

R Installation and Administration

Version 2.11.0 (2010-04-22)

R Development Core Team

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

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

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

Copyright © 2001–2010 R Development Core Team

ISBN 3-900051-09-7

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	4
2.4	Installation	5
2.5	Uninstallation	7
2.6	Sub-architectures	7
2.6.1	Multilib	8
2.7	Testing an Installation	8
3	Installing R under Windows	10
3.1	Building from source	10
3.1.1	Getting the tools	10
3.1.2	Getting the source files	10
3.1.3	Building the core files	11
3.1.4	Building the bitmap files	12
3.1.5	Checking the build	12
3.1.6	Building the manuals	12
3.1.7	Building the Inno Setup installer	12
3.1.8	Building the MSI installer	13
3.1.9	Cross-building on Linux	14
3.2	Testing an Installation	14
3.3	Building R for 64-bit Windows	14
4	Installing R under Mac OS X	16
4.1	Building from source on Mac OS X	16
5	Running R	17
6	Add-on packages	18
6.1	Default packages	18
6.2	Managing libraries	18
6.3	Installing packages	18
6.3.1	Windows	19
6.3.2	Mac OS X	20
6.3.3	Customizing package compilation	20
6.4	Updating packages	20
6.5	Removing packages	20
6.6	Setting up a package repository	21

7	Internationalization and Localization	22
7.1	Locales	22
7.1.1	Locales under Linux	22
7.1.2	Locales under Windows	22
7.1.3	Locales under Mac OS X	23
7.2	Localization of messages	23
8	Choosing between 32- and 64-bit builds	25
8.1	Windows	26
9	The standalone Rmath library	27
9.1	Unix	27
9.2	Windows	28
Appendix A Essential and useful other programs under Unix		30
A.1	Essential programs and libraries	30
A.2	Useful libraries and programs	31
A.2.1	Tcl/Tk	32
A.2.2	Java support	32
A.3	Linear algebra	33
A.3.1	BLAS	33
A.3.1.1	ATLAS	34
A.3.1.2	ACML	34
A.3.1.3	Goto BLAS	34
A.3.1.4	Intel MKL	34
A.3.1.5	Shared BLAS	35
A.3.2	LAPACK	36
A.3.3	Caveats	36
Appendix B Configuration on Unix		38
B.1	Configuration options	38
B.2	Internationalization support	38
B.3	Configuration variables	39
B.3.1	Setting paper size	39
B.3.2	Setting the browsers	39
B.3.3	Compilation flags	39
B.3.4	Making manuals	40
B.4	Setting the shell	40
B.5	Using make	40
B.6	Using FORTRAN	40
B.6.1	Using gfortran	41
B.7	Compile and load flags	41

Appendix C	Platform notes	43
C.1	X11 issues	43
C.2	Linux	44
C.2.1	Intel compilers	45
C.2.2	PGI compilers	45
C.2.3	Sun Studio compilers	46
C.3	FreeBSD	46
C.4	Mac OS X	47
C.4.1	64-bit builds	47
C.4.2	Snow Leopard	47
C.5	Solaris	47
C.5.1	Using gcc	49
C.6	AIX	49
C.7	Cygwin	51
C.8	New platforms	51
Appendix D	The Windows toolset	53
D.1	Perl	54
D.2	L ^A T _E X	54
D.3	The Inno Setup installer	54
D.4	The command line tools	54
D.5	The MinGW compilers	55
D.6	Useful additional programs	55
Function and variable index		56
Concept index		57
Environment variable index		58

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 <http://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 `make` (specifically GNU `make` 3.81) does 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.)

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 (2.x.0) version, ‘r-patched’ will refer to beta/release candidates of the upcoming release, the patched version of the current release being available only via Subversion.)

The tarballs are available from <ftp://ftp.stat.math.ethz.ch/pub/Software/R/>. Download either ‘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 <http://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/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*. 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 required, and that the SSL certificate for the Subversion server of the R project is

```
Certificate information:
- Hostname: svn.r-project.org
- Valid: from Jul 16 08:10:01 2004 GMT until Jul 14 08:10:01 2014 GMT
- Issuer: Department of Mathematics, ETH Zurich, Zurich, Switzerland, CH
- Fingerprint: c9:5d:eb:f9:f2:56:d1:04:ba:44:61:f8:64:6b:d9:33:3f:93:6e:ad
```

¹ e.g. GNU `tar` version 1.15 or later, or that from the ‘libarchive’ (as used on Mac OS 10.6) or ‘Heirloom Toolchest’ distributions.

(currently, there is no “trusted certificate”). You can accept this certificate permanently and will not be asked about it anymore.

Note that retrieving the sources by e.g. `wget -r` or `svn export` from that URL will not work: 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 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 ‘<http://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

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

2 Installing R under Unix-alikes

R will configure and build under a number of common Unix and Unix-alike platforms including ‘cpu-*-linux-gnu’ for the ‘alpha’, ‘arm’, ‘hppa’, ‘ix86’, ‘ia64’, ‘m68k’, ‘mips’, ‘mipsel’, ‘powerpc’, ‘s390’, ‘sparc’, and ‘x86_64’ CPUs, ‘powerpc-apple-darwin’, ‘i386-apple-darwin’, ‘x86_64-apple-darwin’, ‘i386-sun-solaris’, ‘sparc-sun-solaris’, ‘x86_64-*-freebsd’, and ‘powerpc-ibm-aix6*’ as well as perhaps (it is tested less frequently on these platforms) ‘i386-*-freebsd’, ‘i386-*-netbsd’ and ‘i386-*-openbsd’.

In addition, binary distributions are available for some common Linux distributions and for 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.

2.1 Simple compilation

First review the essential and useful tools and libraries in [Appendix A \[Essential and useful other programs under Unix\]](#), [page 30](#), and install those you want or need. Ensure that the environment variable `TMPDIR` is either unset (and ‘/tmp’ exists and can be written in and scripts can be executed from) or points to a valid temporary directory (one from which execution of scripts is allowed).

Choose a place 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 39](#).) Issue the following commands:

```
./configure
make
```

(See [Section B.5 \[Using make\]](#), [page 40](#) if your make is not called ‘make’.)

Then check the built system works correctly by

```
make check
```

Failures are not necessarily problems as they might be caused by missing functionality,¹ 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 17](#)).

More comprehensive testing can be done by

```
make check-devel
```

or

```
make check-all
```

see file ‘tests/README’.

If the command `configure` and `make` commands execute successfully, a shell-script front-end called ‘R’ will be created and copied to ‘*R_HOME*/bin’. You can 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 5](#). 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

¹ for example, if you configured R with ‘--without-recommended’.


```
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 the pathname of the build directory should not contain spaces.)

For those obtaining R *via* Subversion, one additional step is necessary:

```
make vignettes
```

which makes the **grid** vignettes (which are contained in the tarballs): it make take several minutes.

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

2.2 Help options

As from R 2.10.0, 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 continue to build HTML pages at 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,

‘**refman**’ Printed versions of the help pages for the commonly-used packages.

‘**fullrefman**’
Printed versions of all the help pages.

‘**R-FAQ**’ R FAQ

‘**R-intro**’ “An Introduction to R”.

‘**R-data**’ “R Data Import/Export”.

‘**R-admin**’ “R Installation and Administration”, this manual.

‘**R-exts**’ “Writing R Extensions”.

‘**R-lang**’ “The R Language Definition”.

To make these (except ‘**fullrefman**’), use

```
make dvi      to create DVI versions
make pdf      to create PDF versions
make info     to create info files (not ‘refman’).
```

You will not be able to build any of these unless you have `makeinfo` version 4.7 or later installed, and for DVI or 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 T_EX package in re-distributions).

The DVI versions can be previewed and printed using standard programs such as `xdvi` and `dvips`. 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 DVI and PDF versions will be created using the papersize 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 papersize, you should first delete the file `'doc/manual/version.texi'`.)

There are some issues with making the reference manual, and in particular with the PDF version `'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 L^AT_EX Computer Modern fonts. We have provided four alternatives:

- `times` (The default for PDF.) Using standard PostScript fonts. This works well both for on-screen viewing and for printing. The one disadvantage is that the Usage and Examples sections may come out rather wide.
- `lm` Using the *Latin Modern* fonts. These are not often installed as part of a T_EX distribution, but can be obtained from <http://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 <http://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 so readable on-screen as the other three options.
- `ae` (The default for DVI.) 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 variables `R_RD4PDF` and `R_RD4DVI`. (On Unix, these will be picked up at install time.) The default value for `R_RD4PDF` is `'times,hyper'`: omit `'hyper'` if you do not want hyperlinks, e.g. for printing. The default for `R_RD4DVI` is `'ae'`.

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
```

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
```

This causes `make install` to install the R executable 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`. You can install into another directory tree by using

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

at least with GNU `make` (but not e.g. Solaris 8's `make`).

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-2.11.0`.

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

```
make install rhome=/usr/local/lib64/R-2.11.0
```

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

If you have made R as a shared/dynamic 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.

```
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.

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

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

Once again, it is optional to specify *prefix*, *libdir* or *rhome* (the DVI and PDF manuals are installed under the R home directory). (`make install-info` needs Perl installed if there is no command `install-info` on the system.)

More precise control is possible. For info, the setting used is that of `infodir` (default `'prefix/info'`, set by configure option `'--infodir'`). The DVI and 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) should reflect the final destination, and `DESTDIR` should be used: see http://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 DVI, 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.

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 Solaris for different chips (in particular, 32- and 64-bit builds), 64- and 32- bit builds on `'x86_64'` Linux and different CPUs (e.g. `'ppc'`, `'i386'` and `'x86_64'`) under Mac OS ≥ 10.4 .

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 systems 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 63Mb, 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 pkg(s)
```

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/execR_ARCH'`. Normally all installed architectures will run on the platform so the architecture of `Rscript` does not matter.

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

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

to select a sub-architecture to check.

2.6.1 Multilib

On Linux, there is an alternative mechanism for mixing 32-bit and 64-bit libraries known as *multilib*. If a 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 chosen using the `setarch` command. For example, a 32-bit build may be chosen 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).

At present there is a potential problem with packages using Java, as the post-install for a ‘i386’ 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.

2.7 Testing an Installation

Full 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.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 the 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 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 the operating environment has been changed, for example by OS updates or by substituting the BLAS (see [Section A.3.1.5 \[Shared BLAS\]](#), page 35).

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

```
library("tools")
testInstalledBasic("both")
testInstalledPackages("base")
testInstalledPackages("recommended")
```

runs the basic tests and then all the tests on the standard and recommended packages. These tests can be run from anywhere: they write their results in the ‘**tests**’ folder of the R home directory. They run slightly fewer tests than the first approach: in particular they do not test Internet access.

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

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 Windows 2000 or later on ix86 CPUs (including AMD64/Intel64¹ chips and Windows x64).

As from R 2.11.0 there is also a 64-bit version for Windows x64 in the ‘bin/windows64’ directory: this is primarily intended for users who require large amounts of memory and beginning users of R should install the standard version even on a 64-bit Windows OS.

Your file system must allow long file names (as is likely except perhaps for some network-mounted systems).

Installation is *via* the installer ‘R-2.11.0-win32.exe’ (or ‘R-2.11.0-win64.exe’). Just double-click on the icon and follow the instructions. You can uninstall R from the Control Panel or from the (optional) R program group on the Start Menu.

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

See the [R Windows FAQ](#) for more details on the binary installer.

3.1 Building from source

3.1.1 Getting the tools

If you want to build R from the sources, you will first need to collect, install and test an extensive set of tools. See [Appendix D \[The Windows toolset\]](#), [page 53](#) (and perhaps updates in <http://www.murdoch-sutherland.com/Rtools/>) for details.

The ‘Rtools.exe’ executable installer described in [Appendix D \[The Windows toolset\]](#), [page 53](#) also includes some source files in addition to the R source as noted below. You should run it first, to obtain a working `tar` and other necessities. Choose a “Full installation”, and install the extra files into your intended R source directory, e.g. ‘C:/R’. The directory name *should not contain spaces*. We will call this directory *R_HOME* below.

3.1.2 Getting the source files

You need to collect the following sets of files:

- Get the R source code ‘R-2.11.0.tar.gz’ from CRAN. Open a command window (or another shell) at directory *R_HOME*, and run

```
tar xf R-2.11.0.tar.gz
```

to create the source tree in *R_HOME*. **Beware:** do use `tar` to extract the sources rather than tools such as WinZip that do not understand symbolic links.

It is also possible to obtain the source code using Subversion; see [Chapter 1 \[Obtaining R\]](#), [page 1](#) for details.

- If you are not using a tarball you need to obtain copies of the recommended packages from CRAN. Put the ‘.tar.gz’ files in ‘*R_HOME*/src/library/Recommended’ and run `make link-recommended`. If you have an Internet connection, you can do this automatically using

```
make rsync-recommended
```

- Optionally, you can install a version of ATLAS (<http://math-atlas.sourceforge.net/>) tuned to your system for fast linear algebra routines. Pre-built ‘Rblas.dll’ for various CPUs are available in the ‘windows/contrib/ATLAS’ area on CRAN. If you are building R from source, there are macros `USE_ATLAS` and `ATLAS_PATH` in the file ‘MkRules.dist’. Copy

¹ formerly EM64T.

file ‘MkRules.dist’ to ‘MkRules.local’ and edit it to set `USE_ATLAS = YES` and `ATLAS_PATH` to where the ATLAS libraries are located. You will need to make the libraries yourself²: none of the binaries we have seen are compiled for the correct compiler. Since R has its own ‘xerbla’ it is necessary to delete that in ATLAS by

```
ar d /path/to/libf77blas.a xerbla.o
```

There used to be support for AMD’s AMD Core Math Library (ACML) and Kazushige Goto’s BLAS, but neither is currently available for use with the compilers used to compile R.

The following additional items are normally installed by ‘Rtools.exe’. If instead you choose to do a completely manual build you will also need

- The Tcl/Tk support files are in a zip file at <http://www.stats.ox.ac.uk/pub/Rtools/>: unzip this in `R_HOME`, and it will add directory tree ‘`R_HOME/Tcl`’.
- You need libpng and jpeg sources (available, e.g., from <http://www.libpng.org/>, <http://www.ijg.org>, <http://www.libtiff.org/>.³ You will need files ‘libpng-1.2.18.tar.gz’, ‘jpegsrc.v6b.tar.gz’, ‘tiff-3.8.0.tar.gz’ or later.

Working in the directory ‘`R_HOME/src/gnuwin32/bitmap`’, install the libpng and jpeg sources in sub-directories. The jpeg sub-directory for version 8a is named ‘jpeg-8a’; if you use a different version (e.g. ‘jpeg-6b’ or ‘jpeg-7’), copy file ‘`src/gnuwin32/MkRules.dist`’ to ‘`src/gnuwin32/MkRules.local`’ and edit the definition of `JPEGDIR`: the names of the ‘libpng’ and ‘libtiff’ directories can also be set there.

Example:

```
> tar xzf libpng-1.2.40.tar.gz
> mv libpng-1.2.40 libpng
> tar xzf jpegsrc.v8a.tar.gz
> tar xzf tiff-3.9.1.tar.gz
> mv tiff-3.9.1/libtiff .
> rm -rf tiff-3.9.1
```

3.1.3 Building the core files

You may need to compile under a case-honouring file system: we found that a samba-mounted file system (which maps all file names to lower case) did not work.

Open a command window at ‘`R_HOME/src/gnuwin32`’. Look at ‘`MkRules.dist`’ and if settings need to be altered, copy it to ‘`MkRules.local`’ and edit the settings there. Then run

```
make all recommended
```

and sit back and wait while the basic compile takes place.

Notes:

- We have had reports that earlier versions of anti-virus software locking up the machine, but not for several years. However, aggressive anti-virus checking such as the on-access scanning of Sophos can slow the build down several-fold.
- By default Doug Lea’s `malloc` in the file ‘`R_HOME/src/gnuwin32/malloc.c`’ is used for R’s internal memory allocations. You can opt out of this by setting `LEA_MALLOC=NO` in ‘`MkRules.dist`’, in which case the `malloc` in ‘`msvcrt.dll`’ is used. This does impose a considerable performance penalty and has not been tested recently.
- You can run a parallel make by e.g.

² We do this using the Cygwin compilers, often with some difficulty.

³ but the current versions are at <ftp://ftp.remotesensing.org/pub/libtiff/>: we used ‘tiff-3.9.1.tar.gz’.


```
make -j4 all
make -j4 recommended
```

but this is only likely to be worthwhile on a multi-core machine with ample memory, and is not 100% reliable.

3.1.4 Building the bitmap files

The file ‘*R_HOME*/library/grDevices/libs/Rbitmap.dll’ is not built automatically.

Running `make` in ‘*R_HOME*/src/gnuwin32/bitmap’ or `make bitmapdll` in ‘*R_HOME*/src/gnuwin32’ should build ‘Rbitmap.dll’ and install it in ‘*R_HOME*/bin’.

3.1.5 Checking the build

You can test a build by running `make check`. You may need to set `TMPDIR` to the absolute path to a suitable temporary directory: the default is ‘*c:/TEMP*’. (Use forward slashes and do not use a path including spaces. It will be ignored if not set to a directory.)

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 `check-devel` and `check-recommended`.

3.1.6 Building the manuals

The PDF manuals can be made by

```
make manuals
```

If you want to make the info versions (not the Reference Manual), use

```
cd ../../doc/manual
make -f Makefile.win info
```

To make DVI versions of the manuals use

```
cd ../../doc/manual
make -f Makefile.win dvi
```

(all assuming you have `pdftex/pdflatex` or `tex/latex` installed and in your path).

See the [Section 2.3 \[Making the manuals\]](#), [page 4](#) section in the Unix section for setting options such as the paper size.

3.1.7 Building the Inno Setup installer

You need to have the files for a complete R build, including bitmap and Tcl/Tk support and the manuals, as well as the recommended packages and Inno Setup (see [Section D.3 \[The Inno Setup installer\]](#), [page 54](#)).

Once everything is set up

```
make distribution
make check-all
```

will make all the pieces and the installers and put them in the ‘*gnuwin32/cran*’ subdirectory, then check the build. This works by building all the parts in the sequence:

```
Rpwd.exe (a utility needed in the build)
rbuild (the executables, the FAQ docs etc.)
rpackage (the base packages)
```

```

htmldocs (the HTML documentation)
bitmapdll (the bitmap support files)
recommended (the recommended packages)
vignettes (the vignettes in package grid:
    only needed if building from an svn checkout)
manuals (the PDF manuals)
rinstaller (the install program)
crandir (the CRAN distribution directory)

```

The parts can be made individually if a full build is not needed, but earlier parts must be built before later ones. (The ‘Makefile’ doesn’t enforce this dependency—some build targets force a lot of computation even if all files are up to date.) The first four targets are the default build if just ‘make’ is run.

If you want to customize the installation by adding extra packages, replace `make rinstaller` by something like

```
make rinstaller EXTRA_PKGS='pkg1 pkg2 pkg3'
```

An alternative way to customize the installer starting with a binary distribution is to first make a full installation of R from the standard installer (that is, select ‘Full Installation’ from the ‘Select Components’ screen), then add packages and make other customizations to that installation. Then in ‘src/gnuwin32/installer’ run

```
make myR IMAGEDIR=rootdir
```

where ‘rootdir’ is the path to the root of the customized installation (forward slashes and no spaces, please).

Both methods create an executable with the standard name, ‘R-2.11.0-win32.exe’, so please rename it to indicate that it is customized. If you intend to *distribute* a customized installer please do check that license requirements are met – note that the installer will state that the contents are distributed under GPL-2 and this has a requirement for *you* to supply the complete sources (including the R sources even if you started with a binary distribution of R).

The defaults for the startup parameters may also be customized. For example

```
make myR IMAGEDIR=rootdir MDISDI=1
```

will create an installer that defaults to installing R to run in SDI mode. See ‘src/gnuwin32/installer/Makefile’ for the names and values that can be set.

3.1.8 Building the MSI installer

It is also possible to build an installer for use with Microsoft Installer. This is intended for use by sysadmins doing automated installs, and is not recommended for casual use. It is only supported for 32-bit Windows at present.

It makes use of the Windows Installer XML (WiX) toolkit *version 2.0* available from <http://wix.sourceforge.net/>. (This needs the .NET 1.1 framework installed: it ran on a vanilla Windows XP SP2 machine. Unfortunately the file format has been changed within the same version: currently our code works with releases 2.0.4221.0 and 2.0.5805.0 – the latter is now said to be ‘production/stable’.) Once WiX is installed, set the path to its home directory in ‘MkRules.local’.

You need to have the files for a complete R build, including bitmap and Tcl/Tk support and the manuals, as well as the recommended packages. There is no option in the installer to customize startup options, so edit ‘etc/Rconsole’ and ‘etc/Rprofile.site’ to set these as required. Then

```

cd installer
make msi

```

which will result in a file of about 45Mb with a name like ‘R-2.11.0-win32.msi’. This can be double-clicked to be installed, but those who need it will know what to do with it.

Thanks to David del Campo (Dept of Statistics, University of Oxford) for suggesting WiX and building a prototype installer.

3.1.9 Cross-building on Linux

Support for cross-building was withdrawn at R 2.9.0.

3.2 Testing an Installation

The Windows installer contains a set of test files used when building R. These are an optional part of the installation, and if it is desired to run tests on the installation these should be selected as well as the help source files (perhaps most easily by doing a ‘Full Installation’).

The `Rtools` are not needed to run these tests. but more comprehensive analysis of errors will be given if `diff` is in the path.

Once this has been done, launch either `Rgui` or `Rterm`, preferably with ‘`--vanilla`’. Then

```
library("tools")
testInstalledBasic("both")
testInstalledPackages("base")
testInstalledPackages("recommended")
```

runs the basic tests and then all the tests on the standard and recommended packages. These tests can be run from anywhere: they write their results in the ‘`tests`’ folder of the R home directory (as given by `R.home()`).

3.3 Building R for 64-bit Windows

Note that the standard 32-bit build of R runs perfectly well on 64-bit versions of Windows, with a limit on the total memory allocation (and address space) of 4GB. Thus a native 64-bit build will be attractive to those wanting to work with large R objects, or many medium-sized ones.

To build a 64-bit version of R you need a 64-bit toolchain: the only one discussed here is based on the work of the MinGW-w64 project (<http://sourceforge.net/projects/mingw-w64/>, a port of GNU `binutils` and `gcc`), but commercial compilers such as those from Intel and PGI could be used (and have been by R redistributors).

Support for MinGW-w64 was developed in the R sources over the period 2008–10 and was first released as part of R 2.11.0. The binary version available from CRAN was built with the MinGW-w64 toolchain described in this section, using static linking.

Several versions of the MinGW-w64 toolchain are available: use one of them to replace the ‘`MinGW/bin`’ directory of `Rtools` in your path (but still have the `Rtools` ‘`bin`’ directory and Perl in your path). Then copy ‘`src/gnuwin32/MkRules.dist`’ to ‘`src/gnuwin32/MkRules.local`’ and edit it to set `WIN=64` and `BINPREF64` appropriate to your toolchain, then make R in the usual way.

The toolchain we use is technically a cross-compiler: the tools run under 32-bit Windows but produce code to run under 64-bit Windows.⁴ This comes from <http://sourceforge.net/projects/mingw-w64/files/Toolchains%20targetting%20Win64/Automated%20> and has a name like ‘`mingw-w64-1.0-bin_i686-mingw_20100322.zip`’. (Care is needed: that site also has compilers targeting 32-bit Windows and running on other host OSes. You do want the ‘release’ version with `-1.0-` in the name. There was a native x64 toolchain under ‘Personal builds’.) Some versions of that toolchain by default use dynamic linking to the compiler runtimes,

⁴ As the R build process itself runs R, it has to be done under 64-bit Windows. So does installation of source packages.

which causes problems (see <http://www.stats.ox.ac.uk/~ripley/Win64/W64porting.html>), and can be converted to static linking by removing the files like

```
.../x86_64-w64-mingw32/lib64/*.dll.a
.../bin/*.dll
```

where ... is the top-level directory of the toolchain. If do you use a dynamically-linking toolchain be sure to arrange to distribute the required run-time DLLs (and note the requirements that their GPL license imposes on re-distribution).

To build package **tcltk** you need to unpack http://www.stats.ox.ac.uk/pub/Rtools/R64_Tcl_8-5-8.zip at the top of the source tree: otherwise edit 'src/library/tcltk/Makefile.win' to make a stub package.

The default is `options(pkgType="win64.binary")`, supported for R 2.11.x on both CRAN and CRANextras. Packages using Gtk+ (**Cairo**, **RGtk2**, **cairoDevice** and those that depend on them) need the 'bin' directory of the distribution at <http://www.gtk.org/download-windows-64bit.html> in the path: note that this conflicts with the directory of DLLs needed for 32-bit Gtk+ so you need the correct one first in your path if using both.

A package-building service is available at <http://win-builder.r-project.org/>, and some notes on porting packages can be found at <http://www.stats.ox.ac.uk/~ripley/Win64/W64porting.html>.

The default memory limit will be the amount of installed RAM: it can be changed by using (preferably) '`--max-mem-size=`' at startup or `memory.limit` in a running R session.

The personal library is (by default) folder 'R\win64-library\x.y' of your home directory.

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.

4 Installing R under Mac OS X

The ‘`bin/macosx`’ directory of a CRAN site contains binaries for Mac OS X for a base distribution and a large number of add-on packages from CRAN to run on Mac OS X 10.5/6 (with a different build for version 10.4.4 or higher).

The simplest way is to use ‘`R-2.11.0.pkg`’: just double-click on the icon. Note that Tcl/Tk and the Fortran compiler need to be installed separately if needed (unlike for R 2.9.x and earlier).

See the [R for Mac OS X FAQ](#) for more details.

4.1 Building from source on Mac OS X

If you want to build this port from the sources, you can read the FAQ for full details. You will need to collect and install some tools as explained in the document. Then you have to expand the R sources and configure R appropriately, for example

```
tar zxvf R-2.11.0.tar.gz
cd R-2.11.0
./configure --with-blas='-framework vecLib' --with-lapack \
  --with-aqua --enable-R-framework
make
```

and then sit back and wait. The first two options are the default (and strongly recommended), and with some toolsets have been essential.

The second line of options is also default on Mac OS X, but needed only if you want to build R for use with `R.app` Console, and imply ‘`--enable-R-shlib`’ to build R as a shared/dynamic library. These options configure R to be built and installed as a framework called ‘`R.framework`’. The default path for ‘`R.framework`’ is ‘`/Library/Frameworks`’ but this can be changed at configure time by specifying the flag ‘`--enable-R-framework[=DIR]`’ or at install time as

```
make prefix=/where/you/want/R.framework/to/go install
```

(the ‘`R.framework`’ directory should not be included in the path).

For compatibility with the CRAN distribution you may need to specify ‘`--with-included-gettext`’ to avoid linking against a ‘`libintl`’ dynamic library you may have available, for example in ‘`/usr/local/lib`’.

Note that building the ‘`R.app`’ GUI console is a separate project: see the Mac OS X FAQ for details.

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 will have default limits at least as large as these, but apparently NetBSD does not.)

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 39](#)). 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.

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 running 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.

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()`.

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 (and this is done by `Rscript`). But it can also be used to customize R, e.g. for class use.

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 (use ‘`Sys.getenv("R_LIBS_USER")`’ within an R session to see what it is), but only is 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).

6.3 Installing packages

Packages may be distributed in source form or compiled binary form. Installing source packages requires that compilers and 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 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.

If installing packages to be used by other users, ensure that the system `umask` is set to give sufficient permissions (see also `Sys.umask` in R).

¹ unless they were excluded in the build.

² its binding is locked once that files has been read, so users cannot easily change it.

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.

To install packages from source in Unix use

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

The part `-l /path/to/library` can be omitted, in which case the first library in `R_LIBS` is used if set, otherwise the main library `R_HOME/library` is used. (`R_LIBS` is looked for in the environment: note that `.Renviron` is not read by `R CMD`.) Ensure that the environment variable `TMPDIR` is either unset (and `/tmp` exists and can be written in and executed from) or points to a valid temporary directory.

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 set the option `CRAN` to 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 user 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 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 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, Omegahat, local archives, local files, ...). Function `setRepositories()` can select amongst those repositories that the R installation is aware of.

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 Mac OS X) 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) at the list of *binary* versions of packages available for your version of R and downloads the latest versions (if any), although optionally it will also download and install a source package by setting the `type` argument.

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

`R CMD INSTALL` works in Windows to install source packages if you have set up the tools needed (see [Appendix D \[The Windows toolset\]](#), page 53). `R CMD INSTALL --help` will tell you the current options under Windows (which differ from those on a Unix-alike). Installation of binary packages must be done by `install.packages`.

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

We have seen permission problems after unpacking source packages on some Vista/Windows 7/Server 2008 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 <http://win-builder.r-project.org/>.

6.3.2 Mac OS X

On Mac OS X `install.packages` works as it does on other Unix-like systems, but there is an additional type `mac.binary` (the default in the CRAN distribution) that can be passed to `install.packages` in order to download and install binary packages from a suitable repository, and is the default for CRAN builds of R. These Macintosh binary package files have the extension `'tgz'`. The R GUI provides for installation of either binary or source packages, from CRAN or local files.

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'` on Windows), or if that does 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`.

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.

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

```
## for C code
CFLAGS=-g -O
## for C++ code
CXXFLAGS=-g -O
## for Fortran code
FFLAGS=-g -O
## for Fortran 95 code
FCFLAGS=-g -O
```

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. Set the `repos` argument as in the previous section. The `update.packages()` 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.

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, in most installations 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 <http://www.omegahat.org/R/>: this must be an URL scheme that `download.packages` supports (which also includes ‘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. As from R 2.10.0 other forms of compression can be used, e.g. ‘.tar.bz2’ or ‘.tar.xz’ files.
- "win.binary": located at ‘bin/windows/contrib/x.y’ for R versions x.y.z and containing ‘.zip’ files.
- "mac.binary": located at ‘bin/macosx/universal/contrib/x.y’ for R versions x.y.z and containing ‘.tgz’ files. If the repository contains only packages for a specific architecture, the package distribution type can be set to "mac.binary.xxx" where xxx specifies the architecture, replacing `universal` by xxx in the path above. For packages built for the off-CRAN ‘leopard’ binary, replace `universal` by `leopard`.

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 a ‘PACKAGES.gz’ file, a gzip-compressed version of ‘PACKAGES’—as this will be downloaded in preference to ‘PACKAGES’ it should be included for large repositories.

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

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.

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.

Historically R worked in the ISO Latin-1 8-bit character set and so covered English and most Western European languages (if not necessarily their currency symbols). Since R 2.1.0 it has supported (optionally before R 2.10.0) multi-byte character sets such as UTF-8 and others used for Chinese, Japanese and Korean.

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. http://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. http://en.wikipedia.org/wiki/ISO_3166). Sometimes the combination of language and territory is used to specify the encoding, for example to distinguish between traditional and simplified Chinese.
- A charset encoding, which determines both how a byte stream should be divided into characters, and which characters the subsequences of bytes represent.
- Optionally, a modifier, for example to indicate that Austria is to be considered pre- or post-Euro.

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, and most OSes have only one charset each for many languages. Note too the remark above about Chinese.

7.1.1 Locales under Linux

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

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 at

http://msdn.microsoft.com/library/default.asp?url=/library/en-us/vccore98/html/_crt_language_and_country_strings.asp.

It offers only one encoding per language.

7.1.3 Locales under Mac OS X

Mac OS X supports locales in its own particular way, but the R GUI tries to make this easier for users. See <http://developer.apple.com/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’.

Internally Mac OS X uses a form similar to Linux but without specifying the encoding (which is UTF-8). It is based on ICU locales (<http://userguide.icu-project.org/locale>) and not POSIX ones.

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 then¹ 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*² the language you are changing to can be output in the current character set.

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 basic C-level error messages.
- 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.

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. 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.

¹ On some systems setting `LC_ALL` or `LC_MESSAGES` to ‘`C`’ disables `LANGUAGE`.

² If you try changing from French to Russian except in a UTF-8 locale, you will find messages change to English.

Translations in the right language but the wrong charset be 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 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'³ is used.

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

³ with Americanisms.

8 Choosing between 32- and 64-bit builds

Many current CPUs have both 32- and 64-bit sets of instructions: this has long been true for UltraSparc and more recently for MIPS, PPC and ‘x86_64’ (sometimes known as ‘amd64’ and ‘Intel64’: all but a few low-power current ‘ix86’ CPUs support this chipset). Many 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 Mac OS 10.6) 64-bit is.

All current versions of R use 32-bit integers and 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 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).

- 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: and `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* ‘ix86’, 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 Intel Core 2 Duo the vanilla 64-bit version of R is around 10% faster on both Linux and Mac OS X — this can be reduced by tuning the compilation to the chip.

So, for speed you may want to use a 32-bit build, 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 7](#).

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). On all builds of R, the maximum length (number of elements) of a vector is $2^{31} - 1$, about 2 billion, and on 64-bit builds the size of a block of memory allocated is limited to $2^{34} - 1$ bytes (8GB). It is anticipated these will be raised eventually but routine use of 8GB objects is (when this was written in 2005) several years off.

¹ also known as IEC 559 and IEEE 754

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

8.1 Windows

Currently the main Windows build of R from CRAN is a 32-bit executable. This runs happily on Windows x64, but is limited to an address space of 4GB (see the ‘R for Windows FAQ’ for details). As from R 2.11.0 a 64-bit build is also available from CRAN, and a different build has been available from commercial R redistributors since 2009.

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 Unix 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 “2.11.0”.

There is full access to R’s handling of NaNs, 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 will 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

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

Then

```
make
```

will make standalone libraries ‘libRmath.a’ and ‘libRmath.so’: ‘make static’ and make shared will create just one of them.

NB: certain compilers are unable to do compile-time IEEE-754 arithmetic and so cannot compile ‘mlutils.c’ and several other files. The known example is old versions of Sun’s cc (e.g. Forte 6 and 7).

¹ e.g. Bessel, beta and gamma functions

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.

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 up almost all the tools to make R and then run (in a Unix-like shell)

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

For `cmd.exe` 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, 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,dynamic
-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`, 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 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, 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 Unix

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 deb) 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 77 (see [Section B.6 \[Using FORTRAN\]](#), [page 40](#)). Some add-on packages also need a C++ compiler. Your C compiler should be IEC 60059¹, POSIX 1003.1 and C99-compliant if at all possible. R tries to choose suitable flags for the C compilers it knows about, but you may have to set `CC` or `CFLAGS` suitably. For recent versions of `gcc` with `glibc` this means including `‘-std=gnu99’`². If the compiler is detected as `gcc`, `‘-std=gnu99’` will be appended to `CC` unless it conflicts with a setting of `CFLAGS`. (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 of `CFLAGS`.)

Unless you do not want to view graphs on-screen you need `‘X11’` installed, including its headers and client libraries. For recent Fedora distributions it means (at least) `‘libX11’`, `‘libX11-devel’`, `‘libXt’` and `‘libXt-devel’`. On Debian 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 `readline` library available from any GNU mirror: version 4.2 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 for `glibc` but not of most commercial Unixes. However, you can make use of GNU `libiconv` (possibly as a plug-in replacement: see <http://www.gnu.org/software/libiconv/>).

The OS needs to have enough support³ for wide-character types: this is checked at configuration.

A `tar` program is needed to unpack the sources and packages (including the recommended packages). A version⁴ that can automatically detect compressed archives is preferred for use with `untar()`: the configure script looks for `gtar` before `tar`: use environment variable `TAR` to override this.

There need to be suitable versions of the tools `grep`, `sed` and `tr`: 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 `makeinfo` version 4.7 or later installed, and if not some of the HTML manuals will be linked to CRAN. (Version 4.6 is known to create incorrect HTML files.) To make DVI or PDF versions of the manuals you will also need file `‘texinfo.tex’` installed (which is part of the GNU `texinfo` distribution but is often

¹ also known as IEEE 754

² `‘-std=c99’` excludes POSIX functionality, but `‘config.h’` will turn on all GNU extensions to include the POSIX functionality.

³ 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`.

⁴ Such as GNU `tar` 1.15 or later, `bsdtar` (from <http://code.google.com/p/libarchive/>, as used by Mac OS 10.6) or `tar` from the Heirloom Toolchest (<http://heirloom.sourceforge.net/tools.html>).

made part of the \TeX package in re-distributions) as well as `texi2dvi`.⁵ Further, the versions of `texi2dvi` and `'texinfo.tex'` need to be compatible: we have seen problems with older \TeX distributions (TeXLive 2007 and MiKTeX 2.8) used with `texinfo` 4.13.

The DVI and PDF documentation and building vignettes needs `tex` and `latex`, or `pdftex` and `pdflatex`. As from R 2.11.0 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 keep up-to-date.

If you want to build from the R Subversion repository you need both `makeinfo` and `pdflatex`.

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

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 be compiled. These tests will fail if `pkg-config` is not installed, and are likely to fail if `cairo` was built statically (unusual). Most systems with `Gtk+` 2.8 or later installed will have suitable libraries. Mac OS X comes with none of these libraries, but `cairo` support (without `Pango`) has been added to the binary distribution: `pkg-config` is still needed and can be installed from the sources.

For the best font experience with these devices you need suitable fonts installed: Linux users will want the `urw-fonts` package. Another useful set of fonts is the ‘liberation’ truetype fonts available at <https://www.redhat.com/promo/fonts/>,⁶ 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 (http://en.wikipedia.org/wiki/Liberation_fonts).

The bitmapped graphics devices `jpeg()`, `png()` and `tiff()` need the appropriate headers and libraries installed: `jpeg` (version 6b or 7 or 8) or `libpng` (version 1.2.3 or later, including 1.4.0) and `zlib` (version 1.1.3 or later) or `libtiff` (any recent version of 3.x.y – 3.8.2 and 3.9.1 have been tested) respectively.

If you have them installed (including the appropriate headers and of suitable versions), `zlib`, `libbz2` and `PCRE` will be used if specified by `'--with-system-zlib'`, `'--with-system-bzlib'` or `'--with-system-pcre'`: otherwise versions in the R sources will be compiled in. As the latter suffice and are tested with R you should not need to change this. In particular, the version of `zlib` 1.2.3 in the R sources has enhancements to work with large file systems on 32-bit platforms, and `zlib` 1.2.4 is not currently compatible with R.

`liblzma` from `xz-utils` version 4.999 or later will be used if installed: the version in the R sources can be selected instead by configuring with `'--with-system-xz=no'`.

⁵ `texi2dvi` is normally a shell script. Some versions, e.g. that from `texinfo` 4.13a, need to be run under `bash` rather than a Bourne shell as on, say, Solaris.

⁶ `ttf-liberation` in Debian/Ubuntu.

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

Some systems (notably Mac OS X 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 9/10, AIX and Windows, usually as `libicu` or `icu4c`. It will be used by default where available (including on Mac OS X >= 10.4): should a broken version of ICU be found this can be suppressed by ‘`--without-ICU`’.

Code developers will need Perl version 5.8.0 or later, available via <http://www.perl.com/CPAN/>. This is used to package and check packages, for R CMD Rprof⁷ and the rarely used R CMD Sd2Rd: it should be in your PATH when `configure` is run.

The `bitmap` and `dev2bitmap` devices and also `embedFonts()` use `ghostscript` (<http://www.cs.wisc.edu/~ghost>). 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.

A.2.1 Tcl/Tk

The `tcltk` package needs Tcl/Tk >= 8.3 installed: the sources are available at <http://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, setting the paths to the correct config files may be necessary to avoid confusion between them.

Versions of Tcl/Tk up to 8.5.8 have been tested (including most versions of 8.4.x, but not recently).

A.2.2 Java support

`configure` looks for Java support on the host system, and if it finds it sets some settings which are useful for Java-using packages. `JAVA_HOME` can be set during the `configure` run to point to a specific JRE/JDK.

Principal amongst these are setting some library 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-1.6.0-openjdk-1.6.0.0.x86_64/jre
R_JAVA_LD_LIBRARY_PATH=${JAVA_HOME}/lib/amd64/server:${JAVA_HOME}/lib/amd64
```

Note that this unfortunately 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

⁷ for which the R function `summaryProf` is an alternative.

R CMD `javareconf`. The script re-runs Java detection in a manner similar to that of the `configure` script and updates settings in both ‘`Makeconf`’ and ‘`R_HOME/etc/ldpaths`’. See R CMD `javareconf --help` for details.

Another alternative of overriding those setting is to set `R_JAVA_LD_LIBRARY_PATH` (e.g. in ‘`~/.Renviron`’), or use ‘`/etc/ld.so.conf`’ to specify the Java runtime library paths to the system. Other settings are recorded in ‘`etc/Makeconf`’ (or a sub-architecture-specific version), e.g.

```
JAVA = /usr/bin/java
JAVAC = /usr/bin/javac
JAVA_HOME = /usr/java/jdk1.5.0_06/jre
JAVA_LD_LIBRARY_PATH = $(JAVA_HOME)/lib/amd64/server:$(JAVA_HOME)/lib/amd64:\
$(JAVA_HOME)/../lib/amd64:/usr/local/lib64
JAVA_LIBS = -L$(JAVA_HOME)/lib/amd64/server -L$(JAVA_HOME)/lib/amd64
-L$(JAVA_HOME)/../lib/amd64 -L/usr/local/lib64 -ljvm
```

where ‘`JAVA_LIBS`’ contains flags necessary to link JNI programs. Some of the above variables can be queried using R CMD `config`.

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 suitable 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 routines⁸, 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 (`libsunperf` on Sun Sparc⁹, `libessl` on IBM, `vecLib` on Mac OS X). The correct incantation for these is usually 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 more CPU time and even more elapsed time (occasionally dramatically so) than using a similar single-threaded BLAS.

Note that under Unix (but not under Windows) if R is compiled against a non-default BLAS and ‘`--enable-BLAS-shlib`’ is **not** used, then all BLAS-using packages must also be. So if R is re-built to use an enhanced BLAS then packages such as **quantreg** will need to be re-installed.

⁸ unless FORTRAN double complex is not supported on the platform

⁹ Using the Sun Studio `cc` and `f95` compilers

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 usually built as a static library that on some platforms cannot be used with shared libraries such as are used in R packages. Be careful when using pre-built versions of ATLAS (they seem to work on ‘ix86’ platforms, but not on ‘x86_64’ ones).

The usual way to specify ATLAS will be via

```
--with-blas="-lf77blas -latlas"
```

if the libraries are in the library path, otherwise by

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

For systems with multiple processors it is possible to use a multi-threaded version of ATLAS, by specifying

```
--with-blas="-lptf77blas -lpthread -latlas"
```

Consult its file ‘INSTALL.txt’ for how to build ATLAS with position-independent code (at least on version 3.8.0 and later): that file also describes how to build ATLAS as a shared library.

ATLAS can also be used on Windows: see see [Section 3.1.2 \[Getting the source files\]](#), page 10 when building from source, and [R Windows FAQ](#) for adding pre-compiled support to binary versions.

A.3.1.2 ACML

For ‘x86_64’ and ‘i686’ processors under Linux there is the AMD Core Math Library (ACML) <http://www.amd.com/acml>. For the gcc version we could use

```
--with-blas="-lacml"
```

if the appropriate library directory (such as ‘/opt/acml4.4.0/gfortran64/lib’) is in the LD_LIBRARY_PATH. For other compilers, see the ACML documentation. There is a multithreaded Linux version of ACML available for recent versions of gfortran. To make use of this you will need something like

```
--with-blas="-L/opt/acml4.4.0/gfortran64_mp/lib -lacml_mp"
```

See see [Section A.3.1.5 \[Shared BLAS\]](#), page 35 for an alternative (and in many ways preferable) way to use ACML.

A.3.1.3 Goto BLAS

Dr Kazushige Goto has written another tuned BLAS which is available for several processors and OSes. The current version is known as GotoBLAS2.

This has been made available in several formats, but is currently available only as source code. For *academic use only* (after registering) it can be obtained via <http://www.tacc.utexas.edu/tacc-projects>. Once this is built and installed, it can be used by configuring with

```
--with-blas="-lgoto2"
```

See see [Section A.3.1.5 \[Shared BLAS\]](#), page 35 for an alternative (and in many ways preferable) way to use it.

Note that currently a multi-threaded Goto BLAS will be built by default if and only if the building is on a multi-processor system (counting multiple cores and hyperthreading), and at run time the default number of threads is the number of CPUs detected.

A.3.1.4 Intel MKL

For Intel processors under Linux, there is Intel’s Math Kernel Library (<http://www.intel.com/software/products/mkl/>). You are strongly encouraged to read the MKL User’s Guide, which is installed with the library, before attempting to link to MKL.

Version 10 of MKL supports two linking models: the default model, which is backward compatible with version 9 (see below), and the pure layered model. The layered model gives the user fine-grained control over four different library layers: interface, threading, computation, and run-time library support. Some examples of linking to MKL using this layered model are given below. These examples are for GCC compilers on ‘x86_64’. The choice of interface layer is important on ‘x86_64’ since the Intel Fortran compiler returns complex values differently from the GNU Fortran compiler. You must therefore use the interface layer that matches your compiler (`mkl_intel*` or `mkl_gf*`).

R can be linked to a sequential version of MKL by

```
MKL_LIB_PATH=/opt/intel/mkl/10.0.3.020/lib/em64t/
export LD_LIBRARY_PATH=$MKL_LIB_PATH
MKL="-L${MKL_LIB_PATH} -lmkl_gf_lp64 -lmkl_sequential -lmkl_lapack -lmkl_core"
./configure --with-blas="$MKL" --with-lapack
```

The order of the libraries is important. The option ‘`--with-lapack`’ is used since MKL contains a copy of LAPACK as well as BLAS (see [Section A.3.2 \[LAPACK\]](#), page 36).

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

```
MKL="-L${MKL_LIB_PATH} -lmkl_gf_lp64 -lmkl_gnu_thread \
      -lmkl_lapack -lmkl_core -liomp5 -lpthread"
```

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`.

Static MKL may be used with

```
MKL="      -L${MKL_LIB_PATH} \
      -Wl,--start-group \
          ${MKL_LIB_PATH}/libmkl_gf_lp64.a \
          ${MKL_LIB_PATH}/libmkl_gnu_thread.a \
          ${MKL_LIB_PATH}/libmkl_core.a \
      -Wl,--end-group \
      -lgomp -lpthread"
```

(Thanks to Ei-ji Nakama).

The default linking model, which is also used by version 9 of MKL, can be used by

```
--with-blas="-lmkl -lguide -lpthread"
```

This is multi-threaded, but in version 9 the number of threads defaults to 1. It can be increased by setting `OMP_NUM_THREADS`. (Thanks to Andy Liaw for the information.)

A.3.1.5 Shared BLAS

Note that 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 such as is 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 (e.g., ‘-O3’) of compiler optimization are used.

- It is easy to change the BLAS without making 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 dynamic BLAS library (such as ACML or Goto’s) to ‘`R_HOME/lib/libRblas.so`’. For example, just

```
mv R_HOME/lib/libRblas.so R_HOME/lib/libRblas.so.keep
ln -s /opt/acml4.4.0/gfortran64_mp/lib/libacml_mp.so R_HOME/lib/libRblas.so
```

will change the BLAS in use to multithreaded ACML. A similar link works for recent versions of the Goto BLAS and perhaps for MKL (provided the appropriate ‘lib’ directory is in the run-time library path or `ld.so` cache).

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 `libsunperf` on Solaris, `vecLib` on Mac OS X and ACML on ‘`ix86`’/‘`x86_64`’ Linux and Solaris). However, the likely performance gains are thought to be small (and may be negative), and 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. To do so specify something like

```
--with-lapack="-L/path/to/libs -llapack -lcblas"
```

since the ATLAS subset of LAPACK depends on `libcblas`. 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.

Since ACML contains a full LAPACK, if selected as the BLAS it can be used as the LAPACK *via* ‘`--with-lapack`’.

If you do use ‘`--with-lapack`’, be aware of potential problems with bugs in the LAPACK 3.0 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
```

(seen with the Debian ‘`-llapack`’ which was current in late 2002, Fedora Core 4 Extras ‘`-llapack`’ in September 2005 and 64-bit `libsunperf` in Forte 7). Other potential problems are incomplete versions of the libraries: for example `libsunperf` from Sun Forte 6.x was missing the entry point for `DLANGE` and `vecLib` has omitted the BLAS routine `LSAME`. For problems compiling LAPACK using recent versions of `gcc` on ‘`ix86`’ Linux, see [Section C.8 \[New platforms\]](#), [page 51](#).

Please **do** bear in mind that using ‘`--with-lapack`’ is ‘definitely **not** recommended’: it is provided **only** because it is necessary on some platforms.

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 native compilers the flag ‘`-dalign`’ is needed so `libsunperf` can be used.

On some systems it is necessary that an external BLAS/LAPACK was built with the same FORTRAN compiler used to build R: known problems are with R built with **gfortran**, see [Section B.6.1 \[Using gfortran\]](#), page 41.

Appendix B Configuration on Unix

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`).

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.

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

B.2 Internationalization support

Translation of messages is supported via GNU `gettext` unless disabled by the configure option `--disable-nls` or the underlying OS has insufficiently standard C functions to support it. 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.
```

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

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`’ (or failing that² ‘`~/Rconfig`’), 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 and 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.

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 papersize when making DVI and 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 `LD_FLAGS` (for libraries, using ‘`-L`’ flags to be passed to the linker) and `CPP_FLAGS` (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`’ (`LD_FLAGS`, ‘`-L/usr/local/lib64`’ on most 64-bit Linux OSes) and ‘`-I/usr/local/include`’ (`CPP_FLAGS`) to catch the most common cases. If libraries are still not found, then maybe your compiler/linker does not support re-ordering of

² ‘`~/Rconfig`’ was supported (but undocumented) in versions of R prior to 2.10.0.

‘-L’ and ‘-l’ flags (this has been reported to be a problem on HP-UX with the native `cc`). 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` produces worthwhile performance gains. 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 on recent systems ‘-Bdirect,--hash-style=both,-Wl,-O1’ is recommended at <http://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`, `R_RD4DVI` 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. The native `make` has been reported to fail on SGI Irix 6.5 and Alpha/OSF1 (aka Tru64).

To build in a separate directory you need a `make` that uses the `VPATH` variable, for example GNU `make`, or Sun `make` on Solaris 2.7 or later.

`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
```

Note the comment in [Section 2.4 \[Installation\]](#), page 5 about using a parallel `make`.

B.6 Using FORTRAN

To compile R, you need a FORTRAN compiler. The default is to search for `f95`, `fort`, `xlF95`, `ifort`, `ifc`, `efc`, `pgf95` `lf95`, `gfortran`, `ftn`, `g95`, `f90`, `xlF90`, `pgHPF`, `pgf90`, `epcf90`, `g77`, `f77`, `xlF`, `f77`, `pgf77`, `cf77`, `fort77`, `f132`, `af77` (in that order)³, and use whichever is found first; if none is found, R cannot be compiled. However, if `CC` is `gcc`, the matching FORTRAN compiler (`g77` for `gcc` 3 and `gfortran` for `gcc` 4) is used if available.

The search mechanism can be changed using the configure variable `F77` which specifies the command that runs the FORTRAN 77 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 `F77` 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.

³ On HP-UX `fort77` is the POSIX compliant FORTRAN compiler, and comes after `g77`.

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 `COMPLEX*16` variables, which is a Fortran 90 extension. This is checked for at configure time⁴, but you may need to avoid compiler flags⁵ asserting FORTRAN 77 compliance.

For performance reasons⁶ you may want to choose a FORTRAN 90/95 compiler.

It is possible to use `f2c`, the FORTRAN-to-C converter (<http://www.netlib.org/f2c>), via a script. (An example script is given in `'scripts/f77_f2c'`: this can be customized by setting the environment variables `F2C`, `F2CLIBS`, `CC` and `CPP`.) You may need to ensure that the FORTRAN type `integer` is translated to the C type `int`. Normally `'f2c.h'` contains `'typedef long int integer;'`, which will work on a 32-bit platform but not on a 64-bit platform. If your compiler is not `gcc` you will need to set `FPICFLAGS` appropriately.

B.6.1 Using gfortran

`gfortran` is the F95 compiler that is part of `gcc 4.x.y`. There were problems compiling R with the first release (`gcc 4.0.0`) and more with pre-releases, but these are resolved in later versions.

On Linux `'x86_64'` systems there is an incompatibility in the return conventions for double-complex functions between `gfortran` and `g77` which results in the final example in `example(eigen)` hanging or segfaulting under external BLASs built under `g77`. This should be detected by a `configure` test.

The default `FFLAGS` chosen (by `autoconf`) for a GNU FORTRAN compiler is `'-g -O2'`. This seems not to be documented for `gfortran`, and has caused problems (segfaults and infinite loops) on `'x86_64'` Linux (and the optimizer will be shared with other platforms on that CPU type). A maximum optimization of `'-O'` is recommended there.

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

SHLIB_CFLAGS
for shared libraries

FFLAGS debugging and optimization flags, FORTRAN

SAFE_FFLAGS
ditto for source files which need exact floating point behaviour

MAIN_FFLAGS
ditto, for compiling the main program

⁴ as well as its equivalence to the `Rcomplex` structure defined in `'R_ext/Complex.h'`.

⁵ In particular, avoid `g77`'s `'-pedantic'`, which gives confusing error messages.

⁶ e.g., to use an optimized BLAS on Sun/Sparc

SHLIB_FFLAGS

for shared libraries

MAIN_LDFLAGS

additional flags for the main link

SHLIB_LDFLAGS

additional flags for linking the shared libraries

LIBnn

the primary library directory, 'lib' or 'lib64'

CPICFLAGS

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

FPICFLAGS

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

CXXPICFLAGS

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

FCPICFLAGS

special flags for compiling Fortran 95 code to be turned into a shared library

DEFS

defines to be used when compiling C code in R itself

Library paths specified as '-L/lib/path' in LDFLAGS 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 library. Since R as a shared library has about 6200 symbols, if in doubt use the larger version.

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, FFLAGS, CXXFLAGS and LIBS in R_XTRA_CFLAGS (and so on). These are for options which are always required, for example to force IEC 60559 compliance.

Appendix C Platform notes

This section provides some notes on building R on different Unix-like 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 (for R < 2.10.0) contain notes on platforms such as HP-UX, IRIX and Alpha/OSF1 for which we have had no recent reports.

C.1 X11 issues

The ‘X11()’ graphics device is the one started automatically on Unix-alikes when plotting. As its name implies, it displays on a (local or remote) X server, and relies on the services and in particular the fonts provided by the X server. So if you sometimes use R at a console and sometimes remotely from an X11 session running on a Windows machine, you may have to setup the fonts differently for the two usages.

The ‘modern’ version of the ‘X11()’ device is based on ‘**cairo**’ graphics and 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 even higher. 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-truetype
xorg-x11-fonts-Type1
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-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

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

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/X11R6/lib/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.utf8' but not when running in 'en_GB.utf8' on the same machine (although on some systems many UTF-8 X11 locales are aliased to 'en_US.utf8' 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 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 standard compilers.

Remember 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 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).)

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 for 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.

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 8 we used

```
CC="gcc -m32"
CXX="g++ -m32"
F77="gfortran -m32"
FC=${F77}
OBJC=${CC}
LDFLAGS="-L/usr/local/lib"
LIBnn=lib
```

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

There are known problems with several early versions of `gcc` 4, but `gcc` 4.1.0 and later have been tested extensively.

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"
F77="gfortran -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.

C.2.1 Intel compilers

Intel compilers have been used under ‘`ix86`’ and ‘`x86_64`’ Linux and R contains code to set the FPU options suitably. Brian Ripley tried version 9.0 of the compilers for ‘`ix86`’ on Fedora Core 3 *via*

```
CC=icc
F77=ifort
CXX=icpc
ICC_LIBS=/opt/compilers/intel/cc/9.0/lib
IFC_LIBS=/opt/compilers/intel/fc/9.0/lib
LDLFLAGS="-L$ICC_LIBS -L$IFC_LIBS -L/usr/local/lib"
SHLIB_CXXLD=icpc
```

and adding optimization flags failed: at least ‘`src/main/regex.c`’ and ‘`src/modules/lapack/dlamc.f`’ needed to be compiled without optimization. For ‘`x86_64`’ on Fedora Core 5 he used

```
CC=icc
CFLAGS="-g -O3 -wd188 -ip"
F77=ifort
FLAGS="-g -O3"
CXX=icpc
CXXFLAGS="-g -O3"
FC=ifort
FCFLAGS="-g -O3 -mp"
ICC_LIBS=/opt/compilers/intel/cce/9.1.039/lib
IFC_LIBS=/opt/compilers/intel/fce/9.1.033/lib
LDLFLAGS="-L$ICC_LIBS -L$IFC_LIBS -L/usr/local/lib64"
SHLIB_CXXLD=icpc
```

`configure` will add ‘`-c99`’ to `CC` for C99-compliance. This causes warnings with `icc` 10, so use `CC="icc -std=c99"` there. R will add ‘`-mp`’ in `R_XTRA_{C,F,CXX}FLAGS` to maintain correct IEC 60559 arithmetic. 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`.

For some comments on building on an Itanium (‘`ia64`’) Linux system with `gcc` or the Intel compilers see http://www.nakama.ne.jp/memo/ia64_linux/.

Others have reported success with version 10.0.

C.2.2 PGI compilers

Jennifer Lai used the Portland Group compilers on ‘`x86_64`’ to build pre-2.2.0. Updated versions of the settings she used are

```

PG_HOME=/usr/pgi/linux86-64/6.0
CC=pgcc
CFLAGS="-g -O2 -Kieee"
CPPFLAGS="-I$PG_HOME/include -I$PG_HOME/include/CC"
F77=pgf77
FFLAGS="-g -O2 -Kieee"
CXX=pgCC
CXXFLAGS="-g -O2 -Kieee"
FC=pgf95
FCFLAGS="-g -O2 -Kieee"
SHLIB_CXXLDFLAGS=-shared
SHLIB_LDFLAGS=-shared
LDFLAGS="-L$PG_HOME/libso -L/usr/lib64"

```

Note particularly the last, which is needed to ensure that a shared version of `libc` is found. The flag `-Kieee` ensures strict compliance to IEC60659. Also, http://www.amd.com/us-en/assets/content_type/DownloadableAssets/dwamd_PGI_nov603.pdf suggests that `-pc64` may be desirable.

C.2.3 Sun Studio compilers

Brian Ripley tested the Sun Studio 12 (aka Sun Ceres, formerly Sun Forte) compilers (<http://developers.sun.com/sunstudio/index.jsp>) on `'x86_64'` Linux with

```

CC=suncc
CFLAGS="-xO5 -xc99 -xlibmil -nofstore"
CPICFLAGS=-Kpic
F77=sunf95
FFLAGS="-O5 -libmil -nofstore"
FPICFLAGS=-Kpic
CXX="sunCC -library=stlport4"
CXXFLAGS="-xO5 -xlibmil -nofstore -features=tmplrefstatic"
CXXPICFLAGS=-Kpic
FC=sunf95
FCFLAGS=$FFLAGS
FCPICFLAGS=-Kpic
LDFLAGS=-L/opt/sunstudio12.1/rtlibs/amd64
SHLIB_LDFLAGS=-shared
SHLIB_CXXLDFLAGS=-G
SHLIB_FCLDFLAGS=-G
SAFE_FFLAGS="-O5 -libmil"

```

`'-m64'` could be added, but was the default. Do not use `'-fast'`: see the warnings under Solaris. (The C++ options are also explained under Solaris.)

The resulting build of R was not quite as fast as that built with `gcc 4.4.3` at `'-O3'`.

Others have found on at least some versions of `'i386'` Linux that the configure flag `'--disable-largefile'` was needed (since `'glob.h'` on that platform presumed `gcc` was being used).

C.3 FreeBSD

Rainer Hurling has reported success on `'x86_64'` FreeBSD 7.2 (and on 6.x in the past). Since Darwin (the base OS of Mac OS X) is based on FreeBSD we find testing on Darwin tends to pick up any potential problems on FreeBSD problems quite early.

The native BSD `make` suffices to build R but a small number of packages require GNU `make`. (The BSD version is `bsdmake` on Darwin.)

C.4 Mac OS X

You can build R as a Unix application on Mac OS X using the Apple Developer Tools ('Xcode') and `gfortran`. You will also need to install an X sub-system or configure with '`--without-x`'. The X window manager is part of the standard Mac OS X distribution since Mac OS X version 10.3 (Panther), but it is typically not pre-installed prior to 10.5 (Leopard).

For more information on how to find these tools please read the [R for Mac OS X FAQ](#).

The `vecLib` library can be used *via* the (default) configuration options

```
--with-blas="-framework vecLib" --with-lapack
```

to provide higher-performance versions of the BLAS and LAPACK routines. Building R without these options *via*

```
--without-blas --without-lapack
```

can be done (and is provided as an alternative in the binary distribution).

C.4.1 64-bit builds

64-bit builds are supported on 10.5.x (Leopard) and later. All that is needed is to select suitable compiler options, e.g. for recent Intel Macs

```
CC='gcc -arch x86_64'
CXX='g++ -arch x86_64'
F77='gfortran -arch x86_64'
FC='gfortran -arch x86_64'
OBJC='gcc -arch x86_64'
```

in '`config.site`' or on the `configure` command line. (Or on Leopard specify '`gcc-4.2`' etc: those versions are the default on Snow Leopard.)

C.4.2 Snow Leopard

On Snow Leopard you will most likely to need to specify '`-arch i386`' for a 32-bit build and '`-arch x86_64`' for a 64-bit build: the Apple compilers default to 64-bit but the `gfortran` supplied on CRAN defaults to 32-bit.

Another quirk is that people have found that the X11 libraries are not in the default linking path, so something like '`LIBS=-L/usr/X11/lib`' may be required.

C.5 Solaris

R has been built successfully on Solaris 10 (both Sparc and '`x86`') using the (zero cost) Sun Studio (aka Ceres) 12 compilers: there has been some success with `gcc 4/gfortran`, mainly on Sparc. Sun packages for R used to be available from <http://www.sunfreeware.com/> for both architectures, but have not been updated recently. (Recent Sun machines are Opterons ('`amd64`') rather than '`x86`', but 32-bit '`x86`' executables are the default.)

There are also reports of success on OpenSolaris (also known as Solaris Express Community Edition, and sometimes as Solaris 11) on '`x86`'.

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. Also, a version of the preferred GNU `tar` is (if installed) in '`/usr/sfw/bin`', as sometimes are tools like `makeinfo`.

It is necessary to avoid the tools in '`/usr/ucb`': in particular '`/usr/ucb/tr`' causes R CMD `Rdiff` to malfunction. `configure` looks for a POSIX-compliant `tr`, e.g. in '`/usr/xpg4/bin`'.

Modern Solaris systems allow a large selection of Open Source software to be installed from <http://www.opencsw.org> (formerly <http://www.blastwave.org>) via `pkg-get`, by default installed under `/opt/csw`. 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 small number of packages require GNU `make`.

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.

When using the Sun compilers² do *not* specify `-fast`, as this disables IEEE arithmetic and `make check` will fail.

For the Sun Studio compilers a little juggling of paths was needed to ensure GNU `libiconv` is used rather than the Solaris `iconv`: at one time we used

```
CC="cc -xc99"
CPPFLAGS=-I/opt/csw/include
CFLAGS="-O -xlibmieee"
F77=f95
FFLAGS=-O4
CXX="CC -library=stlport"
CXXFLAGS=-O
FC=f95
FCFLAGS=$FFLAGS
LDFLAGS=-L/opt/csw/lib
```

to ensure that the `libiconv` version of `'iconv.h'` was found, and later used self-compiled versions of `libiconv` and `readline` from `/usr/local` (replacing `/opt/csw` in the flags). For a 64-bit target add `-m64` to the compiler macros and use something like `LDFLAGS=-L/opt/csw/lib/sparcv9` or `LDFLAGS=-L/usr/local/lib/amd64` as appropriate.

By default the Sun Studio 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. Also, errors have been reported on `'amd64'` if `-xc99` is omitted.

You can target specific Sparc architectures for (slightly) higher performance: Sun recommend

```
32-bit: -xtarget=ultra3 -xarch=v8plusa
64-bit: -xtarget=ultra3 -xarch=sparcv9
```

(in `CFLAGS` etc.) as a good compromise for recent Sparc chipsets.

On `'x86'` you will get marginally higher performance *via*

```
CFLAGS="-O5 -xc99 -xlibmieee -xlibmil -nofstore -xtarget=native"
FFLAGS="-O5 -libmil -nofstore -xtarget=native"
CXXFLAGS="-O5 -xlibmil -nofstore -xtarget=native"
SAFE_FFLAGS="-libmil -fstore -xtarget=native"
```

The SunStudio compilers provide several implementations of the C++ standard which select both the set of headers and a C++ runtime library. These are selected by the `-library` flag, which as it is needed for both compiling and linking is best specified as part of the compiler. The examples above use `'stlport4'`, currently the most modern of the options: the default (but still needed to be specified as it is needed for linking) is `'Cstd'` which dates from 1998: see http://developers.sun.com/solaris/articles/cmp_stlport_libCstd.html. Note though that most external Solaris C++ libraries will have been built with `'Cstd'` and so an R package using such a package also needs to be. Occasionally the library flag `'-lCrun'` is needed.

Several CRAN packages using C++ need the more liberal interpretation given by adding

² including `gcc` for Sparc from Sun.

```
CXXFLAGS="-features=tmplrefstatic"
```

The Sun performance library `libsunperf` is available for use with the Sun Studio compilers. If selected as a BLAS, it must also be selected as LAPACK via

```
./configure --with-blas=sunperf --with-lapack
```

This has often given test failures in the past, in several different places. For Solaris 10 in early 2010 this worked with 32-bit Sparc and 32-bit Opteron, but failed on both 64-bit Sparc and Opteron in `example(eigen)`.

C.5.1 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. A version of `gcc` optimized for Sparc (using technology from Sun's compilers) is available from Sun.³

Compilation for a 32-bit Sparc target with `gcc 4.3.3` needed

```
CPPFLAGS=-I/opt/csw/include
LDFLAGS="-L/opt/csw/gcc4/lib -L/opt/csw/lib"
```

and for a 64-bit Sparc target

```
CC="gcc -m64"
F77="gfortran -m64"
CXX="g++ -m64"
FC=$F77
CPPFLAGS=-I/opt/csw/include
LDFLAGS="-L/opt/csw/gcc4/lib/sparcv9 -L/opt/csw/lib/sparcv9"
```

Note that paths such as `/opt/csw/gcc4/lib/sparcv9` may need to be in the `LD_LIBRARY_PATH` during configuration.

Tests with `gcc`⁴ on `'x86'` and `'amd64'` have been less successful: `'x86'` builds have failed on tests using complex arithmetic⁵, whereas on `'amd64'` the builds have failed to complete in several different ways, most recently with relocation errors for `'libRblas.so'`.

C.6 AIX

We no longer support AIX prior to 4.2, and `configure` will throw an error on such systems.

Ei-ji Nakama was able to build under AIX 5.2 on `'powerpc'` with GCC 4.0.3 in several configurations. 32-bit versions could be configured with `'--without-iconv'` as well as `'--enable-R-shlib'`. For 64-bit versions he used

```
OBJECT_MODE=64
CC="gcc -maix64"
CXX="g++ -maix64"
F77="gfortran -maix64"
FC="gfortran -maix64"
```

and was also able to build with the IBM `xlc` and Hitachi `f90` compilers by

```
OBJECT_MODE=64
CC="xlc -q64"
CXX="g++ -maix64"
```

³ <http://www.sun.com/download/index.jsp?cat=Application%20Development&tab=3&subcat=Development%20Tools>

⁴ specifically `gcc 4.3.3` from OpenCSW.

⁵ such as those of `fft`: these can often be resolved by turning off optimization


```
F77="f90 -cpu=pwr4 -hf77 -parallel=0 -i,L -O3 -64"
FC="f90 -cpu=pwr4 -hf77 -parallel=0 -i,L -O3 -64"
FLIBS="-L/opt/ofort90/lib -lh90vecmath -lh90math -lf90"
```

Some systems have f95 as an IBM compiler that does not by default accept FORTRAN 77. It needs the flag ‘-qfixed=72’, or to be invoked as `xlfr`.

The AIX native `iconv` does not support encodings ‘latin1’ nor ‘’ and so cannot be used. (As far as we know GNU `libiconv` could be installed.)

Fan Long reports success on AIX 5.3 using

```
OBJECT_MODE=64
LIBICONV=/where/libiconv/installed
CC="xlc_r -q64"
CFLAGS="-O -qstrict"
CXX="xlc_r -q64"
CXXFLAGS="-O -qstrict"
F77="xlfr -q64"
AR="ar -X64"
CPPFLAGS="-I$LIBICONV/include -I/usr/lpp/X11/include/X11"
LDFLAGS="-L$LIBICONV/lib -L/usr/lib -L/usr/X11R6/lib"
```

On one AIX 6.x system it was necessary to use `R_SHELL` to set the default shell to be Bash rather than Zsh.

Kurt Hornik and Stefan Theussl at WU (Wirtschaftsuniversität Wien) successfully built R on a ‘powerpc’ (8-CPU Power6 system) running AIX 6.1, configuring with or without ‘--enable-R-shlib’ (Ei-ji Nakama’s support is gratefully acknowledged).

It helps to describe the WU build environment first. A small part of the software needed to build R and/or install packages is available directly from the AIX Installation DVDs, e.g., Java 6, X11, and Perl. Additional open source software (OSS) is packaged for AIX in ‘.rpm’ files and available from both IBM’s “AIX Toolbox for Linux Applications” (<http://www-03.ibm.com/systems/power/software/aix/linux/>) and <http://www.oss4aix.org/download/>. The latter website typically offers more recent versions of the available OSS. All tools needed and libraries downloaded from these repositories (e.g., GCC, Make, `libreadline`, etc.) are typically installed to ‘/opt/freeware’, hence corresponding executables are found in ‘/opt/freeware/bin’ which thus needs to be in `PATH` for using these tools. Like on other Unix systems one needs GNU `libiconv` as the AIX version of `iconv` is not sufficiently powerful. Additionally, for proper Unicode compatibility one should install the corresponding package from the ICU project (<http://www.icu-project.org/download/>), which offers pre-compiled binaries for various platforms which in case of AIX can be installed via unpacking the tarball to the root file system. For full L^AT_EX support one can install the T_EX Live DVD distribution (<http://www.tug.org/texlive/>): it is recommended to update the distribution using the `tlmgr` update manager. For 64-bit R builds supporting Tcl/Tk this needs to be installed from the sources as available pre-compiled binaries supply only 32-bit shared objects.

The recent WU testing was done using compilers from both the GNU Compiler Collection (version 4.2.4) which is available from one of the above OSS repositories, and the IBM C/C++ (XL C/C++ 10.01) as well as FORTRAN (XL Fortran 12.01) compilers (<http://www14.software.ibm.com/webapp/download/byproduct.jsp#X>).

To compile for a 64-bit ‘powerpc’ (Power6 CPU) target one can use

```
CC ="gcc -maix64 -pthread"
CXX="g++ -maix64 -pthread"
FC="gfortran -maix64 -pthread"
F77="gfortran -maix64 -pthread"
CFLAGS="-O2 -g -mcpu=power6"
```



```
FFLAGS="-O2 -g -mcpu=power6"
FCFLAGS="-O2 -g -mcpu=power6"
```

for the GCC and

```
CC=xlc
CXX=xlc++
FC=xlf
F77=xlf
CFLAGS="-qarch=auto -qcache=auto -qtune=auto -O3 -qstrict -ma"
FFLAGS="-qarch=auto -qcache=auto -qtune=auto -O3 -qstrict"
FCFLAGS="-qarch=auto -qcache=auto -qtune=auto -O3 -qstrict"
CXXFLAGS="-qarch=auto -qcache=auto -qtune=auto -O3 -qstrict"
```

for the IBM XL compilers. For the latter, it is important to note that the decision for generating 32-bit or 64-bit code is done by setting the `OBJECT_MODE` environment variable appropriately (recommended) or using an additional compiler flag (`'-q32'` or `'-q64'`). By default the IBM XL compilers produce 32 bit code. Thus, to build R with 64-bit support one needs to either export `OBJECT_MODE=64` in the environment or, alternatively, use the `'-q64'` compiler options.

It is strongly recommended to install Bash and use it as the configure shell, e.g., via setting `CONFIG_SHELL=/usr/bin/bash` in the environment, and to use GNU Make (e.g., via `(MAKE=/opt/freeware/bin/make)`).

Further installation instructions to set up a proper R development environment can be found in the “R on AIX” project on R-Forge (<http://R-Forge.R-project.org/projects/aix/>).

C.7 Cygwin

The Cygwin emulation layer on Windows can be treated as a Unix-alike OS. This is unsupported, but experiments have been conducted and a few workarounds added for R 2.6.0.

Only building as a shared library works,⁶ so use e.g

```
./configure --disable-nls --enable-R-shlib
make
```

NLS does work if required, although adding `'--with-included-gettext'` is preferable. You will see many warnings about the use of auto-import.

Note that this gives you a command-line application using `readline` for command editing. The `'X11'` graphics device will work if a suitable X server is running, and the standard Unix-alike ways of installing source packages work. There was a bug in the `'/usr/lib/tkConfig.sh'` script in the version we looked at, which needs to have

```
TK_LIB_SPEC='-ltk84'
```

The overhead of using shell scripts makes this noticeably slower than a native build of R on Windows.

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

⁶ Windows DLLs need to have all links resolved at build time and so cannot resolve against `'R.bin'`.

requires a call to (the macro) `fpsetmask(0)` and OSF1 requires 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 Solaris SunPro compilers causes R's NaN to be set incorrectly.

Shared Libraries: There seems to be very little agreement across platforms on what needs to be done to build shared libraries. 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. Set the configure variable `SAFE_FFLAGS` to the flags to be used for this file. If configure detects GNU FORTRAN it adds flag `'-ffloat-store'` to `FFLAGS`. (Other settings are needed when using `icc` on `'ix86'` Linux, for example.)

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 ...

Appendix D The Windows toolset

If you want to build R or add-on packages from source in Windows, you will need to collect, install and test an extensive set of tools. See <http://www.murdoch-sutherland.com/Rtools/> for the current locations and other updates to these instructions. (Most Windows users will not need to build add-on packages from source; see [Chapter 6 \[Add-on packages\]](#), page 18 for details.)

Only building with MinGW `gcc 4.2.1` is supported. It is possible to build with the recently released MinGW `gcc 4.5.0`, but this came too late to be supported for this release of R—beware that this by default uses DLLs for runtimes which need to be distributed with R or packages built with that compiler. There are also differences in the behaviour of `formatC` with later MinGW runtimes.

We have found that the build process for R is quite sensitive to the choice of tools: please follow our instructions **exactly**, even to the choice of particular versions of the tools.¹ The build process for add-on packages is somewhat more forgiving, but we recommend using the exact toolset at first, and only substituting other tools once you are familiar with the process.

This appendix contains a lot of prescriptive comments. They are here as a result of bitter experience. Please do not report problems to the R mailing lists unless you have followed all the prescriptions.

We have collected most of the necessary tools (unfortunately not all, due to license or size limitations) into an executable installer named ‘`Rtools.exe`’, available from <http://www.murdoch-sutherland.com/Rtools/>. You should download and run it, choosing the default “Package authoring installation” to build add-on packages, or the “full installation” if you intend to build R.

You will need the following items to build R and packages. See the subsections below for detailed descriptions.

- Perl (in ‘`Rtools.exe`’)
- The command line tools (in ‘`Rtools.exe`’)
- The MinGW compilers (in ‘`Rtools.exe`’)

For installing simple packages containing data or R source but no compiled code, none of these are needed. Perl is essential to build R, and needed to develop packages.

A complete build of R including PDF manuals, and producing the standalone installer ‘`R-2.11.0-win32.exe`’ will also need the following:

- L^AT_EX
- The Inno Setup installer

It is important to set your `PATH` properly. The ‘`Rtools.exe`’ optionally sets the path to components that it installs.

Your `PATH` may include ‘.’ first, then the ‘`bin`’ directories of the tools, Perl, MinGW and L^AT_EX. Do not use filepaths containing spaces: you can always use the short forms (found by `dir /x` at the Windows command line). Network shares (with paths starting `\\`) are not supported. For example, all on one line,

```
PATH=c:\Rtools\bin;c:\Rtools\perl\bin;c:\Rtools\MinGW\bin;c:\texmf\miktex\bin;
c:\R\bin;c:\windows;c:\windows\system32
```

It is essential that the directory containing the command line tools comes first or second in the path: there are typically like-named tools in other directories, and they will **not** work. The ordering of the other directories is less important, but if in doubt, use the order above.

¹ For example, the Cygwin version of `make 3.81` fails to work correctly.

Our toolset contains copies of Cygwin DLLs that may conflict with other ones on your system if both are in the path at once. The normal recommendation is to delete the older ones; however, at one time we found our tools did not work with a newer version of the Cygwin DLLs, so it may be safest not to have any other version of the Cygwin DLLs in your path.

D.1 Perl

Perl is needed to build R and to use `R CMD build|check|Rprof`: it is no longer needed to install packages.

You will need a Windows port of `perl5` (but only the basic functionality, not any of the third-party Win32 extensions). The Vanilla Perl (version 5.8.x) package is included in `Rtools.exe`. A more full-featured distribution is available from <http://www.activestate.com/Products/ActivePerl/>: other alternatives are listed at <http://win32.perl.org/>.

Beware: you do need a *Windows* port and not the Cygwin one. Users of 64-bit Windows can use a Win64 Perl (such as that from ActiveState) if they prefer.

D.2 L^AT_EX

The ‘MiKTeX’ (<http://www.miktex.org/>) distribution of L^AT_EX includes a suitable port of `pdftex`. The ‘basic’ version of ‘MiKTeX’ almost suffices (the **grid** vignettes need `fancyvrb.sty`), but it will install the 15Mb ‘lm’ package if allowed to (although that is not actually used). The `Rtools.exe` installer does *not* include any version of L^AT_EX.

Please read [Section 2.3 \[Making the manuals\]](#), [page 4](#) about how to make `refman.pdf` and set the environment variables `R_RD4DVI` and `R_RD4PDF` suitably; ensure you have the required fonts installed.

D.3 The Inno Setup installer

To make the installer package (`R-2.11.0-win32.exe`) we currently require Inno Setup 5.3.7 or later from <http://jrsoftware.org/>. We use the Unicode version and recommend that version be used. This is *not* included in `Rtools.exe`. (Earlier versions can be used, but the set of languages will need to be edited in `src/gnuwin32/installer/JRins.pl`, at least to change the location of `Japanese.isl` or remove the reference to it.)

Copy file `src/gnuwin32/MkRules.dist` to `src/gnuwin32/MkRules.local` and edit it to set `ISDIR` to the location where Inno Setup was installed.

D.4 The command line tools

This item and the next are installed by the `Rtools.exe` installer.

If you choose to install these yourself, you will need suitable versions of at least `basename`, `cat`, `cmp`, `comm`, `cp`, `cut`, `date`, `diff`, `echo`, `find`, `grep`, `gzip`, `ls`, `make`, `makeinfo`, `mkdir`, `mv`, `rm`, `rsync`, `sed`, `sh`, `sort`, `tar`, `texindex`, `touch`, `tr` and `uniq`; we use those from the Cygwin distribution (<http://www.cygwin.com/>) or compiled from the sources. You will also need `zip` and `unzip` from the Info-ZIP project (<http://www.info-zip.org/>). All of these tools are in `Rtools.exe`.

Beware: ‘Native’ ports of `make` are **not** suitable (including that called ‘MinGW make’ at the MinGW SourceForge site). There were also problems with several earlier versions of the Cygwin tools and DLLs. To avoid frustration, please use our tool set, and make sure it is at the front of your path (including before the Windows system directories). If you are using a Windows shell, type `PATH` at the prompt to find out.

D.5 The MinGW compilers

This version of R is set up to use gcc 4.2.1 for which MinGW compilers were released in August 2007. The ‘Rtools.exe’ installer currently includes the -sjlj version of 4.2.1 of the MinGW port of gcc from http://sourceforge.net/project/showfiles.php?group_id=2435.

If you would like to install your own copy, we recommend downloading from the URL above, as the download links from <http://www.mingw.org/> have led to obsolete versions. See the notes on <http://www.murdoch-sutherland.com/Rtools/> for updates.

To download the components individually, currently you need

```
mingw-runtime-3.13.tar.gz
w32api-3.10.tar.gz
binutils-2.17.50-20060824-1.tar.gz
gcc-core-4.2.1-sjlj-2.tar.gz
gcc-g++-4.2.1-sjlj-2.tar.gz
gcc-gfortran-4.2.1-sjlj-2.tar.gz
```

(and gcc-objc-4.2.1-sjlj-2.tar.gz if you want Objective C support). Unpack these into the same directory (using `tar xzf tarball_name`). (You may need to copy ‘bin/gcc-sjlj.exe’ to ‘bin/gcc.exe’ and a few badly-written packages need ‘bin/g++-sjlj.exe’ copied to ‘bin/g++.exe’.) This compiler should work on Windows Vista without any workarounds.

Note that mingw-runtime-3.13.tar.gz or later (but not as late as 3.18) and gcc-4.2.1 or later are needed to get a correct build of R itself.

D.6 Useful additional programs

Developers of packages will find some of the ‘goodies’ at <http://www.stats.ox.ac.uk/pub/Rtools/goodies> useful.

There is a version of the `file` command that identifies the type of files, and is used by `Rcmd check` if available.

The file ‘xzutils.zip’ contains the program `xz` which can be used to (de)compress files with that form of compression.

Function and variable index

C

`configure` 3, 5, 6, 39, 40

I

`install.packages` 19

M

`make` 40

R

`R_HOME` 3

`remove.packages` 20

U

`update.packages` 20

Concept index

A

AIX 49

B

BLAS library 33, 41, 47, 49

F

FORTRAN 40

FreeBSD 46

I

Installation 5

Installing under Unix-alikes 3

Installing under Windows 10

Internationalization 22

L

LAPACK library 36, 47, 49

Libraries 18

Libraries, managing 18

Libraries, site 18

Libraries, user 18

Linux 3, 44

Locale 22

Localization 22

M

Mac OS X 3, 16, 47

Manuals 4

Manuals, installing 6

O

Obtaining R 1

P

Packages 18

Packages, default 18

Packages, installing 18

Packages, removing 20

Packages, updating 20

R

Rbitmap.dll 12

Repositories 21

S

Site libraries 18

Solaris 47

Sources for R 1

Subversion 1, 4, 31

U

User libraries 18

V

Vignettes 4, 31

Environment variable index

B

BLAS_LIBS 33

C

CC 41

CONFIG_SITE 39

CPP 41

D

DESTDIR 7, 28

F

F2C 41

F2CLIBS 41

FPICFLAGS 41

J

JAVA_HOME 32

L

LANG 23

LANGUAGE 23, 24

LAPACK_LIBS 36

LC_ALL 23

LC_MESSAGES 23

LD_LIBRARY_PATH 28, 34, 40, 42, 49

O

OBJECT_MODE 51

P

PAPERSIZE 39

PATH 40, 50, 53

R

R_BROWSER 39

R_DEFAULT_PACKAGES 18

R_DISABLE_HTTPD 4

R_JAVA_LD_LIBRARY_PATH 32, 33

R_LIBS 18, 19

R_LIBS_SITE 18

R_LIBS_USER 18

R_PAPERSIZE 5, 17, 39, 40

R_PDFVIEWER 39

R_RD4DVI 5, 40, 54

R_RD4PDF 5, 40, 54

R_SHELL 50

R_USER 17

T

TAR 30

TEMP 17

TMP 17

TMPDIR 3, 12, 17, 19