr/local/clang/bin/clang++ -isysroot
/Library/Developer/CommandLineTools/SDKs/MacOSX.sdk
CXX11 = $CXX
CXX14 = $CXX
CXX17 = $CXX
CXX20 = $CXX
CFLAGS = -g -O2 -Wall -pedantic -Wconversion -Wno-sign-conversion
CXXFLAGS = -g -O2 -Wall -pedantic -Wconversion -Wno-sign-conversion
CXX11FLAGS = $CXXFLAGS
CXX14FLAGS = $CXXFLAGS
CXX17FLAGS = $CXXFLAGS
CXX20FLAGS = $CXXFLAGS
```

(long lines split for the manual only) and `gfortran` by

```
FC = /usr/local/gfortran/bin/gfortran
FLIBS = -L/usr/local/gfortran/lib/gcc/x86_64-apple-darwin18/8.2.0
-L/usr/local/gfortran/lib -lgfortran -lquadmath -lm
```

(with lines broken here for legibility). If that `clang` build supports OpenMP, you can add

```
SHLIB_OPENMP_CFLAGS = -fopenmp
SHLIB_OPENMP_CXXFLAGS = -fopenmp
```

to compile OpenMP-using packages. It will also be necessary to arrange for the `libomp.dylib` library to be found at both install time and run time, for example by copying/linking it somewhere that is searched such as `/usr/local/lib`.

Apple includes many Open Source libraries in macOS but increasingly without the corresponding headers (not even in Xcode nor the Command Line Tools): they are often rather old versions. If installing packages from source using them it is usually easiest to install a statically-linked up-to-date copy of the Open Source package from its sources or from <https://mac.R-project.org/libs-4/>. But sometimes it is desirable/necessary to use Apple’s dynamically linked library, in which case appropriate headers could be extracted from the sources⁵ available *via* <https://opensource.apple.com> – this has been used for OpenSSL and iodbc.

6.3.3 Customizing package compilation

The R system and package-specific compilation flags can be overridden or added to by setting the appropriate Make variables in the personal file `HOME/.R/Makevars-R_PLATFORM` (but `HOME/.R/Makevars.win` or `HOME/.R/Makevars.win64` on Windows), or if that does

⁵ Note that capitalization and versioning may differ from the Open Source project.

not exist, `HOME/.R/Makevars`, where ‘`R_PLATFORM`’ is the platform for which R was built, as available in the `platform` component of the R variable `R.version`. The full path to an alternative personal file⁶ can be specified *via* the environment variable `R_MAKEVARS_USER`.

Package developers are encouraged to use this mechanism 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, or for `clang`.

Note that this mechanism can also be used when it necessary to change the optimization level whilst installing a particular package. For example

```
## for C code
CFLAGS = -g -O -mtune=native
## for C++ code
CXXFLAGS = -g -O -mtune=native
## for fixed-form Fortran code
FFLAGS = -g -O -mtune=native
```

Another use is to override the settings in a binary installation of R. For example, to use a different Fortran compiler on macOS

```
FC = /usr/local/gfortran/bin/gfortran
FLIBS = -L/usr/local/gfortran/lib/gcc/x86_64-apple-darwin16/6.3.0
        -L/usr/local/gfortran/lib -lgfortran -lquadmath -lm
```

(line split for legibility here).

There is also provision for a site-wide `Makevars.site` file under `R_HOME/etc` (in a sub-architecture-specific directory if appropriate). This is read immediately after `Makeconf`, and the path to an alternative file can be specified by environment variable `R_MAKEVARS_SITE`.

Note that these mechanisms do not work with packages which fail to pass settings down to sub-makes, perhaps reading `etc/Makeconf` in makefiles in subdirectories. Fortunately such packages are unusual.

6.3.4 Multiple sub-architectures

When installing packages from their sources, there are some extra considerations on installations which use sub-architectures. These are commonly used on Windows but can in principle be used on other platforms.

When a source package is installed by a build of R which supports multiple sub-architectures, the normal installation process installs the packages for all sub-architectures. The exceptions are

Unix-alikes

where there is an `configure` script, or a file `src/Makefile`.

Windows

where there is a non-empty `configure.win` script, or a file `src/Makefile.win` (with some exceptions where the package is known to have an architecture-independent `configure.win`, or if `--force-biarch` or field ‘`Biarch`’ in the `DESCRIPTION` file is used to assert so).

⁶ using a path containing spaces is likely to cause problems

In those cases only the current architecture is installed. Further sub-architectures can be installed by

```
R CMD INSTALL --libs-only pkg
```

using the path to R or `R --arch` to select the additional sub-architecture. There is also `R CMD INSTALL --merge-multiarch` to build and merge the two architectures, starting with a source tarball.

6.3.5 Byte-compilation

As from R 3.6.0, all packages are by default byte-compiled.

Byte-compilation can be controlled on a per-package basis by the ‘ByteCompile’ field in the DESCRIPTION file.

6.3.6 External software

Some R packages contain compiled code which links to external software libraries. Unless the external library is statically linked (which is done as much as possible for binary packages on Windows and macOS), the libraries have to be found when the package is loaded and not just when it is installed. How this should be done depends on the OS (and in some cases the version).

For Unix-alikes except macOS the primary mechanism is the `ld.so` cache controlled by `ldconfig`: external dynamic libraries recorded in that cache will be found. Standard library locations will be covered by the cache, and well-designed software will add its locations (as for example **openmpi** does on Fedora). The secondary mechanism is to consult the environment variable `LD_LIBRARY_PATH`. The R script controls that variable, and sets it to the concatenation of `R_LD_LIBRARY_PATH`, `R_JAVA_LD_LIBRARY_PATH` and the environment value of `LD_LIBRARY_PATH`. The first two have defaults which are normally set when R is installed (but can be overridden in the environment) so `LD_LIBRARY_PATH` is the best choice for a user to set.

On macOS the primary mechanism is to embed the absolute path to dependent dynamic libraries into an object when it is compiled. Few R packages arrange to do so, but it can be edited⁷ *via* `install_name_tool` — that only deals with direct dependencies and those would also need to be compiled to include the absolute paths of their dependencies. If the choice of absolute path is to be deferred to load time, how they are resolved is described in `man dyld`: the role of `LD_LIBRARY_PATH` is replaced on macOS by `DYLD_LIBRARY_PATH` and `DYLD_FALLBACK_LIBRARY_PATH`. Running `R CMD otool -L` on the package shared object will show where (if anywhere) its dependencies are resolved. `DYLD_FALLBACK_LIBRARY_PATH` is preferred (and it is that which is manipulated by the R script), but as from 10.11 (‘El Capitan’) the default behaviour had been changed for security reasons to discard these environment variables when invoking a shell script (and R is a shell script). That makes the only portable option to set `R_LD_LIBRARY_PATH` in the environment, something like

```
export R_LD_LIBRARY_PATH="'R_RHOME'/lib:/opt/local/lib"
```

The precise rules for where Windows looks for DLLs are complex and depend on the version of Windows. But for present purposes the main solution is to put the directories

⁷ They need to have been created using `-headerpad_max_install_names`, which is the default for an R package.

containing the DLLs the package links to (and any those DLLs link to) on the `PATH`. 64-bit versions of Windows will ignore 32-bit DLLs from 64-bit R and *vice versa*.

The danger with any of the methods which involve setting environment variables is of inadvertently masking a system library. This is less for `DYLD_FALLBACK_LIBRARY_PATH` and for *appending* to `PATH` on Windows (as it should already contain the system library paths).

6.4 Updating packages

The command `update.packages()` is the simplest way to ensure that all the packages on your system are up to date. It downloads the list of available packages and their current versions, compares it with those installed and offers to fetch and install any that have later versions on the repositories.

An alternative interface to keeping packages up-to-date is provided by the command `packageStatus()`, which returns an object with information on all installed packages and packages available at multiple repositories. The `print` and `summary` methods give an overview of installed and available packages, the `upgrade` method offers to fetch and install the latest versions of outdated packages.

One sometimes-useful additional piece of information that `packageStatus()` returns is the status of a package, as "ok", "upgrade" or "unavailable" (in the currently selected repositories). For example

```
> inst <- packageStatus()$inst
> inst[inst$Status != "ok", c("Package", "Version", "Status")]
```

	Package	Version	Status
Biobase	Biobase	2.8.0	unavailable
RCurl	RCurl	1.4-2	upgrade
Rgraphviz	Rgraphviz	1.26.0	unavailable
rgdal	rgdal	0.6-27	upgrade

6.5 Removing packages

Packages can be removed in a number of ways. From a command prompt they can be removed by

```
R CMD REMOVE -l /path/to/library pkg1 pkg2 ...
```

From a running R process they can be removed by

```
> remove.packages(c("pkg1", "pkg2"),
  lib = file.path("path", "to", "library"))
```

Finally, one can just remove the package directory from the library.

6.6 Setting up a package repository

Utilities such as `install.packages` can be pointed at any CRAN-style repository, and R users may want to set up their own. The 'base' of a repository is a URL such as `https://www.stats.ox.ac.uk/pub/RWin/`: this must be an URL scheme that `download.packages` supports (which also includes 'https://', 'ftp://' and 'file://'). Under that base URL there should be directory trees for one or more of the following types of package distributions:

- "source": located at `src/contrib` and containing `.tar.gz` files. Other forms of compression can be used, e.g. `.tar.bz2` or `.tar.xz` files. Complete repositories contain

the sources corresponding to any binary packages, and in any case it is wise to have a `src/contrib` area with a possibly empty `PACKAGES` file.

- "win.binary": located at `bin/windows/contrib/x.y` for R versions `x.y.z` and containing `.zip` files for Windows.
- "mac.binary": located at `bin/macosx/contrib/4.y` for the CRAN builds for macOS for R versions `4.y.z`, containing `.tgz` files.
- "mac.binary.el-capitan": located at `bin/macosx/el-capitan/contrib/3.y` for the CRAN builds for R versions `3.y.z`, containing `.tgz` files.

Each terminal directory must also contain a `PACKAGES` file. This can be a concatenation of the `DESCRIPTION` files of the packages separated by blank lines, but only a few of the fields are needed. The simplest way to set up such a file is to use function `write_PACKAGES` in the `tools` package, and its help explains which fields are needed. Optionally there can also be `PACKAGES.rds` and `PACKAGES.gz` files, downloaded in preference to `PACKAGES`. (If you have a mis-configured server that does not report correctly non-existent files you may need these files.)

To add your repository to the list offered by `setRepositories()`, see the help file for that function.

Incomplete repositories are better specified *via* a `contriburl` argument than *via* being set as a repository.

A repository can contain subdirectories, when the descriptions in the `PACKAGES` file of packages in subdirectories must include a line of the form

Path: *path/to/subdirectory*

—once again `write_PACKAGES` is the simplest way to set this up.

6.7 Checking installed source packages

It can be convenient to run R CMD `check` on an installed package, particularly on a platform which uses sub-architectures. The outline of how to do this is, with the source package in directory `pkg` (or a tarball filename):

```
R CMD INSTALL -l libdir pkg > pkg.log 2>&1
```

```
R CMD check -l libdir --install=check:pkg.log pkg
```

Where sub-architectures are in use the R CMD `check` line can be repeated with additional architectures by

```
R --arch arch CMD check -l libdir --extra-arch --install=check:pkg.log pkg
```

where `--extra-arch` selects only those checks which depend on the installed code and not those which analyse the sources. (If multiple sub-architectures fail only because they need different settings, e.g. environment variables, `--no-multiarch` may need to be added to the `INSTALL` lines.) On Unix-alikes the architecture to run is selected by `--arch`: this can also be used on Windows with `R_HOME/bin/R.exe`, but it is more usual to select the path to the `Rcmd.exe` of the desired architecture.

So on Windows to install, check and package for distribution a source package from a tarball which has been tested on another platform one might use

```
.../bin/i386/Rcmd INSTALL -l libdir tarball --build > pkg.log 2>&1
```

```
.../bin/i386/Rcmd check -l libdir --extra-arch --install=check:pkg.log pkg
```

```
.../bin/x64/Rcmd check -l libdir --extra-arch --install=check:pkg.log pkg
```

where one might want to run the second and third lines in a different shell with different settings for environment variables and the path (to find external software, notably for Gtk+).

R CMD INSTALL can do a i386 install and then add the x64 DLL from a single command by

```
R CMD INSTALL --merge-multiarch -l libdir tarball
```

and `--build` can be added to zip up the installation.

7 Internationalization and Localization

Internationalization refers to the process of enabling support for many human languages, and *localization* to adapting to a specific country and language.

Current builds of R support all the character sets that the underlying OS can handle. These are interpreted according to the current `locale`, a sufficiently complicated topic to merit a separate section. Note though that R has no built-in support for right-to-left languages and bidirectional output, relying on the OS services. For example, how character vectors in UTF-8 containing both English digits and Hebrew characters are printed is OS-dependent (and perhaps locale-dependent).

The other aspect of the internationalization is support for the translation of messages. This is enabled in almost all builds of R.

7.1 Locales

A *locale* is a description of the local environment of the user, including the preferred language, the encoding of characters, the currency used and its conventions, and so on. Aspects of the locale are accessed by the R functions `Sys.getlocale` and `Sys.localeconv`.

The system of naming locales is OS-specific. There is quite wide agreement on schemes, but not on the details of their implementation. A locale needs to specify

- A human language. These are generally specified by a lower-case two-character abbreviation following ISO 639 (see e.g. https://en.wikipedia.org/wiki/ISO_639-1).
- A ‘territory’, used mainly to specify the currency. These are generally specified by an upper-case two-character abbreviation following ISO 3166 (see e.g. https://en.wikipedia.org/wiki/ISO_3166).
- A charset encoding, which determines both how a byte stream should be divided into characters, and which characters the subsequences of bytes represent. Sometimes the combination of language and territory is used to specify the encoding, for example to distinguish between traditional and simplified Chinese.
- Optionally, a modifier, for example to indicate that Austria is to be considered pre- or post-Euro. The modifier is also used to indicate the script (`@latin`, `@cyrillic` for Serbian, `@iqtelif`) or language dialect (e.g. `@saaho`, a dialect of Afar, and `@bokmal` and `@nynorsk`, dialects of Norwegian regarded by some OSes as separate languages, `no` and `nn`).

R is principally concerned with the first (for translations) and third. Note that the charset may be deducible from the language, as some OSes offer only one charset per language.

7.1.1 Locales under Unix-alikes

Modern Linux uses the XPG¹ locale specifications which have the form ‘`en_GB`’, ‘`en_GB.UTF-8`’, ‘`aa_ER.UTF-8@saaho`’, ‘`de_AT.iso885915@euro`’, the components being in the order listed above. (See `man locale` and `locale -a` for more details.) Similar schemes are used by most Unix-alikes: some (including some distributions of Linux) use ‘`.utf8`’ rather than ‘`.UTF-8`’.

¹ ‘X/Open Portability Guide’, which has had several versions.

Note that whereas UTF-8 locales are nowadays almost universally used, locales such as ‘en_GB’ use 8-bit encodings for backwards compatibility.

7.1.2 Locales under Windows

Windows also uses locales, but specified in a rather less concise way. Most users will encounter locales only via drop-down menus, but more information and lists can be found by searching for ‘Windows language country strings’).

It offers only one encoding per language.

Some care is needed with Windows’ locale names. For example, **chinese** is Traditional Chinese and not Simplified Chinese as used in most of the Chinese-speaking world.

7.1.3 Locales under macOS

macOS supports locales in its own particular way, but the R GUI tries to make this easier for users. See <https://developer.apple.com/library/archive/documentation/MacOSX/Conceptual/BPInternational/> for how users can set their locales. As with Windows, end users will generally only see lists of languages/territories. Users of R in a terminal may need to set the locale to something like ‘en_GB.UTF-8’ if it defaults to ‘C’ (as it sometimes does when logging in remotely and for batch jobs: note whether **Terminal** sets the **LANG** environment variable is an (advanced) preference, but does so by default).

Internally macOS uses a form similar to Linux: the main difference from other Unix-alikes is that where a character set is not specified it is assumed to be UTF-8.

7.2 Localization of messages

The preferred language for messages is by default taken from the locale. This can be overridden first by the setting of the environment variable **LANGUAGE** and then² by the environment variables **LC_ALL**, **LC_MESSAGES** and **LANG**. (The last three are normally used to set the locale and so should not be needed, but the first is only used to select the language for messages.) The code tries hard to map locales to languages, but on some systems (notably Windows) the locale names needed for the environment variable **LC_ALL** do not all correspond to XPG language names and so **LANGUAGE** may need to be set. (One example is ‘LC_ALL=es’ on Windows which sets the locale to Estonian and the language to Spanish.)

It is usually possible to change the language once R is running *via* (not Windows) `Sys.setlocale("LC_MESSAGES", "new_locale")`, or by setting an environment variable such as **LANGUAGE**, *provided*³ the language you are changing to can be output in the current character set. But this is OS-specific, and has been known to stop working on an OS upgrade.

Messages are divided into *domains*, and translations may be available for some or all messages in a domain. R makes use of the following domains.

- Domain **R** for the C-level error and warning messages from the R interpreter.

² On some systems setting **LC_ALL** or **LC_MESSAGES** to ‘C’ disables **LANGUAGE**.

³ If you try changing from French to Russian except in a UTF-8 locale, you will most likely find messages change to English.

- Domain `R-pkg` for the R `stop`, `warning` and `message` messages in each package, including `R-base` for the `base` package.
- Domain `pkg` for the C-level messages in each package.
- Domain `RGui` for the menus etc of the R for Windows GUI front-end.

Dividing up the messages in this way allows R to be extensible: as packages are loaded, their message translation catalogues can be loaded too.

R can be built without support for translations, but it is enabled by default.

R-level and C-level domains are subtly different, for example in the way strings are canonicalized before being passed for translation.

Translations are looked for by domain according to the currently specified language, as specifically as possible, so for example an Austrian (`de_AT`) translation catalogue will be used in preference to a generic German one (`de`) for an Austrian user. However, if a specific translation catalogue exists but does not contain a translation, the less specific catalogues are consulted. For example, R has catalogues for `en_GB` that translate the Americanisms (e.g., `gray`) in the standard messages into English.⁴ Two other examples: there are catalogues for `es`, which is Spanish as written in Spain and these will by default also be used in Spanish-speaking Latin American countries, and also for `pt_BR`, which are used for Brazilian locales but not for locales specifying Portugal.

Translations in the right language but the wrong charset are made use of by on-the-fly re-encoding. The `LANGUAGE` variable (only) can be a colon-separated list, for example `se:de`, giving a set of languages in decreasing order of preference. One special value is `en@quot`, which can be used in a UTF-8 locale to have American error messages with pairs of single quotes translated to Unicode directional quotes.

If no suitable translation catalogue is found or a particular message is not translated in any suitable catalogue, ‘English’⁵ is used.

See <https://developer.r-project.org/Translations30.html> for how to prepare and install translation catalogues.

⁴ the language written in England: some people living in the USA appropriate this name for their language.

⁵ with Americanisms.

8 Choosing between 32- and 64-bit builds

Almost all current CPUs have both 32- and 64-bit sets of instructions. Most OSes running on such CPUs offer the choice of building a 32-bit or a 64-bit version of R (and details are given below under specific OSes). For most a 32-bit version is the default, but for some (e.g., ‘x86_64’ Linux and macOS ≥ 10.6) 64-bit is.

All current versions of R use 32-bit integers (this is enforced in the build) and ISO/IEC 60559¹ double-precision reals, and so compute to the same precision² and with the same limits on the sizes of numerical quantities. The principal difference is in the size of the pointers.

64-bit builds have both advantages and disadvantages:

- The total virtual memory space made available to a 32-bit process is limited by the pointer size to 4GB, and on most OSes to 3GB (or even 2GB). The limits for 64-bit processes are much larger (e.g. 8–128TB).

R allocates memory for large objects as needed, and removes any unused ones at garbage collection. When the sizes of objects become an appreciable fraction of the address limit, fragmentation of the address space becomes an issue and there may be no hole available that is the size requested. This can cause more frequent garbage collection or the inability to allocate large objects. As a guide, this will become an issue for 32-bit builds with objects more than 10% of the size of the address space (around 300Mb) or when the total size of objects in use is around one third (around 1Gb).

- Only 64-bit builds support ‘long vectors’, those with 2^{31} or more elements (which needs at least 16GB of storage for each numeric vector).
- Most 32-bit OSes by default limit file sizes to 2GB (and this may also apply to 32-bit builds on 64-bit OSes). This can often be worked around: `configure` selects suitable defines if this is possible. (We have also largely worked around that limit on 32-bit Windows.) 64-bit builds have much larger limits.
- Because the pointers are larger, R’s basic structures are larger. This means that R objects take more space and (usually) more time to manipulate. So 64-bit builds of R will, all other things being equal, run slower than 32-bit builds. (On Sparc Solaris the difference was 15-20%.)
- However, ‘other things’ may not be equal. In the specific case of ‘x86_64’ *vs* ‘i386’, the 64-bit CPU has features (such as SSE2 instructions) which are guaranteed to be present but are optional on the 32-bit CPU, and also has more general-purpose registers. This means that on chips like a desktop Intel i7 the vanilla 64-bit version of R has been around 10% faster on both Linux and macOS. (Laptop CPUs are usually relatively slower in 64-bit mode.)

So, for speed you may want to use a 32-bit build (especially on a laptop), but to handle large datasets (and perhaps large files) a 64-bit build. You can often build both and install them in the same place: See Section 2.6 [Sub-architectures], page 9. (This is done for the Windows binary distributions.)

¹ also known as IEEE 754

² at least when storing quantities: the on-FPU precision is allowed to vary

Even on 64-bit builds of R there are limits on the size of R objects (see `help("Memory-limits")`), some of which stem from the use of 32-bit integers (especially in Fortran code). For example, each dimension of an array is limited to $2^{31} - 1$.

9 The standalone Rmath library

The routines supporting the distribution and special¹ functions in R and a few others are declared in C header file `Rmath.h`. These can be compiled into a standalone library for linking to other applications. (Note that they are not a separate library when R is built, and the standalone version differs in several ways.)

The makefiles and other sources needed are in directory `src/nmath/standalone`, so the following instructions assume that is the current working directory (in the build directory tree on a Unix-alike if that is separate from the sources).

`Rmath.h` contains `'R_VERSION_STRING'`, which is a character string containing the current R version, for example `"4.0.0"`.

There is full access to R's handling of NaN, Inf and -Inf via special versions of the macros and functions

```
ISNAN, R_FINITE, R_log, R_pow and R_pow_di
```

and (extern) constants `R_PosInf`, `R_NegInf` and `NA_REAL`.

There is no support for R's notion of missing values, in particular not for `NA_INTEGER` nor the distinction between `NA` and NaN for doubles.

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 shared library or DLL you may have to use the one supplied, which is the Marsaglia-multicarry with an entry point

```
set_seed(unsigned int, unsigned int)
```

to set its seeds).

The facilities to change the normal random number generator are available through the constant `N01_kind`. This takes values from the enumeration type

```
typedef enum {
    BUGGY_KINDERMAN_RAMAGE,
    AHRENS_DIETER,
    BOX_MULLER,
    USER_NORM,
    INVERSION,
    KINDERMAN_RAMAGE
} N01type;
```

(and `'USER_NORM'` is not available).

9.1 Unix-alikes

If R has not already been made in the directory tree, `configure` must be run as described in the main build instructions.

Then (in `src/nmath/standalone`)

```
make
```

¹ e.g. Bessel, beta and gamma functions

will make standalone libraries `libRmath.a` and `libRmath.so` (`libRmath.dylib` on macOS): ‘`make static`’ and ‘`make shared`’ will create just one of them.

To use the routines in your own C or C++ programs, include

```
#define MATHLIB_STANDALONE
#include <Rmath.h>
```

and link against ‘`-lRmath`’ (and ‘`-lm`’ if needed on your OS). The example file `test.c` does nothing useful, but is provided to test the process (via `make test`). Note that you will probably not be able to run it unless you add the directory containing `libRmath.so` to the `LD_LIBRARY_PATH` environment variable (`libRmath.dylib`, `DYLD_FALLBACK_LIBRARY_PATH` on macOS).

The targets

```
make install
make uninstall
```

will (un)install the header `Rmath.h` and shared and static libraries (if built). Both `prefix=` and `DESTDIR` are supported, together with more precise control as described for the main build.

‘`make install`’ installs a file for `pkg-config` to use by e.g.

```
$(CC) 'pkg-config --cflags libRmath' -c test.c
$(CC) 'pkg-config --libs libRmath' test.o -o test
```

On some systems ‘`make install-strip`’ will install a stripped shared library.

9.2 Windows

You need to set up² almost all the tools to make R and then run (in a Unix-like shell)

```
(cd ../../gnuwin32; make MkRules)
(cd ../../include; make -f Makefile.win config.h Rconfig.h Rmath.h)
make -f Makefile.win
```

Alternatively, in a `cmd.exe` shell use

```
cd ../../include
make -f Makefile.win config.h Rconfig.h Rmath.h
cd ../nmath/standalone
make -f Makefile.win
```

This creates a static library `libRmath.a` and a DLL `Rmath.dll`. If you want an import library `libRmath.dll.a` (you don’t need one), use

```
make -f Makefile.win shared implib
```

To use the routines in your own C or C++ programs using MinGW-w64, include

```
#define MATHLIB_STANDALONE
#include <Rmath.h>
```

and link against ‘`-lRmath`’. This will use the first found of `libRmath.dll.a`, `libRmath.a` and `Rmath.dll` in that order, so the result depends on which files are present. You should be able to force static or dynamic linking *via*

```
-Wl,-Bstatic -lRmath -Wl,Bdynamic
```

² including copying `MkRules.dist` to `MkRule.local` and selecting the architecture.

```
-Wl,-Bdynamic -lRmath
```

or by linking to explicit files (as in the ‘test’ target in `Makefile.win`: this makes two executables, `test.exe` which is dynamically linked, and `test-static.exe`, which is statically linked).

It is possible to link to `Rmath.dll` using other compilers, either directly or via an import library: if you make a MinGW-w64 import library as above, you will create a file `Rmath.def` which can be used (possibly after editing) to create an import library for other systems such as Visual C++.

If you make use of dynamic linking you should use

```
#define MATHLIB_STANDALONE
#define RMATH_DLL
#include <Rmath.h>
```

to ensure that the constants like `NA_REAL` are linked correctly. (Auto-import will probably work with MinGW-w64, but it is better to be sure. This is likely to also work with VC++, Borland and similar compilers.)

Appendix A Essential and useful other programs under a Unix-alike

This appendix gives details of programs you will need to build R on Unix-like platforms, or which will be used by R if found by `configure`.

Remember that some package management systems (such as RPM and Debian/Ubuntu's) make a distinction between the user version of a package and the development version. The latter usually has the same name but with the extension `'-devel'` or `'-dev'`: you need both versions installed.

A.1 Essential programs and libraries

You need a means of compiling C and Fortran 90 (see Section B.6 [Using Fortran], page 57). Your C compiler should be ISO/IEC 60059¹, POSIX 1003.1 and C99-compliant.² R tries to choose suitable flags³ for the C compilers it knows about, but you may have to set `CC` or `CFLAGS` suitably. For versions of `gcc` prior to 5.1 with `glibc`-based Linux this means including `-std=gnu99`⁴. (Note that options essential to run the compiler even for linking, such as those to set the architecture, should be specified as part of `CC` rather than in `CFLAGS`.)

Unless you do not want to view graphs on-screen (or use macOS) you need `'X11'` installed, including its headers and client libraries. For recent Fedora/RedHat distributions it means (at least) RPMs `'libX11'`, `'libX11-devel'`, `'libXt'` and `'libXt-devel'`. On Debian/Ubuntu we recommend the meta-package `'xorg-dev'`. If you really do not want these you will need to explicitly configure R without X11, using `--with-x=no`.

The command-line editing (and command completion) depends on the GNU `readline` library (including its headers): version 6.0 or later is needed for all the features to be enabled. Otherwise you will need to configure with `--with-readline=no` (or equivalent).

A suitably comprehensive `iconv` function is essential. The R usage requires `iconv` to be able to translate between `"latin1"` and `"UTF-8"`, to recognize `" "` (as the current encoding) and `"ASCII"`, and to translate to and from the Unicode wide-character formats `"UCS-[24] [BL]E"` — this is true by default for `glibc`⁵ but not of most commercial Unixes. However, you can make use of GNU `libiconv` (as used on macOS: see <https://www.gnu.org/software/libiconv/>).

¹ also known as IEEE 754

² Note that C11 compilers need not be C99-compliant: R requires support for `double complex` and variable-length arrays which are optional in C11 but are mandatory in C99. C18 (also known as C17) is a 'bugfix release' of C11, clarifying the standard.

³ Examples are `-std=gnu99`, `-std=c99` and `-c99`.

⁴ `-std=c99` excludes POSIX functionality, but `config.h` will turn on all GNU extensions to include the POSIX functionality for R itself: this does not apply to badly-written packages. The default mode for GCC 5.1 and later is `-std=gnu11`, which currently includes the optional features R needs.

⁵ However, it is possible to break the default behaviour of `glibc` by re-specifying the `gconv` modules to be loaded.

The OS needs to have enough support⁶ for wide-character types: this is checked at configuration. Some C99 functions⁷ are required and checked for at configuration. A small number of POSIX functions⁸ are essential, and others⁹ will be used if available.

Installations of **zlib** (version 1.2.5 or later), **libbz2** (version 1.0.6 or later: called **bzip2-libs**/**bzip2-devel** or **libbz2-1.0**/**libbz2-dev** by some Linux distributions) and **liblzma**¹⁰ version 5.0.3 or later are required.

Either PCRE1 (version 8.32 or later, formerly known as just PCRE) or PCRE2 is required: PCRE2 is preferred and using PCRE1 requires `configure` option `--with-pcre1`. Only the 8-bit library and headers are needed if these are packaged separately. JIT support (optional) is desirable for the best performance. For PCRE2 ≥ 10.30 (which is desirable as matching has been re-written not to use recursion and the Unicode tables were updated to version 10)

```
./configure --enable-jit
```

suffices. If building PCRE1 for use with R a suitable `configure` command might be

```
./configure --enable-utf --enable-unicode-properties --enable-jit --disable-cpp
```

The `--enable-jit` flag is supported for most common CPUs. (See also the comments for Solaris.)

Some packages require the ‘Unicode properties’ which are optional for PCRE1: support for this and JIT can be checked at run-time by calling `pcre_config()`.

Library **libcurl** (version 7.28.0 or later¹¹) is required. Information on **libcurl** is found from the `curl-config` script: if that is missing or needs to be overridden¹² there are macros to do so described in file `config.site`.

A **tar** program is needed to unpack the sources and packages (including the recommended packages). A version¹³ that can automagically detect compressed archives is preferred for use with `untar()`: the configure script looks for **gtar** and **gnutar** before **tar** – use environment variable `TAR` to override this. (On NetBSD/OpenBSD systems set this to **bsdtar** if that is installed.)

There need to be suitable versions of the tools **grep** and **sed**: the problems are usually with old AT&T and BSD variants. `configure` will try to find suitable versions (including looking in `/usr/xpg4/bin` which is used on some commercial Unixes).

You will not be able to build most of the manuals unless you have **texi2any** version 5.1 or later installed (which requires **perl**), and if not most of the HTML manuals will be

⁶ specifically, the C99 functionality of headers `wchar.h` and `wctype.h`, types `wctans_t` and `mbstate_t` and functions `mbrtowc`, `mbstowcs`, `wcrtomb`, `wcscoll`, `wcstombs`, `wctrans`, `wctype`, and `iswctype`.

⁷ including `expm1`, `hypot`, `log1p`, `nearbyint` and `va_copy`.

⁸ including `opendir`, `readdir`, `closedir`, `popen`, `stat`, `glob`, `access`, `getcwd` and `chdir` system calls, `select` on a Unix-alike, and either `putenv` or `setenv`.

⁹ such as `realpath`, `symlink`.

¹⁰ most often distributed as part of **xz**: possible names in Linux distributions include **xz-devel**/**xz-libs** and **liblzma-dev**.

¹¹ but not a major version greater than 7 should there ever be one: the major version has been 7 since 2000.

¹² for example to specify static linking with a build which has both shared and static libraries.

¹³ Such as GNU **tar** 1.15 or later, **bsdtar** (from <https://github.com/libarchive/libarchive/>, as used as **tar** by FreeBSD and macOS 10.6 and later) or **tar** from the Heirloom Toolchest (<http://heirloom.sourceforge.net/tools.html>), although the latter does not support **xz** compression.

linked to a version on CRAN. To make PDF versions of the manuals you will also need file `texinfo.tex` installed (which is part of the GNU **texinfo** distribution but is often made part of the \TeX package in re-distributions) as well as `texi2dvi`.¹⁴ Further, the versions of `texi2dvi` and `texinfo.tex` need to be compatible: we have seen problems with older \TeX distributions.

If you want to build from the R Subversion repository then `texi2any` is highly recommended as it is used to create files which are in the tarball but not stored in the Subversion repository.

The PDF documentation (including `doc/NEWS.pdf`) and building vignettes needs `pdftex` and `pdflatex`. We require \LaTeX version 2005/12/01 or later (for UTF-8 support). Building PDF package manuals (including the R reference manual) and vignettes is sensitive to the version of the \LaTeX package **hyperref** and we recommend that the \TeX distribution used is kept up-to-date. A number of standard \LaTeX packages are required (including **url** and some of the font packages such as **times**, **helvetic**, **ec** and **cm-super**) and others such as **hyperref** and **inconsolata** are desirable (and without them you may need to change R's defaults: see Section 2.3 [Making the manuals], page 5). Note that package **hyperref** (currently) requires packages **kvoptions**, **ltxcmds** and **refcount**. For distributions based on TeX Live the simplest approach may be to install collections **collection-latex**, **collection-fontsrecommended**, **collection-latexrecommended**, **collection-fontsextra** and **collection-latexextra** (assuming they are not installed by default): Fedora uses names like **texlive-collection-fontsextra** and Debian/Ubuntu like **texlive-fonts-extra**.

The essential programs should be in your `PATH` at the time `configure` is run: this will capture the full paths.

Those distributing binary versions of R may need to be aware of the licences of the external libraries it is linked to (including 'useful' libraries from the next section). The `liblzma` library is in the public domain and `X11`, `libbzip2`, `libcurl` and `zlib` have MIT-style licences. `PCRE` and `PCRE2` have a BSD-style licence which requires distribution of the licence (included in R's `COPYRIGHTS` file) in binary distributions. GNU `readline` is licensed under GPL (which version(s) of GPL depends on the `readline` version).

A.2 Useful libraries and programs

The ability to use translated messages makes use of `gettext` and most likely needs GNU `gettext`: you do need this to work with new translations, but otherwise the version contained in the R sources will be used if no suitable external `gettext` is found.

The 'modern' version of the `X11()`, `jpeg()`, `png()` and `tiff()` graphics devices uses the `cairo` and (optionally) `Pango` libraries. Cairo version 1.2.0 or later is required. Pango needs to be at least version 1.10, and 1.12 is the earliest version we have tested. (For Fedora users we believe the `pango-devel` RPM and its dependencies suffice.) R checks for `pkg-config`, and uses that to check first that the 'pangocairo' package is installed (and if not, 'cairo') and if additional flags are needed for the 'cairo-xlib' package, then if suitable code can

¹⁴ `texi2dvi` is normally a shell script. Some versions (including those from **texinfo** 5.2 and 6.0-6.6) need to be run under `bash` rather than a Bourne shell, especially on Solaris. Some of the issues which have been observed with broken versions of `texi2dvi` can be circumvented by setting the environment variable `R_TEXI2DVICMD` to the value `emulation`.

be compiled. These tests will fail if `pkg-config` is not installed¹⁵, and are likely to fail if `cairo` was built statically (unusual). Most systems with `Gtk+` 2.8 or later installed will have suitable libraries.

For the best font experience with these devices you need suitable fonts installed: Linux users will want the `urw-fonts` package. On platforms which have it available, the `msttcorefonts` package¹⁶ provides TrueType versions of Monotype fonts such as Arial and Times New Roman. Another useful set of fonts is the ‘liberation’ TrueType fonts available at <https://fedorahosted.org/liberation-fonts/>,¹⁷ which cover the Latin, Greek and Cyrillic alphabets plus a fair range of signs. These share metrics with Arial, Times New Roman and Courier New, and contain fonts rather similar to the first two (https://en.wikipedia.org/wiki/Liberation_fonts). Then there is the ‘Free UCS Outline Fonts’ project (<https://www.gnu.org/software/freefont/>) which are OpenType/TrueType fonts based on the URW fonts but with extended Unicode coverage. See the R help on X11 on selecting such fonts.

The bitmapped graphics devices `jpeg()`, `png()` and `tiff()` need the appropriate headers and libraries installed: `jpeg` (version 6b or later, or `libjpeg-turbo`) or `libpng` (version 1.2.7 or later) and `zlib` or `libtiff` (versions 4.0.[5-10] and 4.1.0 have been tested) respectively. `pkg-config` is used if available and so needs the appropriate `.pc` file (which requires `libtiff` version 4.x and is not available on all platforms for `jpeg` before version 9c). They also need support for either X11 or `cairo` (see above). Should support for these devices **not** be required or broken system libraries need to be avoided there are `configure` options `--without-libpng`, `--without-jpeglib` and `--without-libtiff`. The TIFF library has many optional features such as `jpeg`, `libz`, `lzma`, `jbig` and `jpeg12`, none of which is required for the `tiff()` devices but may need to be present to link the library (usually only an issue for static linking).

Option `--with-system-tre` is also available: it needs a recent version of TRE. (The latest sources are in the git repository at <https://github.com/laurikari/tre/>, but at the time of writing (2016) the resulting build did not pass its checks.).

An implementation of XDR is required, and the R sources contain one which is likely to suffice (although a system version may have higher performance). XDR is part of RPC and historically has been part of `libc` on a Unix-alike. (In principle `man xdr_string` should tell you which library is needed, but it often does not: on Solaris and others it is provided by `libnsl`.) However some builds¹⁸ of `glibc` omit or hide it with the intention that the TI-RPC library be used, in which case `libtirpc` (and its development version) should be installed, and its headers¹⁹ need to be on the C include path or under `/usr/include/tirpc`.

Use of the X11 clipboard selection requires the `Xmu` headers and libraries. These are normally part of an X11 installation (e.g. the Debian meta-package ‘`xorg-dev`’), but some

¹⁵ If necessary the path to `pkg-config` can be specified by setting `PKG_CONFIG` in `config.site`, on the `configure` command line or in the environment. There is a compatible re-implementation of `pkg-config` called `pkgconf` which can be used in the unlikely event that is installed but not linked to `pkg-config`.

¹⁶ also known as `ttf-mscorefonts-installer` in the Debian/Ubuntu world: see also https://en.wikipedia.org/wiki/Core_fonts_for_the_Web.

¹⁷ `ttf-liberation` in Debian/Ubuntu.

¹⁸ This is the default as from `glibc` 2.26 and has been confirmed for Fedora ≥ 28 , which does not mention this on its `man` pages.

¹⁹ R uses `rpc/xdr.h` but that includes `netconfig.h` from the top `tirpc` directory.

distributions have split this into smaller parts, so for example recent versions of Fedora require the ‘libXmu’ and ‘libXmu-devel’ RPMs.

Some systems (notably macOS and at least some FreeBSD systems) have inadequate support for collation in multibyte locales. It is possible to replace the OS’s collation support by that from ICU (International Components for Unicode, <http://site.icu-project.org/>), and this provides much more precise control over collation on all systems. ICU is available as sources and as binary distributions for (at least) most Linux distributions, Solaris, FreeBSD, macOS and AIX, usually as `libicu` or `icu4c`. It will be used by default where available: should a very old or broken version of ICU be found this can be suppressed by `--without-ICU`.

The `bitmap` and `dev2bitmap` devices and function `embedFonts()` use `ghostscript` (<https://www.ghostscript.com/>). This should either be in your path when the command is run, or its full path specified by the environment variable `R_GSCMD` at that time.

At the time of writing a full installation on Fedora Linux used the following packages and their development versions, and this may provide a useful checklist for other systems:

```
bzip2 cairo fontconfig freetype fribidl glib2 harfbuzz libX11 libXext
libXt libcurl libicu libjpeg libpng libtiff libtirpc libxcrypt ncurses
pango pkgconf-pkg-config pcre2 readline tcl tk xz zlib
```

A.2.1 Tcl/Tk

The `tcltk` package needs Tcl/Tk ≥ 8.4 installed: the sources are available at <https://www.tcl.tk/>. To specify the locations of the Tcl/Tk files you may need the configuration options

```
--with-tcltk
    use Tcl/Tk, or specify its library directory

--with-tcl-config=TCL_CONFIG
    specify location of tclConfig.sh

--with-tk-config=TK_CONFIG
    specify location of tkConfig.sh
```

or use the configure variables `TCLTK_LIBS` and `TCLTK_CPPFLAGS` to specify the flags needed for linking against the Tcl and Tk libraries and for finding the `tcl.h` and `tk.h` headers, respectively. If you have both 32- and 64-bit versions of Tcl/Tk installed, specifying the paths to the correct config files may be necessary to avoid confusion between them.

Versions of Tcl/Tk up to 8.5.19 and 8.6.9 have been tested (including most versions of 8.4.x, but not recently).

Note that the `tk.h` header includes²⁰ X11 headers, so you will need X11 and its development files installed.

²⁰ This is true even for the ‘Aqua’ version of Tk on macOS, but distributions of that include a copy of the X11 files needed.

A.2.2 Java support

The build process looks for Java support on the host system, and if it finds it sets some settings which are useful for Java-using packages (such as **rJava** (<https://CRAN.R-project.org/package=rJava>) and **JavaGD** (<https://CRAN.R-project.org/package=JavaGD>): these require a full JDK). This check can be suppressed by configure option `--disable-java`. Configure variable `JAVA_HOME` can be set to point to a specific JRE/JDK, on the `configure` command line or in the environment.

Principal amongst these settings are some paths to the Java libraries and JVM, which are stored in environment variable `R_JAVA_LD_LIBRARY_PATH` in file `R_HOME/etc/ldpaths` (or a sub-architecture-specific version). A typical setting for ‘x86_64’ Linux is

```
JAVA_HOME=/usr/lib/jvm/java-1.8.0-openjdk-1.8.0.265.b01-1.fc32.x86_64/jre
R_JAVA_LD_LIBRARY_PATH=${JAVA_HOME}/lib/amd64/server
```

Unfortunately this depends on the exact version of the JRE/JDK installed, and so may need updating if the Java installation is updated. This can be done by running R CMD `javareconf` which updates settings in both `R_HOME/etc/Makeconf` and `R_HOME/etc/ldpaths`. See R CMD `javareconf --help` for details: note that this needs to be done by the account owning the R installation.

Another way of overriding those settings is to set the environment variable `R_JAVA_LD_LIBRARY_PATH` (before R is started, hence not in `~/.Renvi`), which suffices to run already-installed Java-using packages. For example

```
R_JAVA_LD_LIBRARY_PATH=/usr/lib/jvm/java-1.8.0/jre/lib/amd64/server
```

It may be possible to avoid this by specifying an invariant link as the path when configuring. For example, on that system any of

```
JAVA_HOME=/usr/lib/jvm/java
JAVA_HOME=/usr/lib/jvm/java-1.8.0
JAVA_HOME=/usr/lib/jvm/java-1.8.0/jre
JAVA_HOME=/usr/lib/jvm/jre-1.8.0
```

worked (since the ‘auto’ setting of `/etc/alternatives` chose Java 8 aka 1.8.0).

‘Non-server’ Oracle distributions of Java as from version 11 are of a full JDK. However, Linux distributions can be confusing: for example Fedora 32 had

```
java-1.8.0-openjdk
java-1.8.0-openjdk-devel
java-openjdk
java-openjdk-devel
java-11-openjdk
java-11-openjdk-devel
java-latest-openjdk
java-latest-openjdk-devel
```

where the `-devel` RPMs are needed to complete the JDK. (At the time of writing `java-openjdk` was Java 14.) Debian/Ubuntu use ‘`-jre`’ and ‘`-jdk`’, e.g.

```
sudo apt install default-jdk
```

A.2.3 Other compiled languages

Some add-on packages need a C++ compiler. This is specified by the configure variables `CXX`, `CXXFLAGS` and similar. `configure` will normally find a suitable compiler. It is possible to specify an alternative C++14 compiler by the configure variables `CXX14`, `CXX14STD`, `CXX14FLAGS` and similar (see Section 2.7.3 [C++ Support], page 12). Again, `configure` will normally find a suitable value for `CXX14STD` if the compiler given by `CXX` is capable of compiling C++14 code, but it is possible that a completely different compiler will be needed. (Similar macros are provided for C++17 and C++20.)

For source files with extension `.f90` or `.f95` containing free-form Fortran, the compiler defined by the macro `FC` is used by `R CMD INSTALL`. Note that it is detected by the name of the command without a test that it can actually compile Fortran 90 code. Set the configure variable `FC` to override this if necessary: variables `FCFLAGS` and `FCLIBS_XTRA` might also need to be set.

See file `config.site` in the R source for more details about these variables.

A.3 Linear algebra

A.3.1 BLAS

The linear algebra routines in R can make use of enhanced BLAS (Basic Linear Algebra Subprograms, <http://www.netlib.org/blas/faq.html>) routines. However, these have to be explicitly requested at configure time: R provides an internal BLAS which is well-tested and will be adequate for most uses of R.

You can specify a particular BLAS library *via* a value for the configuration option `--with-blas` and not to use an external BLAS library by `--without-blas` (the default). If `--with-blas` is given with no `=`, its value is taken from the environment variable `BLAS_LIBS`, set for example in `config.site`. If neither the option nor the environment variable supply a value, a search is made for a suitable²¹ BLAS. If the value is not obviously a linker command (starting with a dash or giving the path to a library), it is prefixed by `'-l'`, so

```
--with-blas="foo"
```

is an instruction to link against `'-lfoo'` to find an external BLAS (which needs to be found both at link time and run time).

The configure code checks that the external BLAS is complete (it must include all double precision and double complex routines, as well as `LSAME`), and appears to be usable. However, an external BLAS has to be usable from a shared object (so must contain position-independent code), and that is not checked.

Some enhanced BLASes are compiler-system-specific (`sunperf` on Solaris²², `libessl` on IBM, `Accelerate` on macOS). The correct incantation for these is often found *via* `--with-blas` with no value on the appropriate platforms.

Some of the external BLASes are multi-threaded. One issue is that R profiling (which uses the `SIGPROF` signal) may cause problems, and you may want to disable profiling if you use a multi-threaded BLAS. Note that using a multi-threaded BLAS can result in taking

²¹ The search includes OpenBLAS, ATLAS and a generic `libblas`, plus some platform-specific choices (see below).

²² Using the Oracle Developer Studio `cc` and `f95` compilers

more CPU time and even more elapsed time (occasionally dramatically so) than using a similar single-threaded BLAS. On a machine running other tasks, there can be contention for CPU caches that reduces the effectiveness of the optimization of cache use by a BLAS implementation: some people warn that this is especially problematic for hyperthreaded CPUs.

Note that under Unix (but not under Windows) if R is compiled against a non-default BLAS and `--enable-BLAS-shlib` is **not** used (it is the default on all platforms except AIX), then all BLAS-using packages must also be. So if R is re-built to use an enhanced BLAS then packages such as **quantreg** (<https://CRAN.R-project.org/package=quantreg>) will need to be re-installed; they may be under other circumstances.

R relies on ISO/IEC 60559 compliance of an external BLAS. This can be broken if for example the code assumes that terms with a zero factor are always zero and do not need to be computed—whereas $x*0$ can be NaN. This is checked in the test suite.

External BLAS implementations often make less use of extended-precision floating-point registers (where available) and will almost certainly re-order computations. This can result in less accuracy than using a reference BLAS, and may result in different solutions, e.g. different signs in SVD and eigendecompositions.

Debian/Ubuntu systems provide a system-specific way to switch the BLAS in use. Build R with `-with-blas` to select the OS version of the reference BLAS, and then use `update-alternatives` to switch between the available BLAS libraries. See <https://wiki.debian.org/DebianScience/LinearAlgebraLibraries>.

The URIs for several of these BLAS have been subject to frequent gratuitous changes, so you will need to search for their current locations.

BLAS (and LAPACK) routines may be used inside threaded code, for example in OpenMP sections in packages such as **mgcv**. The reference implementations are thread-safe but external ones may not be (even single-threaded ones): this can lead to hard-to-track-down incorrect results or segfaults.

NOTE: BLAS libraries built with **gfortran** 9 (and later and versions 8.4, 7.5 and later in those series) require calls from C/C++ to handle ‘hidden’ character lengths — R itself does so but many packages do not and some segfault. (This applies also to external LAPACK libraries.)

A.3.1.1 ATLAS

ATLAS (<http://math-atlas.sourceforge.net/>) is a “tuned” BLAS that runs on a wide range of Unix-alike platforms. Unfortunately it is built by default as a static library that on some platforms may not be able to be used with shared objects such as are used in R packages. Be careful when using pre-built versions of ATLAS static libraries (they seem to work on ‘ix86’ platforms, but not always on ‘x86_64’ ones).

ATLAS contains replacements for a small number of LAPACK routines, but can be built to merge these with LAPACK sources to include a full LAPACK library.

Recent versions of ATLAS can be built as a single shared library, either **libsatlas** or **libtatlas** (serial or threaded respectively): these may even contain a full LAPACK. Such builds can be used by one of

```
--with-blas=satlas
```

```
--with-blas=tatlas
```

or, as on ‘x86_64’ Fedora where a path needs to be specified,

```
--with-blas="-L/usr/lib64/atlas -lsatlas"
--with-blas="-L/usr/lib64/atlas -ltatlas"
```

Distributed ATLAS libraries cannot be tuned to your machine and so are a compromise: for example Fedora tunes ‘x86_64’ for CPUs with SSE3 extensions, and separate ‘atlas-sse2’ and ‘atlas-sse3’ ‘i686’ RPMs are available.²³

Note that building R on Linux against distributed shared libraries may need ‘-devel’ or ‘-dev’ packages installed.

Linking against multiple static libraries requires one of

```
--with-blas="-lf77blas -latlas"
--with-blas="-lptf77blas -lpthread -latlas"
--with-blas="-L/path/to/ATLAS/libs -lf77blas -latlas"
--with-blas="-L/path/to/ATLAS/libs -lptf77blas -lpthread -latlas"
```

Consult its installation guide²⁴ for how to build ATLAS as a shared library or as a static library with position-independent code (on platforms where that matters).

According to the ATLAS FAQ²⁵ the maximum number of threads used by multi-threaded ATLAS is set at compile time. Also, the author advises against using multi-threaded ATLAS on hyperthreaded CPUs without restricting affinities at compile-time to one virtual core per physical CPU. (For the Fedora libraries the compile-time flag specifies 4 threads.)

A.3.1.2 OpenBLAS

Dr Kazushige Goto wrote a tuned BLAS for several processors and OSes, which was frozen in mid-2010. OpenBLAS (<http://www.openblas.net/>) is a descendant project with support for some later CPUs.

This can be used by configuring R with something like

```
--with-blas="-lopenblas"
```

See see Section A.3.1.4 [Shared BLAS], page 51, for an alternative (and in many ways preferable) way to use them.

Some platforms provide multiple builds of OpenBLAS: for example Fedora 32 has RPMs²⁶

```
openblas
openblas-threads
openblas-openmp
```

providing shared libraries

```
libopenblas.so
libopenblas.so
```

²³ There were others for earlier versions of ATLAS, and are for non-Intel architectures. The only way to see exactly which CPUs the distributed libraries have been tuned for is to read the `atlas.spec` file: at the time of writing ‘HAMMER64SSE3’ and ‘Corei264AVX’ for ‘x86_64’ Fedora.

²⁴ http://math-atlas.sourceforge.net/atlas_install/

²⁵ <http://math-atlas.sourceforge.net/faq.html#tnum>

²⁶ (and more, e.g. for 64-bit ints and static versions).

`libopenblas.so`

respectively, each of which can be used as a shared BLAS. For the second and third the number of threads is controlled by `OPENBLAS_NUM_THREADS` and `OMP_NUM_THREADS` (as usual for OpenMP) respectively. There is also a Fedora RPM ‘`openblas-Rblas`’ to replace `libRblas.so` in their distribution of R.

Note that building R on Linux against distributed libraries may need ‘`-devel`’ or ‘`-dev`’ packages installed.

For ‘`ix86`’ and ‘`x86_64`’ most distributed libraries contain several alternatives for different CPU microarchitectures with the choice being made at run time.

A.3.1.3 Intel MKL

For Intel processors (and perhaps others) and some distributions of Linux, there is Intel’s Math Kernel Library. You are strongly encouraged to read the MKL User’s Guide, which is installed with the library, before attempting to link to MKL. This includes a ‘link line advisor’ which will suggest appropriate incantations: its use is recommended. Or see <https://software.intel.com/content/www/us/en/develop/articles/intel-mkl-link-line-advisor.html>.

There are also versions of MKL for macOS²⁷ and Windows, but when these have been tried they did not work with the default compilers used for R on those platforms.

The MKL interface has changed several times but has been stable in recent years: the following examples have been used with versions 10.3 to 2020.2, for GCC compilers on ‘`x86_64`’.

To use a sequential version of MKL we used

```
MKL_LIB_PATH=/path/to/intel_mkl/mkl/lib/intel64
export LD_LIBRARY_PATH=$MKL_LIB_PATH
MKL="-L${MKL_LIB_PATH} -lmkl_gf_lp64 -lmkl_core -lmkl_sequential"
./configure --with-blas="$MKL" --with-lapack
```

The option `--with-lapack` is used since MKL contains a tuned copy of LAPACK (often older than the current version) as well as BLAS (see Section A.3.2 [LAPACK], page 52), although this can be omitted.

Threaded MKL may be used by replacing the line defining the variable `MKL` by

```
MKL="-L${MKL_LIB_PATH} -lmkl_gf_lp64 -lmkl_core \
-lmkl_gnu_thread -dl -fopenmp"
```

R can also be linked against a single shared library, `libmkl_rt.so`, for both BLAS and LAPACK, but the correct OpenMP and MKL interface layer then has to be selected via environment variables. With 64-bit builds and the GCC compilers, we used

```
export MKL_INTERFACE_LAYER=GNU,LP64
export MKL_THREADING_LAYER=GNU
```

On Debian/Ubuntu, MKL is provided by package `intel-mkl-full` and one can set `libmkl_rt.so` as the system-wide implementation of both BLAS and LAPACK during installation of the package, so that also R installed from Debian/Ubuntu package `r-base` would use it. It is, however, still essential to set `MKL_INTERFACE_LAYER` and

²⁷ The issue for macOS has been the use of double-complex routines.

`MKL_THREADING_LAYER` before running R, otherwise MKL computations will produce incorrect results. R does not have to be rebuilt to use MKL, but `configure` include tests which may discover some errors such as a failure to set the correct OpenMP and MKL interface layer.

The default number of threads will be chosen by the OpenMP software, but can be controlled by setting `OMP_NUM_THREADS` or `MKL_NUM_THREADS`, and in recent versions seems to default to a sensible value for sole use of the machine. (Parallel MKL has not always passed `make check-all`, but did with MKL 2019.4 and later.)

MKL includes a partial implementation of FFTW3, which causes trouble for applications that require some of the FFTW3 functionality unsupported in MKL. Please see the MKL manuals for description of these limitations and for instructions on how to create a custom version of MKL which excludes the FFTW3 wrappers.

It was reported in 2015 that

```
--with-blas='-mkl=parallel' --with-lapack
```

worked with the Intel 2015.3 compilers on Centos 6.

A.3.1.4 Shared BLAS

The BLAS library will be used for many of the add-on packages as well as for R itself. This means that it is better to use a shared/dynamic BLAS library, as most of a static library will be compiled into the R executable and each BLAS-using package.

R offers the option of compiling the BLAS into a dynamic library `libRblas` stored in `R_HOME/lib` and linking both R itself and all the add-on packages against that library.

This is the default on all platforms except AIX unless an external BLAS is specified and found: for the latter it can be used by specifying the option `--enable-BLAS-shlib`, and it can always be disabled via `--disable-BLAS-shlib`.

This has both advantages and disadvantages.

- It saves space by having only a single copy of the BLAS routines, which is helpful if there is an external static BLAS (as used to be standard for ATLAS).
- There may be performance disadvantages in using a shared BLAS. Probably the most likely is when R's internal BLAS is used and R is *not* built as a shared library, when it is possible to build the BLAS into `R.bin` (and `libR.a`) without using position-independent code. However, experiments showed that in many cases using a shared BLAS was as fast, provided high levels of compiler optimization are used.
- It is easy to change the BLAS without needing to re-install R and all the add-on packages, since all references to the BLAS go through `libRblas`, and that can be replaced. Note though that any dynamic libraries the replacement links to will need to be found by the linker: this may need the library path to be changed in `R_HOME/etc/ldpaths`.

Another option to change the BLAS in use is to symlink a single dynamic BLAS library to `R_HOME/lib/libRblas.so`. For example, just

```
mv R_HOME/lib/libRblas.so R_HOME/lib/libRblas.so.keep
ln -s /usr/lib64/libopenblas.so.0 R_HOME/lib/libRblas.so
```

on 'x86_64' Fedora will change the BLAS used to multithreaded OpenBLAS. A similar link works for most versions of the OpenBLAS (provided the appropriate `lib` directory is in the

run-time library path or `ld.so` cache). It can also be used for a single-library ATLAS, so on ‘x86_64’ Fedora

```
ln -s /usr/lib64/atlas/libsatlas.so.3 R_HOME/lib/libRblas.so
ln -s /usr/lib64/atlas/libtatlas.so.3 R_HOME/lib/libRblas.so
```

can be used with its distributed ATLAS libraries. (If you have the ‘`-devel`’ RPMS installed you can omit the `.0/.3`.)

Note that rebuilding or symlinking `libRblas.so` may not suffice if the intention is to use a modified LAPACK contained in an external BLAS: the latter could even cause conflicts. However, on Fedora where the OpenBLAS distribution contains a copy of LAPACK, it is the latter which is used.

A.3.2 LAPACK

Provision is made for using an external LAPACK library, principally to cope with BLAS libraries which contain a copy of LAPACK (such as `sunperf` on Solaris, `Accelerate` on macOS and ATLAS and MKL on ‘ix86’/‘x86_64’ Linux). At least LAPACK version 3.2 is required. This can only be done if `--with-blas` has been used.

However, the likely performance gains are thought to be small (and may be negative). The default is not to search for a suitable LAPACK library, and this is definitely **not** recommended. You can specify a specific LAPACK library or a search for a generic library by the configuration option `--with-lapack`. The default for `--with-lapack` is to check the BLAS library and then look for an external library ‘`-llapack`’. Sites searching for the fastest possible linear algebra may want to build a LAPACK library using the ATLAS-optimized subset of LAPACK: this is simplest with a dynamic ATLAS library which contains a full LAPACK, when `--with-lapack` suffices.

A value for `--with-lapack` can be set *via* the environment variable `LAPACK_LIBS`, but this will only be used if `--with-lapack` is specified (as the default value is `no`) and the BLAS library does not contain LAPACK.

If you do use `--with-lapack`, be aware of potential problems with bugs in the LAPACK sources (or in the posted corrections to those sources). In particular, bugs in `DGEEV` and `DGESDD` have resulted in error messages such as

```
DGEBRD gave error code -10
```

. Other potential problems are incomplete versions of the libraries, seen several times in Linux distributions over the years.

Please **do** bear in mind that using `--with-lapack` is ‘definitely **not** recommended’: it is provided **only** because it is necessary on some platforms and because some users want to experiment with claimed performance improvements. Reporting problems where it is used unnecessarily will simply irritate the R helpers.

Note too the comments about ISO/IEC 60559 compliance in the section of external BLAS: these apply equally to an external LAPACK, and for example the Intel MKL documentation has said

```
LAPACK routines assume that input matrices do not contain IEEE 754 special values such as INF or NaN values. Using these special values may cause LAPACK to return unexpected results or become unstable.
```

We rely on limited support in LAPACK for matrices with 2^{31} or more elements: it is possible that an external LAPACK will not have that support.

A.3.3 Caveats

As with all libraries, you need to ensure that they and R were compiled with compatible compilers and flags. For example, this has meant that on Sun Sparc using the Oracle compilers the flag `-dalign` is needed if `sunperf` is to be used.

On some systems it has been necessary that an external BLAS/LAPACK was built with the same Fortran compiler used to build R.

LAPACK 3.9.0 has a bug in which the DCOMBSSQ subroutine has a bug that may cause NA to be interpreted as zero. This is fixed in the R 3.6.3 and later sources, but if you use an external LAPACK, you may need to fix it there.

The code (in `dlapack.f`) should read

```
*      ..
*      .. Executable Statements ..
*
      IF( V1( 1 ).GE.V2( 1 ) ) THEN
        IF( V1( 1 ).NE.ZERO ) THEN
          V1( 2 ) = V1( 2 ) + ( V2( 1 ) / V1( 1 ) )**2 * V2( 2 )
        ELSE
          V1( 2 ) = V1( 2 ) + V2( 2 )
        END IF
      ELSE
        V1( 2 ) = V2( 2 ) + ( V1( 1 ) / V2( 1 ) )**2 * V1( 2 )
        V1( 1 ) = V2( 1 )
      END IF
      RETURN
```

(The inner ELSE clause was missing in LAPACK 3.9.0.)

Appendix B Configuration on a Unix-alike

B.1 Configuration options

`configure` has many options: running

```
./configure --help
```

will give a list. Probably the most important ones not covered elsewhere are (defaults in brackets)

```
--with-x    use the X Window System [yes]
```

```
--x-includes=DIR
            X include files are in DIR
```

```
--x-libraries=DIR
            X library files are in DIR
```

```
--with-readline
            use readline library (if available) [yes]
```

```
--enable-R-profiling
            attempt to compile support for Rprof() [yes]
```

```
--enable-memory-profiling
            attempt to compile support for Rprofmem() and tracemem() [no]
```

```
--enable-R-shlib
            build R as a shared/dynamic library [no]
```

```
--enable-BLAS-shlib
            build the BLAS as a shared/dynamic library [yes, except on AIX]
```

You can use `--without-foo` or `--disable-foo` for the negatives.

You will want to use `--disable-R-profiling` if you are building a profiled executable of R (e.g. with `'-pg'`). Support for R profiling requires OS support for POSIX threads (*aka* `'pthreads'`), which are available on all mainstream Unix-alike platforms.

Flag `--enable-R-shlib` causes the make process to build R as a dynamic (shared) library, typically called `libR.so`, and link the main R executable `R.bin` against that library. This can only be done if all the code (including system libraries) can be compiled into a dynamic library, and there may be a performance¹ penalty. So you probably only want this if you will be using an application which embeds R. Note that C code in packages installed on an R system linked with `--enable-R-shlib` is linked against the dynamic library and so such packages cannot be used from an R system built in the default way. Also, because packages are linked against R they are on some OSes also linked against the dynamic libraries R itself is linked against, and this can lead to symbol conflicts.

For maximally effective use of `valgrind`, R should be compiled with `valgrind` instrumentation. 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.) The system headers

¹ We have measured 15–20% on `'i686'` Linux and around 10% on `'x86_64'` Linux.

for `valgrind` can be requested by option `--with-system-valgrind-headers`: they will be used if present (on Linux they may be in a separate package such as `valgrind-devel`). Note though that there is no guarantee that the code in R will be compatible with very old² or future `valgrind` headers.

If you need to re-configure R with different options you may need to run `make clean` or even `make distclean` before doing so.

The `configure` script has other generic options added by `autoconf` and which are not supported for R: in particular building for one architecture on a different host is not possible.

B.2 Internationalization support

Translation of messages is supported via GNU `gettext` unless disabled by the configure option `--disable-nls`. The `configure` report will show NLS as one of the ‘Additional capabilities’ if support has been compiled in, and running in an English locale (but not the C locale) will include

Natural language support but running in an English locale
in the greeting on starting R.

B.3 Configuration variables

If you need or want to set certain configure variables to something other than their default, you can do that by either editing the file `config.site` (which documents many of the variables you might want to set: others can be seen in file `etc/Renviron.in`) or on the command line as

`./configure VAR=value`

If you are building in a directory different from the sources, there can be copies of `config.site` in the source and the build directories, and both will be read (in that order). In addition, if there is a file `~/.R/config`, it is read between the `config.site` files in the source and the build directories.

There is also a general `autoconf` mechanism for `config.site` files, which are read before any of those mentioned in the previous paragraph. This looks first at a file specified by the environment variable `CONFIG_SITE`, and if not is set at files such as `/usr/local/share/config.site` and `/usr/local/etc/config.site` in the area (exemplified by `/usr/local`) where R would be installed.

These variables are *precious*, implying that they do not have to be exported to the environment, are kept in the cache even if not specified on the command line, checked for consistency between two configure runs (provided that caching is used), and are kept during automatic reconfiguration as if having been passed as command line arguments, even if no cache is used.

See the variable output section of `configure --help` for a list of all these variables.

If you find you need to alter configure variables, it is worth noting that some settings may be cached in the file `config.cache`, and it is a good idea to remove that file (if it exists) before re-configuring. Note that caching is turned *off* by default: use the command line option `--config-cache` (or `-C`) to enable caching.

² We believe that versions 3.4.0 to 3.15.0 are compatible.

B.3.1 Setting paper size

One common variable to change is `R_PAPERSIZE`, which defaults to `'a4'`, not `'letter'`. (Valid values are `'a4'`, `'letter'`, `'legal'` and `'executive'`.)

This is used both when configuring R to set the default, and when running R to override the default. It is also used to set the paper size when making PDF manuals.

The configure default will most often be `'a4'` if `R_PAPERSIZE` is unset. (If the (Debian Linux) program `paperconf` is found or the environment variable `PAPERSIZE` is set, these are used to produce the default.)

B.3.2 Setting the browsers

Another precious variable is `R_BROWSER`, the default HTML browser, which should take a value of an executable in the user's path or specify a full path.

Its counterpart for PDF files is `R_PDFVIEWER`.

B.3.3 Compilation flags

If you have libraries and header files, e.g., for GNU readline, in non-system directories, use the variables `LDFLAGS` (for libraries, using `'-L'` flags to be passed to the linker) and `CPPFLAGS` (for header files, using `'-I'` flags to be passed to the C/C++ preprocessors), respectively, to specify these locations. These default to `'-L/usr/local/lib'` (`LDFLAGS`, `'-L/usr/local/lib64'` on most 64-bit Linux OSes) and `'-I/usr/local/include'` (`CPPFLAGS`, but note that on most systems `/usr/local/include` is regarded as a system include directory and so instances in that macro will be skipped) to catch the most common cases. If libraries are still not found, then maybe your compiler/linker does not support re-ordering of `-L` and `-l` flags. In this case, use a different compiler (or a front-end shell script which does the re-ordering).

These flags can also be used to build a faster-running version of R. On most platforms using `gcc`, having `'-O3'` in `CFLAGS` and `FFLAGS` produces worthwhile performance gains with `gcc` and `gfortran`, but may result in a less reliable build (both segfaults and incorrect numeric computations have been seen). On systems using the GNU linker (especially those using R as a shared library), it is likely that including `'-Wl,-O1'` in `LDFLAGS` is worthwhile, and `'-Bdirect,--hash-style=both,-Wl,-O1'` is recommended at <https://lwn.net/Articles/192624/>. Tuning compilation to a specific CPU family (e.g. `'-mtune=native'` for `gcc`) can give worthwhile performance gains, especially on older architectures such as `'ix86'`.

B.3.4 Making manuals

The default settings for making the manuals are controlled by `R_RD4PDF` and `R_PAPERSIZE`.

B.4 Setting the shell

By default the shell scripts such as R will be `'#!/bin/sh'` scripts (or using the `SHELL` chosen by `configure`). This is almost always satisfactory, but on a few systems `/bin/sh` is not a Bourne shell or clone, and the shell to be used can be changed by setting the configure variable `R_SHELL` to a suitable value (a full path to a shell, e.g. `/usr/local/bin/bash`).

B.5 Using make

To compile R, you will most likely find it easiest to use GNU **make**, although the Sun **make** works on Solaris.

To build in a separate directory you need a **make** that supports the **VPATH** variable, for example GNU **make** and Sun **make**.

dmake has also been used. e.g, on Solaris 10.

If you want to use a **make** by another name, for example if your GNU **make** is called '**gmake**', you need to set the variable **MAKE** at configure time, for example

```
./configure MAKE=gmake
```

B.6 Using Fortran

To compile R, you need a Fortran 90 compiler. The current default is to search for **gfortran**, **g95**, **xl95**, **f95**, **fort**, **ifort**, **ifc**, **efc**, **pgfortran**, **pgf95**, **lf95**, **ftn**, **nagfor**, **xl90**, **f90**, **pgf90**, **pghp90**, **epcf90**. (Note that these are searched for by name, without checking the standard of Fortran they support.) The command and flags used should support fixed-form Fortran with extension **.f**: in the unusual case that a specific flag is needed for free-form Fortran with extension **.f90** or **.f95**, this can be specified as part of **FCFLAGS**.

The search mechanism can be changed using the configure variable **FC** which specifies the command that runs the Fortran compiler. If your Fortran compiler is in a non-standard location, you should set the environment variable **PATH** accordingly before running **configure**, or use the configure variable **FC** to specify its full path.

If your Fortran libraries are in slightly peculiar places, you should also look at **LD_LIBRARY_PATH** (or your system's equivalent) to make sure that all libraries are on this path.

Note that only Fortran compilers which convert identifiers to lower case are supported.

You must set whatever compilation flags (if any) are needed to ensure that Fortran **integer** is equivalent to a C **int** pointer and Fortran **double precision** is equivalent to a C **double** pointer. This is checked during the configuration process.

Some of the Fortran code makes use of **DOUBLE COMPLEX** and **COMPLEX*16** variables. This is checked for at configure time, as well as its equivalence to the **Rcomplex** C structure defined in **R_ext/Complex.h**.

gfortran 10 by default gives a compilation error for the previously widespread practice of passing a Fortran array element where an array is expected, or a scalar instead of a length-one array. See https://gcc.gnu.org/gcc-10/porting_to.html.

B.7 Compile and load flags

A wide range of flags can be set in the file **config.site** or as configure variables on the command line. We have already mentioned

CPPFLAGS header file search directory (**-I**) and any other miscellaneous options for the C and C++ preprocessors and compilers

LDFLAGS path (**-L**), stripping (**-s**) and any other miscellaneous options for the linker and others include

CFLAGS debugging and optimization flags, C

MAIN_CFLAGS
ditto, for compiling the main program (e.g. when profiling)

SHLIB_CFLAGS
for shared objects (no known examples)

FFLAGS debugging and optimization flags, fixed-form Fortran

FCFLAGS debugging and optimization flags, free-form Fortran

SAFE_FFLAGS
ditto for source files which need exact floating point behaviour

MAIN_FFLAGS
ditto, for compiling the main program (e.g. when profiling)

SHLIB_FFLAGS
for shared objects (no known examples)

MAIN_LDFLAGS
additional flags for the main link

SHLIB_LDFLAGS
additional flags for linking the shared objects

LIBnn the primary library directory, `lib` or `lib64`

CPICFLAGS
special flags for compiling C code to be turned into a shared object

FPICFLAGS
special flags for compiling Fortran code to be turned into a shared object

CXXPICFLAGS
special flags for compiling C++ code to be turned into a shared object

DEFS defines to be used when compiling C code in R itself

Library paths specified as `-L/lib/path` in **LDLDFLAGS** are collected together and prepended to **LD_LIBRARY_PATH** (or your system's equivalent), so there should be no need for `-R` or `-rpath` flags.

Variables such as **CPICFLAGS** are determined where possible by **configure**. Some systems allows two types of PIC flags, for example `-fpic` and `-fPIC`, and if they differ the first allows only a limited number of symbols in a shared object. Since R as a shared library has about 6200 symbols, if in doubt use the larger version.

Other variables often set by **configure** include `'MAIN_LDFLAGS'`, `'SAFE_FFLAGS'`, `'SHLIB_LDFLAGS'` and `'SHLIB_CXXLDFLAGS'`: see file `config.site` in the sources for more documentation on these and others.

To compile a profiling version of R, one might for example want to use `'MAIN_CFLAGS=-pg'`, `'MAIN_FFLAGS=-pg'`, `'MAIN_LDFLAGS=-pg'` on platforms where `'-pg'` cannot be used with position-independent code.

Beware: it may be necessary to set **CFLAGS** and **FFLAGS** in ways compatible with the libraries to be used: one possible issue is the alignment of doubles, another is the way structures are passed.

On some platforms `configure` will select additional flags for `CFLAGS`, `CPPFLAGS` and `LIBS` in `R_XTRA_CFLAGS` (and so on). These are for options which are always required, for example to force IEC 60559 compliance.

B.8 Maintainer mode

There are several files that are part of the R sources but can be re-generated from their own sources by configuring with option `--enable-maintainer-mode` and then running `make` in the build directory. This requires other tools to be installed, discussed in the rest of this section.

File `configure` is created from `configure.ac` and the files under `m4` by `autoconf` and `aclocal` (part of the `automake` package). There is a formal version requirement on `autoconf` of 2.69 or later, but it is unlikely that anything other than the most recent versions³ have been thoroughly tested.

File `src/include/config.h` is created by `autoheader` (part of `autoconf`).

Grammar files `*.y` are converted to C sources by an implementation of `yacc`, usually `bison -y`: these are found in `src/main` and `src/library/tools/src`. It is known that earlier versions of `bison` generate code which reads (and in some cases writes) outside array bounds: `bison` 2.6.1 was found to be satisfactory.

The ultimate sources for package `compiler` are in its `noweb` directory. To re-create the sources from `src/library/compiler/noweb/compiler.nw`, the command `notangle` is required. Some Linux distributions include this command in package `noweb`. It can also be installed from the sources at <https://www.cs.tufts.edu/~nr/noweb/>⁴. The package sources are only re-created even in maintainer mode if `src/library/compiler/noweb/compiler.nw` has been updated.

³ at the time of revision of this para in late 2018, `autoconf-2.69` from 2012 and `automake-1.16.1` from 2018.

⁴ The links there have proved difficult to access, in which case either point an FTP client at `ftp://www.eecs.harvard.edu/pub/nr/` or grab the copy made available at <https://developer.r-project.org/noweb-2.11b.tgz>.

Appendix C Platform notes

This section provides some notes on building R on different Unix-alike platforms. These notes are based on tests run on one or two systems in each case with particular sets of compilers and support libraries. Success in building R depends on the proper installation and functioning of support software; your results may differ if you have other versions of compilers and support libraries.

Older versions of this manual contain notes on platforms such as HP-UX, IRIX, Alpha/OSF1 (for R < 2.10.0, and support has since been removed for all of these) and AIX (for R <= 3.5.x) for which we have had no recent reports.

C macros to select particular platforms can be tricky to track down (there is a fair amount of misinformation on the Web). The Wiki (currently) at <http://sourceforge.net/p/predef/wiki/Home/> can be helpful. The R sources have used (often in included software under `src/extra`)

```
AIX: _AIX
Cygwin: __CYGWIN__
FreeBSD: __FreeBSD__
HP-UX: __hpux__, __hpux
IRIX: sgi, __sgi
Linux: __linux__
macOS: __APPLE__
NetBSD: __NetBSD__
OpenBSD: __OpenBSD__
Solaris: __sun, sun
Windows: _WIN32, _WIN64
```

C.1 X11 issues

The ‘X11()’ graphics device is the one started automatically on Unix-alikes (except most macOS builds) when plotting. As its name implies, it displays on a (local or remote) X server, and relies on the services provided by the X server.

The ‘modern’ version of the ‘X11()’ device is based on ‘cairo’ graphics and (in most implementations) uses ‘fontconfig’ to pick and render fonts. This is done on the server, and although there can be selection issues, they are more amenable than the issues with ‘X11()’ discussed in the rest of this section.

When X11 was designed, most displays were around 75dpi, whereas today they are of the order of 100dpi or more. If you find that X11() is reporting¹ missing font sizes, especially larger ones, it is likely that you are not using scalable fonts and have not installed the 100dpi versions of the X11 fonts. The names and details differ by system, but will likely have something like Fedora’s

```
xorg-x11-fonts-75dpi
xorg-x11-fonts-100dpi
xorg-x11-fonts-IS08859-2-75dpi
xorg-x11-fonts-Type1
```

¹ for example, X11 font at size 14 could not be loaded.

xorg-x11-fonts-cyrillic

and you need to ensure that the ‘-100dpi’ versions are installed and on the X11 font path (check via `xset -q`). The ‘X11()’ device does try to set a pointsize and not a pixel size: laptop users may find the default setting of 12 too large (although very frequently laptop screens are set to a fictitious dpi to appear like a scaled-down desktop screen).

More complicated problems can occur in non-Western-European locales, so if you are using one, the first thing to check is that things work in the `C` locale. The likely issues are a failure to find any fonts or glyphs being rendered incorrectly (often as a pair of ASCII characters). X11 works by being asked for a font specification and coming up with its idea of a close match. For text (as distinct from the symbols used by `plotmath`), the specification is the first element of the option “X11fonts” which defaults to

```
"-adobe-helvetica-%s-%s-***%d-***-***-***"
```

If you are using a single-byte encoding, for example ISO 8859-2 in Eastern Europe or KOI8-R in Russian, use `xlsfonts` to find an appropriate family of fonts in your encoding (the last field in the listing). If you find none, it is likely that you need to install further font packages, such as ‘xorg-x11-fonts-ISO8859-2-75dpi’ and ‘xorg-x11-fonts-cyrillic’ shown in the listing above.

Multi-byte encodings (most commonly UTF-8) are even more complicated. There are few fonts in ‘iso10646-1’, the Unicode encoding, and they only contain a subset of the available glyphs (and are often fixed-width designed for use in terminals). In such locales *fontsets* are used, made up of fonts encoded in other encodings. If the locale you are using has an entry in the ‘XLC_LOCALE’ directory (typically `/usr/share/X11/locale`), it is likely that all you need to do is to pick a suitable font specification that has fonts in the encodings specified there. If not, you may have to get hold of a suitable locale entry for X11. This may mean that, for example, Japanese text can be displayed when running in ‘ja_JP.UTF-8’ but not when running in ‘en_GB.UTF-8’ on the same machine (although on some systems many UTF-8 X11 locales are aliased to ‘en_US.UTF-8’ which covers several character sets, e.g. ISO 8859-1 (Western European), JISX0208 (Kanji), KSC5601 (Korean), GB2312 (Chinese Han) and JISX0201 (Kana)).

On some systems scalable fonts are available covering a wide range of glyphs. One source is TrueType/OpenType fonts, and these can provide high coverage. Another is Type 1 fonts: the URW set of Type 1 fonts provides standard typefaces such as Helvetica with a larger coverage of Unicode glyphs than the standard X11 bitmaps, including Cyrillic. These are generally not part of the default install, and the X server may need to be configured to use them. They might be under the X11 `fonts` directory or elsewhere, for example,

```
/usr/share/fonts/default/Type1
/usr/share/fonts/ja/TrueType
```

C.2 Linux

Linux is the main development platform for R, so compilation from the sources is normally straightforward with the most common compilers and libraries.²

Recall that some package management systems (such as RPM and deb) make a distinction between the user version of a package and the developer version. The latter usually has

² For example, `glibc`: other C libraries such as `musl` have been used but are not routinely tested.

the same name but with the extension ‘-devel’ or ‘-dev’: you need both versions installed. So please check the `configure` output to see if the expected features are detected: if for example ‘`readline`’ is missing add the developer package. (On most systems you will also need ‘`ncurses`’ and its developer package, although these should be dependencies of the ‘`readline`’ package(s).) You should expect to see in the `configure` summary

```
Interfaces supported:      X11, tcltk
External libraries:       pcre2, readline, curl
Additional capabilities:   PNG, JPEG, TIFF, NLS, cairo, ICU
```

When R has been installed from a binary distribution there are sometimes problems with missing components such as the Fortran compiler. Searching the ‘`R-help`’ archives will normally reveal what is needed.

It seems that ‘`ix86`’ Linux accepts non-PIC code in shared libraries, but this is not necessarily so on other platforms, in particular on 64-bit CPUs such as ‘`x86_64`’. So care can be needed with BLAS libraries and when building R as a shared library to ensure that position-independent code is used in any static libraries (such as the Tcl/Tk libraries, `libpng`, `libjpeg` and `zlib`) which might be linked against. Fortunately these are normally built as shared libraries with the exception of the ATLAS BLAS libraries.

The default optimization settings chosen for `CFLAGS` etc are conservative. It is likely that using `-mtune` will result in significant performance improvements on recent CPUs: one possibility is to add `-mtune=native` for the best possible performance on the machine on which R is being installed. It is also possible to increase the optimization levels to `-O3`: however for many versions of the compilers this has caused problems in at least one CRAN package.

For platforms with both 64- and 32-bit support, it is likely that

```
LDFLAGS="-L/usr/local/lib64 -L/usr/local/lib"
```

is appropriate since most (but not all) software installs its 64-bit libraries in `/usr/local/lib64`. To build a 32-bit version of R on ‘`x86_64`’ with Fedora 32 we used

```
CC="gcc -m32"
CXX="g++ -m32"
FC="gfortran -m32"
OBJC=${CC}
LDFLAGS="-L/usr/local/lib"
LIBnn=lib
```

Note the use of ‘`LIBnn`’: ‘`x86_64`’ Fedora installs its 64-bit software in `/usr/lib64` and 32-bit software in `/usr/lib`. Linking will skip over inappropriate binaries, but for example the 32-bit Tcl/Tk configure scripts are in `/usr/lib`. It may also be necessary to set the `pkg-config` path, e.g. by

```
export PKG_CONFIG_PATH=/usr/local/lib/pkgconfig:/usr/lib/pkgconfig
```

The 32-bit system `libcurl` did not work with the system CA certificates: this is worked around in R’s test suite.

64-bit versions on Linux are built with support for files > 2Gb, and 32-bit versions will be if possible unless `--disable-largefile` is specified.

Note that 32-bit `glibc` before version 2.29 uses a 32-bit `time_t` type, so to pass all the date-time checks needs R built with flag `--with-internal-tzcode`.

Users of ‘ix86’ CPUs with SSE2 support³ may prefer to use the C/C++/Fortran flags

```
-mfpmath=sse -msse2
```

to force floating-point to use the same instructions as ‘x86_64’ builds and hence not make use of 80-bit ‘extended precision’ intermediate results. (NB: this affects more than floating-point operations. For some OSES and versions of gcc it might be necessary to add `-mstackrealign`.)

To build a 64-bit version of R on ‘ppc64’ (also known as ‘powerpc64’) with gcc 4.1.1, Ei-ji Nakama used

```
CC="gcc -m64"
CXX="gxx -m64"
FC="gfortran -m64"
CFLAGS="-mminimal-toc -fno-optimize-sibling-calls -g -O2"
FFLAGS="-mminimal-toc -fno-optimize-sibling-calls -g -O2"
```

the additional flags being needed to resolve problems linking against `libnmath.a` and when linking R as a shared library.

The setting of the macro ‘SAFE_FFLAGS’ may need some help. It should not need additional flags on platforms other than ‘68000’ (not likely to be encountered) and ‘ix86’. For the latter, if the Fortran compiler is GNU (`gfortran` or possibly `g77`) the flags

```
-msse2 -mfpmath=sse
```

are added: earlier versions of R added `-ffloat-store` and this might still be needed if a ‘ix86’ CPU is encountered without SSE2 support. Note that it is a *replacement* for ‘FFLAGS’, so should include all the flags in that macro (except perhaps the optimization level).

Additional compilation flags can be specified for added safety/security checks. For example Fedora adds

```
-Werror=format-security -Wp,-D_FORTIFY_SOURCE=2 -Wp,-D_GLIBCXX_ASSERTIONS
-Fexceptions -fstack-protector-strong -fasynchronous-unwind-tables
-fstack-clash-protection -fcf-protection
```

to all the C, C++ and Fortran compiler flags (even though `_GLIBCXX_ASSERTIONS` is only for C++ in current GCC and `glibc` and none of these are documented for `gfortran`). Use of `_GLIBCXX_ASSERTIONS` will link `abort` and `printf` into almost all C++ code, and R CMD check `--as-cran` will warn.

C.2.1 Clang

R has been built with Linux ‘ix86’ and ‘x86_64’ C and C++ compilers (<http://clang.llvm.org>) based on the Clang front-ends, invoked by `CC=clang CXX=clang++`, together with `gfortran`. These take very similar options to the corresponding GCC compilers.

This has to be used in conjunction with a Fortran compiler: the `configure` code will remove `-lgcc` from `FLIBS`, which is needed for some versions of `gfortran`.

The current out-of-the-box default for `clang++` is to use the C++ runtime from the installed `g++`. Using the runtime from the `libcxx` project (<http://libcxx.llvm.org/>, Fedora RPM `libcxx-devel`) via `-stdlib=libcxx` has also been tested.

³ Likely all since 2005, including Pentium 4 and all ‘x86_64’ CPUs with 32-bit compilers.

Recent versions have (optional when built) OpenMP support.⁴

There is a project called `flang` (<https://github.com/flang-compiler/flang>) to develop a Fortran compiler similar to `clang` but based on the Portland Group's front end. This needs something like

```
FC=/usr/local/flang/bin/flang
LDFLAGS="-L/usr/local/flang/lib -L/usr/local/lib64"
```

Note that `flang` accepts all the flags which `clang` does (the driver is a modified version of `clang`, and `flang` is a symbolic link to `clang`), but does not implement all of them for Fortran compilation: it also accepts most PGI-style flags such as `-mp` for OpenMP. It currently produces few diagnostics even with `-Wall -pedantic`.

`flang`'s Fortran runtime is compiled against OpenMP and it seems this conflicts with using OpenMP in R. So it may be necessary to disable the latter by configuring R using `--disable-openmp`.

It is not clear what architectures `flang` intends to support: our experiments were done on 'x86_64'. At the time of writing binary 'releases' were available for that platform (called by them 'x86') and 'ppc64le'.

C.2.2 Intel compilers

Intel compilers have been used under 'ix86' and 'x86_64' Linux. Brian Ripley used version 9.0 of the compilers for 'x86_64' on Fedora Core 5 with

```
CC=icc
CFLAGS="-g -O3 -wd188 -ip -mp"
FC=ifort
FLAGS="-g -O3 -mp"
CXX=icpc
CXXFLAGS="-g -O3 -mp"
ICC_LIBS=/opt/compilers/intel/cce/9.1.039/lib
IFC_LIBS=/opt/compilers/intel/fce/9.1.033/lib
LDFLAGS="-L$ICC_LIBS -L$IFC_LIBS -L/usr/local/lib64"
SHLIB_CXXLD=icpc
```

It may be necessary to use `CC="icc -std=c99"` or `CC="icc -c99"` for C99-compliance. The flag `-wd188` suppresses a large number of warnings about the enumeration type 'Rboolean'. Because the Intel C compiler sets '`__GNUC__`' without complete emulation of `gcc`, we suggest adding `CPPFLAGS=-no-gcc`.

To maintain correct IEC 60559 arithmetic you most likely need add flags to `CFLAGS`, `FFLAGS` and `CXXFLAGS` such as `-mp` (shown above) or `-fp-model precise -fp-model source`, depending on the compiler version.

Others have reported success with versions 10.x and 11.x. Björn-Helge Mevik reported success with version 2015.3 of the compilers, using (for a SandyBridge CPU on Centos 6.x)

```
fast="-fp-model precise -ip -O3 -opt-mem-layout-trans=3 -xHost -mavx"
CC=icc
CFLAGS="$fast -wd188"
FC=ifort
```

⁴ This also needs the OpenMP runtime which has sometimes been distributed separately.

```
FFLAGS="$fast"
CXX=icpc
CXXFLAGS="$fast"
```

It is possible that 32-builds need to force the use of SSE2 instructions in `SAFE_FFLAGS`, e.g. by

```
SAFE_FFLAGS=-axsse2
```

C.3 macOS

(‘macOS’ was known as ‘OS X’ from 2012–2016.)

The instructions here are for Intel 64-bit (‘x86_64’) builds on 10.13–10.15 (High Sierra, Mojave and Catalina)⁵. R can in principle be built for 10.6–10.12, although this is little tested and it may be necessary to install later versions of software such as `libcurl` and an older version⁶ of `gfortran`.

C.3.1 Prerequisites

The following are essential to build R:

- Apple’s ‘Command Line Tools’: these can be (re-)installed by running `xcode-select --install` in a terminal. (If you have a fresh OS installation, running e.g. `make` in a terminal will offer the installation of the command-line tools. If you have installed Xcode, this provides the command-line tools. The tools may need to be reinstalled when macOS is upgraded, as upgrading may partially remove them.)

The Command Line Tools provide C and C++ compilers derived from LLVM’s `clang`.

- A Fortran compiler. An installer for `gfortran` is available at <https://github.com/fxcoudert/gfortran-for-macOS/releases/download/8.2/gfortran-8.2-Mojave.dmg>⁷. This installs into `/usr/local/gfortran`.
- Binary components `pcre2` and `xz` (for `liblzma`) from <https://mac.R-project.org/libs-4/>. For example

```
curl -OL https://mac.r-project.org/libs-4/pcre2-10.34-darwin.17-x86_64.tar.gz
sudo tar -xvzf pcre2-10.34-darwin.17-x86_64.tar.gz -C /
curl -OL https://mac.r-project.org/libs-4/xz-5.2.4-darwin.17-x86_64.tar.gz
sudo tar -xvzf xz-5.2.4-darwin.17-x86_64.tar.gz -C /
```

(Messages like ‘`usr/local/: Can’t restore time`’ should be ignored.)

and desirable

- GNU `readline` from <https://mac.R-project.org/libs-4/>,⁸. If `readline` is not present, the emulation in Apple’s version of `libedit` (aka `editline`) will be used: if you wish to avoid that, configure with `--without-readline`.

⁵ Details for the upcoming macOS 11 ‘Big Sur’ remain to be confirmed: no Fortran compiler was available for ‘Apple Silicon’ at the time of writing although there are ports under development.

⁶ For example, from <https://github.com/fxcoudert/gfortran-for-macOS/releases>.

⁷ This is said to be for Mojave and later but works well enough on High Sierra.

⁸ For licence reasons this is version 5.2 of `readline`: for those who want a more recent version it is straightforward to compile it from its sources.

- Components `jpeg`, `libpng`, `pkgconfig`, `tiff` and `pkgconfig` from <https://mac.R-project.org/libs-4/>, for the full range of bitmapped graphics devices.
- An X sub-system unless configuring using `--without-x`: see <https://xquartz.macosforge.org/>. R's `configure` script can be told to look for X11 in XQuartz's main location of `/opt/X11`, e.g. by

```
--x-includes=/opt/X11/include --x-libraries=/opt/X11/lib
```

although linked versions under `/usr/X11` will be found. (It seems that for recent versions of macOS/XQuartz `/usr/X11` may exist but be a link to a non-existent directory: reinstalling XQuartz might help.)

- An Objective-C compiler, as provided by `clang` in the Command Line Tools: this is needed for the `quartz()` graphics device.

Use `--without-aqua` if you want a standard Unix-alike build: apart from disabling `quartz()` and the ability to use the build with `R.APP`, it also changes the default location of the personal library (see `?libPaths`).

- Support for `cairo` (without `Pango`) will be enabled if `pkg-config`⁹ and XQuartz are available. Make sure the XQuartz's `pkg-config` files are found first on the configuration path: for example by setting

```
export PKG_CONFIG_PATH=/opt/X11/lib/pkgconfig:/usr/local/lib/pkgconfig:/usr/lib/pkgconfig
```

or appending that variable to the `configure` command (as below).

- A TeX installation. See Section C.3.3 [Other libraries], page 68.
- `texi2any` from a 'texinfo' distribution, which requires `perl` (currently version 5.18 is a default part of macOS but it has been announced that it may not be in future). A version of `texi2any` has been included in the binary distribution of R and there is a `texinfo` component at <https://mac.R-project.org/libs-4/>.

For C/C++ compilers supporting OpenMP, see the subsection below.

To build R itself from the sources with the compilers in the Command Line Tools (or Xcode) and `gfortran` from the installer mentioned above, use a file `config.site` containing

```
CC=clang
OBJC=$CC
FC=/usr/local/gfortran/bin/gfortran
CXX=clang++
```

and configure by something like

```
./configure -C \
--enable-R-shlib --enable-memory-profiling \
--x-includes=/opt/X11/include --x-libraries=/opt/X11/lib \
--with-tcl-config=/usr/local/lib/tclConfig.sh \
--with-tk-config=/usr/local/lib/tkConfig.sh \
PKG_CONFIG_PATH=/opt/X11/lib/pkgconfig:/usr/local/lib/pkgconfig:/usr/lib/pkgconfig
```

(See below for Tcl/Tk, including other options.)

To install packages using compiled code one needs the Command Line Tools and appropriate compilers, e.g. Fortran and the C/C++ compilers from those tools. Some packages have further requirements such as `pkg-config`.

⁹ available from <https://mac.R-project.org/libs-4/> as `pkgconfig`.

A subversion client, `svn` version 1.10, used to be part of macOS¹⁰. A more recent client can be obtained from <https://mac.r-project.org/tools/>, for example by

```
curl -OL https://mac.r-project.org/tools/subversion-1.14.0-darwin15.6.tar.gz
tar xf subversion-1.14.0-darwin15.6.tar.gz
sudo cp subversion-1.14.0-darwin15.6/svn /usr/local/bin
```

For some distributions of Command Line Tools version 12, `-Wno-implicit-function-declaration` needs to be included in `CFLAGS` to configure R 4.0.2 and earlier as well as some external software.

C.3.1.1 Note for Catalina users

The default security settings for Catalina can make it difficult to install Apple packages built after 2019-06-01 which have not been ‘notarized’¹¹ by Apple. And not just packages, as this has been seen for executables contained in tarballs/zipfiles (for example, for `pandoc`).

There are workarounds, including in some cases not installing the latest version. Usually one can use ‘Open With’ (Control/right/two-finger-click in Finder), then select ‘Installer’ and ‘Open’ if you get a further warning message.

This applies also to some ‘nightly builds’ from <https://mac.R-project.org/>.

If you run into problems with ‘quarantine’ for tarballs downloaded in a browser, consider using `curl -OL` to download (as illustrated above) or `xattr -c` to remove extended attributes.

C.3.2 Other C/C++ compilers

Other distributions of `clang` are available from <https://releases.llvm.org/>. In particular, these include support for OpenMP which Apple builds of `clang` do not, and may contain more bugfixes and new features (such as better diagnostics) than that in Command Line Tools/Xcode.

Suppose one of these distributions is installed under `/usr/local/clang`. Use a file `config.site` containing

```
CC=/usr/local/clang/bin/clang
OBJC=$CC
FC=/usr/local/gfortran/bin/gfortran
CXX=/usr/local/clang/bin/clang++
LDFLAGS="-L/usr/local/clang/lib -L/usr/local/lib"
R_LD_LIBRARY_PATH=/usr/local/clang/lib:/usr/local/lib
```

The care to specify library paths is to ensure that the OpenMP runtime library, here `/usr/local/clang/lib/libomp.dylib`, is found when needed. If this works, you should see the line

```
checking whether OpenMP SIMD reduction is supported... yes
```

in the `configure` output. Also, ‘`R_LD_LIBRARY_PATH`’ needs to be set to find the latest version of the C++ run-time libraries rather than the system ones.

¹⁰ but as a stub invoking an executable provided by the Command Line Tools and the latter is not in versions 11.5 and later of the Command Line Tools.

¹¹ See https://developer.apple.com/documentation/xcode/notarizing_macos_software_before_distribution.

For Catalina, the system paths need to be specified *via*

```
CC="/usr/local/clang/bin/clang -isysroot
/Library/Developer/CommandLineTools/SDKs/MacOSX.sdk"
CXX="/usr/local/clang/bin/clang++ -isysroot
/Library/Developer/CommandLineTools/SDKs/MacOSX.sdk"
```

(Long lines split for the manual, only.) Should the location of the SDK change, it can be found by running `xcrun -show-sdk-path`.

The approach in the previous paragraph also works for Mojave, but for versions 10.x of the Command Line Tools an alternative is to install the headers to `/usr/include`: from a Terminal run

```
sudo installer -pkg \
/Library/Developer/CommandLineTools/Packages/macOS_SDK_headers_for_macOS_10.14.pkg \
-target /
```

(This will need to be re-run if the OS decides to update the Command Line Tools.)

It is usually possible to build R with GCC (built from the sources, from a `gfortran` distribution, from Homebrew, ...). When last tested¹² it was not possible to use `gcc` to build the `quartz()` device, so `configure --without-aqua` may be required.

C.3.3 Other libraries

Pre-compiled versions of many of the Section A.2 [Useful libraries and programs], page 43, are available from <https://mac.R-project.org/libs-4/> or <https://mac.R-project.org/libs/>.

The `Accelerate` library¹³ can be used *via* the configuration options

```
--with-blas="-framework Accelerate"
```

to provide potentially higher-performance versions of the BLAS and LAPACK routines.¹⁴ This also includes a full LAPACK which can be used *via* `--with-lapack`: however, the version of LAPACK it contains is often seriously old.

Threading in `Accelerate` is controlled by ‘Grand Central Dispatch’ and is said not to need user control. Test `nls.R` in package `stats` usually fails with the `Accelerate` BLAS.

Looking at the top of `/Library/Frameworks/R.framework/Resources/etc/Makeconf` will show the compilers and configuration options used for the CRAN binary package for R: at the time of writing the non-default options

```
--enable-memory-profiling --enable-R-framework
--x-libraries=/opt/X11/lib --x-includes=/opt/X11/include
```

were used. (`--enable-R-framework` implies `--enable-R-shlib`.)

Configure option `--with-internal-tzcode` is the default on macOS, as the system implementation of time zones does not work correctly for times before 1902 or after 2037 (despite using a 64-bit `time_t`).

¹² with `gcc` 10.2 on High Sierra.

¹³ <https://developer.apple.com/documentation/accelerate>.

¹⁴ It was reported that for some non-Apple toolchains `CPPFLAGS` needed to contain `-D__ACCELERATE__`: not needed for `clang` from releases.llvm.org though.

The \TeX implementation used by the developers is MacTeX¹⁵ (<https://www.tug.org/mactex/>): the full installation is about 7GB, but a smaller version (‘Basic TeX’) is available at <https://www.tug.org/mactex/morepackages.html> to which you will need to add some packages to build R, e.g. for the 2020 version we needed to add¹⁶ **cm-super**, **helvetic**, **inconsolata** and **texinfo** which brought this to about 380MB. ‘TeX Live Utility’ (available *via* the MacTeX front page) provides a graphical means to manage \TeX packages. This is documented to require High Sierra or later: for earlier versions see the instructions on the MacTeX front page.

Checking packages thoroughly requires **ghostscript** (part of the full MacTeX distribution or separately from <https://www.tug.org/mactex/morepackages.html>) and **qpdf** (from <https://mac.R-project.org/libs-4/>, a version of which is in the `bin` directory of a binary installation of R, usually `/Library/Frameworks/R.framework/Resources/bin/qpdf`).

One macOS quirk is that the default path has `/usr/local/bin` after `/usr/bin`, contrary to common practice on Unix-alikes. This means that if you install tools from the sources they will by default be installed under `/usr/local` and not supersede the system versions.

C.3.4 Tcl/Tk headers and libraries

If you plan to use the `tcltk` package for R, you need to install a distribution of Tcl/Tk. There are two alternatives. If you use R.APP you will want to use X11-based Tcl/Tk (as used on other Unix-alikes), which is installed under `/usr/local/lib` as part of the CRAN binary for R.¹⁷ This may need `configure` options

```
--with-tcltk=/usr/local/lib
or
--with-tcl-config=/usr/local/lib/tclConfig.sh
--with-tk-config=/usr/local/lib/tkConfig.sh
```

Note that this requires a matching XQuartz installation.

There is also a native (‘Aqua’) version of Tcl/Tk which produces widgets in the native macOS style: this will not work with R.APP because of conflicts over the macOS menu, but for those only using command-line R this provides a much more intuitive interface to Tk for experienced Mac users. Earlier versions of macOS came with an Aqua Tcl/Tk distribution but these were often not at all recent versions of Tcl/Tk (8.5.9 in High Sierra, which is not even the latest patched version in that series). (Catalina included a partial distribution which lacks the files needed to compile against Tcl/Tk.) It is better to install Tcl/Tk 8.6.x from the sources or a binary distribution from <https://www.activestate.com/activetcl/downloads>. For the latter, configure R with

```
--with-tcl-config=/Library/Frameworks/Tcl.framework/tclConfig.sh
--with-tk-config=/Library/Frameworks/Tk.framework/tkConfig.sh
```

If you need to find out which distribution of Tk is in use at run time, use

```
library(tcltk)
```

¹⁵ An essentially equivalent \TeX installation can be obtained by the Unix TeX Live installation scripts: MacTeX also contains GhostScript.

¹⁶ E.g. *via* `tlmgr install cm-super helvetic inconsolata texinfo`.

¹⁷ Just that component can be selected from the installer for R: at the ‘Installation Type’ screen select ‘Customise’ and then just the ‘Tcl/Tk 8.6.6’ component.

```
tclvalue(.Tcl("tk windowingsystem")) # "x11" or "aqua"
```

C.3.5 Java

macOS no longer comes with an installed Java runtime (JRE) and a macOS upgrade may remove one if already installed: it is intended to be installed at first use. Check if a JRE is installed by running `java -version` in a **Terminal** window: if Java is not installed¹⁸ this should prompt you to install it. You can also install directly a recent Java from Oracle (currently from <https://www.oracle.com/java/technologies/javase-downloads.html>). Builds of OpenJDK may also be available, e.g. from <https://adoptopenjdk.net> or <http://jdk.java.net/>. We recommend you install a version with long-term support, e.g. 8 or 11 (but not 9, 10, 12–16 which have/had a 6-month lifetime).

Binary distributions of R are built against a specific version (e.g. 11.0.6) of Java so `sudo R CMD javareconf` will likely be needed before using Java-using packages.

To see what compatible versions of Java are currently installed, run `/usr/libexec/java_home -V -a x86_64`. If needed, set the environment variable `JAVA_HOME` to choose between these, both when R is built from the sources and when R CMD `javareconf` is run.

Configuring and building R both looks for a JRE and for support for compiling JNI programs (used by packages **rJava** (<https://CRAN.R-project.org/package=rJava>) and **JavaGD** (<https://CRAN.R-project.org/package=JavaGD>)); the latter requires a JDK (Java SDK) and not just a JRE¹⁹.

The build process tries to fathom out what JRE/JDK to use, but it may need some help, e.g. by setting `JAVA_HOME`. A Oracle JDK can be specified explicitly by something like

```
JAVA_HOME=/Library/Java/JavaVirtualMachines/jdk-11.jdk/Contents/Home
JAVA_CPPFLAGS="-I/${JAVA_HOME}/include -I/${JAVA_HOME}/include/darwin"
JAVA_LD_LIBRARY_PATH="${JAVA_HOME}/lib/server"
JAVA_LIBS="-L/${JAVA_HOME}/lib/server -ljvm"
```

or

```
JAVA_HOME=/Library/Java/JavaVirtualMachines/jdk1.8.0_191.jdk/Contents/Home
JAVA_CPPFLAGS="-I/${JAVA_HOME}/include -I/${JAVA_HOME}/include/darwin"
JAVA_LD_LIBRARY_PATH="${JAVA_HOME}/jre/lib/server"
JAVA_LIBS="-L/${JAVA_HOME}/jre/lib/server -ljvm"
```

in `config.site`.

To use the builds of OpenJDK from <https://adoptopenjdk.net> (Apple package) to

```
JAVA_HOME=/Library/Java/JavaVirtualMachines/adoptopenjdk-11.jdk/Contents/Home
```

Note that it is necessary to set the environment variable `NOAWT` to 1 to install many of the Java-using packages.

¹⁸ In the unlikely event that the version reported does not start with 1.8.0, 11 or higher you need to update your Java.

¹⁹ As from Java 11, there is no separate client JRE distribution from Oracle, but there may be from OpenJDK.

C.3.6 Frameworks

The CRAN build of R is installed as a framework, which is selected by the option

```
./configure --enable-R-framework
```

(This is intended to be used with an Apple toolchain: others may not support frameworks correctly but those from releases.llvm.org do.)

It is only needed if you want to build R for use with the R.APP console, and implies `--enable-R-shlib` to build R as a dynamic library. This option configures R to be built and installed as a framework called `R.framework`. The default installation path for `R.framework` is `/Library/Frameworks` but this can be changed at configure time by specifying the flag `--enable-R-framework[=DIR]` (or `--prefix`) or at install time *via*

```
make prefix=/where/you/want/R.framework/to/go install
```

Note that installation as a framework is non-standard (especially to a non-standard location) and Unix utilities may not support it (e.g. the `pkg-config` file `libR.pc` will be put somewhere unknown to `pkg-config`).

C.3.7 Building R.app

Note that building the R.APP GUI console is a separate project, using Xcode. Before compiling R.APP make sure the current version of R is installed in `/Library/Frameworks/R.framework` and working at the command-line (this can be a binary install).

The current sources can be checked out by

```
svn co https://svn.r-project.org/R-packages/trunk/Mac-GUI
```

and built by loading the `R.xcodeproj` project (select the `R` target and a suitable configuration), or from the command-line by e.g.

```
xcodebuild -target R -configuration Release
```

See also the `INSTALL` file in the checkout or directly at <https://svn.r-project.org/R-packages/trunk/Mac-GUI/INSTALL>.

R.APP does not need to be installed in any specific way. Building R.APP results in the R.APP bundle which appears as one R icon. This application bundle can be run anywhere and it is customary to place it in the `/Applications` folder.

C.3.8 Building binary packages

CRAN macOS binary packages are distributed as tarballs with suffix `.tgz` to distinguish them from source tarballs. One can `tar` an existing installed package, or use `R CMD INSTALL --build`.

However, there are some important details.

- Current CRAN macOS distributions are targeted at High Sierra so it is wise to ensure that the compilers used generate code that will run on High Sierra or later. With the recommended compilers we can use

```
CC="clang -mmacosx-version-min=10.13"
CXX="clang++ -mmacosx-version-min=10.13"
FC="/usr/local/gfortran/bin/gfortran -mmacosx-version-min=10.13"
```

- Using the flag `-Werror=partial-availability` can help trigger compilation errors on functionality not in High Sierra.
- Check that any compiled code is not dynamically linked to libraries only on your machine, for example by using `otool -L` or `objdump -macho -dylibs-used`. This can include C++ and Fortran run-time libraries under `/usr/local`: one can use `install_name_tool` to point these at system versions or those shipped with R, for example

```
install_name_tool -change /usr/local/clang/lib/libc++.1.dylib \
  /usr/lib/libc++.1.dylib \
  pkg.so
```

```
install_name_tool -change /usr/local/gfortran/lib/libgfortran.5.dylib \
  /Library/Frameworks/R.framework/Versions/4.0/Resources/lib/libgfortran.5.dylib \
  pkg.so
```

(where the details depend on the compilers and CRAN macOS R release).

- For C++ code there is the possibility that calls will be generated to entry points not in the system `/usr/lib/libc++.1.dylib`. The previous step allows this to be tested against the system library on the build OS, but not against earlier ones. It may be possible to circumvent that by static linking to `libc++.a` and `libc++abi.a` by something like

```
SHLIB_CXXLD = /usr/local/clang/bin/clang
PKG_LIBS = /usr/local/clang/lib/libc++.a /usr/local/clang/lib/libc++abi.a
```

in `src/Makevars`. It would also be possible to static link the Fortran runtime libraries `/usr/local/gfortran/lib/libgfortran.a` and `/usr/local/gfortran/lib/libquadmath.a` should the Fortran compiler have later versions (but `gfortran` 8, 9 and 10 all have version 5.0.0).

The CRAN binary packages are built with the Apple compiler on High Sierra, which avoids the first two and any issues with C++ libraries.

C.4 Solaris

R has been built successfully on Solaris 10 using the (zero cost) Oracle Developer Studio²⁰ compilers: there has also been success with `gcc/gfortran`. (Recent Sun machines are AMD Opterons or Intel Xeons (`'amd64'`) rather than `'x86'`, but 32-bit `'x86'` executables are the default.) How these compilers identify themselves is slightly confusing: commands `CC -V` with Developer Studio 12.5 and 12.6 report as versions 5.14 and 5.15. We will only consider Developer Studio versions 12.5 (May 2016) and 12.6 (July 2017).

There have been few reports on Solaris 11, with no known extra issues.

Solaris was last tested on Sparc machines in June 2017.

The Solaris versions of several of the tools needed to build R (e.g. `make`, `ar` and `ld`) are in `/usr/ccs/bin`, so if using those tools ensure this is in your path. A version of the preferred GNU `tar` is (if installed) in `/usr/sfw/bin`. It may be necessary to avoid the tools in `/usr/ucb`: POSIX-compliant versions of some tools can be found in `/usr/xpg4/bin` and `/usr/xpg6/bin`.

²⁰ Oracle Solaris Studio prior to 2016, and previously Sun Studio.

A large selection of Open Source software can be installed from <https://www.opencsw.org>, by default installed under `/opt/csw`. Solaris 10 ships with `bzlib` version 1.0.6 (sufficient) but `zlib` version 1.2.3 (too old): OpenCSW has 1.2.8. (Note from 2019: updating of OpenCSW has slowed or stopped.)

At least when compiling with Oracle compilers, Solaris uses far more stack space than other platforms. This makes it desirable to build PCRE1 (and PCRE2 < 10.30) with the option `--disable-stack-for-recursion`: the OpenCSW distribution was at the time of writing.

The Oracle compilers are unusual in not including `/usr/local/include` in the default include search path: R's default `CPPFLAGS=-I/usr/local/include` remedies this. If you rely on OpenCSW software you may need `CPPFLAGS=-I/opt/csw/include` (or both).

You will need GNU `libiconv` and `readline`: the Solaris version of `iconv` is not sufficiently powerful.

The native `make` suffices to build R but a number of packages require GNU `make` (some without declaring it as 'SystemRequirements' in the DESCRIPTION file).

The support for the C99 `long double` type on Sparc hardware uses quad-precision arithmetic, and this is usually slow because it is done by software emulation. On such systems the `configure` option `--disable-long-double` can be used for faster but less accurate computations.

The Solaris time-zone conversion services seem to be unreliable pre-1916 in Europe (when daylight-savings time was first introduced): most often reporting in the non-existent DST variant. Using `configure` option `--with-internal-tzcode` is recommended, and required if you find time-zone abbreviations being given odd values (as has been seen on 64-bit builds without it).

When using the Oracle compilers do *not* specify `-fast`, as this disables IEEE arithmetic and `make check` will fail.

A little juggling of paths was needed to ensure GNU `libiconv` (in `/usr/local`) was used rather than the Solaris `iconv`:

```
CC="cc -xc99"
CFLAGS="-O -xlibmieee"
FC=f95
FFLAGS=-O
CXX=CC
CXXSTD="-std=c++11 -library=stdcpp,CrunG3"
CXX11STD="-std=c++11 -library=stdcpp,CrunG3"
CXX14STD="-std=c++14 -library=stdcpp,CrunG3"
CXXFLAGS=-O
R_LD_LIBRARY_PATH="/opt/developerstudio12.6/lib:/usr/local/lib:/opt/csw/lib"
```

The Oracle compilers do not by default conform to the C99 standard (appendix F 8.9) on the return values of functions such as `log`: use `-xlibmieee` to ensure this.

A peculiarity of some versions of the Fortran compiler has been that when asked to link a shared object they did not link against all the Fortran 9x runtime libraries, hence

```
FCLIBS_XTRA="-lfsu /opt/developerstudio12.6/lib/libfui.so.2"
```

has been needed.

Using `-xlibmil` in `CFLAGS` or `FFLAGS` allows more system mathematical functions to be inlined.

On ‘x86’ you will get marginally higher performance *via*

```
CFLAGS="-x05 -xlibmieee -xlibmil -nofstore -xtarget=native"
FFLAGS="-x05 -xlibmil -nofstore -xtarget=native"
CXXFLAGS="-x05 -xlibmil -nofstore -xtarget=native"
SAFE_FFLAGS="-O -xlibmil -fstore -xtarget=native"
```

but the use of `-nofstore` can be less numerically stable, and some packages have in the past failed to compile at optimization level 5.

The Oracle compilers provide several implementations of the C++ standards which select both the set of headers and a C++ runtime library. One of those is selected by the `-library` flag, which as it is needed for both compiling and linking is best specified as part of the compiler or standard. Current R expects a C++11 compiler, for which the choice given above is the only possibility. Although version 12.5 accepted the flag `-std=c++14`, it did not pass `configure`’s conformance tests: version 12.6 does.

The performance library `sunperf` is available for use with the Oracle compilers. If selected as a BLAS, it must also be selected as LAPACK *via*

```
./configure --with-blas='-library=sunperf' --with-lapack
```

This has often given test failures in the past, in several different places.²¹

Parsing very complex R expressions needs a lot of stack space when the Oracle compilers are used: several packages require the stack increased to at least 20MB.

Some people have reported that the Solaris `libintl` needs to be avoided, for example by using `--disable-nls` or `--with-included-gettext` or using `libintl` from OpenCSW. (On the other hand, there have been many successful installs which automatically detected `libintl` from OpenCSW or selected the included `gettext`.)

It has been reported that some Solaris installations need

```
INTERNET_LIBS="-lsocket -lnsl"
```

on the `configure` command line or in file `config.site`; however, there have been many successful installs without this.

C.4.1 64-bit builds

On both ‘x86’ and ‘Sparc’ platforms the compilers default to 32-bit code.

For a 64-bit target add `-m64` to the compiler macros and use something like `LDFLAGS=-L/usr/local/lib/amd64` or `LDFLAGS=-L/usr/local/lib/sparcv9` as appropriate (and other 64-bit library directories if used, e.g. `-L/opt/csw/lib/amd64`). It will also be necessary to point `pkg-config` at the 64-bit directories, e.g. by something like

```
PKG_CONFIG_PATH= /usr/local/lib/amd64/pkgconfig:/opt/csw/lib/64/pkgconfig:/usr/lib/64/
```

and to specify a 64-bit Java VM by e.g.

```
JAVA_CPPFLAGS="-I${JAVA_HOME}/../include -I${JAVA_HOME}/../include/solaris"
JAVA_LD_LIBRARY_PATH=${JAVA_HOME}/lib/amd64/server
JAVA_LIBS="-L${JAVA_HOME}/lib/amd64/server \
-R${JAVA_HOME}/lib/amd64/server -ljvm"
```

²¹ When last checked it failed in `tests/reg-BLAS.R`, and on some builds, including for ‘amd64’, it failed in `example(eigen)`.

C.4.2 Using gcc

If using `gcc`, ensure that the compiler was compiled for the version of Solaris in use. (This can be ascertained from `gcc -v`.) `gcc` makes modified versions of some header files, and several reports of problems were due to using `gcc` compiled on one version of Solaris on a later version. Note that this can even apply to OS patches: some 2016 patches to Solaris 10 changed its C header files in way incompatible²² with the modified versions included with OpenCSW's binary distribution.

The notes here are for `gcc` set up to use the Solaris linker: it can also be set up to use GNU `ld`, but that has not been tested. The tests were for compilers from the OpenCSW repository: Solaris systems often come with much older compilers installed under `/usr/sfw/bin`. One of `-m32` or `-m64` will be the default and could be omitted, but it is not easy to find out which. (For OpenCSW it is `-m32`.)

Compilation for an 'x86' target with `gcc` 5.2.0 needed

```
CC="gcc -m32"
CPPFLAGS="-I/opt/csw/include -I/usr/local/include"
FC="gfortran -m32"
CXX="g++ -m32"
LDFLAGS="-L/opt/csw/lib -L/usr/local/lib"
```

For an 'amd64' target we used

```
CC="gcc -m64"
CPPFLAGS="-I/opt/csw/include -I/usr/local/include"
FC="gfortran -m64"
CXX="g++ -m64"
LDFLAGS="-L/opt/csw/lib/amd64 -L/usr/local/lib/amd64"
```

Note that paths such as `/opt/csw/lib`, `/usr/local/lib/amd64` and `/opt/csw/lib/amd64` may need to be in the `LD_LIBRARY_PATH` during configuration.

C.5 FreeBSD

There have been few recent reports on FreeBSD: there is a 'port' at <https://www.freebsd.org/ports/math.html>. Recent versions of FreeBSD use Clang and the `libc++` C++ headers and runtime, but the 'port' is configured to use GCC.

Use of ICU for collation and the `configure` option `--with-internal-tzcode` are desirable workarounds.

C.6 OpenBSD

Ingo Feinerer installed R version 3.2.2 on OpenBSD 5.8 arch 'amd64' (their name for 'x86_64'). Details of the build (and patches applied) are at <http://cvsweb.openbsd.org/cgi-bin/cvsweb/ports/math/R/>. (Downgrading the `zlib` requirement to 1.2.3 is against the advice of the R developers.)

²² In particular, header `cmath` in C++11 mode includes `math.h` and `iso/math_c99.h` and `gcc` had 'fixed' an earlier version of the latter.

C.7 Cygwin

The 32-bit version never worked well enough to pass R's `make check`, and residual support from earlier experiments was removed in R 3.3.0.

The 64-bit version was never supported.

C.8 New platforms

There are a number of sources of problems when installing R on a new hardware/OS platform. These include

Floating Point Arithmetic: R requires arithmetic compliant with IEC 60559, also known as IEEE 754. This mandates the use of plus and minus infinity and NaN (not a number) as well as specific details of rounding. Although almost all current FPUs can support this, selecting such support can be a pain. The problem is that there is no agreement on how to set the signalling behaviour; Sun/Sparc, SGI/IRIX and 'ix86' Linux require no special action, FreeBSD requires a call to (the macro) `fpsetmask(0)` and OSF1 required that computation be done with a `-ieee_with_inexact` flag etc. On a new platform you must find out the magic recipe and add some code to make it work. This can often be done via the file `config.site` which resides in the top level directory.

Beware of using high levels of optimization, at least initially. On many compilers these reduce the degree of compliance to the IEEE model. For example, using `-fast` on the Oracle compilers has caused R's NaN to be set incorrectly, and gcc's `-ffast-math` and clang's `-Ofast` have given incorrect results.

Shared Objects: There seems to be very little agreement across platforms on what needs to be done to build shared objects. there are many different combinations of flags for the compilers and loaders. GNU libtool cannot be used (yet), as it currently does not fully support Fortran: one would need a shell wrapper for this). The technique we use is to first interrogate the X window system about what it does (using `xmkmf`), and then override this in situations where we know better (for tools from the GNU Compiler Collection and/or platforms we know about). This typically works, but you may have to manually override the results. Scanning the manual entries for `cc` and `ld` usually reveals the correct incantation. Once you know the recipe you can modify the file `config.site` (following the instructions therein) so that the build will use these options.

It seems that gcc 3.4.x and later on 'ix86' Linux defeat attempts by the LAPACK code to avoid computations entirely in extended-precision registers, so file `src/modules/lapack/dlamc.f` may need to be compiled without optimization or with additional flags. Set the configure variable `SAFE_FFLAGS` to the flags to be used for this file.

If you do manage to get R running on a new platform please let us know about it so we can modify the configuration procedures to include that platform.

If you are having trouble getting R to work on your platform please feel free to use the 'R-devel' mailing list to ask questions. We have had a fair amount of practice at porting R to new platforms ...

Function and variable index

C

`configure` 3, 4, 7, 55, 57

I

`install.packages` 24

M

`make` 57

R

`R_HOME` 3

`remove.packages` 30

U

`update.packages` 30

Concept index

B

BLAS library 47, 68, 74

F

Fortran 57
FreeBSD 75

I

Installation 7
Installing under Unix-alikes 3
Installing under Windows 17
Internationalization 33

L

LAPACK library 52, 68, 74
Libraries 23
Libraries, managing 23
Libraries, site 23
Libraries, user 23
Linux 3, 61
Locale 33
Localization 33

M

macOS 3, 19, 65
Manuals 5
Manuals, installing 8

O

Obtaining R 1
OpenBSD 75

P

Packages 23
Packages, default 23
Packages, installing 24
Packages, removing 30
Packages, updating 30

R

Repositories 30

S

Site libraries 23
Solaris 72
Sources for R 1
Subversion 1, 43

U

User libraries 23

V

Vignettes 43

Environment variable index

B

BLAS_LIBS 47

C

CONFIG_SITE 55

D

DESTDIR 9, 39

J

JAVA_HOME 46

L

LANG 34

LANGUAGE 34, 35

LAPACK_LIBS 52

LC_ALL 34

LC_COLLATE 16

LC_MESSAGES 34

LD_LIBRARY_PATH 39, 57, 58, 75

P

PAPERSIZE 56

PATH 43, 57

R

R_ARCH 10

R_BROWSER 56

R_DEFAULT_PACKAGES 23

R_DISABLE_HTTPD 5

R_GSCMD 45

R_INSTALL_TAR 25

R_JAVA_LD_LIBRARY_PATH 46

R_LIBS 23

R_LIBS_SITE 23

R_LIBS_USER 23

R_PAPERSIZE 5, 22, 56

R_PDFVIEWER 56

R_RD4PDF 6, 56

R_SCRIPT_DEFAULT_PACKAGES 23

R_USER 22

T

TAR 42

TAR_OPTIONS 1

TEMP 22

TMP 22

TMPPDIR 3, 22, 24